CHEP2009 March 26, 2009

Offline Computing for the MINERVA Neutrino Experiment – Heidi Schellman

MINERνA Experimental Set-up

Surface 100 m

Side HCAL: 116 tons

Side ECAL Pb: 0.6 tons

Cryotarget

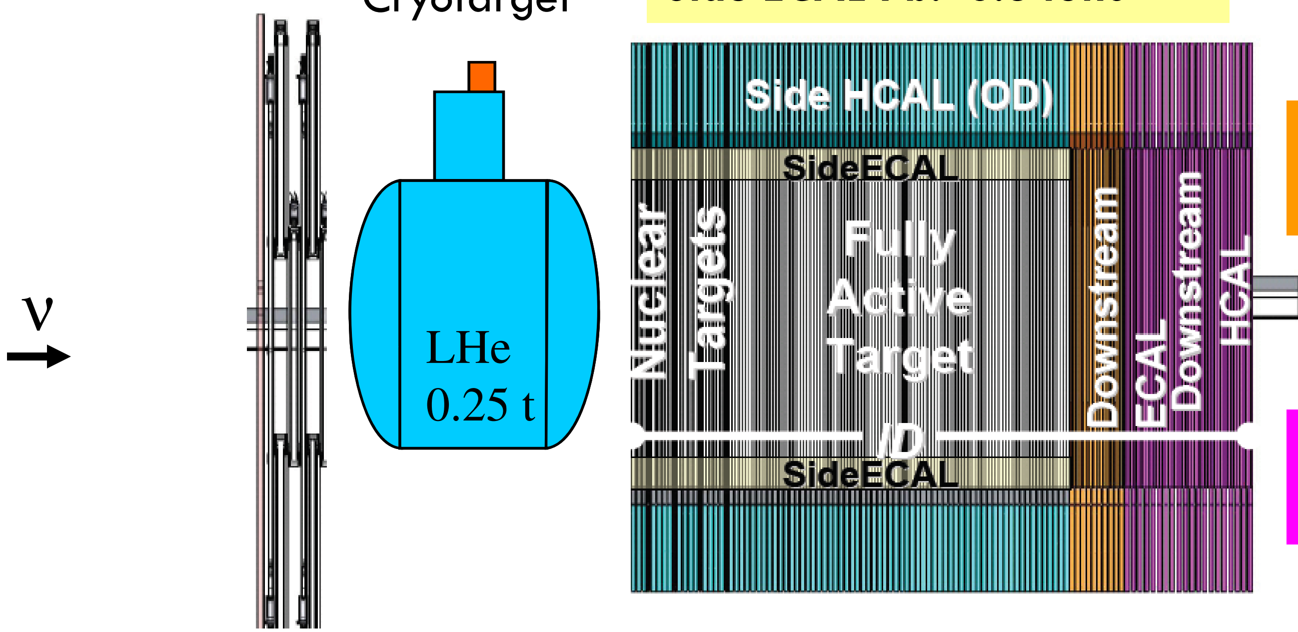

DS ECAL: 15 tons

DS HCAL: 30 tons

VetoWall

Nuclear Targets: 6.2 tons (40% scint.)

Fully Active Target: 8.3 tons

MINERvA – precision neutrino cross sections

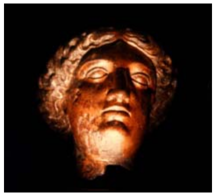

Beam flux

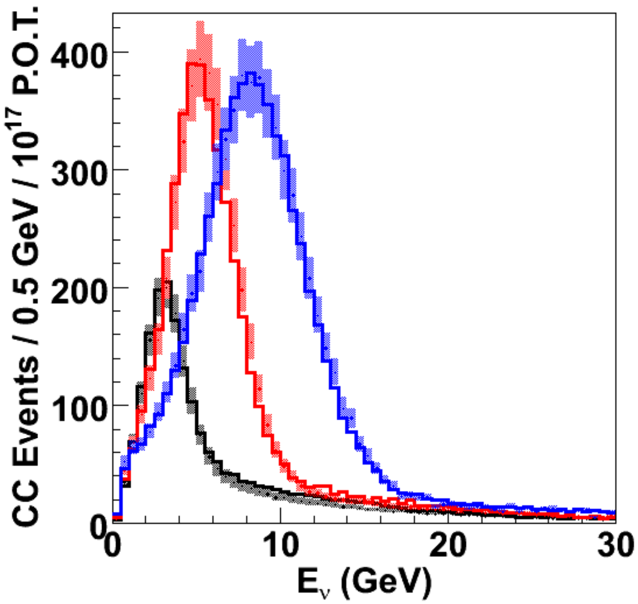

Event rates in 2010-2014

- 8.6 M ν events in CH
- 1.4 M ν events in C
- 2.9 M ν events in Fe
- 2.9 M ν events in Pb

