Report from CHEP 2003
UCSD, La Jolla, March 24 – 28

Outline

Introductory Remarks
Observations and Trends
Underrepresented
Open Questions, Concerns, Challenges
My Favourites
General Remarks
Conclusions

Introductory Remarks

- many (up to 9) sessions in parallel
- more than 300 talks in 5 days
- impossible to even listen to all “interesting” talks
- rely on session summaries
  - which were of very different quality, of course
- a personal selection ... starting from the conference summary talk by Torre Wenaus ...

Program

- 16 plenary talks
- 286 talks in 11 parallel sessions plus 3 subsessions
- 61 presentations in poster session
- 12 summary talks
- most talks can be found at
  - www.chep2003.org
  - chep03.ucsd.edu

- Parallel Sessions
  1. Grid Architecture, Infrastructure, and Middleware (28 talks)
     a. Grid Security (10)
     b. Monitoring (10)
  2. Grid Applications, Testbeds, and Demonstrations (32 + 15 posters)
  3. Computing Systems and Infrastructure (28)
  4. High Performance Networking (5)
  5. Data Acquisition, Triggers, and Controls (33 + 16 posters)
     a. First Level Triggers and Trigger Hardware (5)
  6. Lattice Gauge Computing (10)
  7. HENP Software Architecture and Software Engineering (26)
  8. Data Management and Persistency (29)
  9. Data Analysis Environment, Algorithms, and Visualization (33)
  10. Simulation and Modeling (25)
  11. Colaborative Tools and Information Systems (12)

Observations and Trends

OO and C++ the accepted paradigm

- no major OO/C++ migration or usage angst, it is done and widely accepted
  - offline and online: "*triumph of C++ in HEP DAQ confirmed*" (G.Dubois-Felsman, DAQ summary)
- reports on Nth generation C++ software
  - "*every subsystem has been rewritten at least once*" - "*in 2003 the only Fortran is in generators and other outside products. Very few want to go back to f77*" (L.Sexton-Kennedy, CDF)
  - "*many more talks about redesign than about design*" (J.Becla, Data Management Summary)
- reports on maturation and emergence of tools as broad standards, after years of developments and refinements (e.g. GEANT4, ROOT)

Java

- Java ... successful to a degree
  - important limitations being addressed, e.g. manageable C++ interoperability (JACE autogeneration of interface)
  - it is easier than C++ (and less error-prone)
  - used for GUIs in trigger and DAQ along with C++ (e.g. run control GUIs) (G.Dubois-Felsman)
    - similarly in many other areas
  - some examples: JAS (FreeHEP Java Library in general), NLC software, IceCube, ...
  - but not broadly competing with C++ in usage so far

today's computer centers

for example: CERN

PC Farms or commodity computing

- *"triumph of commodity computing hardware and operating systems"* (G.Dubois-Felsman, DAQ Summary)
  - large farms of small boxes model confirmed
- *"CPU power is plentiful and cheap, networking is adequate, local resources (memory, disk) are sufficient"* (S.Wolbers)
- *"don't underestimate farm installation and operations"* (R.Divia)
- *"commodity computing is a challenge"* (S.Wolbers)
  - much variety, constant change
  - big issues are power and cooling and space
  - Watts/$ steadily rising (R.Mount)

Analysis Computing

- moved almost completely from big SMPs to PC based systems
  - driven by cost, performance, capabilities
- *"tough problem"* (S.Wolbers)
  - many clients, unpredictable behaviour, short time scales (conferences), large peaks of load
- *"last problem to be addressed"* (R.Brun)
  - data acquisition, online, production, Monte Carlo all dealt with first
- *"competes to some extent with production services"* (S.Wolbers)
  - tape drives, network, staging systems, ...
- is distributed world-wide
  - the challenge (see later)

RooT

- used for a variety of applications
  - for analysis, I/O (persistency), and much else
  - now fully supported at CERN (EP/SFT)
  - close interaction with experiments and other groups on new developments
    - CDF, D0, ALICE, LCG, BaBar, ...
    - foreign classes (e.g. RooFit - roofit.sourceforge.net), PROOF, geometry modeler, grid integration, ...
- mentioned in about 50 talks at this conference
- certainly one of the success stories ...
- ... I will come back to this ...

