Report on CLAS12 Software Workshop*

- Goals:
  - Broad view of the state-of-the-art in offline analysis.
  - Status of the CLAS12 software program.
  - Training for new users/developers.
- Held at the University of Richmond, May 25-26, 2010.
- Travel funds provided for students and post-docs.
- Approximately 50 people attended the plenary session on May 25.
- Twenty-five people (plus instructors) attended the tutorials on May 26 which ran until 1 pm.

*Supported by the JSA/SURA Initiatives Fund.

Website: `http://conferences.jlab.org/CLAS12Software/index.html`

Program for CLAS12 Software Workshop

Tuesday, May 25, 2010 - Gottwald Center for the Sciences Auditorium (Chair: David Heddle, CNU)		
9:00-9:20	Welcome and Summary of the 12 GeV Upgrade	Jerry Gilfoyle (Richmond)
9:20-10:10	Big Science in an Era of Large Datasets - BNL Experience and Perspective in Nuclear and Particle Physics Data Analysis	Michael Ernst (BNL)
10:10-10:50	Scientific Computing at Jefferson Lab: Petabytes, Petaflops, and GPUs	Chip Watson (JLab)
10:50-11:10	**Break**	
11:10-12:00	The Offline Analysis Framework at CDF	Elena Gerchtein (FNAL)
12:00-12:30	JLab Data Analysis Coordination, Planning and Funding	Graham Heyes (JLab)
12:30-2:00	**Lunch**	
2:00-2:20	Overview of the CLAS12 Software Enterprise	Dennis Weygand (JLab)
2:20-2:50	CLARA: Service oriented architecture based PDP application development framework	Vardan Gyurjyan (JLab)
2:50-3:10	The JANA Reconstruction Framework	Dave Lawrence (JLab)
3:10-3:30	Event Reconstruction - The Big Picture	Mac Mestayer (JLab)
3:30-3:50	Socrat - CLAS12 Electron Reconstruction	Sebastien Procureur (Saclay)
3:50-4:10	**Break**	
4:10-4:30	Reconstructing CLAS12 events using JANA	Maurizio Ungaro (UConn
4:30-4:45	gemc - A Modern Simulation for CLAS12	Maurizio Ungaro (UConn
4:45-5:00	Simulation Results for CLAS12 From gemc	Jerry Gilfoyle (Richmond)
5:00-5:20	CED12 - Seeing Tracks Through Thick and Thin	Dave Heddle (CNU)
5:20-5:40	Code Sharing and the EVIO Package	Elliott Wolin (Jlab)
5:40-6:00	Data Storage Format for CLAS using HDF5	Gagik Gavalian (ODU)
6:00-8:00	Reception in the Gottwald Atrium	

Program for CLAS12 Software Workshop

Tuesday, May 25, 2010 - Gottwald Center for the Sciences Auditorium (Chair: David Heddle, CNU)		
9:00-9:20	Welcome and Summary of the 12 GeV Upgrade	Jerry Gilfoyle (Richmond)
9:20-10:10	Big Science in an Era of Large Datasets - BNL Experience and Perspective in Nuclear and Particle Physics Data Analysis	Michael Ernst (BNL)
10:10-10:50	Scientific Computing at Jefferson Lab: Petabytes, Petaflops, and GPUs	Chip Watson (JLab)
10:50-11:10	**Break**	
11:10-12:00	The Offline Analysis Framework at CDF	Elena Gerchtein (FNAL)
12:00-12:30	JLab Data Analysis Coordination, Planning and Funding	Graham Heyes (JLab)
12:30-2:00	**Lunch**	
2:00-2:20	Overview of the CLAS12 Software Enterprise	Dennis Weygand (JLab)
2:20-2:50	CLARA: Service oriented architecture based PDP application development framework	Vardan Gyurjyan (JLab)
2:50-3:10	The JANA Reconstruction Framework	Dave Lawrence (JLab)
3:10-3:30	Event Reconstruction - The Big Picture	Mac Mestayer (JLab)
3:30-3:50	Socrat - CLAS12 Electron Reconstruction	Sebastien Procureur (Saclay)
3:50-4:10	**Break**	
4:10-4:30	Reconstructing CLAS12 e	
4:30-4:45	gemc - A Modern Simulatic	
4:45-5:00	Simulation Results for CLA	
5:00-5:20	CED12 - Seeing Tracks T	
5:20-5:40	Code Sharing and the EVI	
5:40-6:00	Data Storage Format for C	
6:00-8:00	Reception in the Gottwald	

Wednesday, May 26, 2010, Room D-115, Gottwald Center for the Sciences		
9:00-12:00	Morning (lab session) Tutorials on CLAS12 Software.	
	1. Building and using gemc for CLAS12 simulation.	Jerry Gilfoyle
	2. Building and using Socrat/JSocrat for event reconstruction.	Maurizio Ungaro Dennis Weygand
	3. Running CLARA.	Vardan Gyurjyan
	4. Running CED12.	Dave Heddle

Big Science in an Era of Large Datasets

Michael Ernst, Director, RHIC and Atlas Computing Facility

- Relativistic Heavy-Ion Collider (RHIC).
  - Formation of quark-gluon plasma (QCD at high temperatures).
  - Spin program using spin-polarized protons.
  - Four experimental areas.

- Large Hadron Collider (LHC).
  - Will collide protons at 7 TeV or lead nuclei at 574 TeV to search for the Higgs, SUSY particles, ...
  - Seven experiments.

- Data rates and sets.

