



Hall D Computing

JLab Computing Review

Mark Ito

Jefferson Lab

December 16, 2021



Outline

- Introduction
- Hall D Online Skim System
- Calibration
- Reconstruction
- Analysis
- Simulation
- Artificial Intelligence and Machine Learning
- GPUs and Amplitude Analysis
- Reconstruction on the Open Science Grid
- Areas for Improvement
- Conclusions

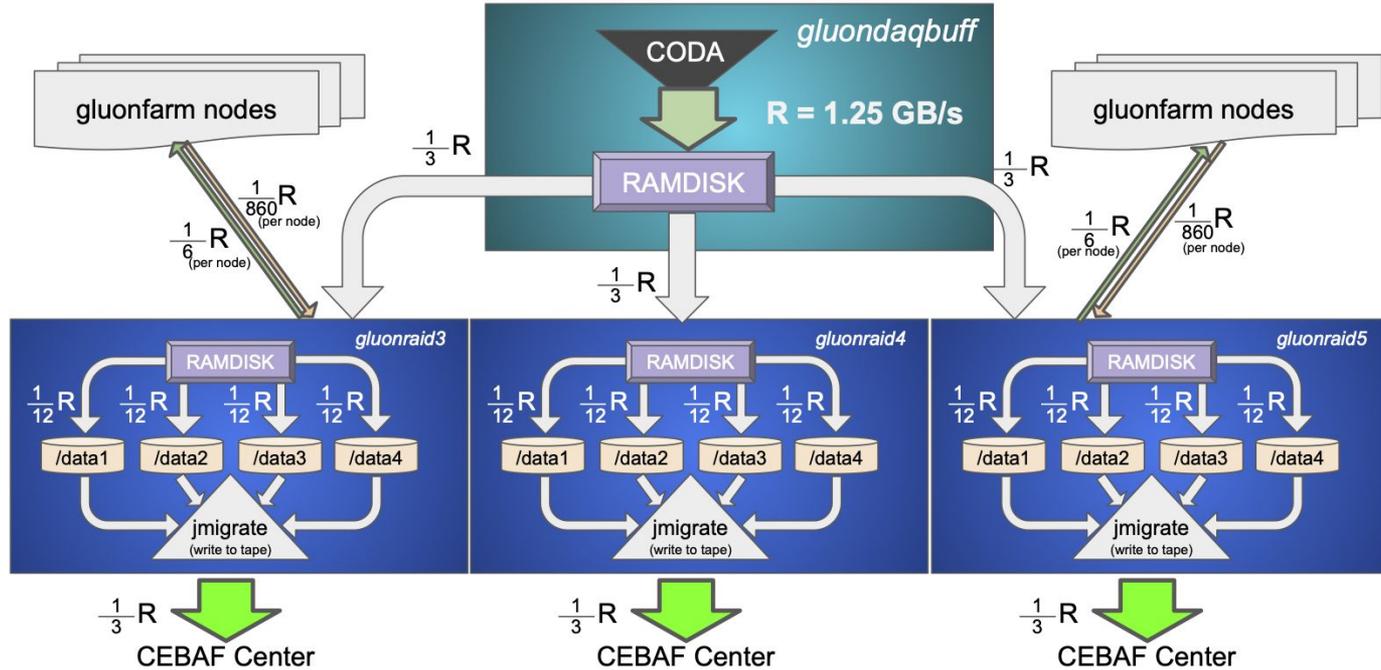
Introduction

- GlueX designed to search for exotic mesons (hybrids)
- Needs high-statistics, large data sets
- Aim for central management of software effort as much as possible
- Use standard interfaces to facilitate interaction with the data
- Standards enable economies of scale, reduce errors, conserve computing resources
- Exploit new opportunities from developments in machine learning

Hall D Online Skim System (HOSS)

- GlueX Phase I running: 400 MB/s data rate
- GlueX Phase II running: 1.3 GB/s
- HOSS: a new system
 - Deal with increased data rate
 - Facilitate online skimming of calibration data
- Front end: RAM disk
- Back-end: three independent RAID servers each with its own RAM disk input buffer
- Skimming calibration events saves tape latency incurred in offline environment

The HOSS System



Calibration

- Online calibrations an on-ongoing effort
 - Speeds up time to physics results
 - Allows online data quality checks to have better sensitivity
- Recent running: key physics signals extracted within hours of data recording
- HOSS facilitates the online calibration process
- Still challenged by charged track reconstruction for all but a sample of events
- AI-based systems will help in the future (see section on FPGAs)

