Supernova Neutrino Burst Data Streaming from DUNE FD DAQ

Getting good enough pointing information out fast enough

Brett Viren

2026-11-05

Context

- BNL developed a conceptual design for a small DUNE DAQ improvement to stream a tiny amount of Supernova Neutrino Burst (SNB) triggered readout to a computing with low latency. (2022 SciDAC proposal).
- Reconstructing the pointing of this data can take on various forms and should not add more than 3 minutes latency (competitive with Super-Kamiokande today).
- BNL has SPNG, GPU accelerated SP for high quality input to pointing reco.
- SPNG plans to run on DUNE-controlled GPU hardware which can be immediately switched to running pointing reco on the SNB stream when the SNB trigger is raised.

For this to work we want an expectation that the SNB data rich enough for pointing data can be extracted from the DAQ fast enough to be timely.

The scheme for prompt, SNB data streaming from DUNE DAQ

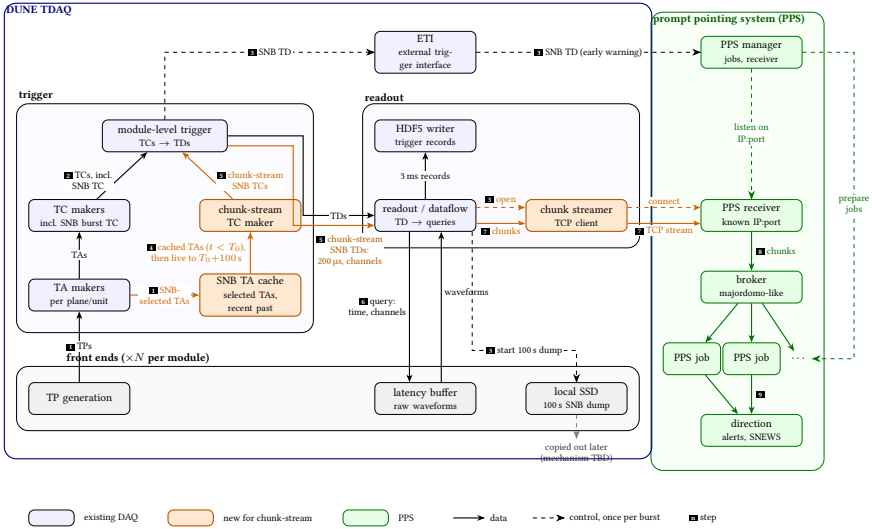

TN Fig. 1: The chunk-stream scheme. After the SNB TD (3), new trigger elements turn cached, then live, SNB-selected TAs into a sequence of chunk-stream TDs (4, 5); the readout fetches 200 µs chunks from the front-end buffers (6) and streams them over TCP to the prompt pointing system, warned and prepared through the ETI (7–9).

The study's questions

Assuming the conceptual design for DUNE FD TDAQ to produce a **minimal data stream** that is **highly enriched with SNB pointing information** we want to answer:

1. What would be the efficiency and purity to acquire SNB nu interactions?
2. What would be the resulting data rate and latency?
3. How does *pointing resolution* depend on the design choices?

Out of scope for now: actual SNB pointing algorithms and the subsystem that runs them.

- Here, want to **maximize pointing information** and **minimize data rate** input to reco.
- Latency goals: input to reco: $\mathcal{O}(1s)$, final pointing: **3 minutes**.

Outline of the study:

- Produce a full-chain simulation covering FDVD and FDHD.
  - Boost signal and background rates for efficient computing, reduced back in results.
- Implement trigger algorithms from DUNE-DAQ and `dunetrigger` (larsoft) in Wire-Cell Toolkit.
- Develop and evaluate selection algorithms reasonable for implementing in DAQ.
- Examine efficiency, purity, a pointing resolution *proxy* metric and data rate as function of selections.

Details on implementation are in the backups.

- I'll go through results.