Big Science in an Era of Large Datasets

Michael Ernst, Director, RHIC and Atlas Computing Facility

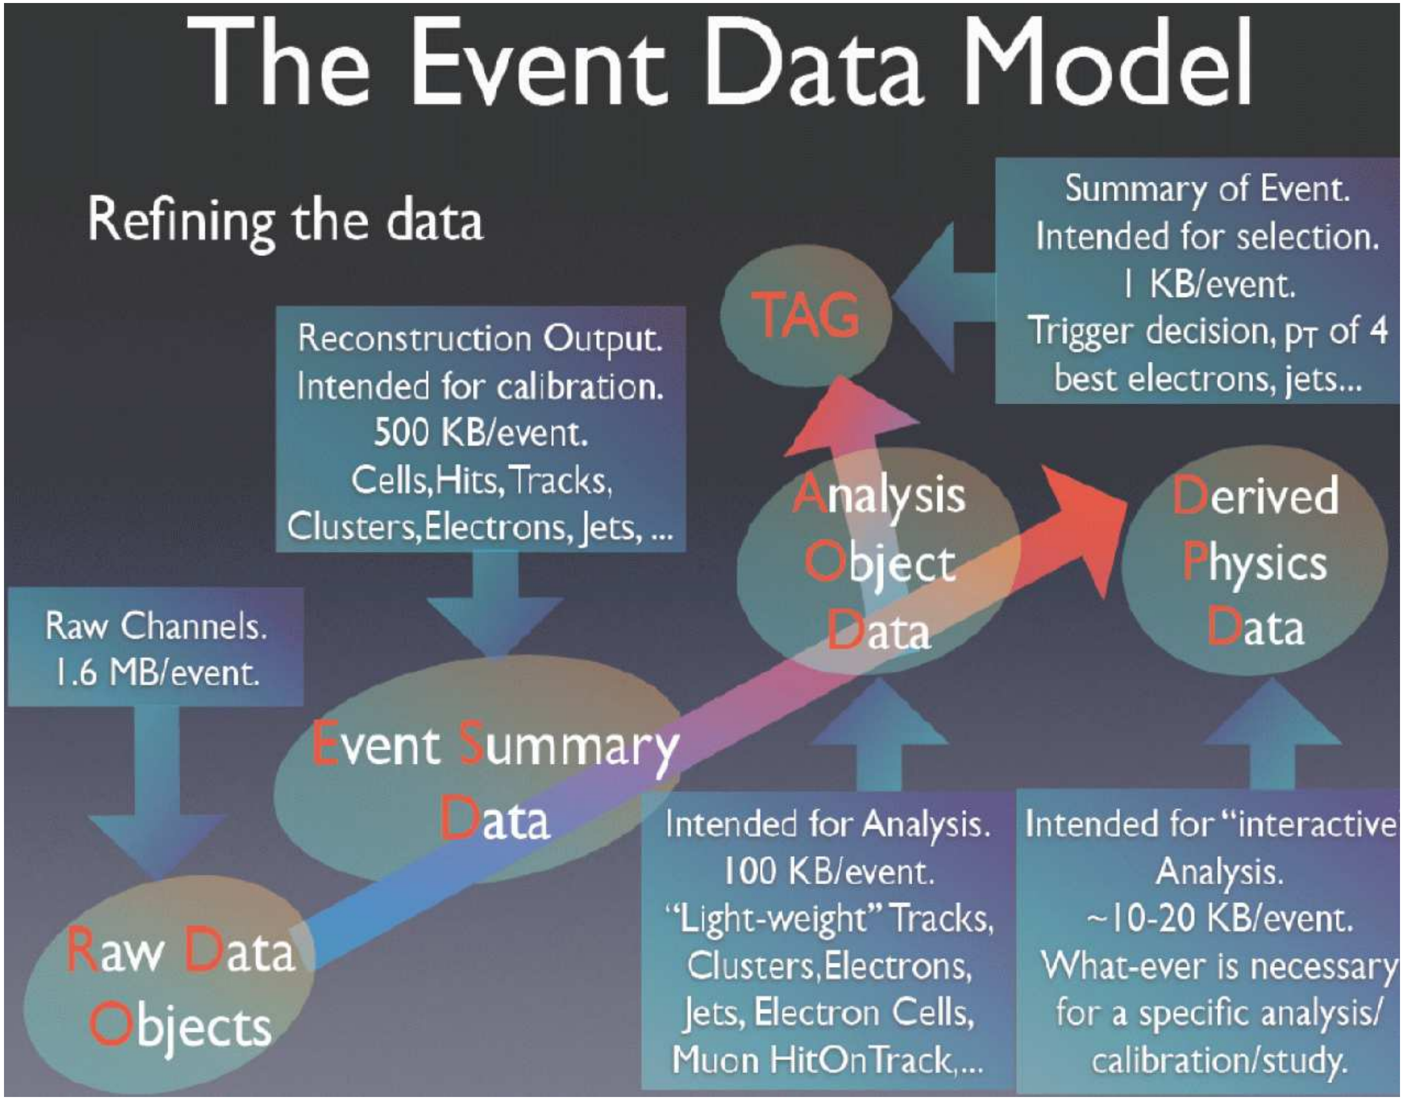

Data Formats

- Raw data (~ 200 Hz; 1.6 MB/ev) are written out in $4-7$ (RAW) physics streams based on trigger signature
- At the Tier-0 (CERN) reconstruction software runs on the raw data streams and produces
  - Event Summary Data (ESD) of about 500 kB/ev (target size – currently 50% more)
    - ESDs allow re-running reconstruction, particle identification, jet-finding, track re-fitting etc.
  - Analysis Object Data (AOD) of about 100 kB/ev (target size – currently 70% more)
    - AODs are made from ESDs and allow most common analyses
  - Derived Physics Data (DPD) of about 10 kB/(AODev) (assuming 10 largely disjunct primary DPD's with 10% skimming/thinning efficiency)
    - DPDs are filtered AODs with reduced information and optionally added analysis data
  - Tag Data (TAG) of about 1 kB/ev
    - TAGs are event-level metadata allowing fast selection of events with certain signatures

Software Framework:

- 'Simpler is better!' - Modular, common tools, validation.
- athena
  - Based on LHCb framework for reconstruction and processing.
  - Does RAW $\rightarrow$ ESD/AOD/DPD.
  - Based on C^{++} and Python.
  - Distributed as shared libraries.
  - Linux only now; Mac later.
- Additional frameworks:
  - AthenaROOTAccess: athena classes available to ROOT; provides access to RAW data.
  - Several other frameworks built on athena/ROOT/Python.

- Services: handle infrastructure.
- Algorithms: invoked by scheduling service in event loop.
- Tools: invoked by Algorithms.

Big Science in an Era of Large Datasets

Michael Ernst, Director, RHIC and Atlas Computing Facility

- Parallel processing.
  - Done with athenaMP.
  - Single events distributed to single cores.
  - Memory is shared among athenaMP workers on each core.

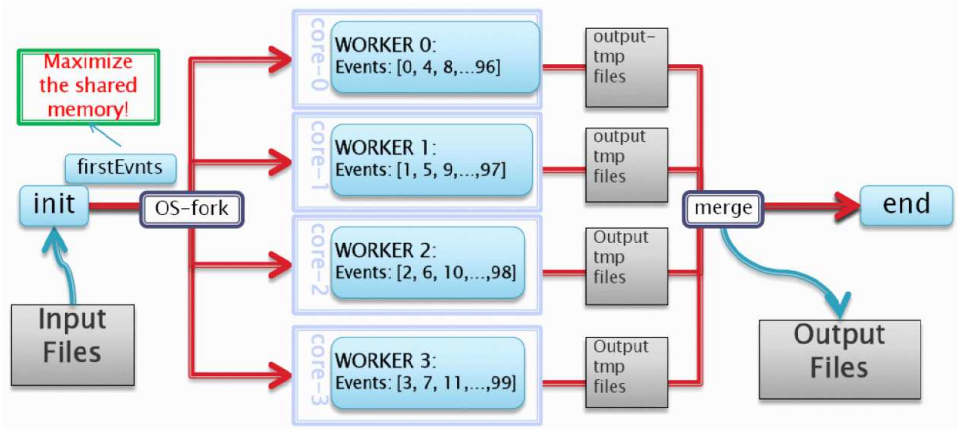

- The Grid.
  - PanDA - used by Atlas since 8/05; in ATLAS-wide production use since early 2008.
  - Rapid increase in recent use.

User counts	1mo	3mo	6mo
Feb-09	90	352	473
Mar-10	580	780	965

Scientific Computing Resources

- 1400 compute servers; 200 in analysis farm; multi-threading, multi-GPU.
- 50% increase in available compute servers for summer ‘bump’ (loaners for LQCD farm).
- Future upgrades driven by compute requirements.
- Annual budget $\approx \$50k$ each for compute servers, disk space, and tape storage.
- 64-bit part of farm is not fully utilized!

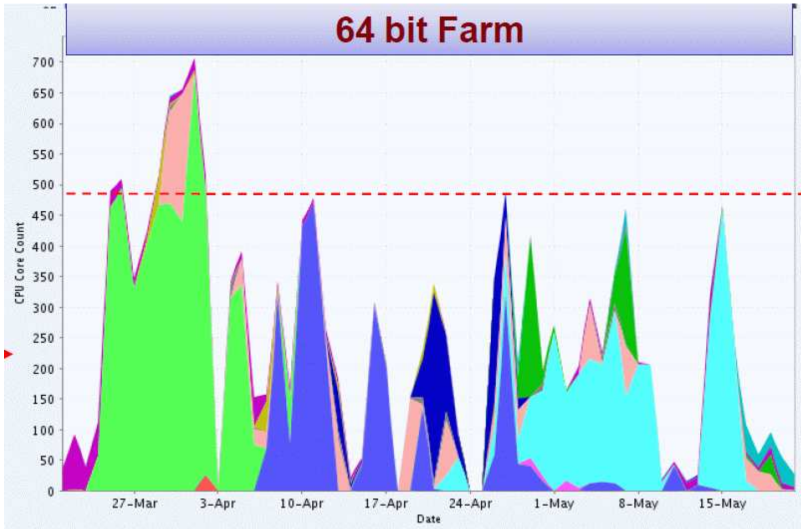

Computing Trends

- Computing power continues to grow; clock speeds no longer improving (chips got to hot) so number of cores is increasing.
- GPUs to the rescue!!??
  - Graphical Processing Units can have hundreds of cores, high memory bandwidth $\rightarrow$ large increase in speed.
  - Improve $\$/MFlops$ by a factor of 10.
  - Only works for the right problem:
    1. high data parallelism.
    2. high flop count on a small kernel.
  - Requires lots of new, clever programming (manual optimization, ...).
  - 530 GPUs at JLab by July.

