Hall D Software Overview

12 GeV Software Review

Mark M. Ito

Jefferson Lab

June 7, 2012

Outline

1. Basic Components: Processors/Data-Formats/Databases/Management
2. Staffing
3. Computing Resource Requirements
4. Other Topics
    - Grid
    - Calibration and Alignment
    - Amplitude Analysis and GPU's
    - Data Challenges
5. Conclusions

Flow Diagram

- boxes: processors
- ovals: data formats
- punch cards: databases

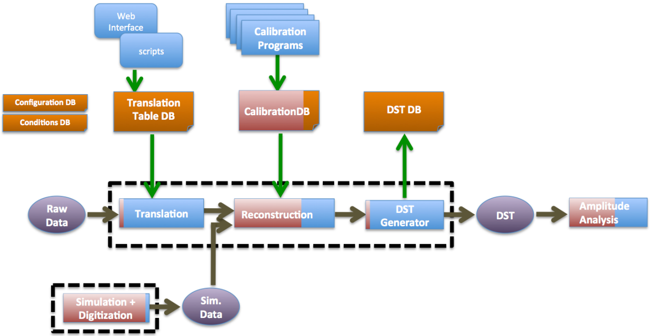

Simulation/Digitization

- GEANT3-based (mature)
  - Geometry defined in configuration files, not source code
    - Hall D Detector Specification (HDDS), an XML implementation
  - Electromagnetic background included (as accidentals)
  - Resolution/digitization introduced in separate process (mcsmear)
- Geant4 Conversion
  - Started as a background task
  - Geometry from same HDDS files as for GEANT3 implementation
  - Hit generation: no new algorithms, re-use same core C code as before

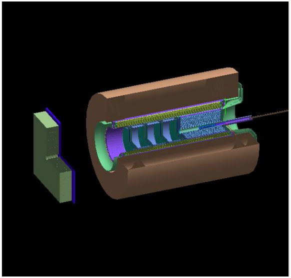

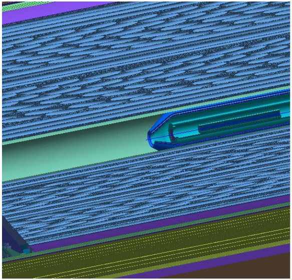

Reconstruction

- Based on JANA framework (mature)
  - C++
  - multi-threaded
  - factory model
- Provides full feature set
- Factories attached to framework current area of development
  - *e. g.*, cluster finding, track fitting
  - Direction driven by simulation based analysis

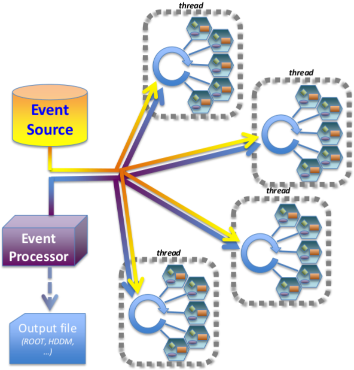

Efficiency, Resolution, and Background

- Reconstruction code has improved significantly over the past year
- Remains major area of effort right now
- Track reconstruction in a non-uniform magnetic field
  - curling tracks
  - areas of reduced efficiency/resolution
  - reconstruction speed
- Photon reconstruction
  - split-offs
  - merged clusters
  - effort to lower thresholds
  - hadronic contamination