The SNB signal sim

- **MARLEY 1.2.1** ν_e CC on ^{40}Ar (Bhattacharya 2009), pinched thermal spectrum, $\langle E \rangle = 11$ MeV, $\alpha = 2.5$; run as a separate step, JSON events.
- Elastic scattering (ES) samples from MARLEY too, for pointing.
- `LowEDepoSource`: for each depo set: Poisson count, uniform in time and volume, random direction; particles go to `ToyTransport`.
- Reference burst: 0.0075 Hz/t (∼3000 events in 40 kt over 10 s at 10 kpc); boosted sim rates are divided back for analysis.

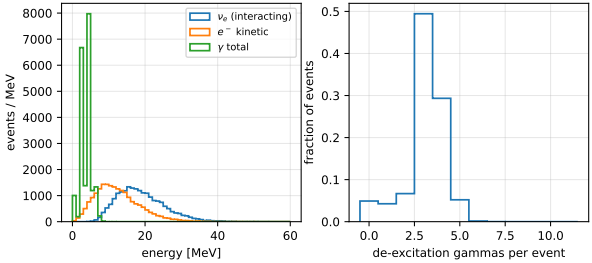

TN Fig. 6: MARLEY ν_e CC events (20000 events).

The SNB background sim

- Decays in the argon bulk: ^{39}Ar (1.01 Bq/kg, known), ^{85}Kr, ^{42}Ar/^{42}K, ^{222}Rn chain with the ^{214}Bi–^{214}Po coincidence.
- Beta shapes (allowed, forbidden), simplified schemes with gammas.
- MeV-scale sources boosted for tail statistics.
- Neutron captures (6.1 MeV γ): stress test, arbitrary rate.
- Missing: rock neutrons, external γ, surface U/Th, ^{220}Rn.

	HD APA			VD CRU		
source	rate [Hz]	boost	simulated	rate [Hz]	boost	simulated
snb	1.03	967	12788	0.339	1469	8504
Ar39	1.39×10^5	1	1781304	4.59e+04	1	784297
Ar42	19.5	1	250	5.5	1	94
K42	17.7	50	11344	5.92	50	5064
Kr85	1.59e+04	1	203218	5.24e+03	1	89623
Rn222	139	1	1778	46.4	1	794
Po218	144	1	1847	43.9	1	751
Pb214	137	5	8789	45.3	5	3876
Bi214	138	10	17711	45.7	10	7821
Po214	132	10	16942	44.3	10	7568
nCapture	68.7	10	8789	22.4	10	3822

TN Table 2: Sources per simulated unit: physical rate, boost and number of interactions simulated.

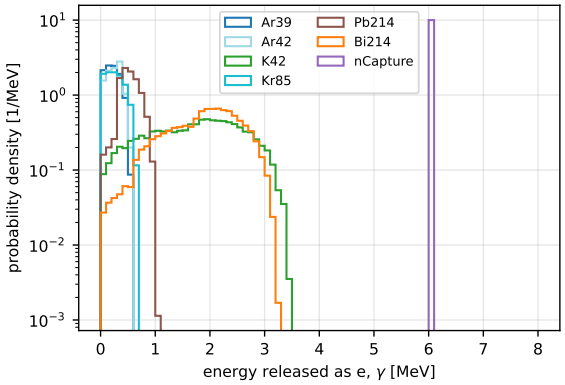

TN Fig. 7: Energy released to electrons and gammas per decay (alphas excluded).

Particle passage sim

`ToyTransport`: **no Geant4**; particles to depos in infinite LAr. Aims at what matters for TPs: how far MeV gammas carry energy, electron track lengths, energy sharing.

- $e^{\pm}$**:** 0.5 mm steps at the collision stopping power (ESTAR-like), bremsstrahlung above 100 keV, Highland multiple scattering; e^{+} annihilate.
- γ**:** Klein-Nishina Compton (exact), photoelectric, pair; exponential paths.
- **p,** α**, ions:** deposited at once (Bragg-Kleeman range for recombination).
- **n,** ν**, nuclei:** not transported.
- Charge from WCT modified box recombination.
- Unit tests: energy conservation, ranges (2%), attenuation (1%), γ spread.

