

TPG rate issues

Artur Sztuc

a.sztuc@ucl.ac.uk

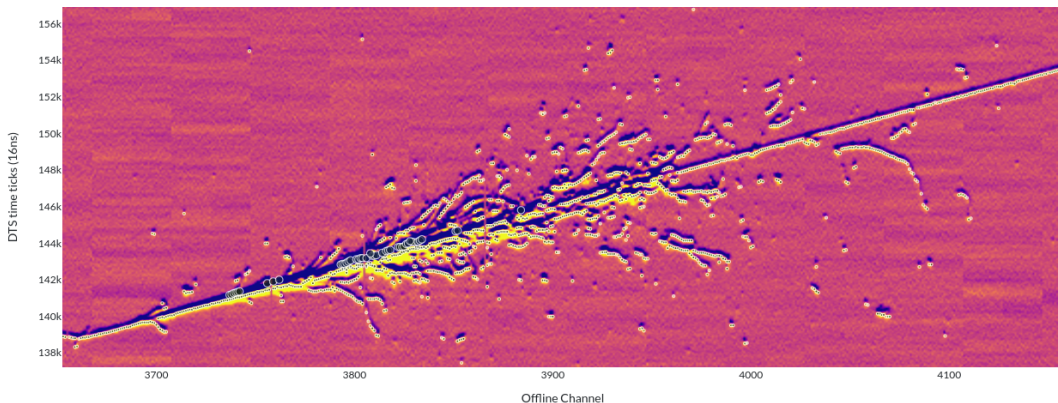
University College London

3 Jun 2024



DEEP UNDERGROUND
NEUTRINO EXPERIMENT

- TP rates far higher than initially expected, causing the readout to fail.
- Tested many solutions, eventually enabling to run TPG on one plane/APA.
- Another solution ready to test, that will enable TPG on all planes.
- Even with TPG on one plane/APA, collecting good-looking data!




Mega TPGeneration: the problem

- Started as with the CBs: ADC scan, all planes SimpleThreshold.
- Ivana calculated ADC threshold from raw data $RMS+5\sigma$: 100.
- We hit 'start_run', and all the readout apps failed immediately!
- TP rates higher than anything we've dealt with before.
- Tried higher threshold (500), increasing queue capacity etc.
 - At that point, even TPG on single plane was failing.




ER...	2024-05-20 10:12:14	runp04srv029eth0	tp_link_1: Unable to push within timeout period (timeout period was 0 milliseconds)	iomanager::QueueTimeoutExpired	np04-srv-029
ER...	2024-05-20 10:12:14	runp04srv029eth0	tp_link_1: Unable to push within timeout period (timeout period was 0 milliseconds)	iomanager::QueueTimeoutExpired	np04-srv-029
INFO	2024-05-20 10:12:02	trigger	Start of run 26299	trigger::TriggerStartOfRun	np04-srv-018

Was failing around $\sim 650\text{kHz}/\text{APA}$.



RawDataHandler


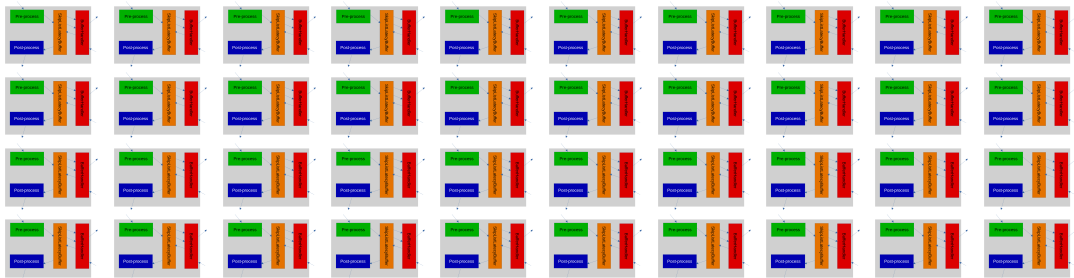


TPDataHandler

- ReadoutModel gets some type of data, places it into pre-processor.
- Next, places the data into a buffer, which sorts.
- With some delay, it will request sorted data from the buffer and place it into post-processing task.
- It will then send processed data out.
- BufferHandler blocks buffer with mutex:
 - External data request.
 - Periodic cleanup of the buffer.

Mega TPGeneration: the explanation

RawDataHandler(s)




- Post-processing from tp-handler task receives less data, Buffer blocked!
- Queue to the TP Data Handler gets filled. Raw Data Handlers cannot push!
- We think it's a combination:
 - BufferHandler occupying the buffer with mutexes.
 - Could be from: frequent buffer cleanups, large data requests etc.
 - 0 timeout on pushing into queue to TPDataHandler.
- Tried a few possible solutions!

We tested a few proposed solutions:

1. Increased the size of the buffer to decrease cleanup frequency.
2. FollyQueue between RawDataHandlers & TPDataHandler non-blocking.
3. Adding non-0 timeout on pushing onto that queue.
 - Seemed to help a little – could run at slightly higher thresholds, fewer warnings.
4. Placing all the TPDataHandler's threads on one NUMA node for shared CPU cache.
 - So BufferHandler, cleanup task & post-processing task each would sit on their own core & hypercore.
 - This made the largest difference.

In the end, a combination of 3. and 4. allowed us to run TPG stably on one plane, with thresholds we require, without crashing/errors/warnings.



- An elegant solution from Giovanna & Alessandro: one TPhandler per plane!
- Splits the workload into three.
- But Trigger will by default make three TriggerActivityMakers per APA?
- Which is fine!
 - TriggerCandidateMakers could merge not just between APAs, but planes too.
 - And it split TriggerActivityMaker's processing into three threads.
- Yet to test it: other work ongoing.

- We have reached the limit of what readout could do.
- We pushed that limit and can now run TPG stably on one plane (all APAs).
 - Huge thanks to Roland, Wes, Giovanna and Alessandro.
- There's more elegant solution waiting to be tested.
 - Hopefully will end the “too much for readout” saga.
 - And most likely expose another limit/bottleneck (trigger/dataflow/tpwriter?)!
- The TP rates we see on the surface detector are far higher than underground.
 - And although FD will be bigger, volume per APA will not.
 - The processing is done on per-APA-level.
 - This is not going to be an issue for FD.