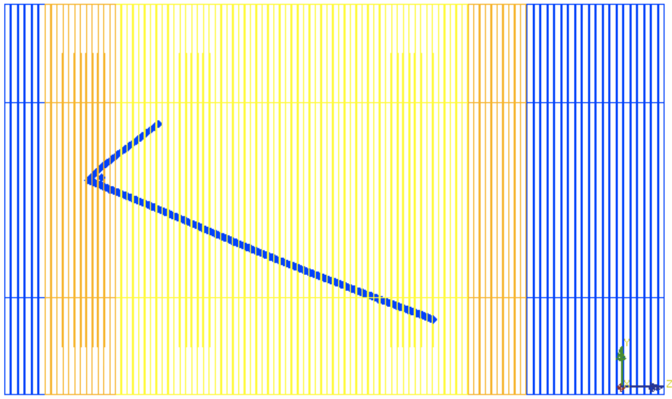

Quasielastic event

$$\nu_\mu\, n \to \mu^-\, p$$

Simulated events and particles

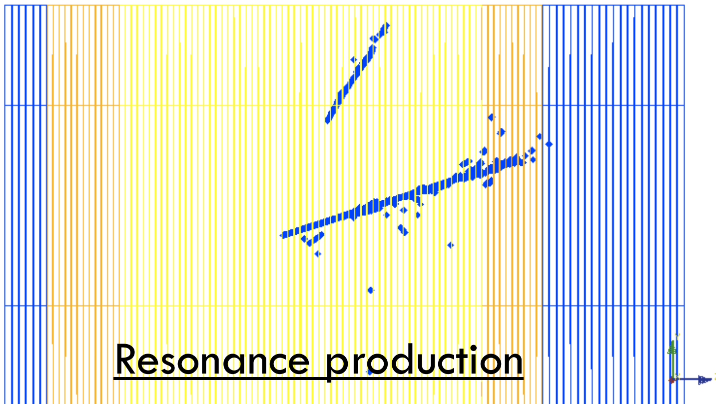

$\nu_\mu\, p \rightarrow \nu_\mu\, \Delta^+ \rightarrow \mu^-\, p\, \pi^0$

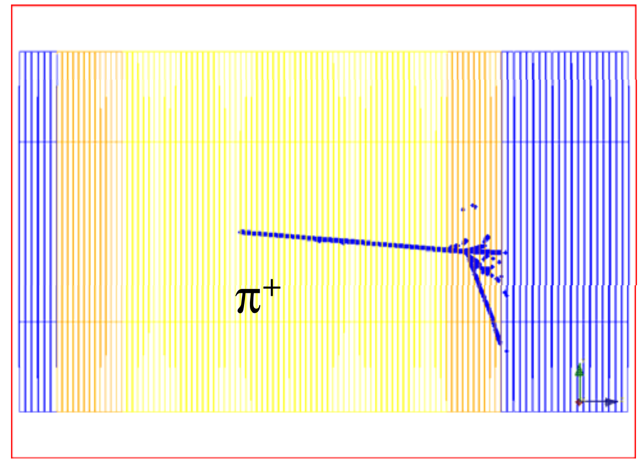

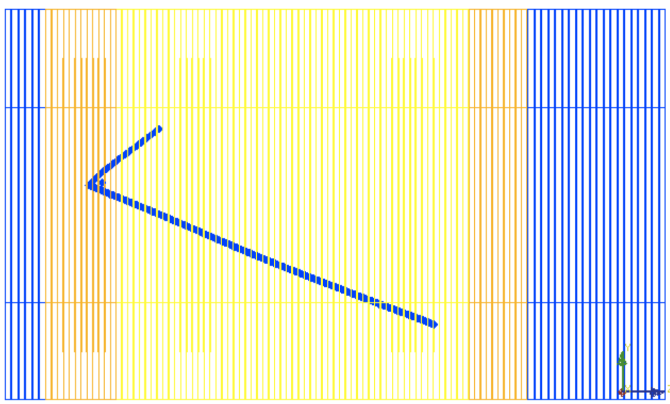

Quasielastic event

$\nu_\mu\, n \rightarrow \mu^-\, p$

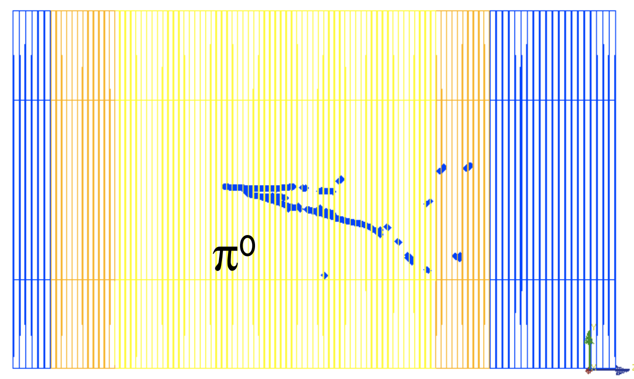

Minerva 24 module prototype

Currently:

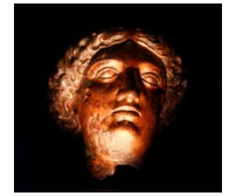

- Tracking Prototype ~ 20% of the final detector
  - Running on cosmics on the surface
  - Will move to neutrino beam next week
  - Current readout is 100kB/event no zero suppression
- Final detector – 2010-2014
  - ~32,000 channels with 3 gains/timing each
  - Multiple readouts over 10μs gate every ~2 sec
  - Raw event size is ~ 1 MB before zero suppression
  - Rate is 1 Hz (data + monitoring) → 1MB/sec

