The protoDUNE-SP experiment and its prompt processing system

Maxim Potekhin (Brookhaven National Laboratory) for the DUNE Collaboration

EPS Conference on High Energy Physics

July 6 2017



DUNE

- Primary Science Program
 - Precision measurement of neutrino oscillation parameters
 - Search for proton decay
 - Detection and measurement of the neutrino flux from supernovae



- DUNE detectors and subsystems
 - for details please see a presntation by Nicola McConkey in this session
 - a massive 40kT Liquid Argon modular time-projection chamber deployed as a far neutrino detector 1,300 km from FNAL and 1.5km underground



Single-Phase LArTPC (DUNE and its prototypes)

- One of the *two design options* for the modules of the DUNE Far Detector TPC
- Liquid Argon serves as both the target and the sensitive medium.
- Wire Sensors: two *induction planes* at a **stereo angle** (35.7°) and one *collection* plane. Implemented as Anode Plane Assemblies (APA).
- Two coordinates (in the plane) are determined via stereo projections on three planes, and the third (along the drift) via the time measurement



DUNE Single-Phase LAr TPC prototypes

- Experiment with a 35t Single-Phase prototype at FNAL
 - Construction in 2015, data taking completed in early 2016 (cosmic rays)
- Larger scale prototype: protoDUNE-SP
 - Under construction at CERN
 - Major structural elements of the TPC are of the same size as in DUNE
 - Beam and Cosmic Ray data will be taken in 2018
 - Detailed detector characterization for different particle species at different momenta

The 35t prototype at FNAL

- View of Anode Plane Assemblies installed inside the cryostat
- Image below shows the field cage
- Integrated Photon Detector
- Cold electronics (analog and digital)
- ~100 scintillation paddles outside (for triggers)





DUNE

M.Potekhin | protoDUNE-SP

The 35t experiment data

- Took cosmic ray data in January-March 2016
- Please see a dedicated presentation in this conference (M.Wallbank)
- Example: three projection views of an electron shower
- Encountered more electronics noise than expected, problems ultimately diagnosed and understood
- Automated reconstruction shown to work on actual data



The protoDUNE program at CERN

- The protoDUNE program:
 - large-scale engineering prototypes of Single- and Dual-Phase Liquid Argon TPCs for DUNE (CERN designation NP04 and NP02 correspondingly)
 - an international project made possible by the <u>CERN Neutrino Platform</u> Organization
- Test-beam facility is under construction in the extension (EHN1) of the CERN North Experimental Area Hall, with a purpose-built tertiary beam from the SPS (H4) to provide various particle types



Goals of protoDUNE-SP

- Prototype the production and installation procedures for the Single-Phase Far Detector design.
- Validate the design from the perspective of basic detector performance.
- Accumulate large samples of test-beam data to understand/calibrate the response of the detector to different particle species.
- Demonstrate the long-term operational stability of the detector as part of the risk mitigation program ahead of the construction of the first 10kt Far Detector module.





The Liquid Argon TPC of protoDUNE-SP

- The protoDUNE detector includes the following principal DUNE components
 - 6 Anode Plane Assemblies (APA)
 - 6m H × 2.4m W
 - 18 Cathode Plane Assemblies (CPA)
 - $2m H \times 1.2m W$
 - Field Cage (FC)
 - Photon Detector (PD)
 - Cold Electronics: inside the cryostat
 - Electron drift distance: 3.6m
 - 500V/cm drift field



- Warm Interface Boards are outside of the cryostat volume and receive data from cold electronics, transmit to DAQ via optic fiber
- Data is received by Board Readers, then assembled from fragments by Event Builders and written into high-performance storage appliance (Online Buffer)



The Anode Plane Assembly Construction

Anode Plane Assembly: 6×2.4m





Other protoDUNE-SP subsytems



- Beam Intrumentation: TOF × 2, Cherenkov, Fiber Tracker (for beam particles)
- Cosmic ray tagger (in red)
- Some of these detectors will have data streams separate from those of the TPC will need to monitor separately and merge offline

protoDUNE-SP: readout parameters

- Please see the BACKUP section for more detail
- TPC channel count: 15,360
- Digitization frequency: 2MHz
- Nominal readout window: 5ms
- Nominal beam trigger rate: 25Hz
- Single readout size: 230MB
- Lossless compression factor: 4
- Post-compression peak data rate: 1.4GB/s, Average: 0.300GB/s
- Contingencies may lead to higher rate, the benchmark of 3GB/s is assumed
- Up to 3PB of data to be collected during the run



protoDUNE-SP: data flow



Online Monitoring vs Data Quality Monitoring

- Two complementary types of data processing providing actionable information
- Online Monitoring
 - close to real time
 - high bandwidth
 - fixed CPU resources located in the vicinity of DAQ
 - data streaming out of DAQ is the sole source of input
- Data Quality Monitoring (DQM)
 - processing time: up to an hour, allowing for more time-consuming algorithms
 - low bandwidth (only a small fraction of data is processed)
 - flexible access to a variety of resources (CERN Tier-0, local ad hoc clusters etc)
 - can use data which is not part of the DAQ stream e.g. the Beam Instrumentation data or any other type of data which will need to be merged downstream