RooT

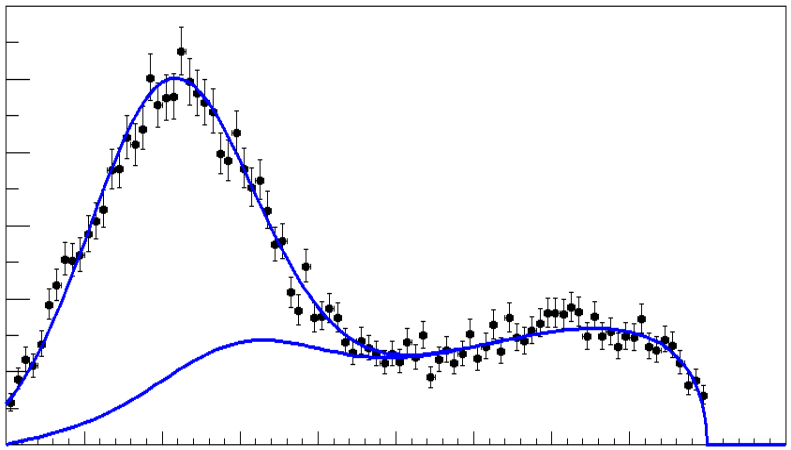

RooFit (roofit.sourceforge.net)

RooT downloads

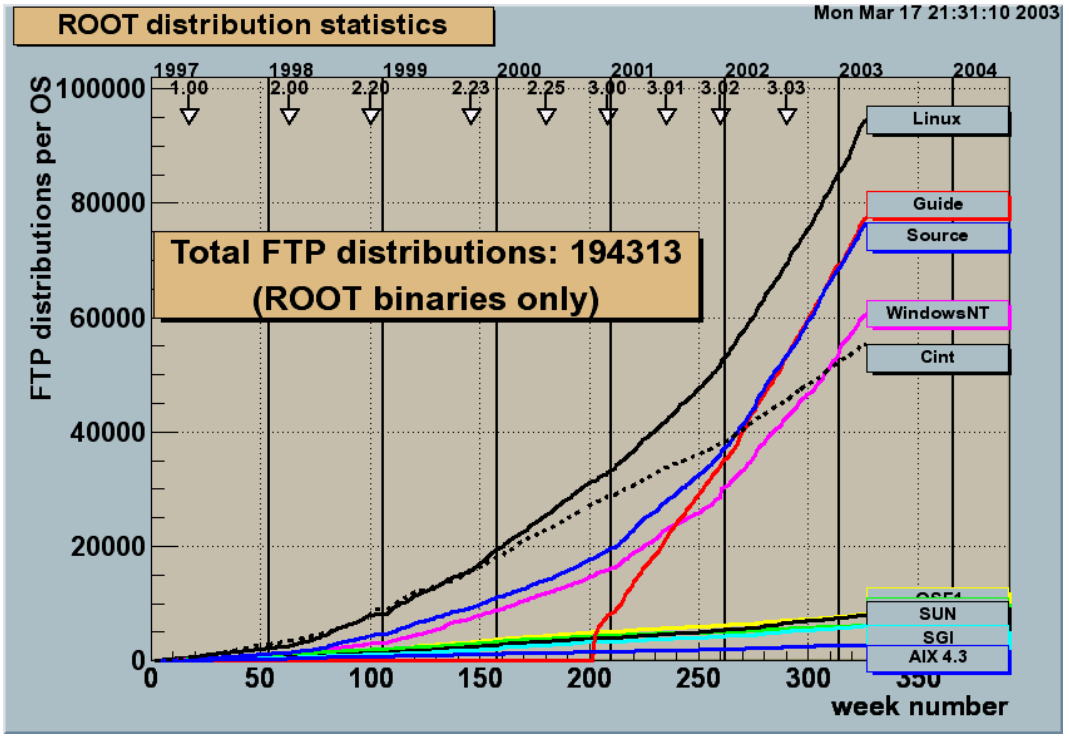

RDBMS revival

- commercial and open source databases (MySQL, Postgres, Oracle ...)
  - metadata, distributed computing, conditions, ...
  - together with ROOT I/O for persistency (e.g. POOL project; see below)
  - open source dominating but also Oracle
- MySQL mentioned in 37 talks, Postgres in 8, Oracle in 27

online – offline continuum

- similar Linux farm environments
- at least mostly the same language (C++) and the same OS (Linux)
- use of offline code in high-level software triggers
  - same appication framework
  - or even same offline reconstruction code (V.Boisvert, ATLAS)
- benefits are
  - greatly simplifies incorporation of trigger algorithms in simulations
  - eases development and validation of trigger algorithms
- several current experiments and most future ones are doing this
- stringent performance/robustness requirements on software

common projects

- *"joint projects are hard, but worth it"* (and one of the CDF/D0 successes) (S.Wolbers)
- Grid projects (e.g. LHC Computing Grid project)
- laudable but difficult and increasingly forced by the circumstances
  - require coordination with (many) other parties
  - resouce constraints and increasing scale and complexity makes go-it-alone N times too costly
- *"somewhat less success in online where it is even harder than offline"* (G.Dubois-Felsman, DAQ summary)
- related is software reuse ...
  - respect what we know about long software development time scales ...