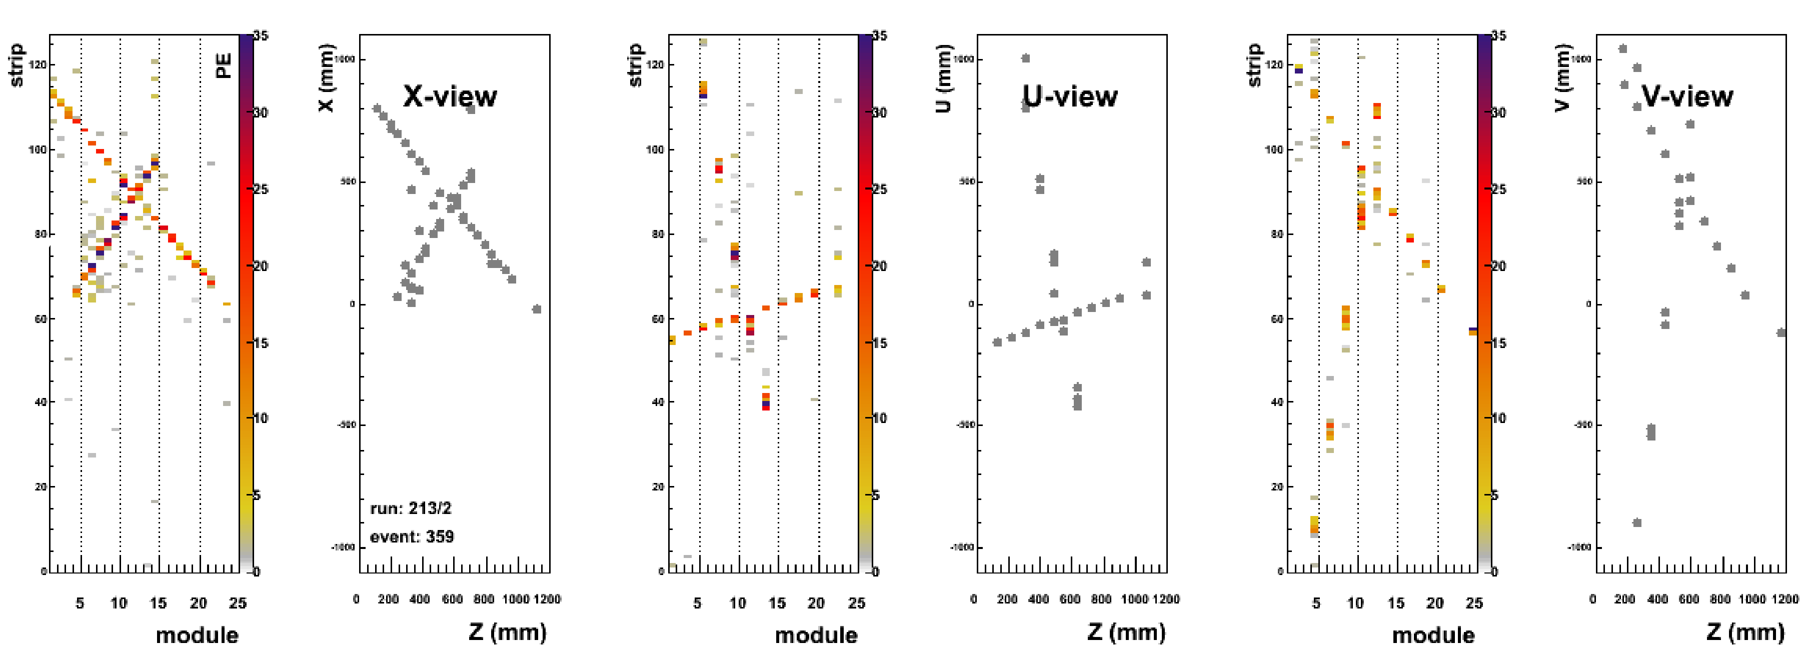

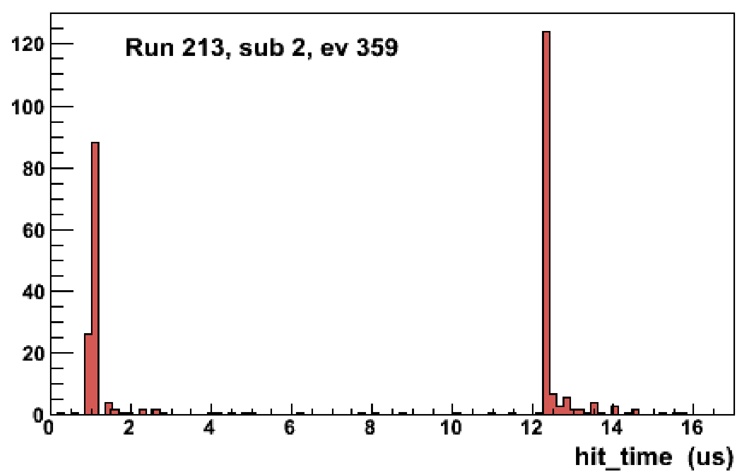

Event with 2 cosmics in the gate..