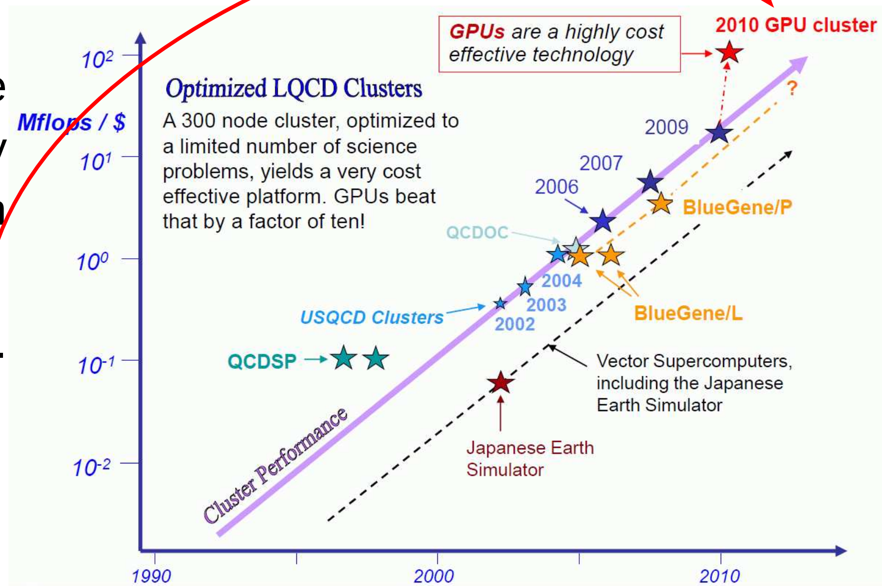

Disruptive technology!

12-GeV Computing Model

Assumptions

- Moore's Law will hold for another 4 years (in core counts)
- Your software will become multi-threaded to use the cores
- A 10x growth in computing power is free, and matches the expected growth from the 6 GeV era to the 12 GeV era
- While computing is (approximately) free, people are not (currently 80% of farm related budget goes into people)
- The 12 GeV operations budget will allow an additional doubling in compute capacity (hardware) with no discussion (i.e. upon demand); this is only a 20% budget impact
- All raw, reconstructed, and simulation data will be held at Jlab, with redundant copies for all raw data

12-GeV Computing Model

Assumptions

Issues

- Tape Storage:
  - 6 GeV - $80k/year
  - 12 GeV - $300k/year.
- Simulations: Generate-On-Demand versus Store-and-Retrieve.
- Parallel computing: Grid versus Central Computing
  - Grid requires a person.
  - 80% of farm budget is people.
  - More cost effective to put money into compute servers, ...
  - Sensitive to physics computing requirements.

The Offline Analysis Framework at CDF

Elena Gerchtein - Chief of Offline Operations, CDF, FermiLab

- Collider Detector at FermiLab.
  - For proton-proton collisions up to center-of-mass energy of 2 TeV.
  - Built around large solenoid magnet.
  - Discovered the top quark.

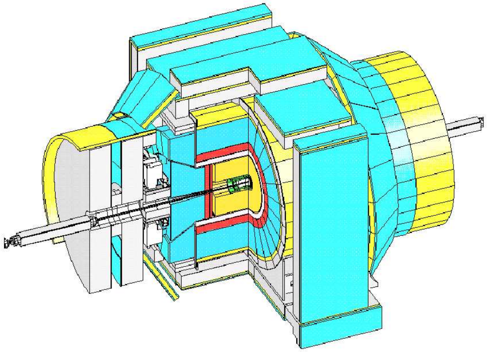

- CDF data flow.
- Computing resources - GRID enabled farms: CDFGrid (5K), NamGRID (300 MIT + 150 KISTI)

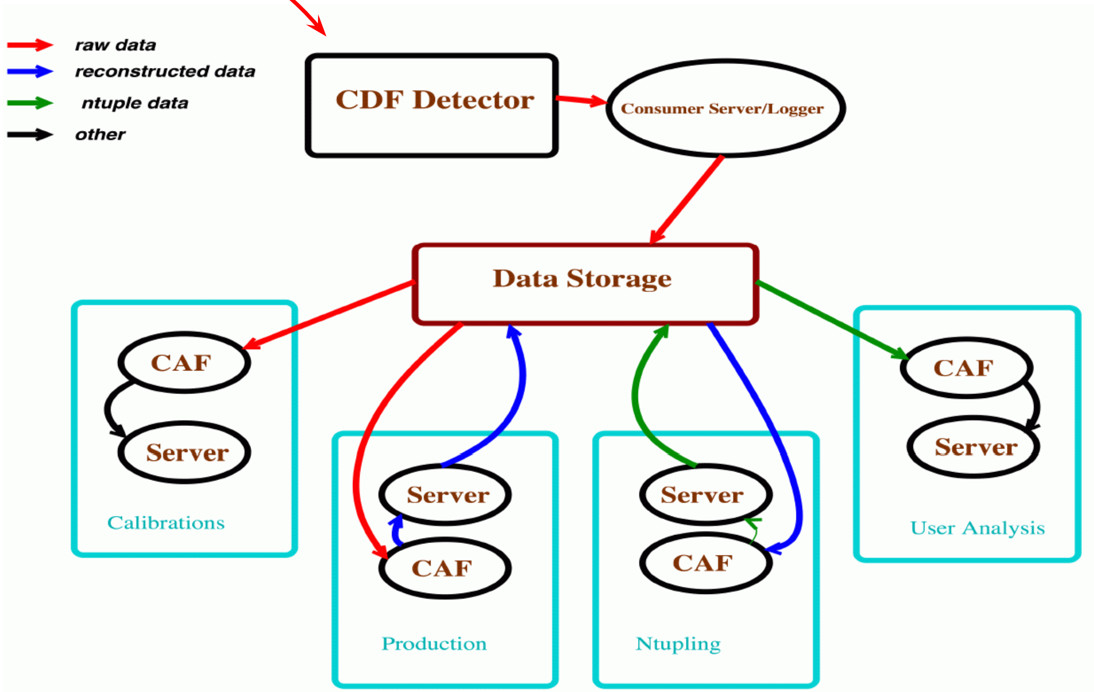

The Offline Analysis Framework at CDF

Elena Gerchtein - Chief of Offline Operations, CDF, FermiLab

Software Framework - AC^{++}

- Does it all: reconstruction, post-reconstruction analysis, online triggering, calibration, and monitoring.
- Modular.
- Configured with tcl.
- Based on C^{++}.

Offline Processing

- Releases: bimonthly during development; yearly during operations.
- Calibrations: $\approx 1/5$ of data reconstructed for calibrations; results stored in database.
- Reconstruction: one executable *ProductionExe*, only higher level objects kept.
- Ntupling: AC^{++} processes reconstructed data to make ROOT ntuples.

Pro	- dataset size reduced 70%. - simplified format. - reconstruction corrections done as needed.	Con	- different physics groups run the same algorithms to produce different datasets: data redundancy and excessive CPU usage.

CDF software: lessons learned

- **Code:**
  - **Need to work harder to enforce better coding practice**
  - **critical design pieces should be reviewed by experts**
  - develop and maintain validation plan/code in parallel with major algorithms
  - Plan for maintenance: documentation, transitions, training
- **Algorithms**
  - **have speed/size target for each algorithm, subject to review and negotiation**
- **Operations**
  - **offline software shifts (run validation, checking reported software errors, do bug fixes) were helpful**

Data Storage: lessons learned

- need better plan for resource allocation/recycling.
- tapes may be need recycling in the nearest future.
- disk pool for temporary output is 50% junk data but no tools/policy on how to clean up.
- jobs waiting for files from tape and wasting CPU.

Simulation

- Simulation Framework (SF) is integrated into AC++ application framework.
- Mixture of generic programming and OOP $\longrightarrow$ SF is easily extensible and time efficient.
- CDF uses the same geometry for reconstruction, simulation, and visualization.
- The same data objects are used for simulation and real data reconstruction.
- Software provided to submit large scale MC jobs to GRIDs according to complicated run dependent job plan.
- Generators: Need to be able to plug in ANY event generator with minimum efforts by demand; read arbitrary event generator output and pass it through detector simulation.
  - Eleven different generators (see talk on webpage).
  - Les Houches Accords: universal interface to MC. Implemented in PYTHIA and GRAPPA
- Tracking of particles through matter is based on GEANT3 tracking.
- Fast parametric simulation available for some sub-detectors.
- Transparent to user: Configure with tcl scripts and submit to Grid.