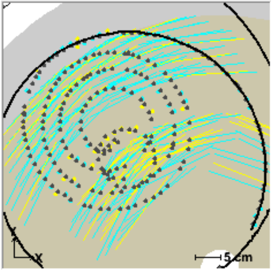

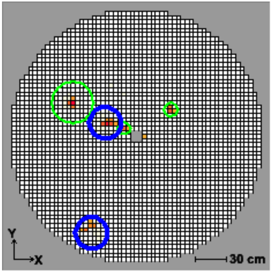

Flow Diagram

- boxes: processors
- ovals: data formats
- punch cards: databases

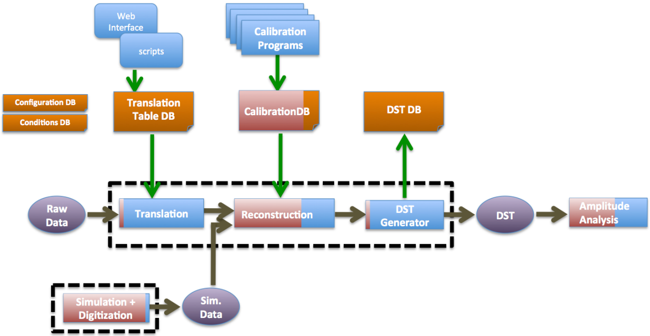

Serialized data formats

- Raw data
  - EVIO: CEBAF Online Data Acquisition (CODA) format
- Simulated data
  - HDDM: Hall D Data Model, self-documenting on two levels
    1. data itself XML-like (compressed)
    2. each file contains complete mini-schema
  - support utilities for conversion
- DST data (reconstructed)
  - HDDM
  - others possible

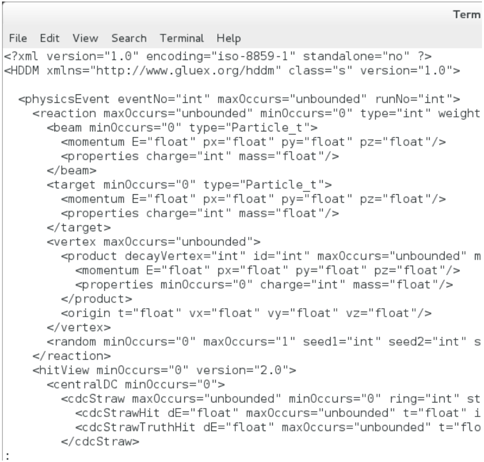

Flow Diagram

- boxes: processors
- ovals: data formats
- punch cards: databases

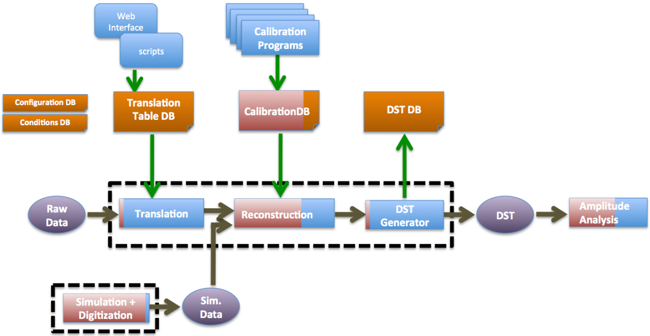

Databases

- Calibration and Conditions Database (CCDB)
  - developed by Hall D
  - in beta testing (Hall B is using it too)
  - put into production mid-July, 2012
- Other Database Applications
  - Online Conditions/Parameters
  - Translation Tables
  - DST

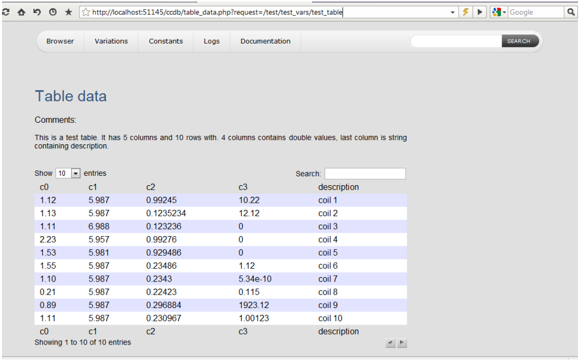

c0	c1	c2	c3	description
1.12	5.987	0.99245	10.22	coil 1
1.13	5.987	0.1235234	12.12	coil 2
1.11	6.988	0.123236	0	coil 3
2.23	5.957	0.99276	0	coil 4
1.53	5.981	0.929486	0	coil 5
1.55	5.987	0.23486	1.12	coil 6
1.10	5.987	0.2343	5.34e-10	coil 7
0.21	5.987	0.22423	0.115	coil 8
0.89	5.987	0.296884	1923.12	coil 9
1.11	5.987	0.230967	1.00123	coil 10
c0	c1	c2	c3	description