Raw data and tracking clusters are shown

The MINERvA Collaboration

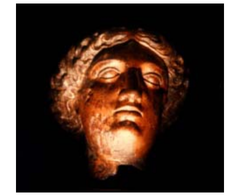

~ 20 active user/developers

- Angelidakis, P. Stamoulis, G. Tzanakos
  University of Athens, Athens, Greece
- C. Castromonte, G.A. Fiorentini, H. da Motta, M. Vaz, J.L. Palomino
  Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, Brazil
- D. Casper, C. Simon, B. Ziemer
  University of California, Irvine, California, USA
- E. Paschos
  University of Dortmund, Dortmund, Germany
- M. Andrews, B. Baldin, D. Boehnlein, R. DeMaat, C. Gingu, N. Grossman, D. A. Harris#, J. Kilmer, J.G. Morfin*, J. Olsen, A. Pla-Dalmau, P. Rubinov, D. Schmitz, P. Shanahan
  Fermi National Accelerator Laboratory, Batavia, Illinois, USA
- J. Grange, J. Mousseau, B. Osmanov, H. Ray
  University of Florida, Gainesville, Florida, USA
- J. Castorena, J. Felix, R. Gutierrez, A. Higuera, G. Moreno, M. Reyes, Z. Urrutia, G. Zavala
  Universidad de Guanajuato -- Instituto de Fisica, Guanajuato, Mexico
- M. E. Christy, C. E. Keppel, T. Walton, L. Zhu
  Hampton University, Hampton, Virginia, USA
- A. Butkevich, S. Kulagin
  Institute for Nuclear Research, Moscow, Russia
- I. Niculescu. G. Niculescu
  James Madison University, Harrisonburg, Virginia, USA
- W.K. Brooks@, R. Ent, D. Gaskell, D. Meekins, W. Melnitchouk, S. Wood
  Jefferson Lab, Newport News, Virginia, USA
- E. Maher
  Massachusetts College of Liberal Arts, North Adams, Massachusetts, USA
- R. Gran, C. Rude
  University of Minnesota-Duluth, Duluth, Minnesota, USA
- A. Jeffers, D. Buchholz, B. Gobbi, A. Loveridge, J. Hobbs, V. Kuznetsov, L. Patrick, H. Schellman
  Northwestern University, Evanston, Illinois, USA
- N. Tagg *Otterbein College, Westerville, Ohio, USA*
- L. Aliaga, C. Araujo, J. Bazo, A. M. Gago, C. E. Perez
  Pontificia Universidad Catolica del Peru, Lima, Peru
- S. Boyd, S. Dytman, I. Danko, B. Eberly, D. Naples, V. Paolone
  University of Pittsburgh, Pittsburgh, Pennsylvania, USA
- S. Avvakumov, A. Bodek, R. Bradford, H. Budd, J. Chvojka, M. Day, R. Flight, H. Lee, S. Manly, K. McFarland*, A. McGowan, A. Mislivec, J. Park, G. Perdue
  University of Rochester, Rochester, New York, USA
- R. Gilman, G. Kumbartzki, R. Ransome#, E. Schulte, B. Tice
  Rutgers University, New Brunswick, New Jersey, USA
- S. Kopp, L. Loiacono, M. Proga
  University of Texas, Austin, Texas, USA
- H. Gallagher, T. Kafka, W.A. Mann, W. Oliver
  Tufts University, Medford, Massachusetts, USA
- A. Chamorro, K. Hurtado, C. Romero, C.J. Solano Salinas
  Universidad Nacional de Ingenieria, Lima, Peru
- M. Kordosky, A. Krajeski, A.G. Leister, J.K. Nelson, USA
  The College of William and Mary, Williamsburg, Virginia

Software and computing

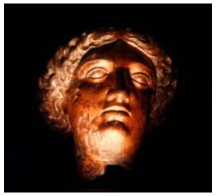

- ~ 20-40 developers/users
- International – US, Europe, Latin America
- Very young – mainly postdocs and students
  - No IT personnel assigned
  - Physicist liaison from FNAL CD shared with other neutrino experiments
  - Rely on FNAL CD for specific tasks
  - Use HEP standard codes as much as possible

Software Framework

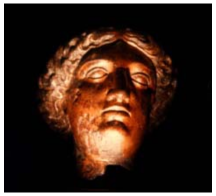

- MINERvA uses the GAUDI framework developed by the LHCb collaboration and used by ATLAS
- MINERvA is a fairly small experiment with limited manpower and the decision was made some time ago to take advantage of the available GAUDI framework
  - CERN had already spent a decade developing it
  - had been successfully used by other small experiments
  - with LHCb and ATLAS using it, will continue to be supported well beyond lifetime of MINERvA
- We also take advantage of many tools from LHCb

Software Framework

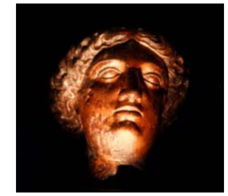

- GAUDI is built on a set of applications and external software products available in a bundle from LCG