Omits: Landau/delta-ray fluctuations, recombination fluctuations, track shape details, neutrons. So selections may rely only on **total energy** and **length scales**, with margins.

Expected improvements with Geant4

Replacing the toy by Geant4 (edep-sim, same truth ids) would:

- **Populate tails:** energy-loss and recombination fluctuations push radiological decays past their end points. The working-point cuts sit just above them: their margins get tested.
- **Real track shapes:** multiple scattering, delta rays, brems. Lets us *measure* direction reconstruction and replace the assumed pointing models.
- **Neutrons:** transport and capture, in SNB events and from rock and materials. At ~ 1 Hz/module, captures cost < 1 MB/s even at loose cuts; needs a physics list and capture data that get argon right.
- **Gamma cascades** with proper correlations: tests the neighborhood sum's 30 cm radius.

Expect: somewhat lower purity at fixed cut and a realistic answer to "which cut", not a different method.

Detector simulation

Standard WCT chain, 500 ns sampling, 3 ms frames: drift, diffusion, lifetime, 2D field + electronics response, empirical noise, digitizer. Some parameters from older designs.

- **HD:** one APA, both faces, 2560 channels, 14-bit, ProtoDUNE-HD noise.
- **VD:** one top-drift CRU (2 CRMs), 1536 channels, 12-bit, ProtoDUNE-VD top responses and noise.
- Thresholds about 5× noise; VD36: the `dunetrigger` VD values (36 ADC).

	HD	VD	VD36
geometry	1x2x6 APA 0	v6 top CRU 160	v6 top CRU 160
argon mass [t]	137.9	45.4	45.4
channels	2560	1536	1536
ADC bits	14	12	12
electron lifetime [ms]	8	8	8
noise RMS U/V/W [ADC]	13.3/13.3/10.4	3.0/3.0/3.0	3.0/3.0/3.0
TP threshold U/V/W [ADC]	70/70/60	15/15/15	36/36/36
threshold / noise (W)	5.8	5.1	12.1
ADC per visible MeV (W)	3637	471	357
units per 14 kt active	102	309	309

TN Table 3: Simulated detector units. HD: APA 0 of `dune10kt-1x2x6`. VD: top CRU 160 of the VD v6 geometry. VD36: same CRU, `dunetrigger` VD thresholds. Noise: robust RMS of the simulated ADC.

Example events

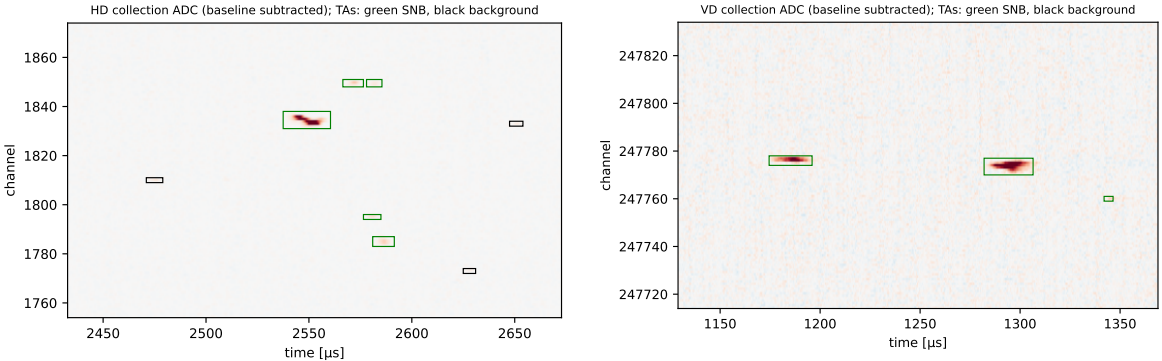

TN Fig. 9: Collection ADC around an SNB interaction in a frame, HD (left) and VD (right), with the TAs found.