Rene's time to develop plot

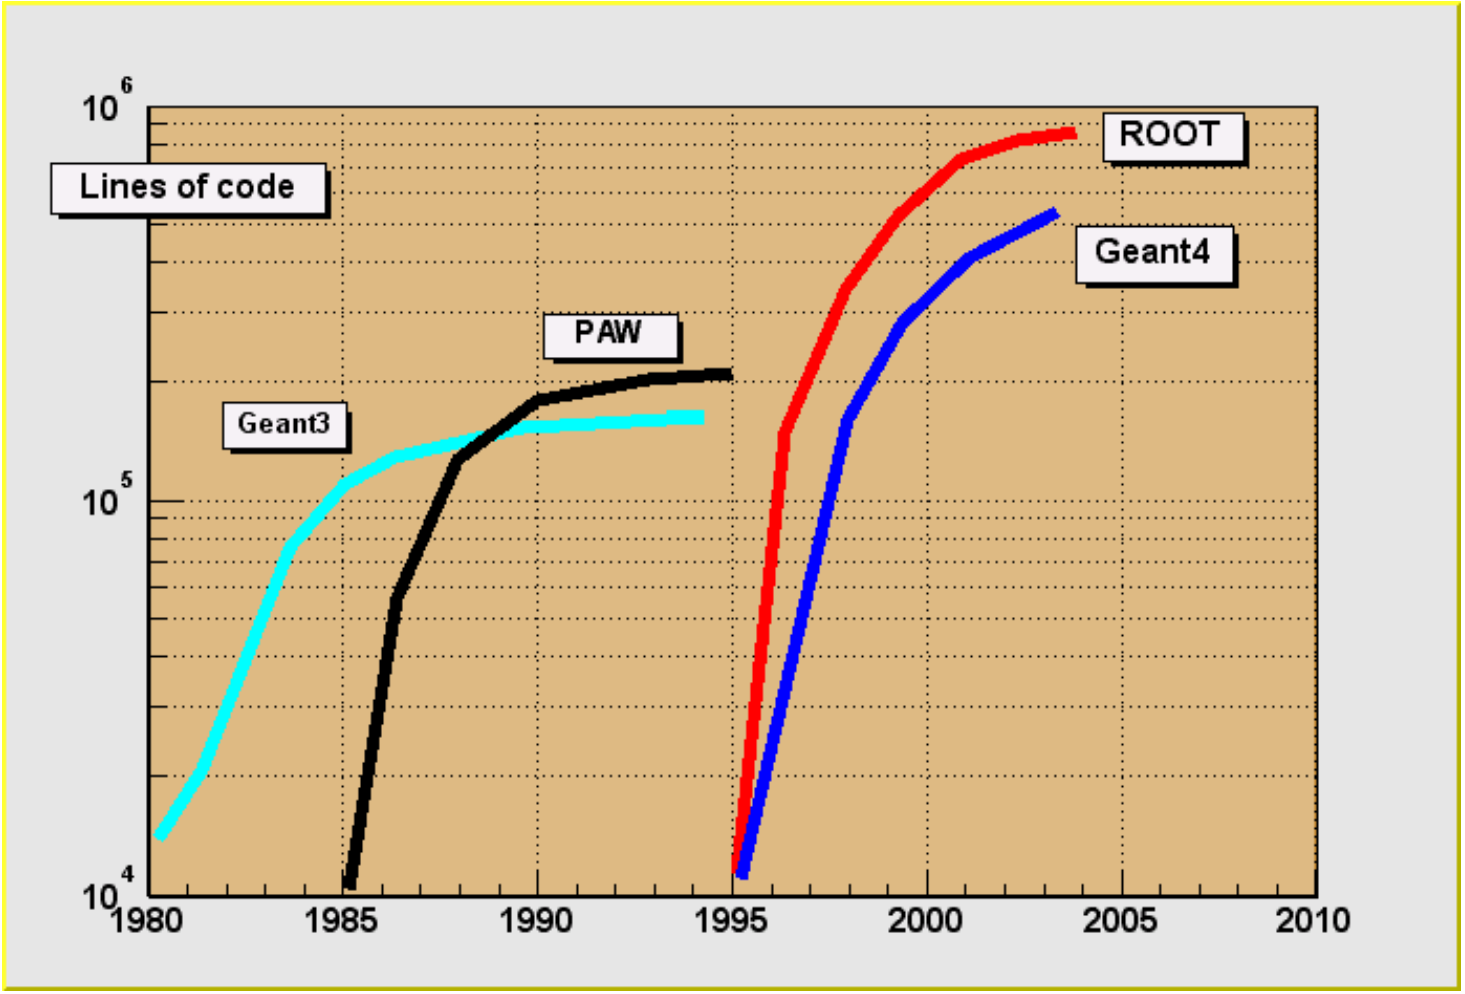

XML

- obviously the choice for small scale structured data storage and exchange
- *“ new: widespread use of XML for non-event information (configuration, monitoring)” (G.Dubois-Felsman, DAQ Summary)*
- many talks in many areas involving XML applications
  - detector description, conditions info, configuration, monitoring, graphics, object modules, data/object interchange, dictionary generation, layered applications, ...
- XML mentioned in 37 talks
- but XML itself does not define common format/schema, and much divergence and duplication exists in how XML is used
  - e.g. detector description

Open Source

- *"Open Source, please. Your interests rarely in commercial vendor's interest"* (M.Purschke, PHENIX)
- *"Live free or die"* (F.Carminati)
- in the CDF/D0 success column (S.Wolbers)
  - similarly all over
- DBs, Qt, utility libraries ... and of course Linux
- extraordinary capability and quality
- ... but *"the world of Open Software licenses is very complicated"* (F.Carminati)
- Open Source has two aspects
  - using open source software (the easy part?)
  - delivering open source products (which products are of common interest?) - discussions concerning Geant 4, ROOT, ...

Simulation Engines

- Geant 4 made great strides towards a production status (www.cern.ch/geant4)
  - in production in BaBar
  - ATLAS on the way to completing Geant4 transition (after 2 years of physics validation)
  - CMS, LHCb, Linear Collider also transitioning
  - ALICE still counts on G3 and Fluka
  - applications in many areas (e.g. accelerator studies, ...)
- FLUKA established and widely used, but new integration efforts as a detector simulation engine (www.fluka.org)
  - FLUGG interface to G4 geometry
  - *"joint CERN-INFN project to develop, maintain, and distribute Fluka initiated"* (2003) (A.Fasso)