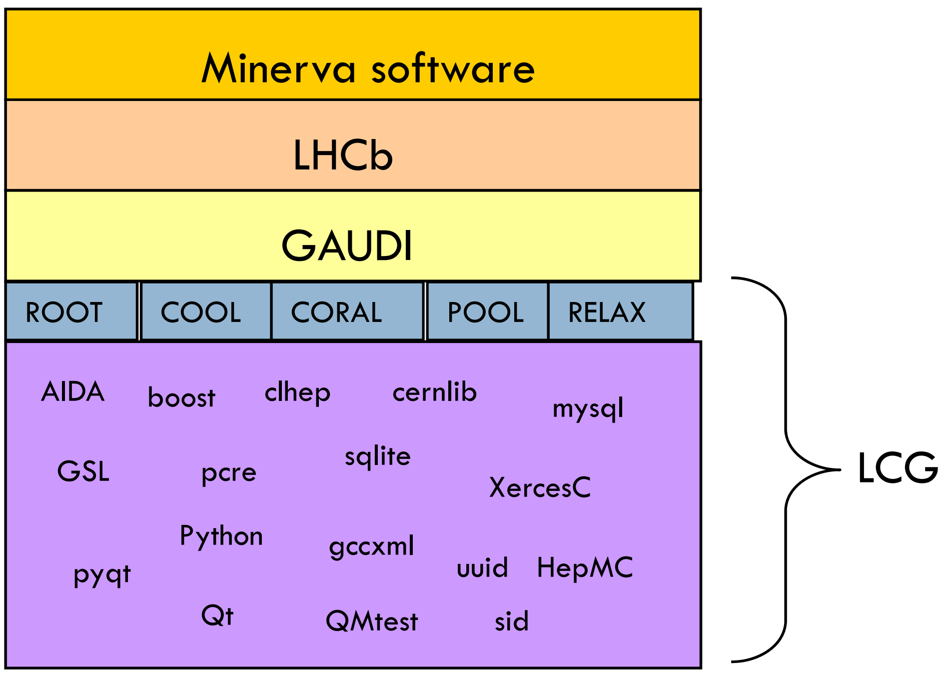

LHCb (v26r1)

↓ GAUDI (v20r4)

↓ LCG v55c

so we currently use binary distributions available for SL4 and build our software on top

SL5 available soon from LHCb

MAC OSX planned

Software Framework

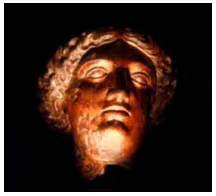

- We also use some general tools from LHCb :
  - GaudiObjDesc – define data model in xml format
  - DetDesc – xml geometry and materials
  - GiGa – interfaces GEANT4 to framework
  - Panoramix – geometry visualization and event display
- but because we use binaries, have not extracted these components from rest of LHCb base ☹ working on it now.
- but if running SL4, very easy to install. source an install script, go for a coffee – or maybe dinner.

Framework

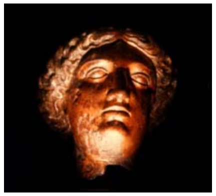

- GAUDI allows reconstruction, simulation and data analysis to be nicely integrated into a single framework
  - We also use GAUDI for data taking and online monitoring…
- POOL ➔ automatic I/O of entire event model at any stage of processing (Persistency for arbitrary transient C++ objects. No need to write converters)
- Built in histogram and ntuple services makes doing high level analysis within the framework reasonable
  - Can avoid many independent ROOT-based analyses which each build analysis tools from scratch
  - (We have created a DSTWriter (flat root tree) for doing studies of data, making plots, developing analysis, etc.)