SNB TA energy and count spectra

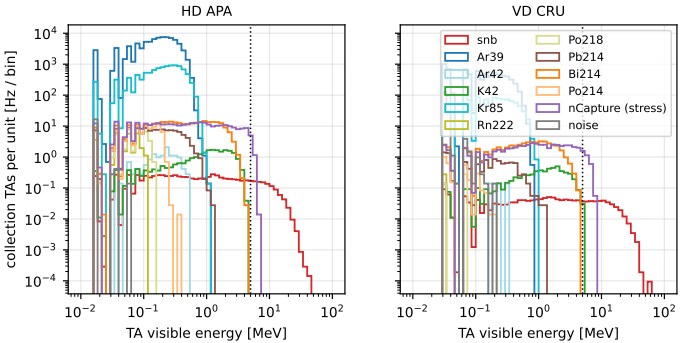

TN Fig. 11: Visible energy of collection TAs by dominant source, at physical rates (the supernova at the reference burst rate). The dotted line is the single-TA working point.

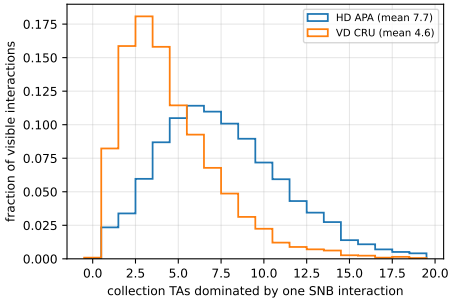

TN Fig. 12: Collection TAs per visible supernova interaction.

TA selection

Two simple methods, cuts in visible MeV (MeV_{vis}) of collection TAs:

- **Single-TA charge** (Tx): select a TA with $\geq x$ MeV_{vis}.
  - Working point T5: just above the radiological end points (^{42}K 3.5 MeV shows near 4.5 MeV_{vis}).
  - Cost: an SNB interaction's energy is split over several TAs.
- **Neighborhood sum** (Sx): a seed TA ($\geq$1 MeV_{vis}) is selected if the TAs ($\geq$0.3 MeV_{vis} each) within R =30 cm (channels and drift time, same face) sum to $\geq x$. In HD this box spans the whole 6 m wire.
  - Working point S8: recovers the gamma energy, can sit above a 6.1 MeV neutron capture.
  - Includes ^{39}Ar pile-up; cheap, a TA-level clustering.

Both rely only on energy and length scales the toy gets right.

TA-seeded readout

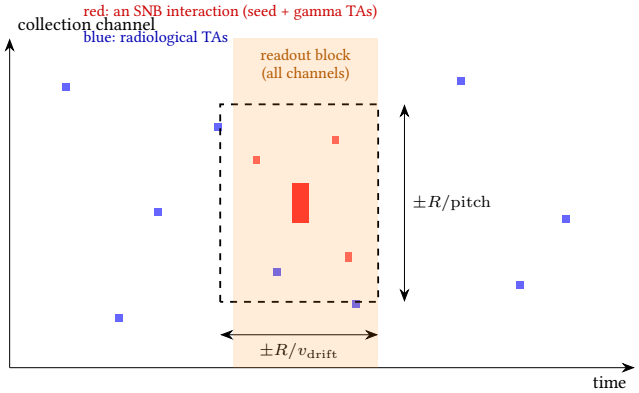

TN Fig. 14: The neighborhood sum around a seed TA and the readout block it seeds.

- Each selected TA reads a block of 200 µs (chosen to match response kernel) over **all channels** of its unit.
- 1.79 MB per HD APA, 922 kB per VD CRU.
- Overlapping blocks merged.
- Data rate = background TA rate × block, plus SNB blocks during a burst.

