Hall D and IT

at Internal Review of IT in the 12 GeV Era

Mark M. Ito

Hall D

May 20, 2011

Hall D in a Nutshell

- search for exotic mesons in the 1.5 to 2.0 GeV region
- 12 GeV electron beam
  - coherent bremsstrahlung photon beam
  - coherent peak at 9 GeV
  - photon tagger
- 4π detector: GlueX
  - charged tracking
  - calorimetry
  - particle ID (time of flight)
- amplitude analysis (a.k.a. partial wave analysis (PWA)) necessary

Topics: Aspects of Computing

- Requirements
- Status of the Software
- Planning, Tests, System Engineering

Requirements

- attempt to capture all components of offline computation resources
  - Calibration
  - Reconstruction
  - Skimming
  - Analysis
  - Simulation
- will not cover online computing
  - data acquisition software
  - software trigger
- not all plans are fully formed
- welcome comments on holes in planning

Raw Data

- strategy: take entire hadronic cross-section through level 1 trigger
- start at photon intensity of 1×10^7/s in coherent peak
- ramp to 1×10^8/s after "two" years
- also ramp software trigger to give factor of 10 rejection
- net effect: constant event rate to tape
- 20 kHz when running
- study: event size = 15 kB, instantaneous data rate = 300 MB/s
- 35 weeks of running a year, 50% running efficiency
- average data rate: 200 billion events/year or 3.2 PB/year

Calibration

- assume detectors can be calibrated using 5% of the raw data
  - gross simplification
  - for estimating purposes
- assume that calibrations will have to be done twice

Reconstruction

- turning detector hits into particles
- bulk of offline processing
- 133 ms per event
- assume output data 1/5 size of input data
- assume that it needs to be done twice
- do this at JLab
  - avoids shipping around raw data

Streaming

- dividing reconstructed events into separate "streams" for specific analyses
- based on event topology
- assume CPU load 1/10 that of reconstruction
- assume that number of events output 1/10 of that input
- assume five streams need to be produced
- must be done twice (like reconstruction)

Analysis

- extraction of physics signals from reconstructed events
- assume that CPU required is 1/10 of that needed to reconstruct (including repetitions factors)
- multiple physics analysis, assume 10 of them
- assume each analysis needs 20 TB of data on disk
- statistical results, negligible storage requirements for output
- PWA
  - highly CPU intensive
  - suitable for off-site (modest data requirements)
  - may require special resources (branch un-intensive: GPU farms)
  - not included in this estimate

Simulation

- create simulated data
- assume reconstruction time same as that for raw data
- generation time half of that to reconstruct
- number events needed assumed to be twice raw data
  - want statistical error to be small compared to that of data, factor of 10 more
  - more selective generation, factor of 5 less
- assume that it needs to be done twice
- use of off-site resources is an option
  - collaborating institutions have farms
  - grid paradigm established (OSG)
  - other paradigms possible
- simulation studies outside this estimate
  - ideas often generated during analysis (creativity in science!)
  - large demand for resources with desire for quick turn around
  - grid-type environment well-suited for these cases

Summary of Requirements

Process	CPU (kCores[a])	Disk (TB)	Tape (PB/y)
Raw Data	–	–	3.2
Calibration	0.09	–	0.06
Reconstruction	1.8	–	1.3[b]
Streaming	0.9	–	0.6
Analysis	0.9	200	–
Simulation	5.4[c]	–	2.5[b]
Total	9	200	8

[a] single thread on a 2.8 GHz Nehalem machine

[b] roughly half may be able to be recycled

[c] significant amount may be done off-site

Status of the Software

- Geometry
- Simulation
- Reconstruction
- Partial Wave Analysis
- Calibration Database
- Data Format
- Utilities

Geometry

- implemented in XML
- Hall D Detector Specification (HDDS)
  - XML elements and attributes closely follow GEANT defined shapes and their parameters
  - mature
- goal: keep the geometry in one place, use in
  - simulation (fully implemented)
  - reconstruction (partially implemented)
  - event display (to be implemented)

Simulation

- GEANT3-based: HDGEANT
  - geometry information auto-coded into FORTRAN code from HDDS information
  - hits (*i. e.*, digitization) coded separately
  - output in Hall Data Description Model format (HDDM, see slide below)
  - mature
- experimental resolution added in separate stage: mcsmear
  - HDDM in, HDDM out
  - in use
  - development continues
- effort started to transition to GEANT4

Reconstruction

- JANA
  - multi-threaded: each thread a separate event stream
  - algorithms for different detectors implemented as "factories"
- ROOT used for some general utilities
- hooks for user code
  - user's class inherits from abstract base class
  - must be registered with the framework
  - multiple user classes possible
- plug-in mechanism
  - e. g., define user class at run time
- mature

Partial Wave Analysis (PWA)