Software Framework

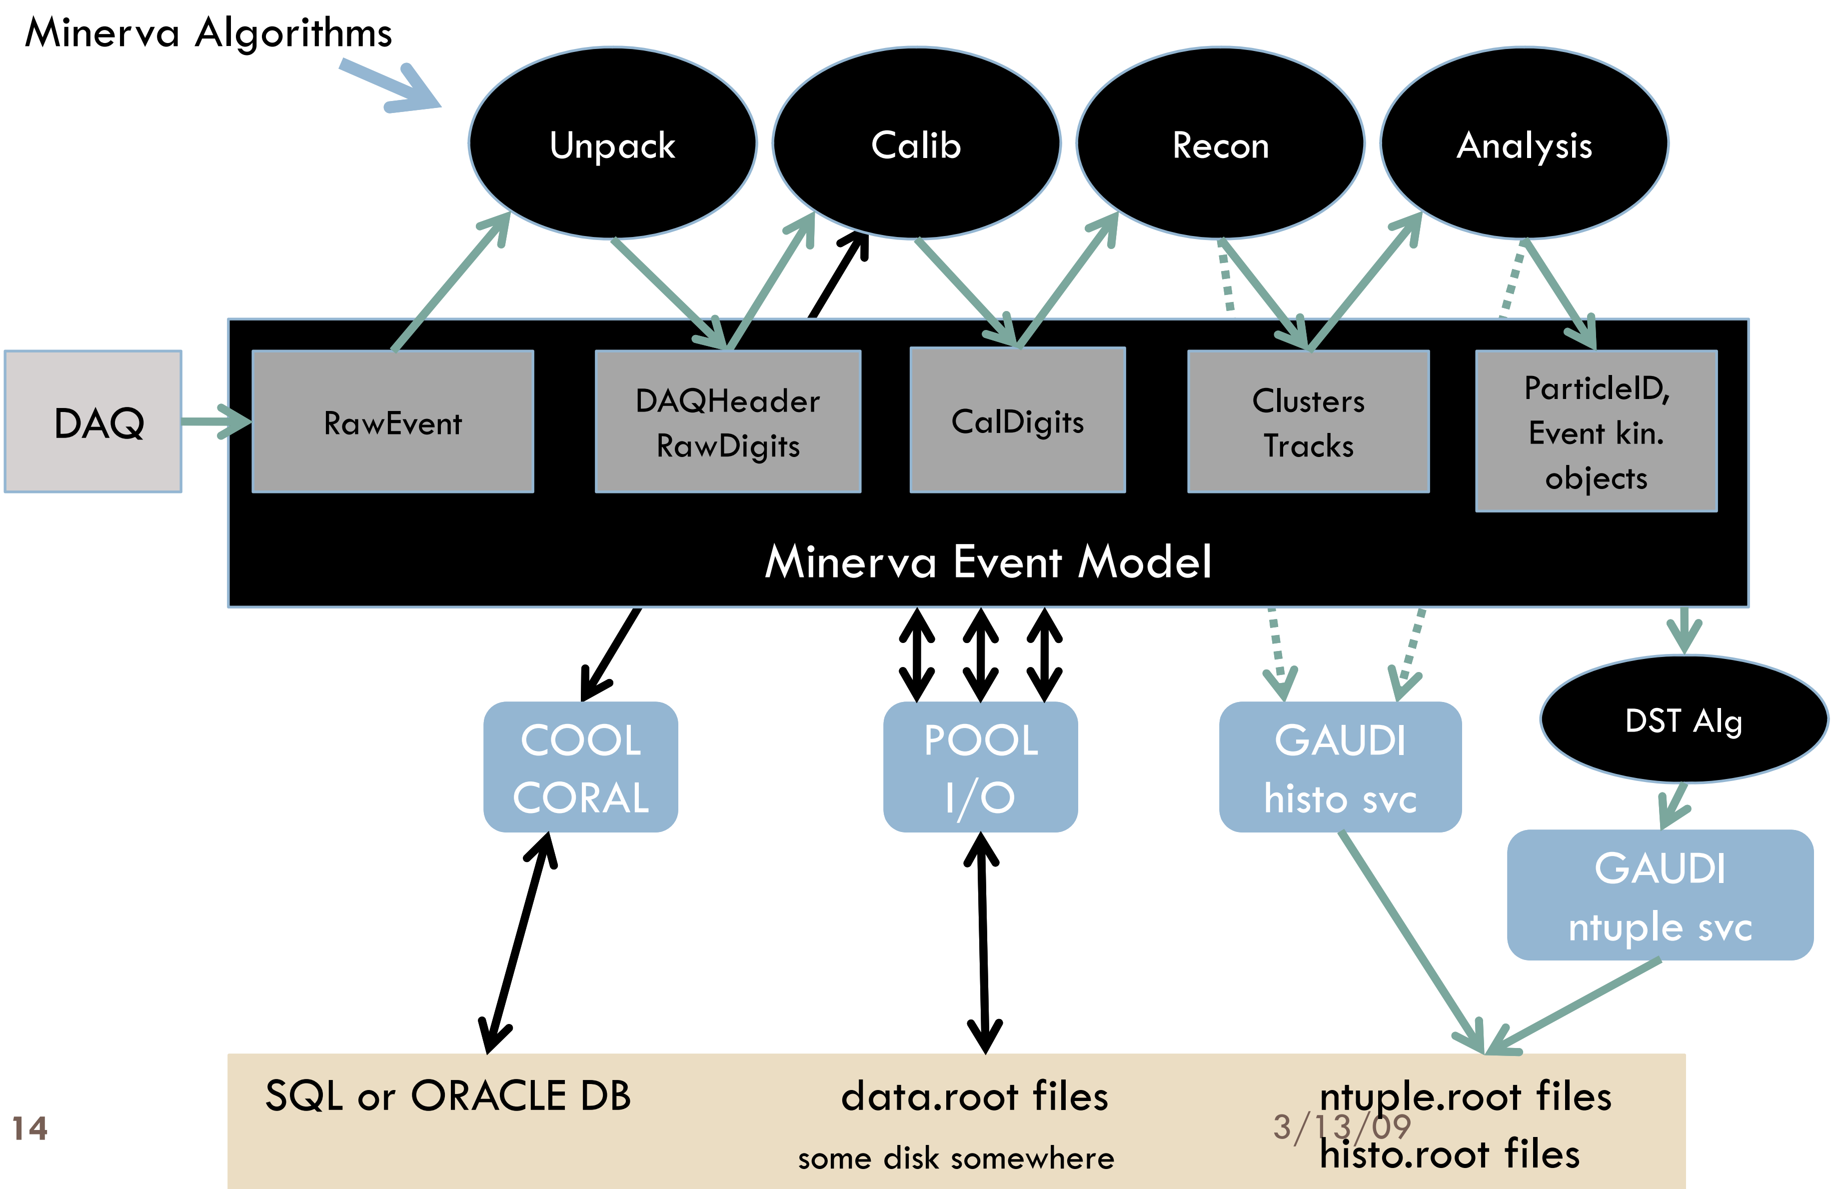

Example: COOL db

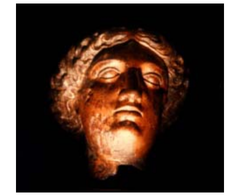

- We started exploring COOL for our conditions database last summer.
- Implement in SQLlite and MYSQL
- We started taking real data in the Tracking Prototype
  - Gains, calibrations→ COOL
  - Our unpacking algorithms use them (1 week effort to implement!)
- Small effort by 2-3 experts led to a nice working system very quickly
- Now considering move to ORACLE with MYSQL or SQLlite extracts.