TA spectra

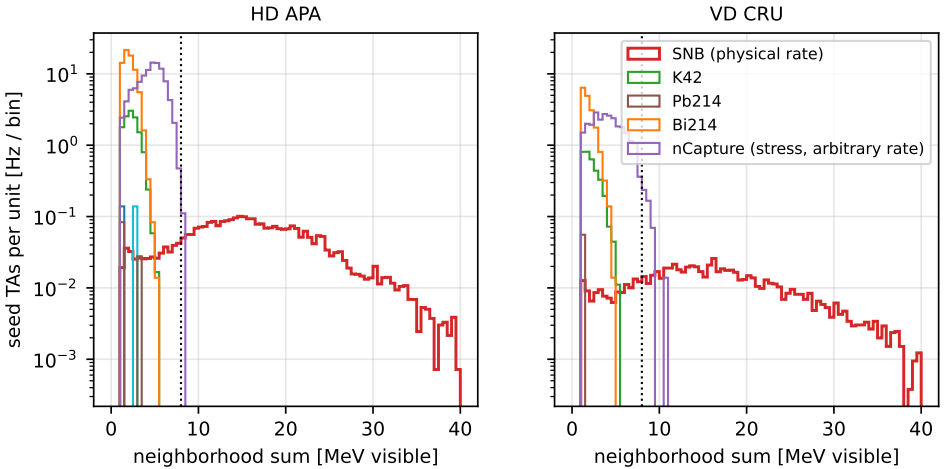

TN Fig. 15: Neighborhood sum of seed TAs, by source, at physical rates. The dotted line is the working point.

Background TA rate vs selection cut

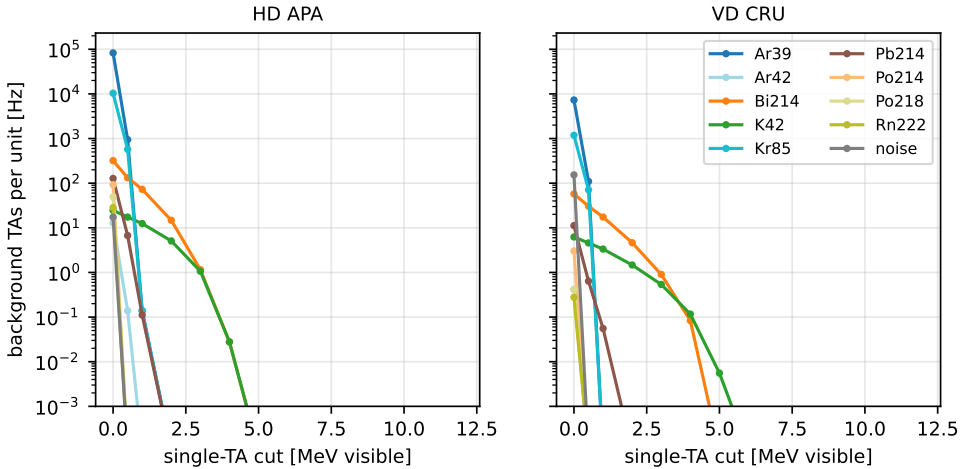

TN Fig. 13: Background TA rate by source against the single-TA cut.

Pointing resolution proxy - pointing power vs reco assumption

- ES electrons point forward and dominate; CC nearly isotropic (included, adds little).
- True SNB dir $\hat{n}$, per event dir $\hat{u}$, energy bin g:
  - mean: $\mu_g = \langle \hat{u} \cdot \hat{n} \rangle$
  - variance: $s_g = (1 - \langle (\hat{u} \cdot \hat{n})^2 \rangle)/2$.
- For N_g, the angular error $\sigma^{-2} = \sum_g N_g \, \mu_g^2 / s_g$, and half-angle cone with 68% of directions: $\theta_{68} = 1.51\,\sigma$.
- 10 kpc burst, one 14 kt module: 1050 CC, 196 ES.
- Reconstruction of e^- direction (assumed):
  - *ideal*: none
  - *nominal*: $\sigma = 25^\circ (10\,\mathrm{MeV}/T)$, 20% flip
  - *flat*: 25° at all T, 20% flip
- Nominal and flat bracket the cost of a cut.
- Optimistic: perfect ES/CC tagging and no ES-like background (NC $^{40}\mathrm{Ar}^*$ single γ, not in MARLEY; external γ).

