# DRAFT

# DUNE CERN prototype Computing: data rate estimates and requirements

Maxim Potekhin
Brett Viren
*Brookhaven National Laboratory*

CERN integration meeting
TBA

# Intro

- According to the proposal submitted to SPSC, the data collected with the DUNE prototype at CERN (provisional designation DUNE-PT) is to be committed to tape at CERN as the primary custodial copy.
- Data will be replicated from a disk buffer at CERN to a few locations in the United States. There will be a full replica of the data at FNAL used for production, and partial copies at other DUNE institutions.
- Bulk of the processing will be performed at the US computing sites. An effort will be made to establish "express processing streams" for a portion of the data for QA and debugging purposes.
- Recording the data collected with the prototype is obviously critical to the experiment but it also constitutes one of the central points of interaction and integration between DUNE-PT and CERN.
- Disclaimer: we are still in early stages of planning and developing relevant information, so any and all numbers quoted in this presentation should be considered as preliminary.

# Overview of the data characteristics

- One significant factor that defines the characteristics of LAr TPC in the context of the test beam is the combination of relatively long electron drift time (nominal ~2.25ms) and occupancy from cosmic ray muons (cf. the top face of the detector is ~50m$^2$).
- As a result, 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.
- The run plan calls for a total of 5M triggers in a few event categories. Since the data will be dominated by signals from cosmic ray muons, it is possible to estimate of the total data volume just by looking at the data due to "cosmic muon overlay".
- Signal processing strategies such as Zero Suppression (ZS) and Region of Interest (ROI) are being considered to reduce the volume of the data to be transmitted by DAQ and committed to mass storage. This is work in progress and at this time we aim to provide a range of parameters to help planning process.

# Estimating the Data Volume

- At this point we have a range of values for the scale of zero-suppressed TPC data due to a "nominal" 4GeV muon track from 0.1MB to 1.2MB. This obviously depends on data design, formats and algorithms.
- Data from the photon detector is of substantially smaller volume so it won't be considered here.
- Refining the algorithms currently available to us and improving the estimate will take some time, for now we must keep with the higher end of the range to ensure adequately provisioned resources and scalability in DAQ and other components.
- Based on the parameters quoted above, and subject to the run plan and DAQ setup we arrive to the following estimates

  ~100MB per triggered event (TPC readout).

  ~500TB for the projected volume of raw data to be recorded

- We would like to have an opportunity to expand certain datasets if such need would arise during the data taking, and for that purpose have a safety margin of x2.
- The baseline request is then for 1PB of tape storage at CERN as the primary custodial copy of the data.

# Assumptions about the beam

- 45s SPS cycle
- Two spills per cycle, each of 4.5 duration
- Debunching
- Low intensity of the beam to minimize pile-up (i.e. <200Hz during the spill)