CDF Offline: Lessons Learned

- Plan for scalability more then you currently expect.
- CPU gets faster while loads get bigger (more jobs are submitted).
- Systems run longer then expected.
- Datasets get bigger.
- Scalability in DB access, job control, data access.
- Frequent system downtimes unexpected, not all services are easily restored (major downtimes for security upgrades every 60 days).
- Performance improvement over time.

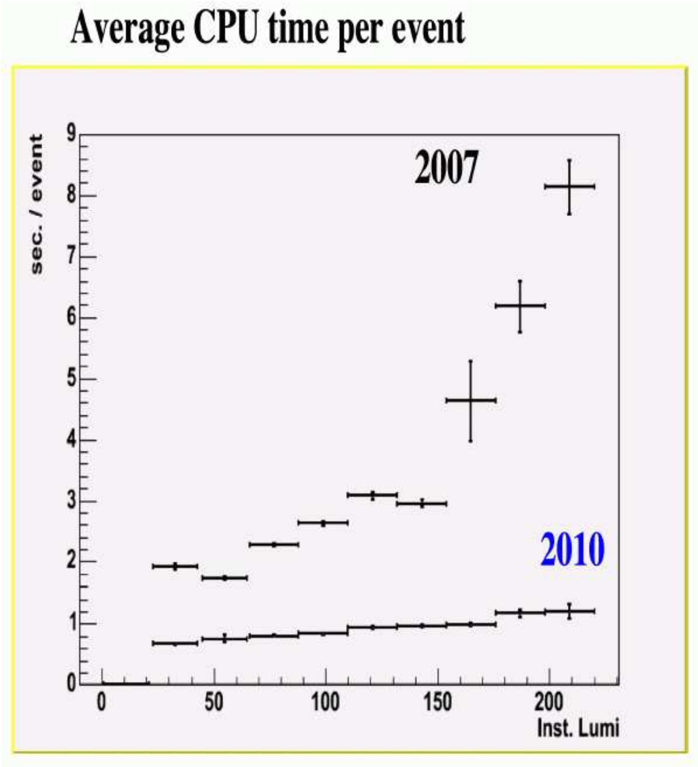

Data Analysis Coordination, Planning and Funding

Graham Heyes - Data Acquisition and Analysis group, Physics division, JLab

- Discussed details of organization, budgeting, and procedures for obtaining funding (see the website).
- Comparison of proposed CLAS12 requirements and current farm.

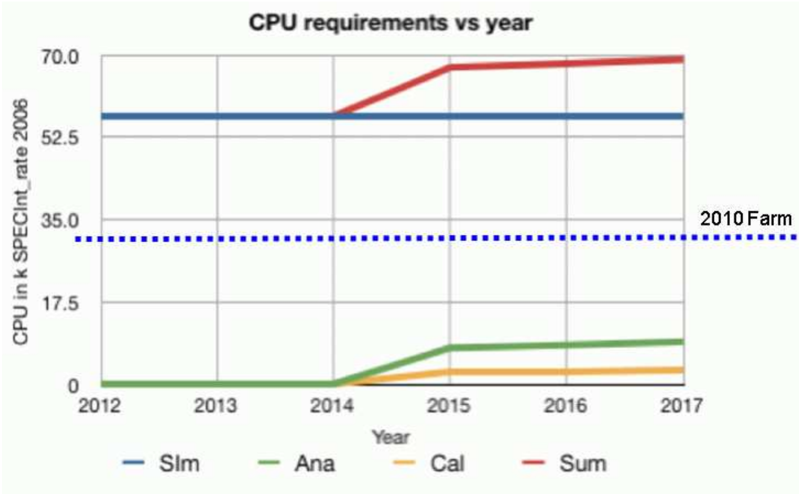

- The CPU requirements until 2015 can be met by the current budget of \$50k per year for nodes. A bump to \$200k is needed in 2015 for GLUEX calibration and analysis.
- Disk budget appears to be in good shape.

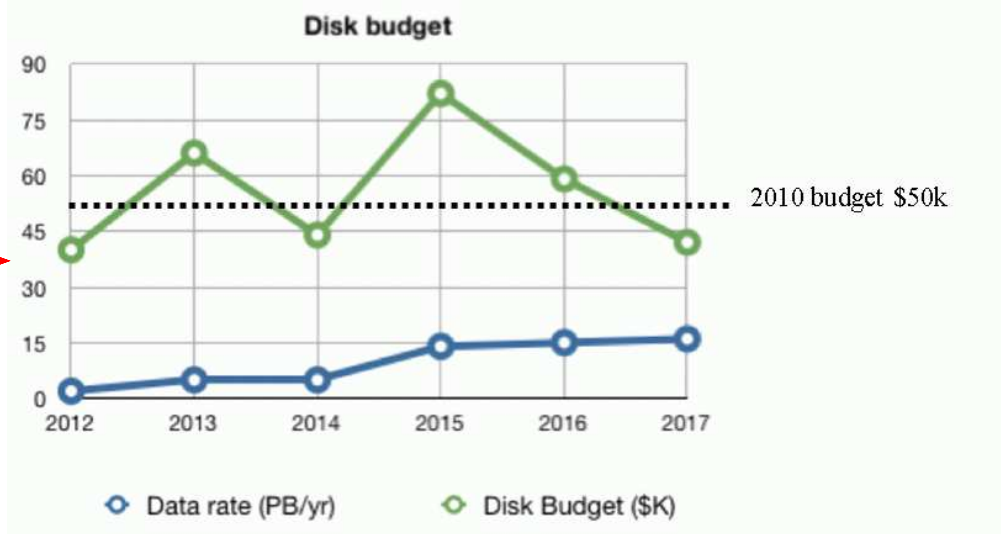

Data Analysis Coordination, Planning and Funding

Graham Heyes - Data Acquisition and Analysis group, Physics division, JLab

- Given the current requirements from CLAS12 and GLUEX and current trends in cost and performance we appear to be in good shape.
  - Numbers should be checked! Factors of 2 either way are not good.
- A budget of \$50k for nodes, \$50k for disk and \$50k for other items gets us very close to where we need to be even as late as 2014.
- An additional \$200k is needed in FY15 to purchase the additional nodes and disk needed to transition to processing raw data and detector calibration.
- The tape costs level out at about \$500k per year in 2016/17 but should start to fall after as the \$/TB tape cost improves.
  - Media costs are traditionally paid out of the hall operation budget.
- At 15 petabytes per year the tape library fills up very quickly.
  - The current library outfitted with next generation tape drives would only hold 20 Petabytes!
  - The funding plan must include a \$100k library expansion in 2014.
- Even so we must have a good plan for quickly processing the raw data and ejecting tapes from the library to make space for new data.

Data Analysis Coordination, Planning and Funding

Graham Heyes - Data Acquisition and Analysis group, Physics division, JLab

- Given the current requirements from CLAS12 and GLUEX and current trends in cost and performance we appear to be in good shape.
  - Numbers should be checked! Factors of 2 either way are not good.
- A budget of $50k for nodes, $50k for disk and $50k for other items gets us very close to where we need to be even as late as 2014.
- An additional $200k … additional nodes and disk needed to trans… ector calibration.
- The tape costs level … 6/17 but should start to fall after as the $/TE…
  - Media costs are traditionally paid out of the hall operation budget.
- At 15 petabytes per year the tape library fills up very quickly.
  - The current library outfitted with next generation tape drives would only hold 20 Petabytes!
  - The funding plan must include a $100k library expansion in 2014.
- Even so we must have a good plan for quickly processing the raw data and ejecting tapes from the library to make space for new data.