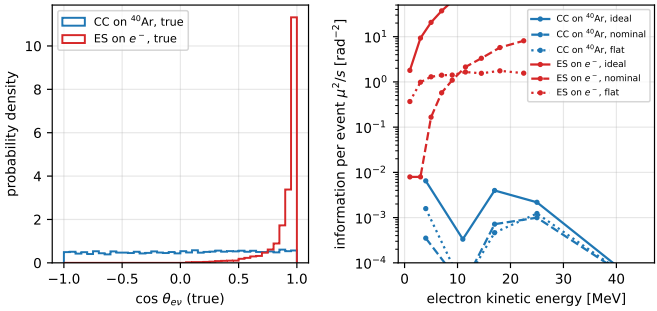

TN Fig. 22: Left: angle between the electron and the neutrino (MARLEY, HD samples). Right: information per event μ^2/s against the electron kinetic energy, for the ideal and nominal reconstruction.

Efficiency and pointing resolution

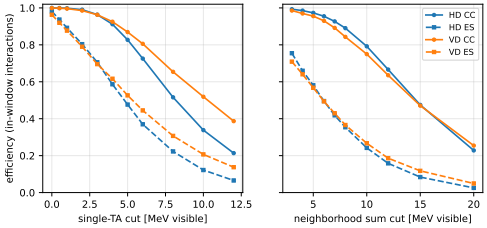

TN Fig. 23: Efficiency for CC (solid) and ES (dashed) interactions arriving in the readout, against the cut of each method.

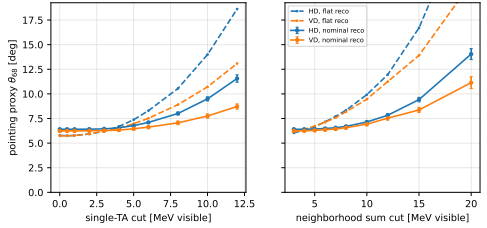

TN Fig. 24: Pointing proxy θ_{68} for a 10 kpc burst in one module, against the cut of each method: nominal reconstruction with bootstrap errors (solid) and flat reconstruction (dashed).

HD/VD data rates vs selection cut with 7.2 seconds / 2400 frames sim.

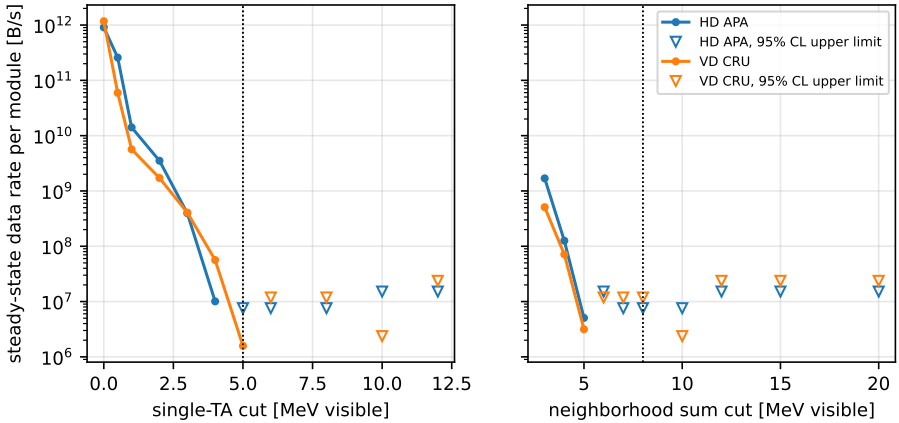

TN Fig. 21: Steady-state data rate per module (14 kt) with 200 µs readout blocks, against the cut of each method (the cuts of the efficiency and pointing figures). Open triangles are 95% CL upper limits where no background TA passes, set by the boosted statistics of the sources that could pass; dotted lines mark the working points.