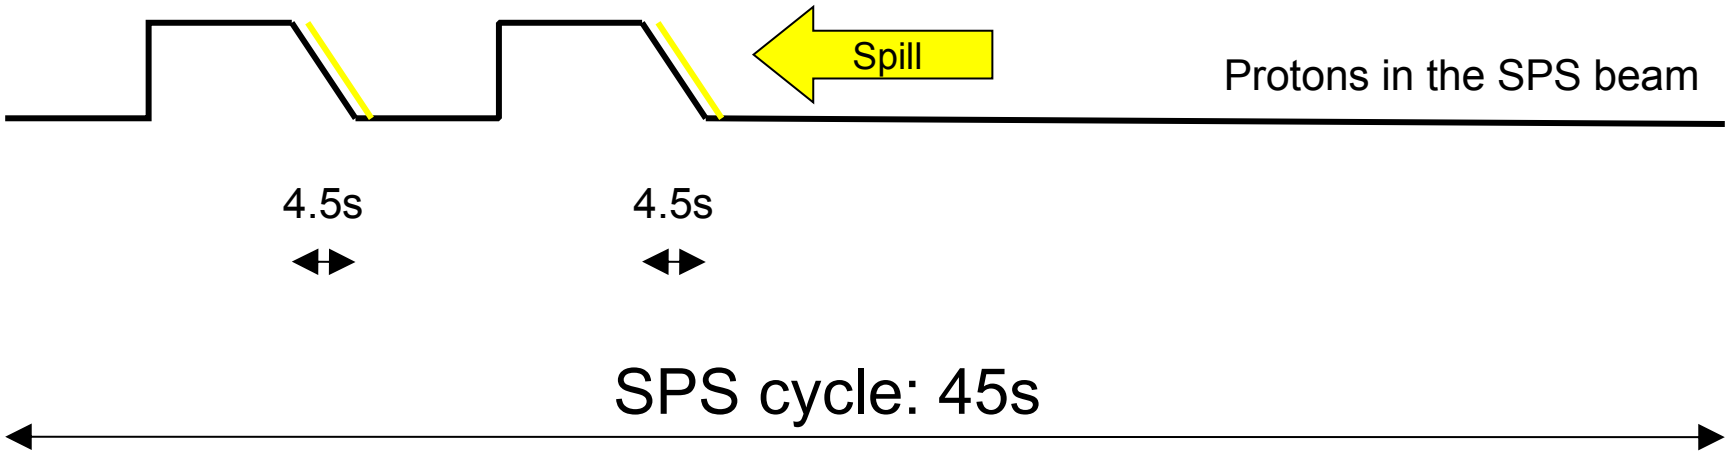

# Estimating the Data Rates

- The top trigger rate for DUNE-PT will likely be determined by throughput of the DAQ system and its attached storage (both peak and sustained) - but there are a few other parameters.
- Assuming the nominal 200Hz rate of beam partifcles during the spill, the spill duration allows for 2.4 drift time intervals on average to be read out for each beam particle. As indicated above, there are two spills per cycle, each of approx. 4.5s duration). This translates DAQ peak rate of ~10GB/s and into 2.0GB/s sustained data rate averaged over the SPS spill.
- That number obviously falls into the higher part of the range currently handled by CERN experiments and it poses challenges to every component of the data handling chain, from internal DAQ bandwidth to network connection out of CERN (cf. ~20gbps from CERN to ATLAS Tier-1 at BNL). The number is conservative in that it is likely to be revised downward depending on progress with better zero suppression and localized readout of ROI, and other similar measures.
- At the same time, we can establish the lower practical limit for the data rate requirement. At 100MB/s the measurement program can be completed in 2 months (with 100% efficiency). This may cut into our statistics, as it does not allow for troubleshooting time as well as other interruptions inevitable in a test enivironment. For that reason, it is reasonable to define the lowest possible requirement for the bandwidth as 200MB/s.

# Capturing the Data from DAQ

- LHC experience - both ATLAS and CMS transmit data from their respective DAQ to EOS, with subsequent creation of a custodial copy in CASTOR.
- Buffering in DAQ: several buffer nodes may be needed to absorb the instantenous rate coming out of DAQ.
- DAQ to EOS: multiple connections via xrdcp.
- For DUNE-PT we do not anticipate the need for Tier-0, so EOS would be used to offload data from DAQ and buffer it before it's committed to CASTOR.
- There are multiple tools under consideration to transport data from CERN to data centers in the US, one possibility is to leverage existing expertise with Spade (IceCube, Daya Bay) which provides ease of installation, redundancy and monitoring.
- Procedures will need to be put in place to restore data from tape (stored at CERN) in case replicas are corrupted.