- achieving physics goals of GlueX depends critically on PWA.
- "Collaborative Research: Open Access Amplitude Analysis on a Grid"
  - NSF-funded effort
  - Carnegie Mellon, Indiana, Connecticut
- AmpTools
  - PWA toolkit
  - Indiana University
  - GPU-based implementation
- Ruby-PWA
  - PWA toolkit
  - Carnegie Mellon University
- plan to use off-site resources

Calibration Database

- relational database
- based on CLAS experience (Hall B, JLab) with improvements
- complete version history, with version choice at API level
- facility for private versions, with history
- tagging facility
- code base exists
- alpha testing next

Data Format

- raw data: EVIO
  - native CODA format
  - mature
- simulation output: HDDM
  - Hall Data Description Model
  - a compressed XML
  - retains schema-like template at beginning of each file (uncompressed)
  - C-based API, mature
  - C++ API, in testing
- reconstruction output
  - options:
    - HDDM
    - EVIO
    - ROOT trees
  - need to finalize plans

Utilities

- XML parsing: Xerces
- source code management: subversion
- source code documentation: doxygen
- building scripts: GNU Make
- database: MySQL
- general documentation
  - GlueX Notes: DocDB
  - webpages: mediawiki

Planning, Tests, System Engineering

- Project Management
- Simulation and Reconstruction Testing
- Code Review–Repository Control
- End-To-End Offline Test
- Documentation
- Communications
- Off-Site Computing Facilitation

Project Management

- history
  - part of BIA
  - captured in formal PM system
- leverage this effort going forward

Labor and Costs: offline_BL10_04_CR

WBS	Task	Labor (FTE-wks)	Sci'ist (%)	Contrib (%)	Cost (USD)
5533010	MC Studies for Detector Optimization	56.5	50.00	50.00	74930.30
5533015	Reconstruction Framework	18.83	100.00	0.00	49944.69
5533020	Integration of Slow Controls	33	100.00	0.00	87529.20
5533025	DAQ to Detector Translation Table	44	9.09	90.91	10609.60
5533030	Micro DST Writer	22	9.09	90.91	5304.80
5533045	Track Finding	54.98	100.00	0.00	145828.95
5533050	Track Fitting	54.98	100.00	0.00	145828.95
5533055	BCal Reconstruction	44	9.09	90.91	10609.60
5533060	FCal Reconstruction	33	9.09	90.91	7957.20
5533070	TOF Reconstruction	33	9.09	90.91	7957.20
5533080	Tagger Reconstruction	33	9.09	90.91	7957.20
5533085	Start Counter Reconstruction	22	9.09	90.91	5304.80
5533090	Particle ID	44	9.09	90.91	10609.60
5533095	Kinematic Fitter	44	9.09	90.91	10609.60
5533100	Integration/QC	44	100.00	0.00	116705.60
5533105	Calibration Database	33	66.67	33.33	58352.80
5533110	CDC Calibration	33	9.09	90.91	7957.20
5533115	FDC Calibration	33	9.09	90.91	7957.20
5533120	BCal Calibration	33	9.09	90.91	7957.20
5533125	FCal Calibration	33	9.09	90.91	7957.20
5533140	Tagger Calibration	33	9.09	90.91	7957.20

Simulation and Reconstruction Testing

- collaborators exercising code and generating feedback (now and forever)
- analysis of simulated data underway
- systematic reconstruction integrity
  - traditional approach: generate standard histograms
    - weekly simulation/reconstruction suite running in cron job
    - only exotic meson channel simulated
  - to add: tests of individual software components
    - pinpoint problem areas
    - major effort: coding the tests, generating appropriate test vectors

Code Review–Repository Control

- problem area
- worry that bad code gets checked in
- worry that restrictions inhibit productivity/creativity

End-To-End Offline Test

- start in ET ring with raw data
- calibration, reconstruction, analysis
- need raw data format from DAQ group
- alternately, use HDDM surrogate for raw data
- resource-use-system development: need software just to use resources
- incremental development of test
- in conceptual stages

Documentation

- major challenge
  - no one likes to write it
  - critical to have it
- outline of formal system exists
- on-going discussion

Communications

- first-rate video conference capability
  - Hall D already a heavy user
  - reflects investment of collaborating institutions
- remote viewing/inspection of experimental features
  - online/DAQ plots
  - reading voltages
  - *et cetera*
- actual control/change of parameters likely done by people physically in the counting room
- locations at JLab:
  - counting house
  - non-accelerator site locations

Off-Site Computing Facilitation

- need a way to ship data to and from JLab
- will require resources
- robust storage resource manager (SRM) is planned
- grid utilities already installed at JLab

Summary and Conclusions

- simulation and reconstruction is done routinely
- positioned to exploit multi-core technology
- computing model developed; will be iterated
- many major components in place, others identified

...but...

- reconstruction still needs refinement
- calibration not seriously addressed
- adequate manpower a challenge