Keys to Success:

1. Realistic requirements.
2. Robust communication.

CLAS12 Offline Software

D.P. Weygand (JLab)

- Service-Oriented Architecture (SOA) for distributed physics data processing.
- An architecture (NOT a technology) based on well defined, reusable components: services.
- Services are loosely coupled - minimizes dependencies between services; only requires they maintain an awareness of each other.
- Abstraction - Hide logic from the outside world.
- Reusability - Logic is divided into services with the intention of promoting reuse.
- Composability - Collections of services can be coordinated and assembled to form composite services
- Discoverability - Outwardly descriptive so that they can be found and assessed via available discovery mechanism.

CLAS12 Offline Software

D.P. Weygand (JLab)

The overriding design feature of CLAS12 software is flexibility

The traditional monolithic software design of HEP is constraining, and inhibits innovative algorithms

Within CLAS6 the online and offline demarcation was a wall that was not crossed, in the present design it is intended that successful services will be employed both by the DAQ/online and offline analysis

Service-based CLAS12 full track reconstruction application at ~750microsec per event relative processing time on the ClaRA platform using modest 10 JLAB farm nodes.

Successful implementation of the SOA architecture will permit large scale distributed computing resources to be deployed on a single input stream, eg the DAQ.

It is practical to reconstruct CLAS12 data as it is being acquired.

To fully exploit this opportunity, concentration of effort should be focussed on calibration, typically a year long process in CLAS6.

The SOA paradigm brings with it a whole new class of problems: specifically the level of granularity- but we have several years to work out details of the implementation.

CLAS12 Reconstruction and Analysis Framework (CLARA)

V. Gyurjyan (JLab), S. Paul (CNU), S. Heddle (CNU)

- CLARA is a Service Oriented Architecture (SOA) for physics data processing.
- Multi-threaded, distributed.
- Data-centric : data flow defines the services.
- Deployed on JLab and CNU machines.
- Full reconstruction algorithm in CLARA.
- Results from reconstruction service multithreading.

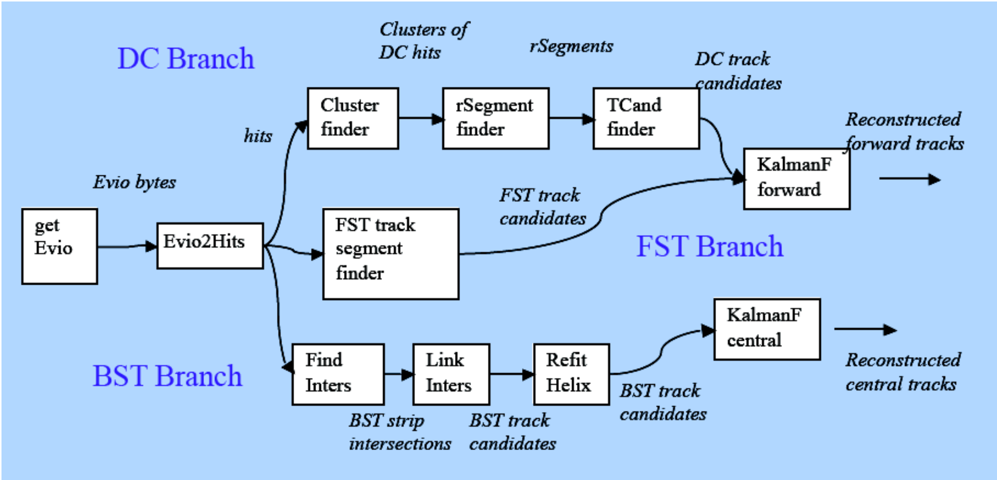

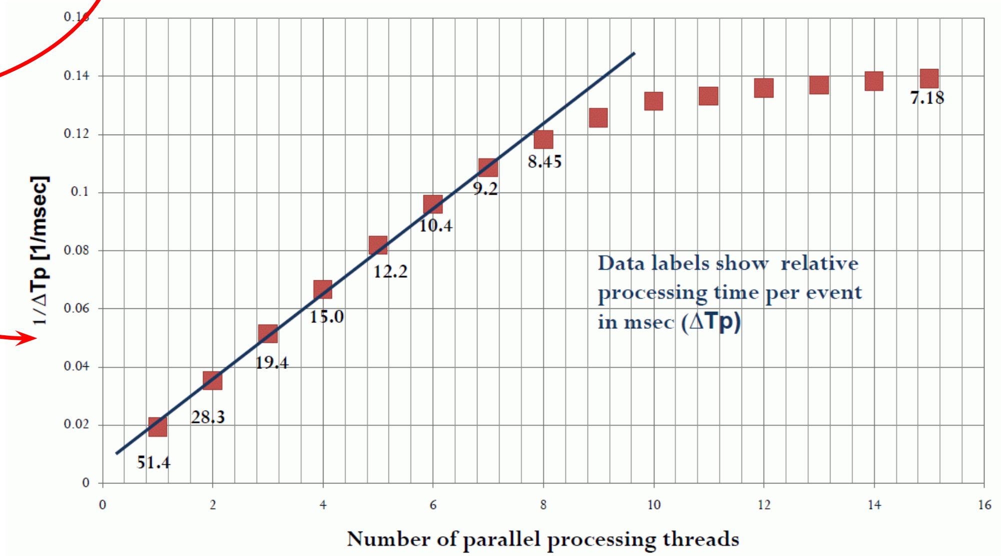

2x4 Xeon, 3 GHz, 8 GByte

The JANA Reconstruction Framework

David Lawrence (JLab)

- JLab ANAlysis framework (more accurately, a reconstruction) framework.
- C++ framework that formalizes the organization of algorithms and data transfer for event based processing. Designed to optimize CPU usage.
- True multi-threaded event processing.
- Data-on-demand model.
- Numerous features custom for experimental Nuclear Physics.
- Support for alternative algorithms and plugins make the framework highly configurable to end users as well as developers.
- Built-in features (janadot) allow for both profiling and easily documenting.

The JANA Reconstruction Framework

David Lawrence (JLab)