- CCDB Features:
  - relational database
  - standard look up by run
  - hierarchical calibration type structure
  - history kept, previous versions selectable at run time
  - branches/private/tagged versions supported

Software Management Tools

- Source code version control: Subversion
- Regular tagged releases of simulation and reconstruction software ("sim-recon"): about every 6 weeks
- Nightly builds of sim-recon
  - all Lab-supported Linux flavors
  - Doxygen documentation generated
- Semi-weekly tests: histograms generated, archived
  1. single-charged particles events
  2. multi-track, multi-photon events (5π)

$$
\begin{aligned}
\gamma p \rightarrow & X p \\
& \hookrightarrow b_1^{\pm}\pi^{\mp} \\
& \quad \hookrightarrow \omega\pi^{\pm} \\
& \quad\quad \hookrightarrow \pi^+\pi^-\pi^0 \\
& \quad\quad\quad \hookrightarrow \gamma\gamma
\end{aligned}
$$

- Mantis bug tracker for issue tracking

Staffing Requirements and Project Progress

	Budgeted Labor Units (MW)	FTE-year	% complet	Responsible Institution	Responsible Persons	fraction of project
GEANT 3 simulation	88	2.0	100%	UConn	Richard Jones	5.3%
GEANT 4 simulation	88	2.0	33%	UConn	Richard Jones	5.3%
DAQ to Detector Translation Tab	44	1.0	5%	JLab		2.7%
Reconstruction	495	11.3	66%			30.0%
Reconstruction Framework	44	1.1	81%	JLab	David Lawrence	
CDC Reconstruction	33	0.9	78%	JLab	David Lawrence	
FDC Reconstruction	33	1.1	73%	JLab	Simon Taylor	
Track Finding	66	2.0	75%	JLab/CMU	Simon Taylor/David Lawrence	
Track Fitting	66	3.0	67%	JLab/CMU	S. Taylor/D. Lawrence/P. Mattione	
BCal Reconstruction	44	1.0	50%	IU/Regina	Matt Shepherd/Zisis Papandreou	
FCal Reconstruction	33	0.8	75%	IU/UConn	Matt Shepherd/Richard Jones	
TOF Reconstruction	33	0.8	50%	FSU	Paul Eugenio	
Tagger Reconstruction	33	0.8	0%	UConn/CUA	Richard Jones	
Start Counter Reconstruction	22	0.5	50%	FIU	Simon Taylor/Werner Boeglin	
Particle ID	44	1.0	75%	CMU/JLab	Paul Mattione	
Kinematic Fitter	44	1.0	95%	MIT/CMU	Mike Williams	
Calibration	242	5.5	23%			14.7%
Calibration Database	33	0.8	80%	MEPHI/JLab	Dmitry Romanov	
CDC Calibration	33	0.8	5%	CMU	Naomi Jarvis	

- Complete task list with labor estimates and fraction completed
- Reconstruction 66% complete, calibration 23% complete
- Total effort: 38 FTE-years, complete: 51%, remaining work: 19 FTE-years

Staffing Resources

- Survey of estimated staffing resources available for software infrastructure
- Different categories weighted differently
- Total for 2012-2014: 23 FTE-years
- Good match with requirements, but little safety factor

Institution	Names/Categories	2012	2013	2014	3-year Totals
Arizona State University	got email from Barry, 2/3				
	grad student	0.00	0.50	0.50	
	totals	0.00	0.50	0.50	1.00
University of Athens	- sent an additional email to Christina, 2/2				
Carnegie Mellon University	- got email from Curtis, 1/24				
	Curtis Meyer	0.10	0.10	0.10	
	Paul Matione	0.75	0.75	0.75	
	Will Levine/student-to-be-named	0.56	0.75	0.75	
	totals	1.41	1.60	1.60	4.61
Catholic University	- got email from Franz, 1/24				
	Franz Klein	0.12	0.25	0.25	
	grad student	0.10	0.10	0.10	
	totals	0.23	0.35	0.35	0.92