Results at chosen cut points

unit	cut	ϵ_{CC}	ϵ_{ES}	θ_{68} [deg]	bkg [Hz]	purity	data/unit	data/module
HD	T5	0.822	0.477	6.8±0.1	<0.042	>0.96	<74.7 kB/s	<7.58 MB/s
HD	S8	0.874	0.354	6.7±0.1	<0.042	>0.99	<74.7 kB/s	<7.58 MB/s
VD	T5	0.861	0.526	6.5±0.1	0.00556	0.983	5.12 kB/s	1.58 MB/s
VD	S8	0.831	0.366	6.6±0.2	<0.042	>0.94	<38.4 kB/s	<11.8 MB/s
VD36	T5	0.847	0.488	6.6±0.2	0.0833	0.801	76.8 kB/s	23.7 MB/s
VD36	S8	0.766	0.340	6.9±0.2	<0.083	>0.86	<76.8 kB/s	<23.7 MB/s

TN Table 1: Results at the working points, per unit and per module (14 kt), at physical rates. Tx: single collection TA $\geq x$ MeV$_{\mathrm{vis}}$; Sx: seed TA whose 30 cm neighborhood sum is $\geq x$ MeV$_{\mathrm{vis}}$. ϵ_{CC}, ϵ_{ES}: efficiencies for ν_e CC and ES interactions; θ_{68}: pointing proxy, 10 kpc, one module, nominal reconstruction. "bkg": selected background TA rate ("<": 95% CL limit); each reads 200 µs over all channels. VD36: VD CRU with the `dunetrigger` 36 ADC thresholds.

- Pointing degrades much less than ES efficiency: cuts remove soft, poorly-pointing electrons first.
- Loose sum S5 is best per byte here, but cavern-γ pile-up and ^{208}Tl (not simulated) may cost $\sim$ 100 MB/s/module; S6 is safer.
- Missing backgrounds (Chao's estimates): external γ floor $\sim$ 0.3–3 MB/s/module at T5, S8; with NC $^{40}\mathrm{Ar}^{*}$, $\theta_{68} \sim$ 7–9°.

Summary and possible next steps

This is a quick (vibed) study but realistic enough to conclude that the scheme is viable.

- A reasonably small chunk-stream data rate can retain majority of pointing information.

Future study may include:

- Introducing Geant4, perhaps utilizing upcoming FDVD production sim.
- The full DUNE background model (external γ, surface U/Th, $^{220}\mathrm{Rn}$) and NC $^{40}\mathrm{Ar}^{*}$ as ES-like background in the pointing proxy.
- Propagating data rate and actual sim through SPNG to determine compute costs and latency.
  - SPNG needs $\mathcal{O}(1)$s for 3ms of APA, do we see linear speed up to 200 μs?
  - Is there a new bottleck with many small readouts?
- More realistic reco or a *reco pointing proxy* developed.
- Of course, test with real pointing reco when available.

Outreach, in progress:

- Reestablishing ties with DAQ DS/PP group.
- Engaging with SNB pointing Genesis group.

FIN

Backups

The `trig` package

Added a Wire-Cell Toolkit `trig/` sub-package, WCT core dependencies only.

- **Data:** trigger *primitive*, *activity* and *candidate* objects (TP, TA, TC)
- **Algorithm cores:** ports of DAQ `tpglibs` (TPG) and `triggeralgs` (Supernova, SWIFT TA; Supernova TC), checked by unit tests against verbatim upstream code; a new *burst* TC.
- **WCT nodes:** `TPFinder`, `TAFinder`, `TASetFanin`, `TCFinder`.
- **Study tools:** `LowEDepoSource` (SNB + radiologicals), `ToyTransport`, `TATruth` (truth tables of TAs and interactions).
- **Etc**: configs, tests, run and analysis scripts and docs ("tech note" PDF).