DQM types of processing

- A summary/plots of ADC-level data e.g. mean/RMS values at channel level and as a statistics over various groupings requires data decompression.
- A summary/plots of the ADC-level data in frequency space (FFT) on channel waveforms. This largely provides measures of noise and its evolution.
- A summary of the data after signal processing (uses FFT). It includes
 - "stuck code" mitigation (ADC anomalies)
 - coherent noise removal important for catching more subtle noise issues
 - noise subtraction and filtering
 - deconvolution of the response function (please see the backup slides)
 - calculation of signal correlations for diagnostic purposes
- Visualization e.g. histograms, plots, or a basic 2D event display before and after signal processing
- Beam Instrumentation monitor, validation of the trigger vis-a-vis the TPC data

An example of the DQM Workflow



protoDUNE prompt processing system (p3s)

- "Express streams" and "prompt processing systems" exist in many experiments as complementary to main production systems
- p3s provides more agility, automation and monitoring than is possible by using the batch system alone, for "nearline/prompt" job execution
 - minimalistic design and straightforward installation
 - runs with or without Grid
 - simple automation of workflows: modeled as DAGs and use templates
- Flexible prioritization (and de-prioritization) scheme to assure timely execution of critical workflows
- Automatic injection of workflows triggered by fresh data
- Distributed data handling is not a part of the system
- Code and documentation on GitHub: https://github.com/DUNE/p3s



p3s: technologies used

- pilot-based job dispatch, can be implemented with or without a batch system
- Utilising well established, popular, reliable components
 - Django (Python)
 - Apache
 - PostgreSQL (may use other RDBMS)
- CERN Openstack VM instances to host Web and DB servers
- JSON
- XML: GraphML schema for describing workflows as graphs
 - compatible with a number of third party editing and visualization packages
- NetworkX (Python)
- CVMFS for software provisioning
- EOS Storage at CERN
- Fermi-FTS to get access to data and to transmit results for storage
- Most of the payload code is based on LArSoft/art developed and maintained at FNAL, and is expected to be portable between Online Monitoring and DQM



p3s - workload parameters

- FFT per event ~5 min
- Full reco per event ~30 min
- With the nominal 25Hz trigger and projected SPS spill profile: 300 events/min
- An assumption only a few O(1) events per minute will need to be fully processed i.e. result in event displays before and and after signal processing
- A small number of events from each file will be used to ensure the results are most current
- The projected footprint of p3s in CERN Tier-0 is about 300 cores which will handle all kinds of DQM payloads, for example
 - 300 cores keep up with 20% of the event rate just for the FFT stage
 - ...or 3% for RECO
- The exact composition and prioritization of p3s jobs is yet to be determined
- Memory use is a concern as vast majority of available nodes have 2GB/core

Summary

- The DUNE Collaboration has established an ambitious prototyping program which started with a 35t detector at FNAL, and which now involves two largescale Liquid Argon detectors under construction at CERN.
- Both the protoDUNE-SP detector and its projected data are large scale.
- To facilitate comissioning of the detector and to ascertain its optimal operation and efficient use of beam time a Data Quality Monitoring (DQM) system will be put in place.
- The protoDUNE prompt processing system "p3s" has been developed to support DQM, it is portable and flexible due to utilization of standard technologies and components which are easy to deploy.
- p3s has been tested at CERN, with realistic jobs and at scale
- Service deployed in the CERN Cloud (Openstack)

Backup slides



protoDUNE-SP: readout parameters detail

- TPC channel count: 15360
- Digitization frequency: 2MHz
- Nominal readout window: 5ms
- Nominal beam trigger rate: 25Hz
- Single readout size: 230MB
- Lossless compression factor: 4
- Post-compression peak data rate: 1.4GB/s
- Average data rate: 0.300GB/s
- Contingencies
 - higher trigger rate (50-100Hz)
 - lesser compression (if the noise levels are above predicted)
 - larger number of cosmic triggers (if necessary for certain measurements)
- To size up the data handling system, the benchmark of 3GB/s rate is assumed
- Nominal 20Gbps network bandwidth from the experiment to CERN central storage
- Up to 100TB daily data volume
- ~3PB of data to be collected during the run

Signal formation in the Single-Phase LArTPC

Drifting charge Signal on wire

Shockley-Ramo theorem — instantaneous electric current induced by a charge moving in the vicinity of an electrode: $i=E_yqv$

ADC input

- Electronics response (e.g. amplifier) is • convoluted with the wire signal shape
- Bipolar signals on induction wires •





Figure 2. Simulated digitized signals from a central wire in each plane for a minimum-ionizing particle track traveling parallel to the wire plane and perpendicular to each wire plane orientation. The number of drifting ionization electrons is assumed to be 4600/mm of track. The gain and peaking time of the cold electronics are assumed to be 14.7 mV/fC and 2 μ s, respectively.