Example: Event display and tracking

Ben Ziemer - Riverside

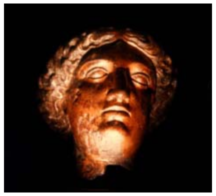

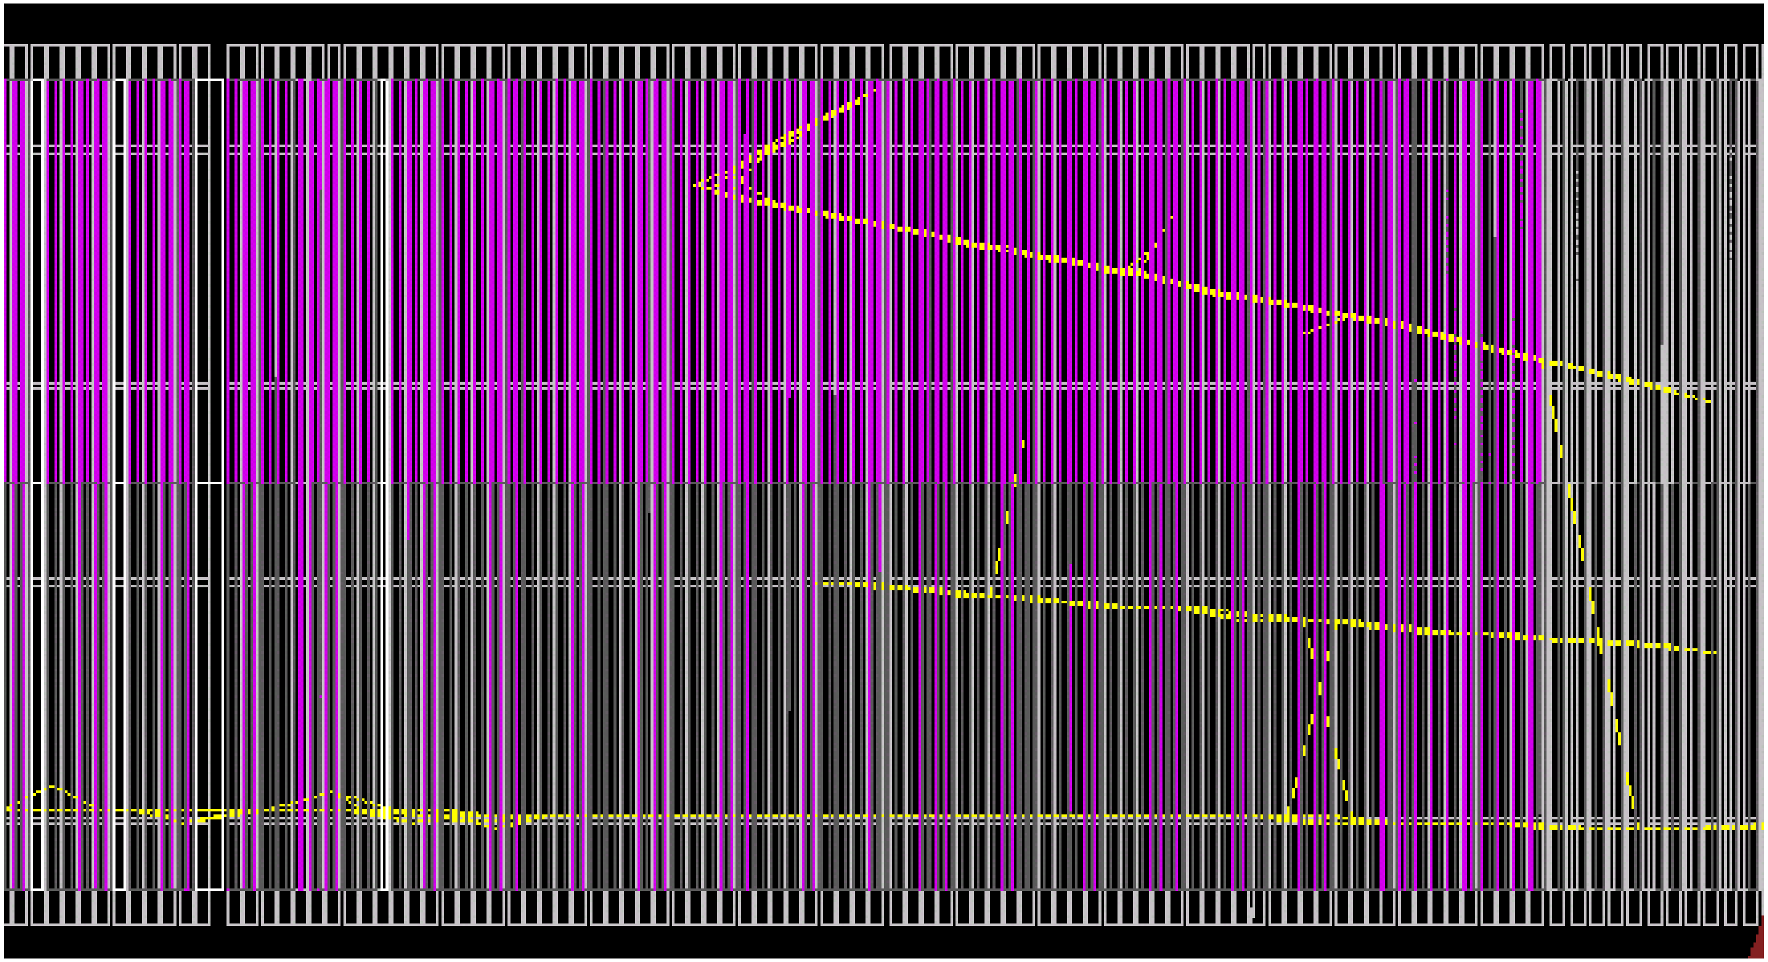

Present concerns

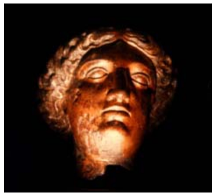

- No local super-expert
  - Seem to be doing ok even so
  - Gaudi-talk and other experts worldwide are very helpful
  - What do we bring to the effort?
- Reliance on LCG/LHCb for binaries, FNAL/FSL for OS
  - Potential to get stuck if OS/code diverge

MINERvA conclusions

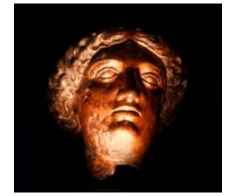

- MINERvA uses GAUDI, LCG tools and some LHCb code
- Not trivial to set up/run but..
  - Lots of valuable pre-existing tools
  - Many helpful people worldwide
- We're taking data and have first versions of almost all tools we need for the full detector run in 2010.

- Thanks to GAUDI/LCG/ROOT/SL … developers

The bigger picture – neutrinos (originating) at FNAL

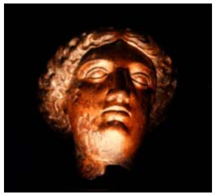

- MiniBooNE 2004-- ? (Oscillations)
- MINOS 2004 --? (Oscillations)
- SciBooNE 2007-2008 (Cross section)
- MINERvA 2009 –2014 (Oscillations)
- Argoneut 2008-2010 (test)
- NOvA 2012 --? (Oscillations/Astro?)
- MicroBooNE 2011 --? (Oscillations/test)
- DUSEL H_2O 2020 ? (Oscillations/Astro/Proton decay)
- DUSEL LAr 2020 ? (Oscillations/Astro)

Recent workshop on neutrino computing at Fermilab

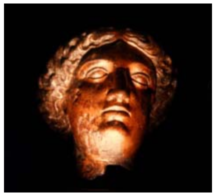

- Representatives from (almost) all the experiments
- Discussions of hardware infrastructure
- Discussions of frameworks/methods

- Typical experiment has ~ 100 "authors", 20-60 active users
- > = 5 different frameworks
  - FNAL, CMSlite, Gaudi, FMWK, LOON, T2K

Sizes of the computing problems

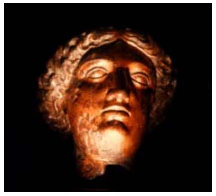