# Moving the data – initial thoughts?

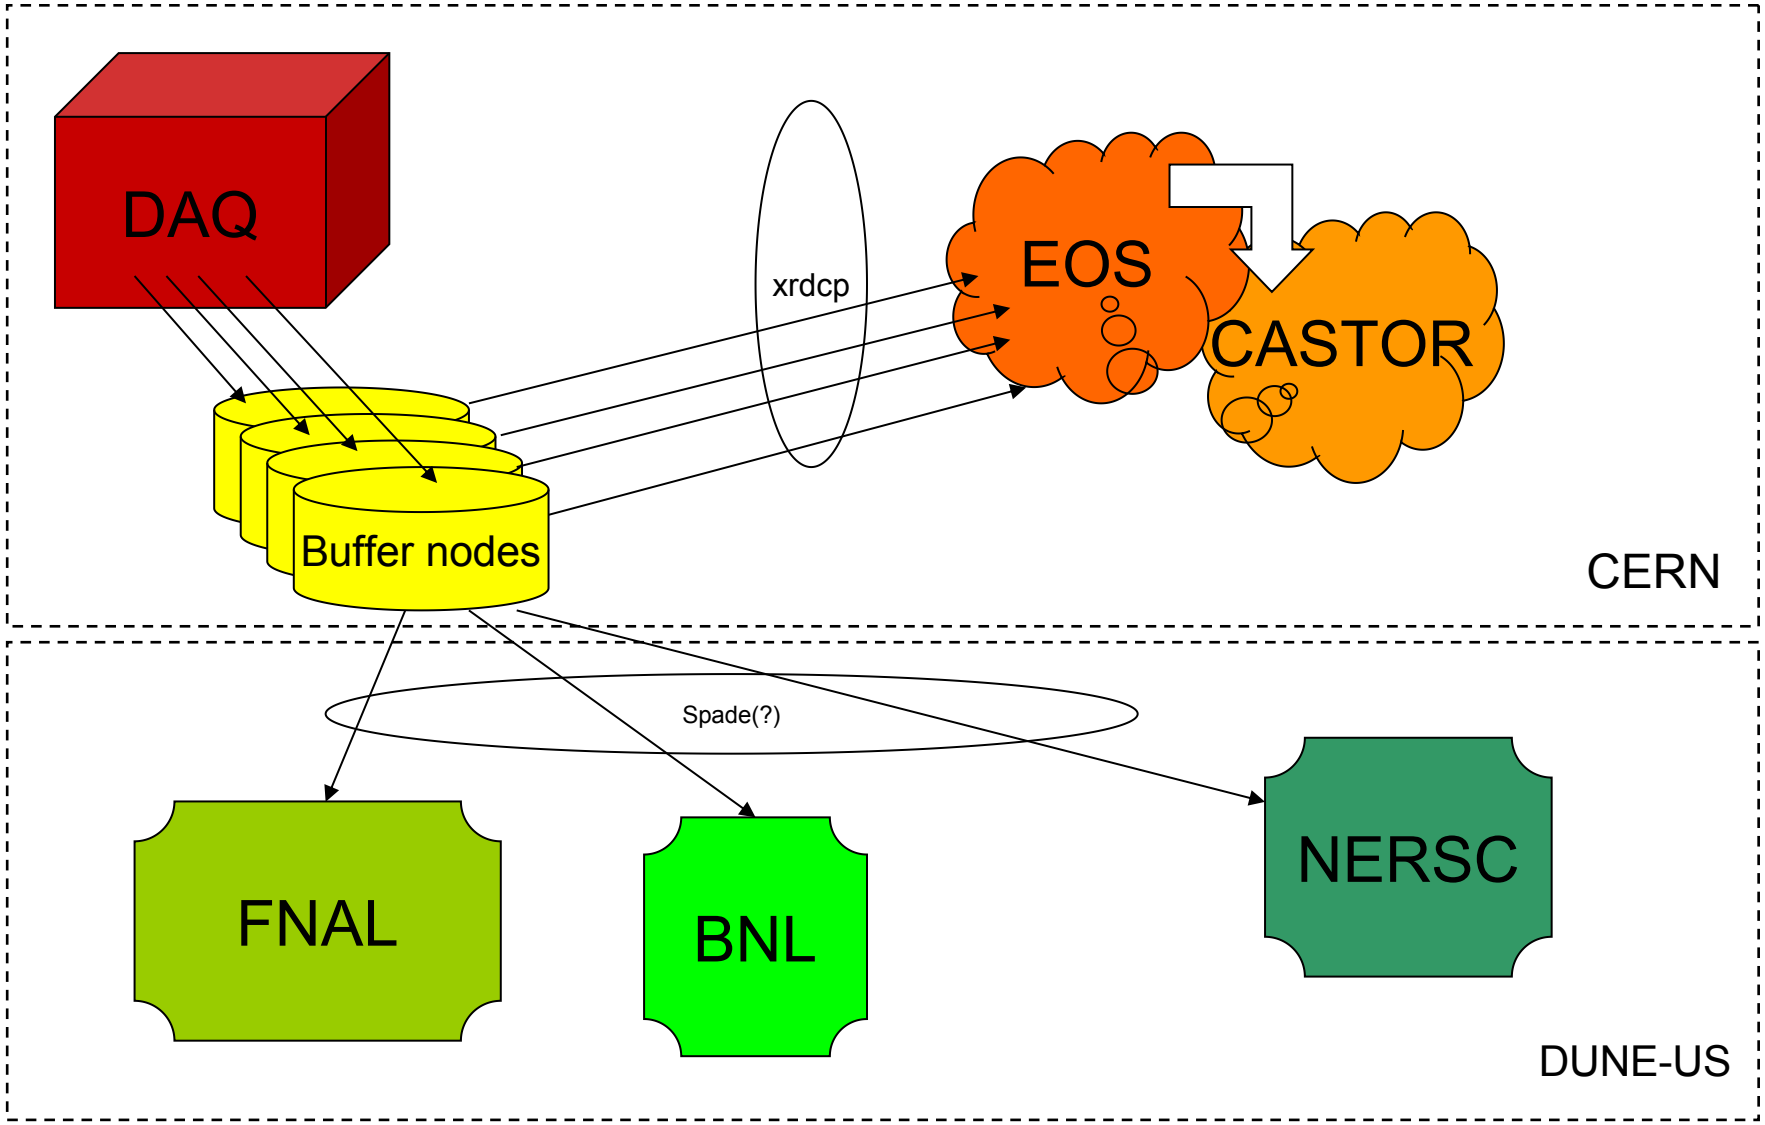

# Summary

- The exact size of the DUNE-PT event will depend on R&D done in both online and offline areas, for now we are using a conservative estimate of 100MB per event.
- Initial conservative estimate is 1PB of tape storage at CERN as the primary copy of the data.
- The DUNE-PT DAQ will need a network connection of at least 200MB/s and preferably twice that.