Reconstruction

- Most compute intensive activity
- Use the JLab Farm
- **And** use multiple High Performance Computing (HPC) sites
 - NERSC
 - Pittsburgh Supercomputing Center (I and II)
 - BigRed 3 (Indiana University)
- 40/60 split between JLab and HPC
- Next frontier: reconstruction on the Open Science Grid

Reconstruction Statistics

Run period	Length of processing	Fraction of events at each site	CPU used at each site in millions of core-hours	Number of jobs
2017-01	1 Month	JLab (100%)	6	42165
2018-01	4 Months 1 Month	NERSC (81%) JLab (19%)	20.4 2.3	77603 16279
2018-08	2 Months 1 Month 1 Month	NERSC (52%) PSC (Bridges) (26%) JLab (22%)	9.25 0.81 2.1	24669 6990 13358
2019-11	4 Months 2 Months 2 Months 4 Months 4 Months	NERSC PSC (Bridges) PSC (Bridges2) BigRed3 JLab (57%)	12.4 2.33 4.03 3.56 23.9	45236 17752 19694 16392 119397

Analysis

- Total volume of reconstructed data (REST) too large to keep on disk at JLab
- Post-reconstruction process developed to make ROOT trees, the “Analysis Launch”
- Specific reactions requested via web form
- Pass made over reconstructed data producing multiple independent set of ROOT trees, one set per reaction
- Multiple analysis launches on a given reconstructed data set
 - with new reactions or
 - same reactions with different ROOT tree generation algorithms
- Output for a given reaction small enough to ship to collaborating institution
- Tape latency a bottleneck for Analysis Launches
 - REST data on disk would improve throughput

Analysis Launch Web Form

Please fill out your reaction below:

Use add/remove particle to add/remove a particle from the products side of the reaction.

Each product comes as a set of three objects:

- 1) the main selector where you can select the product.
- 2) a tri-state button to let you flag the particle as "m" (missing) or "M" (NOT Mass constrained) as desired.
- 3) a checkbox to indicate the product decays

B (Beam Bunches): 3 T (Extra Charged Tracks): 3 F (Fit Type): P4 and Vertex U (unused tracks):

Initial Particles ----> Final State Particles



LEVEL 1



```
Reaction1: 1_14_7_17_14
Reaction1: Decay1 17_7_8_9
Reaction1: Flags B3_M17
```



Simulation

- Standard Tools Developed for
 - event generation
 - detector simulation (interactions with matter)
 - smearing (adding detector resolution and accidentals)
- Post simulation processing identical to that for real data
- Transition from HDGeant (CERNLIB-based) to HDGeant4 (Geant4-based) complete
 - multi-year process
 - validation of HDGeant4 results against those from HDGeant
 - Now HDGeant4 is used in production simulation exclusively

Simulation with MCwrapper

- Web-based simulation job submission
- Menu presented for configuration choices
 - event generator
 - real data runs to simulate
 - reaction choice (for ROOT tree generation)
 - curated set of other options
 - command line interface also available
- Jobs submitted to the Open Science Grid (principally)
- Statistics kept on all projects

MCwrapper Web Form

halld_recon version: recon-2019_11-ver01.0

halld_sim version: 4.34.0

version Set: recon-2019_11-ver00_2.xml

Run Number: 11366 Number of Events: 1000000

Output Directory Name: My_MC

Generator: bggen

Full Path to Generator Config: full path must be reachable by tbrinton from an ifarm r

Flux to Generate: ccdb cobrems

Min Photon E: 3.0 [GeV] Max Photon E: 11.6 [GeV]

Post-Processing: None

Geant Version: Geant3 Geant4

Geant Secondaries?

Background: None

ReactionFilter: Reactions Set: Reaction1_14_7_17_14

analysis version Set: None

recon-2019_11-ver00_2.xml

- **created:** 2021-11-29
- **description:** Modified from 4.24.0 to get modern simulation. Includes Sean's backport of tagger fixes.

