

Ideas for data logging at DUNE

Kurt Biery

DAQ Coordination Meeting

03 June 2019

Background

- Recent discussions with offline folks about how to handle long TPC readout windows (~ 1 sec) at protoDUNE SP
 - Full events tax memory limits for grid jobs (& other issues)
 - A short-term help at protoDUNE is geographic grouping
- Choices of A) classic event building and B) writing data in art/ROOT format intended to be helpful, but maybe they are not helpful for DUNE FarDet
- Informal discussions among DAQ folks before the collaboration meeting about this and other topics
- Giovanna suggested that a move away from (A) and (B) might be more helpful

An idea

- What if each ‘*artdaq* BoardReader’ wrote each data fragment (one fragment per trigger per APA) to a file on a distributed file system in a binary format...
 - The parallel parts of the initial offline processing could access manageable chunks of the full raw event
 - (Of course, this defers the assembling the data from the whole detector into one unit until later in the offline processing, but at that point the data would be smaller)
 - The 150 files per trigger per module could be written to a trigger-specific subdirectory, and they could be packaged with a simple Linux utility like *tar*
 - The job of the classic EventBuilder would become one of book-keeping – making DB entries for the trigger metadata and writing appropriate metadata file(s) on disk

Handling of SNB triggers

- Initial discussions suggested time-based splitting to keep "files" at a reasonable size for files

Another idea

- An alternative to writing the raw data to disk files would be to write it to a key-value store.
 - The translation to something that is ingestible by the offline would need to be considered in this case

Possible to-do items

If we choose to pursue this, some of the work items:

- Evaluate the performance of distributed file systems to see if they would be sufficient
- Same for whatever key-value stores are available
- Consider online use cases to see if a distributed logging of the data would cause unacceptable issues elsewhere
 - (it could make some parts of the dataflow a little easier)
- What else?