Weighting Field of a U Wire

M.Potekhin | protoDUNE-SP

p3s: web-based UI

թઝઁջ

Sites	DAGs Wor	<u>kflows</u>	8	Jobs F	lots	5	Ē	Data	-	Data	Туре	s						
	Summary		System															
Object 🗠		Number 🗠		Attribute 🗠	e 🗠 Value 🔿													
Pilots: total/idle/running/stopped		75/1/74/0		Current time	05/04/17 21:16:40 EST													
Jobs: total/defined/running/finished 1075/:			jobs	(total in DB: 5)														
Workflows	: total/-/-/-	0/-/-/-	Job St	ates (select one):														
			• (• (Subm	Running Finished Pilot Timed Out														
			ID 🛆 U	Juid 🛆	User 🛆	Name 🛆	Wfuuid 🛆	type 🗠	Payload	Params 🛆	Pri. 🛆 t.li	imit 🛆 Sta	ate 🛆 d	lefined 🛆	started 🗅	stopped 🗠	Pid 4	Err
			7 6	44ee502-1112-11e7-9f74-001d096774dd	mxp	sleep1000	-	sleep	/mnt/nas01/users /mxp/projects /p3s/inputs /jobs/sleep100.sh	-	1 10	10 rur	nning 21	0170411 9:55:51	20170411 19:55:57	-	15952	-
			3 k	3378608-1f06-11e7-b642-001d0967b087	mxp	sleep1000	-	sleep	/home/maxim /projects /p3s/inputs /jobs/sleep100.sh	-	1 10	10 no	nstarter 1	0170411 8:32:09	20170411 18:32:12	-	_	_
			4 3	08007ca-1107-11e7-a220-001d0967b087	mxp	sleep1000	-	sleep	/mnt/nas01/users /mxp/projects /p3s/inputs /jobs/sleep100.sh	_	1 10	10 fini	shed 21	0170411 8:35:39	20170411 18:35:43	20170411 18:37:24	11401	0
			5 6	2636506-1112-11e7-9fb8-001d096774dd	mxp	sleep1000	-	sleep	/mnt/nas01/users /mxp/projects /p3s/inputs /jobs/sleep100.sh	-	1 10	10 rur	nning 21	0170411 9:55:47	20170411 19:55:57	-	4051	_
									/									

p3s: GraphML for workflow description

```
<data key="d104">100</data>
  <data key="d105">1</data>
</node>
<node id="reco">
  <data key="d101">/bin/grep -c TEST $OUTTXT > $FINTXT</data>
 <data key="d102">{"P3S EXECMODE":"SHELL"}</data>
  <data key="d103">reco</data>
 <data key="d104">100</data>
 <data key="d105">1</data>
</node>
<node id="NOOP2"> <!-- Dummy node is a placeholder since out-data is final -->
  <data key="d101"></data>
  <data key="d102">{}</data>
  <data key="d103">noop</data>
  <data key="d104">0</data>
 <data key="d105">0</data>
</node>
<!-- ++++++ DATA ++++++ -->
<edge source="NOOP1" target="reduction">
 <data key="d1">input.txt</data>
 <data key="d2">/home/maxim/p3sdata/</data>
  <data key="d3">TXT</data>
  <data key="d4">INPTXT</data>
  <data key="d5">Simulated input</data>
</edge>
```

p3s service: Running in CERN Cloud

CERN Accelerating science									Si	gned in as: mpotekhi Sign (out Directory	
🖸 openstack	■ DUNE →										Тоо	s
Project	Instance	S										
Compute												
Overview			INSTANCE NAME = +					A LAUNCH INST	MORE ACTION	MORE ACTIONS +		
Instances	Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions	
Volumes	p3s-content	CC7 - x86_64 [2017-04-06]	188.185.77.722001:1458:d00:9::100:142	m2.medium	-	Active	cern-geneva-a	None	Running	3 days	CREATE SNAPSHOT	
Access & Security	p3s-db	CC7 - x86_64 [2017-04-06]	188.185.85.2052001:1458:d00:f::100:1c7	m2.medium	-	Active	cern-geneva-a	None	Running	3 days, 1 hour	CREATE SNAPSHOT	ų
Container Infra	p3s-web	CC7 - x86_64 [2017-04-06]	188.185.85.1752001:1458:d00:f::100:1a9	m2.medium	-	Active	cern-geneva-a	None	Running	3 days, 1 hour	CREATE SNAPSHOT	ų.
Orchestration	Displaying 3 items											