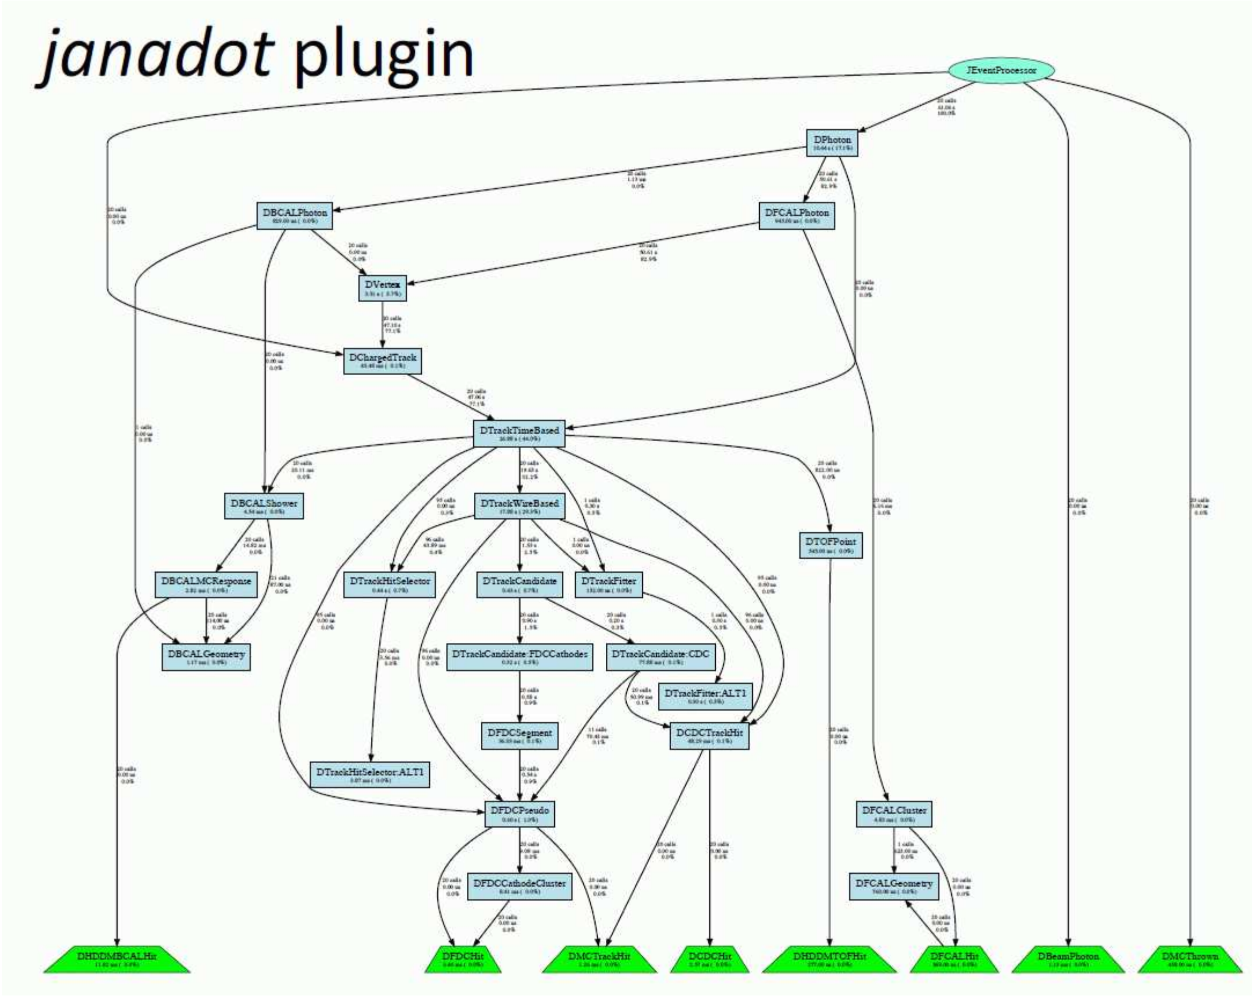

A closer look at *janadot*

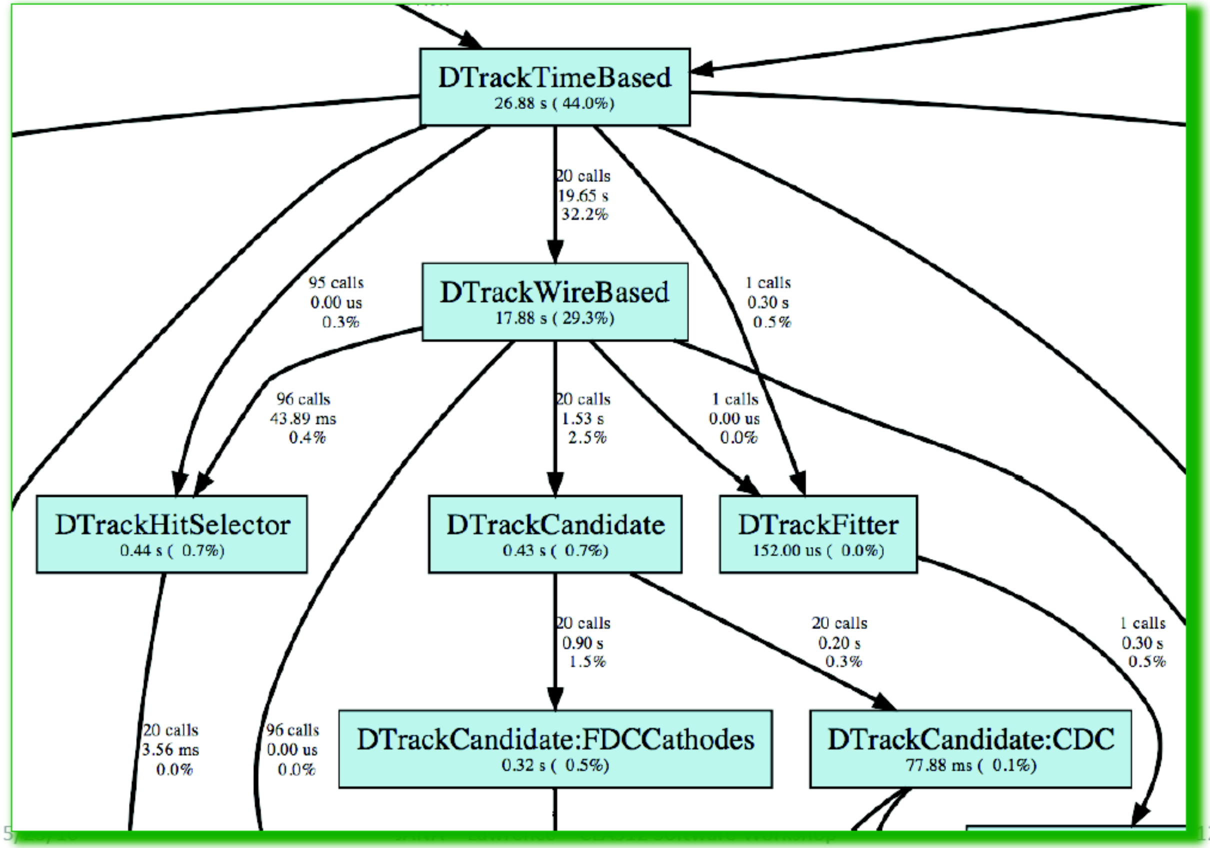

What's the Object(ive) Of Tracking?

Mac Mestayer (JLab)

- Things: hits, wires, clusters, track segments, tracks, trajectories, outer detector matching hits,...
- Actions: find clusters of hit wires in superlayer, find possible track segments in cluster, link segments, ...
- Order matters:
  - identify clusters of hits in superlayers
  - find paired clusters in neighboring superlayers
  - link clusters to find track candidates and trim to viable segments
  - convert time distance for hits in segments
  - do local l-r ambiguity space hits
  - fit space hits to a track
- Full list of objects (see website for details): Signal, WireHit, TimeHit, TrackTimeHit, TrackSpaceHit, TrackSegment, Track.

SOCRAT: SOftware for Clas12 Reconstruction And Tracking

S. Procureur (CEA-Saclay)

- Reconstruction is based on the Kalman Filter algorithm.
- Charged particle reconstruction in the Forward DC (24k sense wires), and Central Tracker (3x2 Silicon + 3x2 Micromegas).
- For accurate vertex reconstruction of forward tracks the Forward Vertex Tracker (FVT) is also include.
- Tracking studies in Central Detector, Forward Tracker, track matching in FVT.

→ Helped the design of several detectors:
  - strip layout of the BST (zigzag to graded stereo angle)
  - mini-stagger in DC Region 3 (300 microns)
  - geometry of the Micromegas FVT at small acceptance (in progress)

→ Code available on svn

→ See CLAS note 2008-015 for more details

→ Runs with Geant4 events (or internal track generator)

→ User friendly script to run Gemc+Socrat on the batch (Simu.pl)

Reconstructing CLAS12 events using JANA

M. Ungaro (UConn)

Introducing Scrat

Software for CLAS12 Reconstruction And Tracking

- Original Socrat code from S.Procureur migrated to C++ objects and factories.
- EVIO input/output.
- Bank list Filtering.
- JANA Multi-Threading, benchmarking.
- New Banks: DC_CLUSTER, DC_SEGM, DC_RSEGM, DC_TCANDIDATE.
- New banks can be read by CED12.

JANA reconstruction using factory methods allows change of modules.

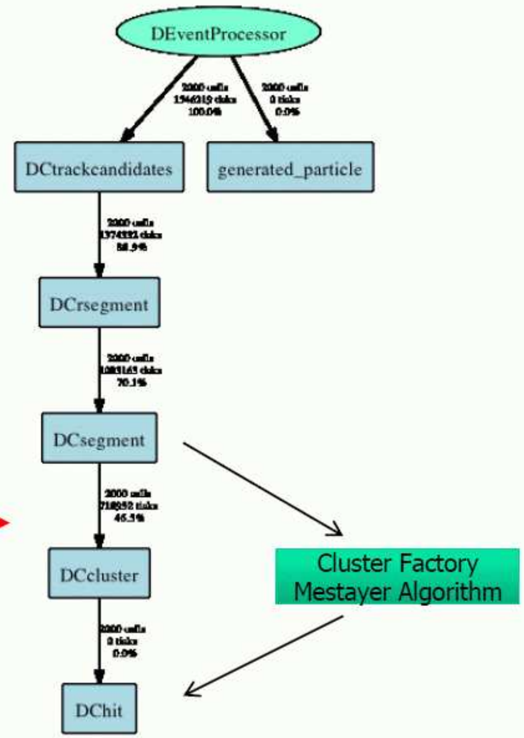

http://www.jlab.org/~ungaro/maureepage/proj/scrat/doxy

gemc - Simulation Results for CLAS12

G.P.Gilfoyle (Richmond) and M. Ungaro (UConn)

- Essential tool for design and analysis (*e.g.* acceptance calculations).
- Quality of the results may be limited by systematic uncertainties (not statistics).
- Will need about four times as much Monte Carlo data as CLAS12 collects.