Job

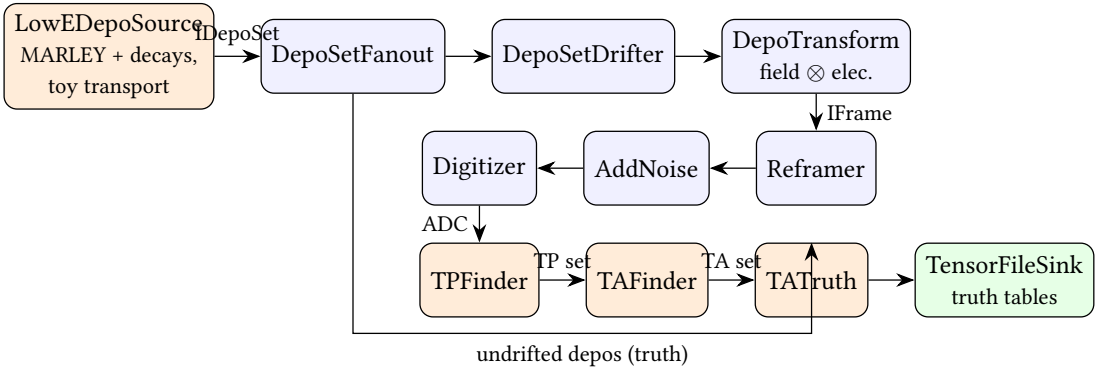

TN Fig. 2: Data flow of the study job. Orange nodes are from `trig`. The undrifted depos go to the truth node along with the TAs found.

Trigger Primitive (TP)

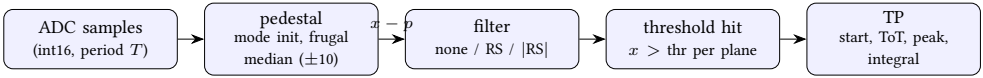

TN Fig. 3: Trigger primitive generation per channel.

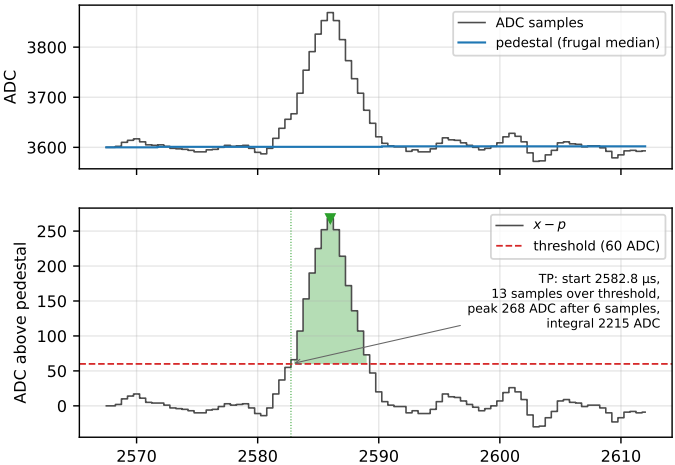

TN Fig. 4: The TP chain on one HD collection channel crossed by a supernova interaction. Top: ADC samples and the frugal median pedestal p, started at the mode of the frame. Bottom: $x - p$ against the threshold; the shaded run of samples over threshold makes one TP (start, samples over threshold, samples to peak, peak and integral of $x - p$ over the run).

Trigger Activity (TA)

A TA groups TPs from one plane to make a (channel $\times$ tick) rectangle.

- An energy measure $\mathrm{MeV}_{\mathrm{vis}}$ can be made from its ADC sum.

Two algorithms are ported from DUNE-DAQ:

- **Supernova TA maker**: greedy clustering in channel and tick.
- **SWIFT TA maker**: TPs in fixed windows. Available, not used for these results.

Only **collection-plane** TAs are used (DAQ can also use induction planes).

Trigger Candidate (TC) and Decision (TD) and Record (TR)

Made from TAs, indicates the nature of some activity, may be local to one anode or span many.

Relevant to this study:

- **Supernova TC** (DAQ, ported): count TAs with >2 TPs in an 8 s sliding window, TC when >3.
- **Burst TC** (new): count *selected* TAs (algorithm, TP count, ADC, spans) in a sliding window W; TC when the count reaches N, then hold off. Tolerates late TAs. Not yet used in the results.

The SNB *trigger decision* (TD) is essentially a promotion of the SNB TC by the *module level trigger*.

- The SNB pointing stream is **predicated** on the SNB TD being raised.

A TD leads to Trigger Record (TR) holding front-end buffer readouts.

- The 100 seconds of readout from an SNB TD is one TR.
- The SNB streaming transcends TC and TD levels by having no files involved.