DUNE CERN prototype: evolution of data rate estimates and requirements

Maxim Potekhin
Brett Viren
Brookhaven National Laboratory

CERN single-phase meeting
July 24 2015

Overview

- Please look at our previous meetings' material for the context detail, I won't reproduce it here for brevity's sake.
- Looking back at the past two years, there has been a pretty steady flow of information and discussion regarding the size of data for an "event" or "track" in various scenarios.
- Without the final design of the data structures and algorithms (DAQ etc) it was hard to get a reliable estimate for this sort of number, so we had to obtain guidance from MC data, while complementing with "first principles" estimates based on parameters such as ADC clock, number of wires activated etc .
- Zero suppression is a crucial component in data handling in DUNE. As a rule of thumb, the often quoted factor in data reduction due to ZS was 10. We understand this is a ballpark estimate.
- In dealing with the event size estimates (e.g. from MC) the following factors need to be included:

  "Bloat", i.e. exessive data which is due purely to the design of the data structures.

  ZS.

  "Native" compression in ROOT.

- This makes these estimates not as straightforward as we would like them to be.
- CERN prototype will provide expertise and experience in this area, with real data.

A brief history of the single track data dize

- One number quoted in the past was 10MB per track, based on MC data (file size etc).
- The "best guess" data in the proposal as of late April was 5MB per track.
- This was revisited in the process of developing requisite information for the CDR (credits: Brett Viren) in May 2015, and also in the continued editing of the Computing section of the proposal. Current event size estimates are calculated using a spreadsheet uploaded to this indico item.
- More scrutiny in July 2015, apparent editing error on page 50 of the proposal where non-ZS data is added to a prior ZS estimate.
- As a reminder:

  it is estimated that there will be ~68 cosmic muon tracks (or track segments) in addition to the "main" triggered beam event. This includes additional "padding", i.e. readout time intervals just before and just after the trigger which are necessary to fully characterize the beam event. These extra tracks produce significantly more ionization (and correspondingly more data) than the beam trigger events and make the "single track data size" issue a prominent one.

  NB: The current readout scheme in 35t is 2.4 drift times total, not 3.

  The run plan calls for a total of 5M triggers in a few event categories.

"Padded" Readout vs DC mode

- The "nominal" 200Hz trigger rate is meant to reduce beam particle pile-up in the relatively slow detector (cf. Poisson).
- At 200 Hz, reading out more than 2.2 drift times (still less than 35t's 2.4 drift times) the DAQ will oversample the data, i.e. record overlapping portions of it.
- Therefore, it would be more economical to run in a "DC" mode where data is continually acquired throughout the spill.
- For the sparse Kaon beams, a beam-PID trigger can still be used to veto non-Kaon contamination.

Estimating the Data Volume and Rate

- At this point the most trusted number we have for the scale of ZS data due to a "nominal" 4GeV muon track is ~0.1MB (taking into account ROOT compression). This is now considered more realistic than the higher value of 1MB quoted two weeks ago, and is consistent with an estimate used in the proposal (minus the unintended addition of non-ZS data in later versions).
- The number won't be much different for other energies since a large fraction of it is "bloat" i.e. artifacts of the format used.
- We won't be able to take much data without ZS since the rate will be overwhelming.
- Based on the parameters quoted above, and subject to the run plan and DAQ setup we arrive to the following estimates
  - ~5MB per triggered event
  - instantaneous DAQ rate of ~1GB/s
  - sustanained average rate of 0.20GB/s
- Considerations:
  - disk buffer (or multiple buffer nodes) will need to be designed to handle the instantaneous rate
  - we'll likely require network connectivity of ~2gbps bandwidth
  - close coordination with CERN central services will be needed to operate at these rates
- The revised estimates leave us with more breathing space and the run plan may potentially be augmented to take advantage of that.

Moving the data – initial thoughts?

(this is an illustration, not a design)