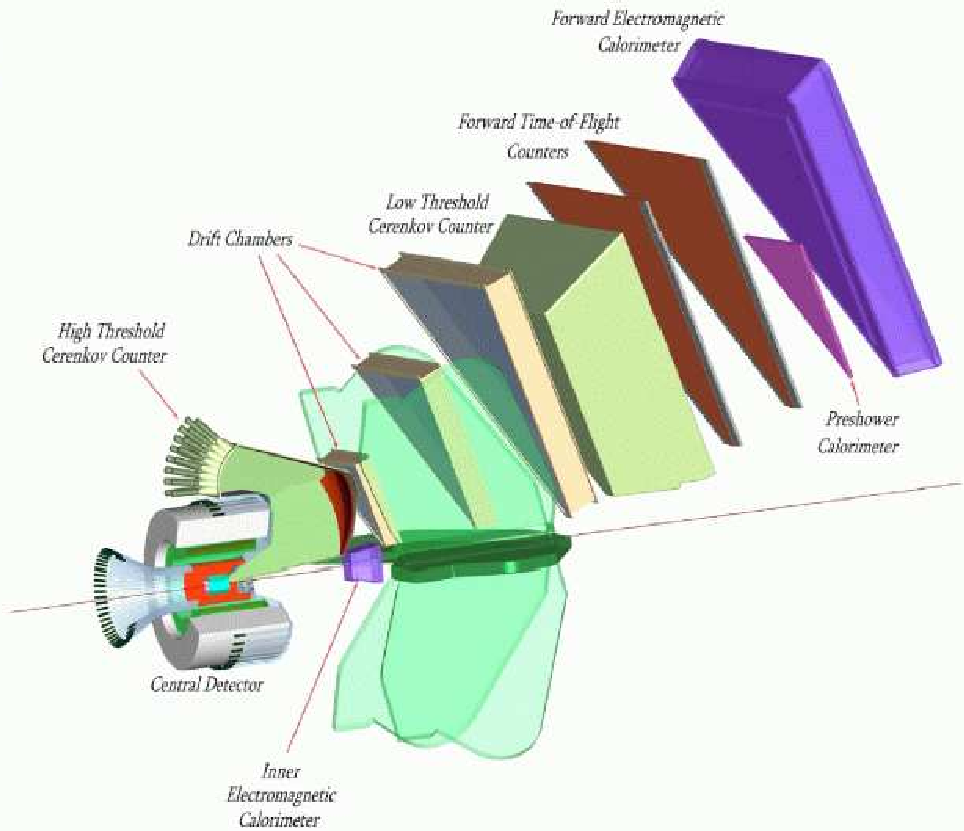

The CLAS12 detector.

gemc

- ◇ Modern, object-oriented, Geant4-based simulation.
- ◇ Still in development stage.
- ◇ Needed subsystems: ⟶
  - ♣ Pre-shower calorimeter (PCal).
  - ♣ Cerenkov counter (CC).

gemc - Simulation Results for CLAS12

G.P.Gilfoyle (Richmond) and M. Ungaro (UConn)

1. CLAS6 results from EC NIM paper (NIM A460 (2001) 239-265) using GSIM.
2. Black curve in each plot based on ionization energy loss in materials preceding the EC.
3. Caveats: (1) electrons passed through different material, (2) $\vec{B}$ field off in the CLAS12 simulation and $\vec{B}$ field on in the CLAS6 simulation, (3) tracking used in CLAS6 simulation while Monte Carlo information used for CLAS12, (4) energy loss out the back, ...
4. Note different horizontal ranges.

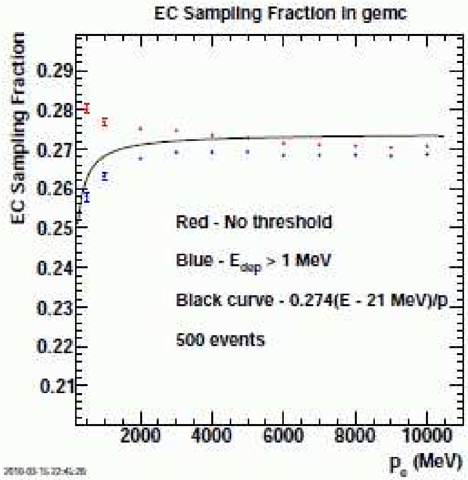

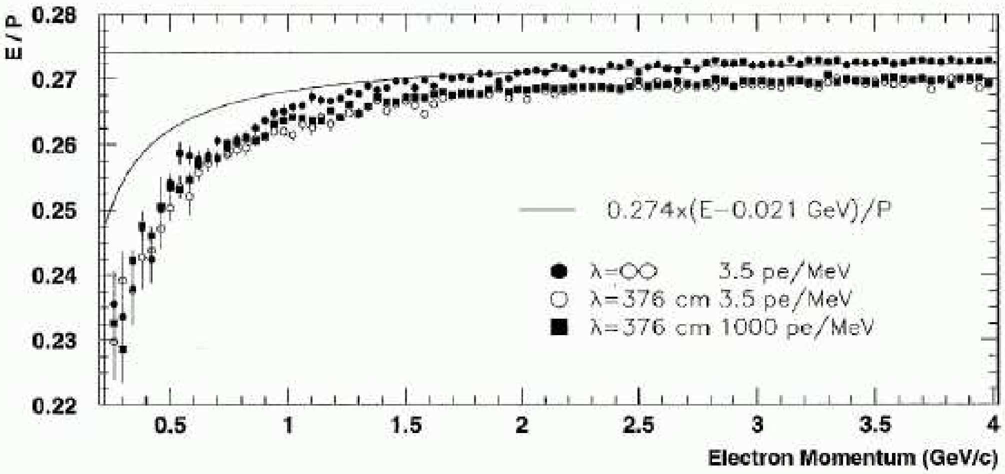

gemc and GSIM sampling fractions are consistent to 10-15%.

gemc - Simulation Results for CLAS12

G.P.Gilfoyle (Richmond) and M. Ungaro (UConn)

- Preliminary simulation of TOF NDE measurement consistent with CLAS6 results.
- EC simulation consistent with CLAS6 measurements.
- Simulation of DVCS events and protons and mesons in the SVT validate design.
- Background studies show that solenoid field will suppress background events and extend SVT lifespan.
- Studies of forward tagger are encouraging and a full proposal will be forthcoming.
- To do:
  - PCal (Mike Wood)
  - Cerenkov counters.
  - Full simulation of tagged neutrons for NDE ($ep \rightarrow e'\pi^+n$).

ced_{12} - Seeing Tracks Through Thick and Thin (Clients that is.)

David Heddle, Andrew Blackburn, George Ruddy (CNU)

- New event display for CLAS12 descended from CED for CLAS6.
- Clients:
  - Thick: Traditional, full-featured Desktop Application.
  - Thin: Same full-featured Web 2.0 application delivered in a browser?
  - We will deliver both, using (approximately) the same code base.
- Java-based, will use Adobe FLEX (2004) FLASH player as virtual machine (VM) (97% percent penetration across all platforms).
  - Browser delivers VM and provides real estate.
  - Compiled application runs in vendor VM.
  - VM, not browser, renders. Browser's primary role has changed! It is a VM container.
- Using same code base for events displays in Halls B and D. Lots of new and interesting features.
- Availability: (1) svn checkout https://clas12svn.jlab.org/repos/trunk/clas12/cedExport, (2) run script (NO BUILD!).
- Development continues: SOA deployed at CNU, more geometries, Web 2.0 version, 3D views, ...

ced_{12} - Seeing Tracks Through Thick and Thin (Clients that is.)