- ALICE Virtual Monte Carlo as uniform interface to multiple engines (FLUKA, Geant4, Geant3)
  - interest from other experiments – joint LCG project starting

Geant 4

- Geant 4 maturation
  - EM validation in hand, hadronic beginning (extensive physics validation in general)
  - robust and reasonably fast
  - will be used for ATLAS/CMS/LHCb data challenges starting late this summer
- new developments
  - significant developments in EM and hadronic physics models
  - improvements in geometry modeling (e.g. GDML 1.0 released)
  - interface for using EGS4 physics processes in Geant 4
- ... we should join the effort
  - for LC detector studies
  - for machine studies, e.g. radiation protection, ...

Automation in Software Development/Management

- reports on several automated tools for code building and testing, release integration and tag management, configuration management (Software Architecture and Software Engineering Session)
- two examples:
  - Athena Startup Kit (ASK, ATLAS): users need to know only few commands to jump into analysis
  - OVAL, BOA (CMS): tools for testing and automated builds
- popular new software Web portal at CERN LCG/SPI
  - www.cern.ch/savannah
- automated textual and statistical analysis of test outputs

- ... we should maybe have a look ...

Lattice Gauge Computing

- *"exciting new generation of specialized lattice gauge computers"* (B.Sugar)
  - two tracks
    - QCD on a chip
    - e.g. QCDOC, a "technical marvel", project with IBM
      - $1M/Tflop, aiming at 10+ Tflop at BNL in 2004
    - optimized commodity clusters
      - Pentium4, Myrinet/Gbit Ethernet
      - e.g. 10+ Tflop at FNAL and JLAB by 2006
  - US activities: SciDAC Lattice Gauge Computing (www.lqcd.org)
  - German (+Austrian,...) activities: LATFOR
  - ... DESY has a word here ...