⋮

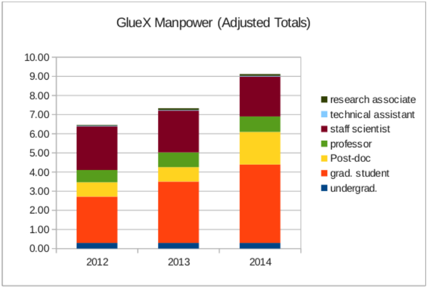

Data Analysis Model

- From detector to online buffer disk in Counting House
- Over network to Tape Library in Computer Center
- Reconstruction on JLab batch farm, reduction of volume by factor of 5
- Resulting DST data written to Tape Library
- Mini-DST's ("4-vectors") in several streams
- Similiar scenario for simulation except "raw" events not kept
- Analysis engines access this data:
  - JLab batch farm
  - Individual work stations
  - Collaborating institutions
  - Amplitude analysis resources

Data/Computing Model

Trigger rates:

- Phase III: beam rate 10^7 γ/s in coherent peak
- Full hadronic cross section $\Rightarrow$ 20 kHz trigger rate
- Beyond Phase III: use level 3 software trigger to keep rate roughly constant
- Start with "monitoring farm", upgrade to "trigger farm"

Assumptions (generic):

- 20 kHz off detector
- 15 kB events
- run 35 weeks year, 50% running efficiency
- 133 ms to reconstruct an event (measured)
- 2 Monte Carlo events per data event (on average)
- 67 ms to generate Monte Carlo events (reconstruction time comparable to data)
- factor for multiple iterations: 2
- Other loads:
  - calibration processing
  - skims/mini-DST production
  - physics analysis

CPU and Tape Requirements

Process	CPU (kCores)	Tape (PB/y)
Raw Data	–	3.2
Calibration	0.09	0.06
Reconstruction	1.8	1.3
Skims/mini-DST	0.9	0.6
Analysis	0.9	–
Simulation	5.4	2.5
Total	9	8

CPU represents amount of computing power required to keep up with average rate off detector. Sets the scale; should be viewed as a minimum requirement.

Another View of Requirements: Wait Time

How long do we wait for results?

Assume a 10 kCore farm, unloaded, nominal running efficiency, and one iteration only. Other assumptions the same as above.

	Phase I	**Phase II**	**Phase III**
Year of run	2014	2015	2016
Days of running	60	60	120
Trigger rate (kHz)	2	20	20
Number of events	5.18×10^{9}	5.18×10^{10}	1.04×10^{11}
Reconstruction time (days)	0.8	8	16
Simulation time (gen.+recon.) (days)	2.4	24	48
Recon.+Sim. time (days)	3.2	32	64
Total data to tape (PB)	0.2	1.8	3.6

Disk Requirements

Data Type	Phase II (TB)	Phase III (TB)
Calibration disk	62	124
Coherent-peak skim DST Reconstructed data, from coherent bremsstrahlung peak	25	50
Inclusive background simulation DST Simulation, minimum bias	265	531
Individual analysis skims Skims of data and simulation for individual analyses, 10 analyses	207	415
Mini-DST's for amplitude analysis 4-vectors	7	15
Total	567	1134

- Disk to support analysis activities, not large-scale reconstruction/simulation
- In addition: work/scratch/staging disk of 300 TB
- **Grand Total**: 2.0 PB

Open Science Grid (OSG)

- Plan to use Grid resources to augment those at JLab
  - Monte Carlo generation and reconstruction (no raw data transfer)
  - Opportunistic usage: possible to have quick turn-around on specific tasks
- Have established GlueX as a Virtual Organization (VO) with the OSG
- Nodes from UConn have been contributed
- Significant analysis has been performed using the grid (3π and 5π analyses)
- GlueX consumption has roughly matched contribution
- Grid tools have been installed at JLab (no plan for JLab to contribute nodes)
- Plan to use Grid Storage Resource Manager (SRM) to move data on and off JLab site