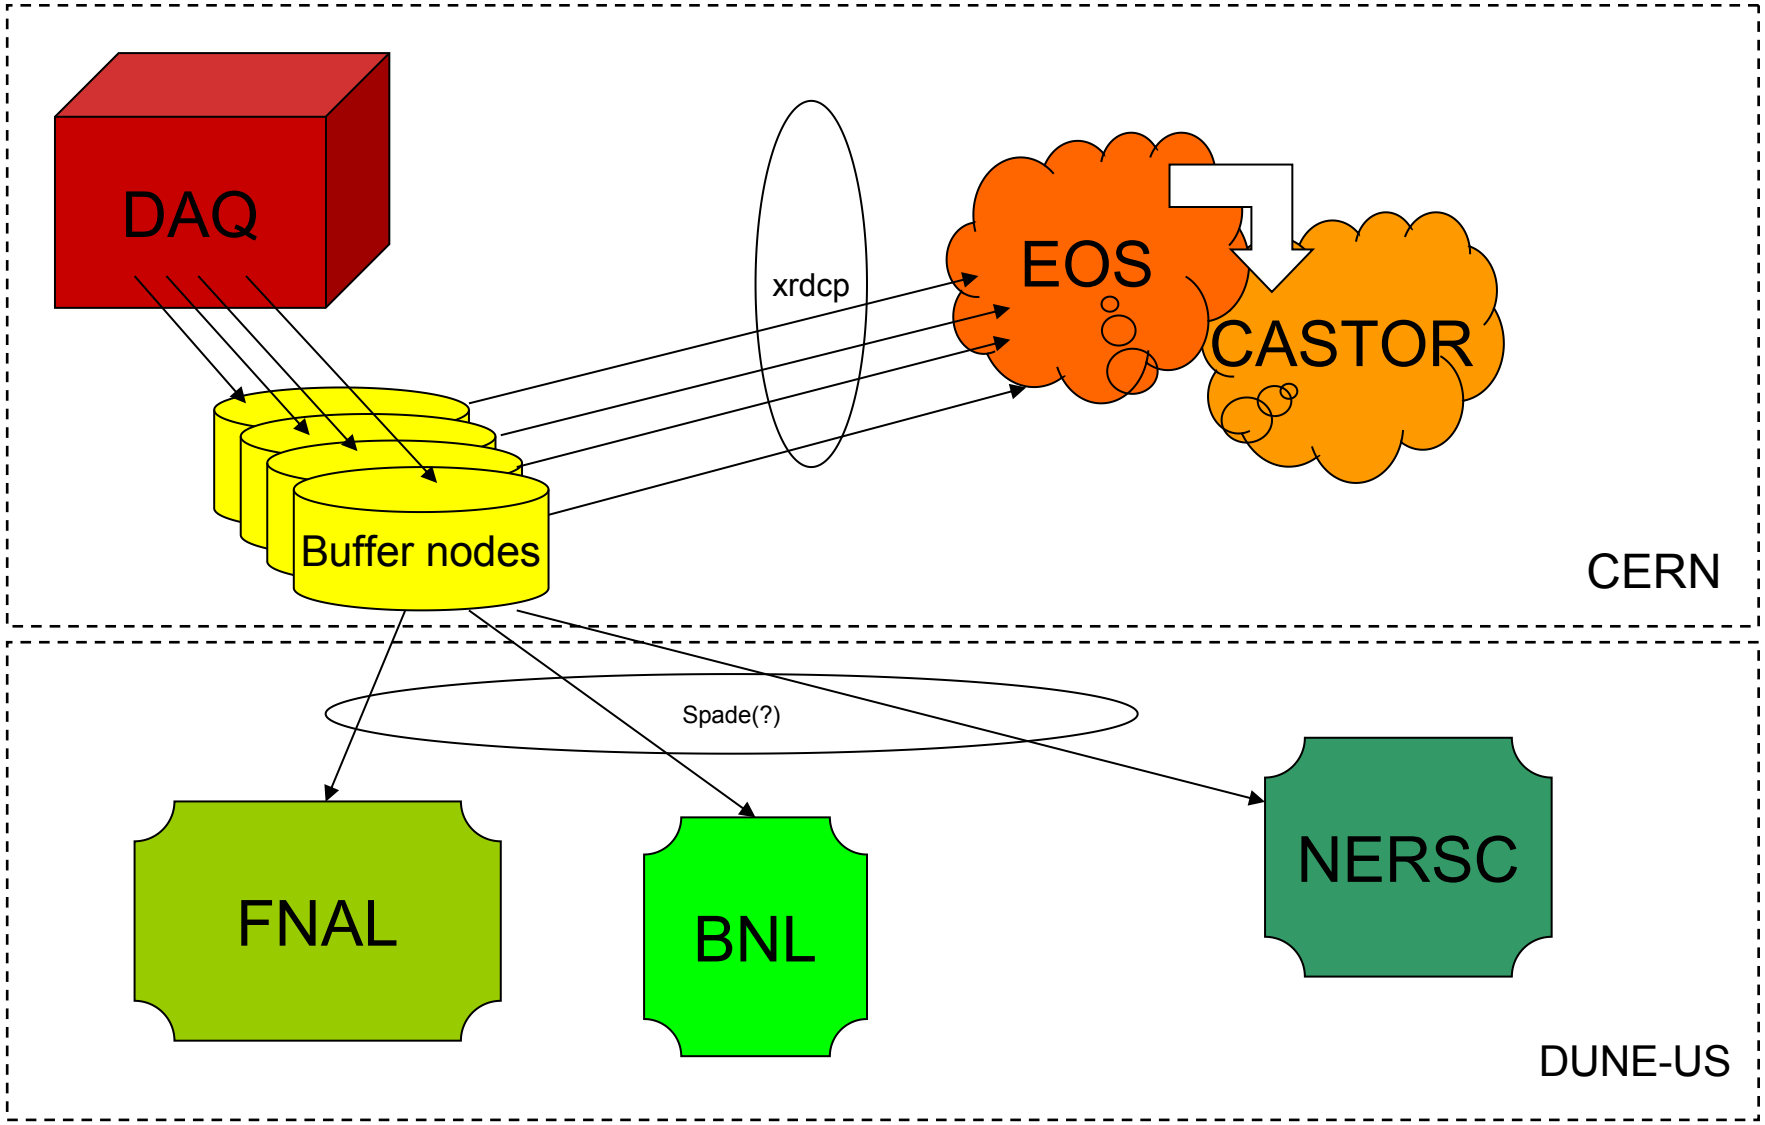

Particle Range

- Since the data load is mostly due to cosmic ray muons, reading out only the part of the detector which contains a stopped beam particle could help reduce the data rate and volume. Need comments from the measurement program team.