Package	Version	Directory Tag	Debug Level
amptools	0.12.2		
ccdb	1.06.07		
cernlib	2005		
diracxx	2.0.1		
evio	4.4.6		
evtgen	01.07.00		
geant4	10.04.p02		
gluex_MCwrapper	v2.5.0		
gluex_root_analysis	1.20.0	rec191110	
halld_recon	recon-2019_11-ver01.0		
halld_sim	4.34.0	rec191110	
hdds	4.10.0	x323	
hdgeant4	2.30.0	rec191110	
hd_utilities	1.36		
hepmc	2.06.10		
jana	0.7.9p1	x323	
lapack	3.6.0		
photos	3.61		
rcdb	0.06.00		
root	6.08.06	bs221	
sqlitecpp	2.2.0	bs130	

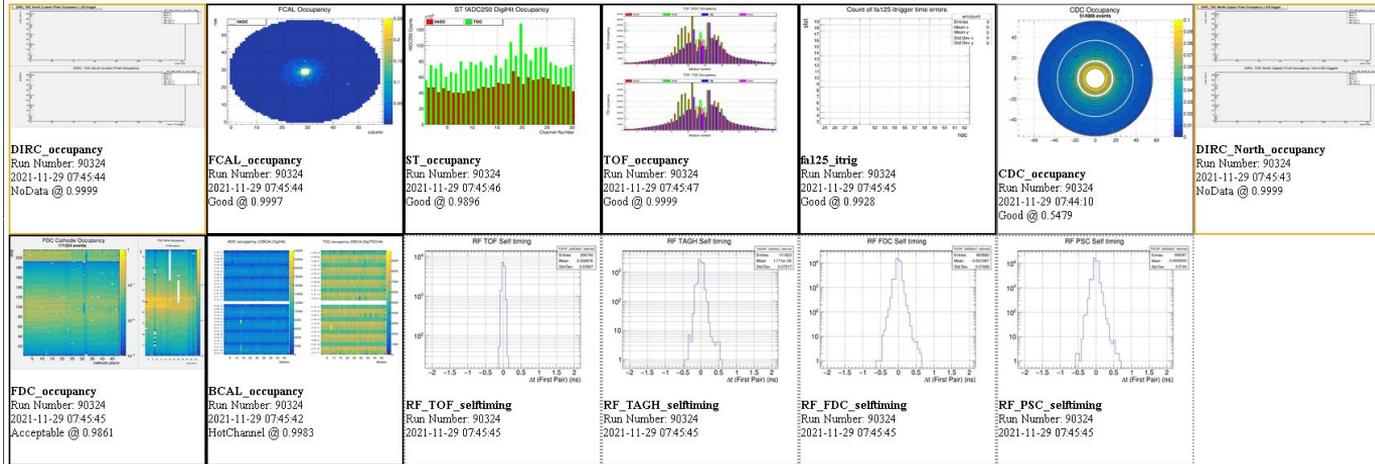
Artificial Intelligence and Machine Learning

- Forward calorimeter shower shape analysis: photons vs. hadronic interaction split-offs
- Running conditions predicting central drift chamber calibration constants
- FPGAs for Online Processing
- Hydra (online monitoring histogram problem detection)
- Bethe-Heitler lepton pair identification: e^+e^- vs. $\mu^+\mu^-$
- Charged particle identification

Hydra Display Screen

90324

Last Updated: 14.27 second(s) ago



GPUs and Amplitude Analysis

- Amplitude analysis (or partial-wave analysis) key technique in the GlueX program
- Standard tool developed at Indiana University: [AmpTools](#)
- Used for about a decade at GlueX
- Feature: can be run on GPUs
- JLab GPU-enabled nodes being used in this context
- New high-memory GPUs would allow scaling of data set size and physics-model complexity

Reconstruction on the Open Science Grid

- So far, only simulation has been done on the OSG
 - No large input data files to stage
 - Simulation can be broken into small pieces to run opportunistically
- Demonstration project has been developed and used
 - Processes 20 GB raw data files in 60 to 70 pieces with jobs that run for 2 to 3 hours
 - Maintains access to opportunistic resources
 - Incurs the need to merge 60 to 70 output files, both reconstructed data and diagnostic ROOT histograms
- Deployment at scale a project for the coming year

Areas for Improvement

- Data Catalog
- Work Flow Management
- Maintenance/development of legacy systems
- Improve Continuous Integration
- Improve Software Testing
- Unit Testing
- Documentation

Conclusions

- Hall-D/GlueX in full production performing all functions from recording raw data to producing physics results.
- Wide participation from the GlueX Collaboration in many and varied areas of activity.
- Plan to add ability to do reconstruction on the OSG
- Areas of improvement have been identified.
- Would benefit from more human resources deployed at the interface of data analysis and software engineering.