Contribution and Consumption

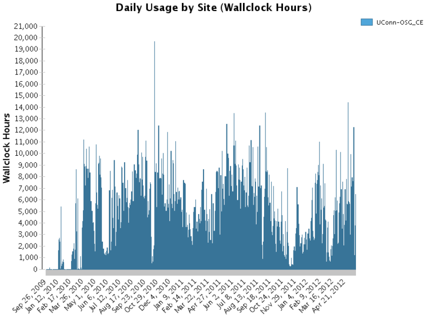

OSG usage at UConn (contributed)

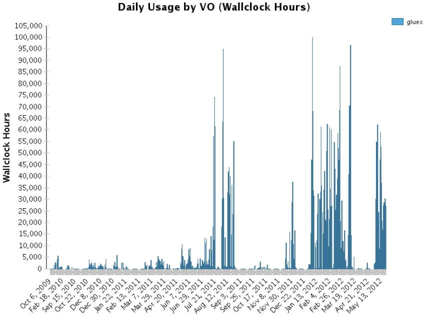

GlueX usage on the OSG (consumed)

Calibration/Alignment

- Developing plans for calibration
  - Each working group tasked to list methods, data requirements, compute requirements
  - Examples:
    - π^0 calibration of forward calorimeter
    - run plan for tracking chamber alignment
    - time walk corrections for time-of-flight using plane-to-plane information
- Hardware and software have reached the point where we can turn to calibration
- Goal: have most calibration systems available in a year
  - fits with detector installation time frame

Amplitude Analysis on GPU's

- Challenge of amplitude analysis largely met by GPU technology
- Calculation of log likelihoods for individual events
  - independent for each event
  - computationally expensive
  - not branch intensive
  - I/O to do computation is small
  - results are to be summed
- Technique starting to be used by many experiments
- One of leading implementations developed by GlueX collaborator: AmpTools
- Large reduction in scale of compute task (1-2 orders of magnitude)
- Proof-of-principle established (3π, 5π)
- Implement a user-friendly "production" system by Spring 2013

Data Challenges

- Simulate raw data to DST chain on a large scale
- Need to develop job management tools

Data Challenge Timeline

	Task	**Date**
1	deploy calibration database	2012-07-15
2	deploy translation database	2012-09-01
3	complete raw data format specification	2012-10-01
4	complete specification of reconstructed data format	2012-10-15
5	data challenge: simulated raw data to DST data (one day at 10^7 γ/s)	2012-12-01
6	complete mini-DST writer	2013-02-01
7	create mini-DST data samples from data challenge events	2013-03-01

Next step: one week's worth of data

Collaboration with IT Division's Scientific Computing Group

- Looking forward to continued collaborative efforts in JLab's 12 GeV era
- Physicists want to better understand the computing resources being provided
  - un-steepen the learning curve for new collaborators
  - increase efficiency in use of resources
- Possible areas of development:
  - Enhanced reporting of farm job priority/status/disposition
  - Central support of key scientific software packages (ROOT, CERNLIB, Geant4, CLHEP)

Summary/Conclusions

Accomplishments

- End-to-end solution for simulation and reconstruction in hand.
  - realistic event generators
  - hit-level reconstruction of photons and charged particles
  - major components mature: simulation, reconstruction framework
  - amplitude analysis capability demonstrated
- Software management systems in place
  - Tagged releases
  - Automated software builds
  - Automated reconstruction tests

Work Ahead

- Put calibration database into production (2012-07-01)
- Deploy translation table database (2012-09-01)
- Large-scale data challenge (2012-12-01)
- Calibration software systems in full development (2013-06-01)
- Production GPU system (2013-06-20)
- Reconstruction quality at near-publication level (2014-03-31)

Backup Slides

$Id: software_overview_halld_MMI.tex 9207 2012-06-06 17:56:31Z marki $