David Heddle, Andrew Blackburn, George Ruddy (CNU)

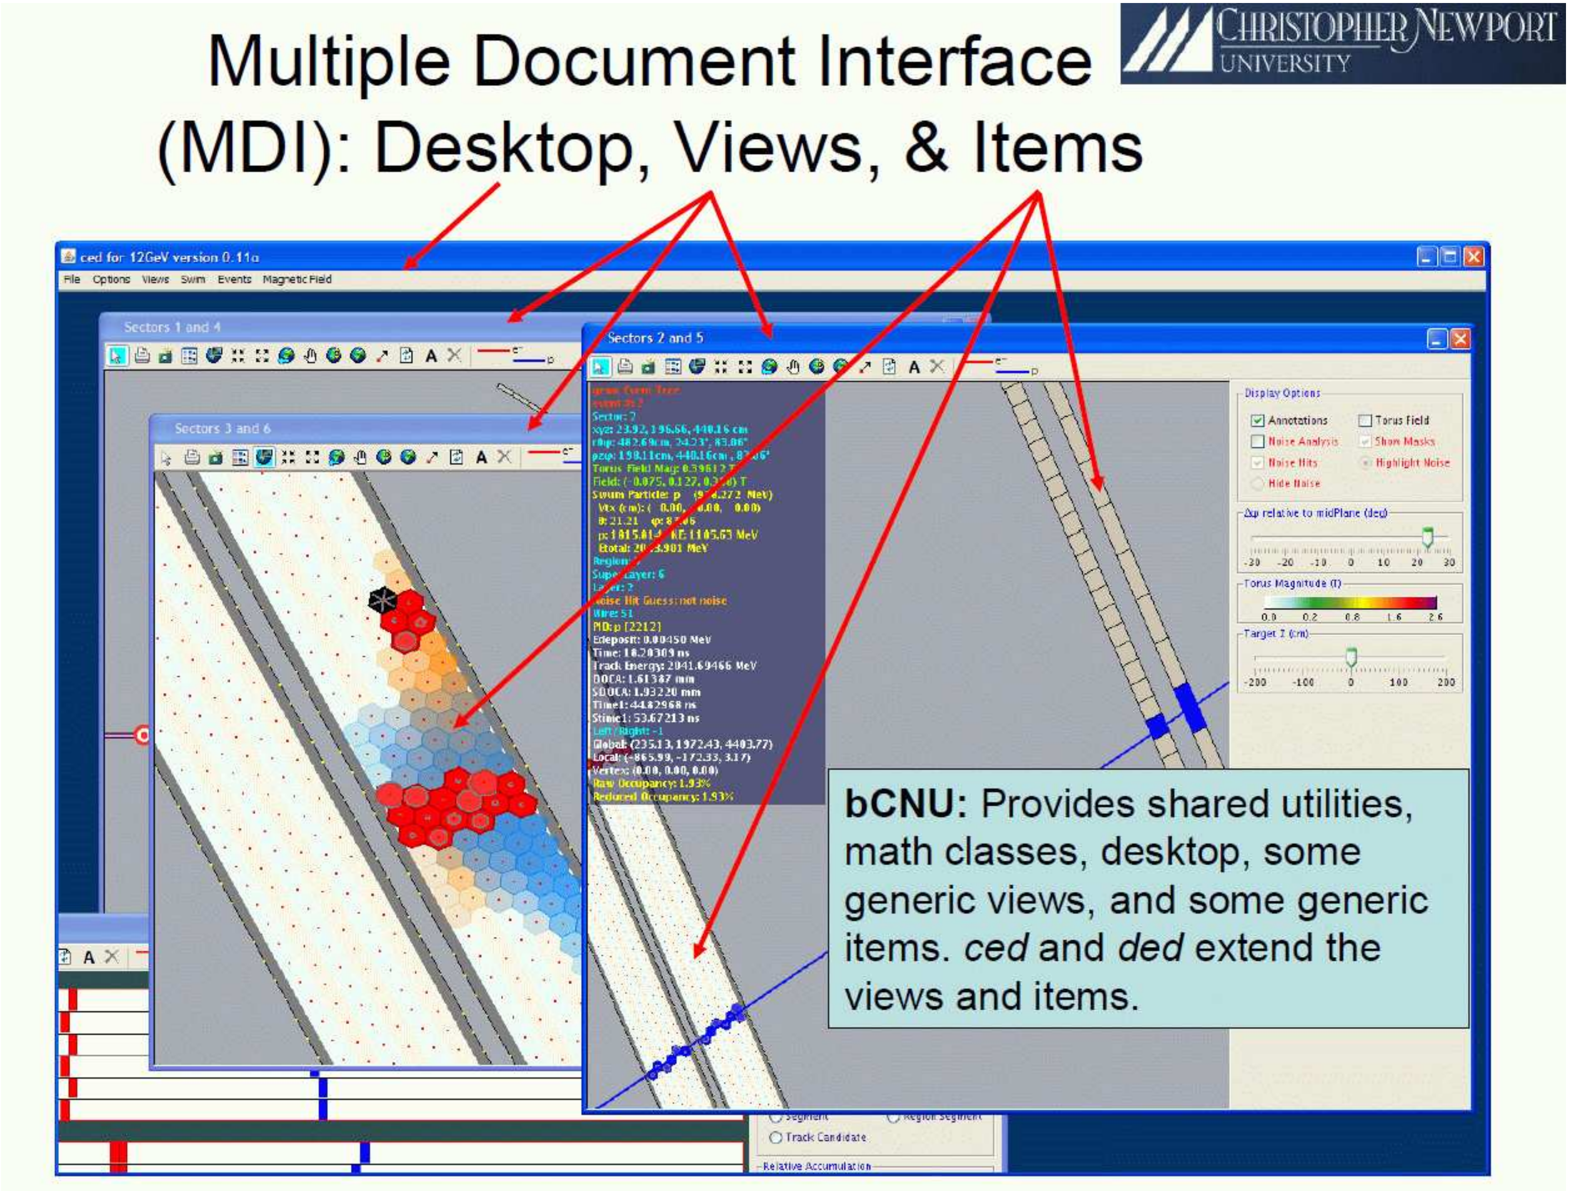

Code Sharing and the EVIO Package

Elliott Wolin (JLab)

- Code Sharing between CLAS12 and Hall D being driven by nearly identical online requirements and similar offline ones.
- Shared projects: CLARA, JANA, EVIO, Event Display, cMsg (JLab messaging package), ...
- Current joint development projects: Event display and RootSpy (display distributed Root histograms.
- EVIO -
  - Original C package just did binary buffer I/O; now implements in-memory object model.
  - Machine/architecture independent, automatically handles endian conversions.
  - Used for CODA raw event I/O, serialize objects to binary array for cMg and disk storage, Geant4 output, and event display input.
- Easily modified to meet future CLAS12 and Hall D needs.

CLAS data format (HDF5)

G.Gavalian (ODU)

- For data-mining project we need unified DST structure for all the run periods (acceptances,run information, flux , etc.).
- HDF - Hierarchical Data Format; HDF5 is the most recent implementation.
- HDF5 is a data model, library, and file format for storing and managing data.
- Made comparisons with BOS, EVIO, and other data formats.
- Bank structure stored with the DST provides backward compatibility (no re-compile, no ddl).
- Custom embedded objects in the DST file.

CLAS12 Tutorials

M.Ungaro, D.Heddle, and V.Gyurjyan

- Fedora 11 was loaded on 12 teaching lab computers at Richmond and gemc installed via yum.
- About 25 students, postdocs, faculty, and staff members attended starting at 9 am on Wednesday, May 26 and lasting until about 1 pm.
- gemc - Participants were able to running gemc interactively, view 3D renderings of CLAS12 and simulate events. Instructor: Maurizio Ungaro.
- ced_{12} - Downloaded and installed on the same teaching machines. Used to read sample events and to explore different views. Instructor: Dave Heddle.
- CLARA - Demonstrations and discussion of use of CLARA. Instructor: Vardan Gyurjyan.