- Technology sets data size
  - Phototube based detectors (MINOS, SciBooNE, MINIBooNE, MINERvA, NOvA , DUSEL H_2O) have similar scales < 1 MB/spill
  - Liquid Argon – whole different ballgame in terms of storage/reconstruction? ~ MicroBooNE → 40 MB/spill
- Science sets live time and data rates
  - Cross sections/oscillations – rate set by machine cycle .5-10 Hz , 0.01% live
  - Proton decay/Astrophysics – 100% live

	PMT	**LAr**
Beam	MINOS SciBooNE MiniBooNE MINERvA ~ 10 TB/yr	Argoneut MicroBooNE 1 PB/yr
100% live	NOvA? DUSEL H_2O 100 TB/yr	DUSEL LAr > 100 PB/yr…?

Example: CPU needs from MINOS

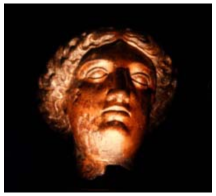

	Pre-2009	2009	2010	2011	2012	2013	2014
Running	?	?	?	?			
Reconstruc	200	200	300	400			
Calibration	small	small	small	small			
Skimming	100	100	100	100			
Analysis	400	500	600	600			
Simulation	small	small	small	small			

Please use CPU-years on a current machine
e.g. # events * time per event in sec * 3 x10^7 * reprocessing factor

Analysis computing needs are driven by advanced analysis techniques: ensemble tests and matrix-element type methods.

Conclusions from workshop

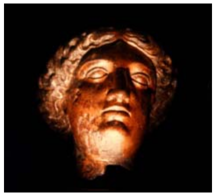

- Experiments within a class (PMT vs. LAr) are pretty similar in terms of needs - we should be able to share similar hardware architecture
- Each experiment's software is unique
  - Most are very happy with their system
- All use common tools
  - ROOT
  - GEANT
  - GENIE/NUGEN
  - SL4/5
- Try to work together more in future

Summary: How will computing people see the neutrino program?

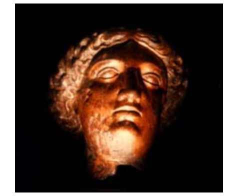

- A coherent program?
- Many annoying gnats…

2009 Super Bowl winners

2007 AL Playoffs