Objectivity and ODBMS in general

- *"jury still out"* at CHEP 2000 (P.Sphicas), but now clear
- Objectivity dropped or being phased out by LHC experiments, COMPASS, BaBar event store
  - in PHENIX "becoming a liability" (compiler issue); augmented with RDBMSs
  - not due to technical failure but a mix of technical problems. commercial concerns, manpower costs, availability of an alternative
- its replacements are not other ODBMSs but files (often RooT) plus RDBMS (MySQL, Oracle, ...) for metadata (e.g. POOL)
  - this seems to be the community's consensus now
- Objectivity still used for BaBar conditions database and in some other experiments ...

Commercial Software

- *" the HEP community dislikes the paper work involved in paying money, the licensing issues, and the going-out-of-business risks"* (M.Fischler)
- some in decline (Objectivity, LHC++, ...) but some are indispensable or superior (e.g. control software – PVSS) and new prospects are opening (IBM, Sun, MS, ...) in Grid
- Open source now has an important commercial element
  - Red Hat, MySQL, Qt, ...

... and of course
“ The Grid”

... and the Grid ...

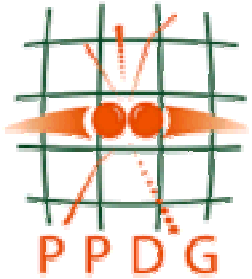

- the excitement remains
- but there is not "The Grid" but rather many of them
  - ... and compatibility is becoming an issue
  - ... there is light at the far end of the tunnel
    - GGF (www.ggf.org)
    - common tools (e.g. VDT + GLUE)
- people are talking about deployment rather than ideas, organization, structures, ...
  - ... but there is still a long way to go

Grid in General

- the central importance of distributed computing to future (increasingly present) HENP is long known
- " the Grid" as the means to that is now established
  - major broad successes in funding and in attracting collaboration with Computer Science
    - *" HEP has set a model for integration, focus, coordination"* (F.Berman, SDSC)
  - progress in applying Grid software and infrastructure to real problems (deployment)
    - e.g. batch production
  - clearly the chosen path – success to be proven, but has promise and broad commitment
- *" fundamental to a grid connected facility is the ability to support multiple experiments at a minimum and ideally multiple disciplines"* (I.Fisk) – that's what we have at DESY!

Grids

- Grids on the horizon
  - must be useful, usable, stable, supported (F.Berman)
    - *"still very manpower intensive – when the support team goes on holiday, so does the Grid"* (M.Livny)
  - more cooperative than competitive
    - requires *"true collaboration, open minds"* (F.Berman)
    - not always the case
  - application are key to success
    - Grid killer application: a focus on data
    - that's the field of HENP
  - much to do and improve but important progress, e.g.
    - VDT (Virtual Data Toolkit) as standard middleware suite
    - OGSA (Open Grid Service Architecture) as essential ingredient in Globus toolkit V3 (www.globus.org) (I.Foster)

LHC Computing Grid Project (LCG)

- major new internationally supported effort to build the distributed computing environment of the LHC which encompasses
  - the distributed computing facility
    - site fabrics (facilities), middleware selection, integration, testing, deployment at distributed sites, operations and support, ...
  - the common physics applications software
    - persistency, core libraries and services, physics analysis interfaces, simulation and other frameworks, all in a distributed environment
- first testbed deployment is this summer (LCG-1)
  - including the first major applications deployment, POOL persistency framework (RooT I/O plus MySQL hybrid)

LCG – my view

- very much focused on LHC, of course
  - ... what about the (us) others ?
- the “magic word” at CERN
- viewed as the prerequisite to the success of LHC computing
- CERN is now putting everything under the cloak of LCG
  - persistency, simulation (G4), application libraries, ...
  - re-organization of CERN's IT activities
- ... this is not necessarily what everybody else might want ...

LCG – SEAL

- SEAL (Shared Environment for Applications at LHC)
  - provide the software infrastructure, basic frameworks, libraries, and tools that are common among the LHC experiments
  - select, integrate, develop, and support foundation and utility class libraries (e.g. mathematical libraries, ...)
  - develop a coherent set of basic framework services to facilitate the integration of LCG and non-LCG software
  - the (ultimate) successor of good old CERN libraries?
  - www.cern.ch/seal

LCG – POOL

- POOL (Pool Of persistent Objects for LHC)
  - a technology neutral API
    - abstract C++ interfaces
  - RooT I/O for object streaming
    - complex data, simple consistency model (write once)
  - RDBMS for consistent meta data handling
    - simple data
    - transactional consistency
  - lcgapp.cern.ch/project/persist

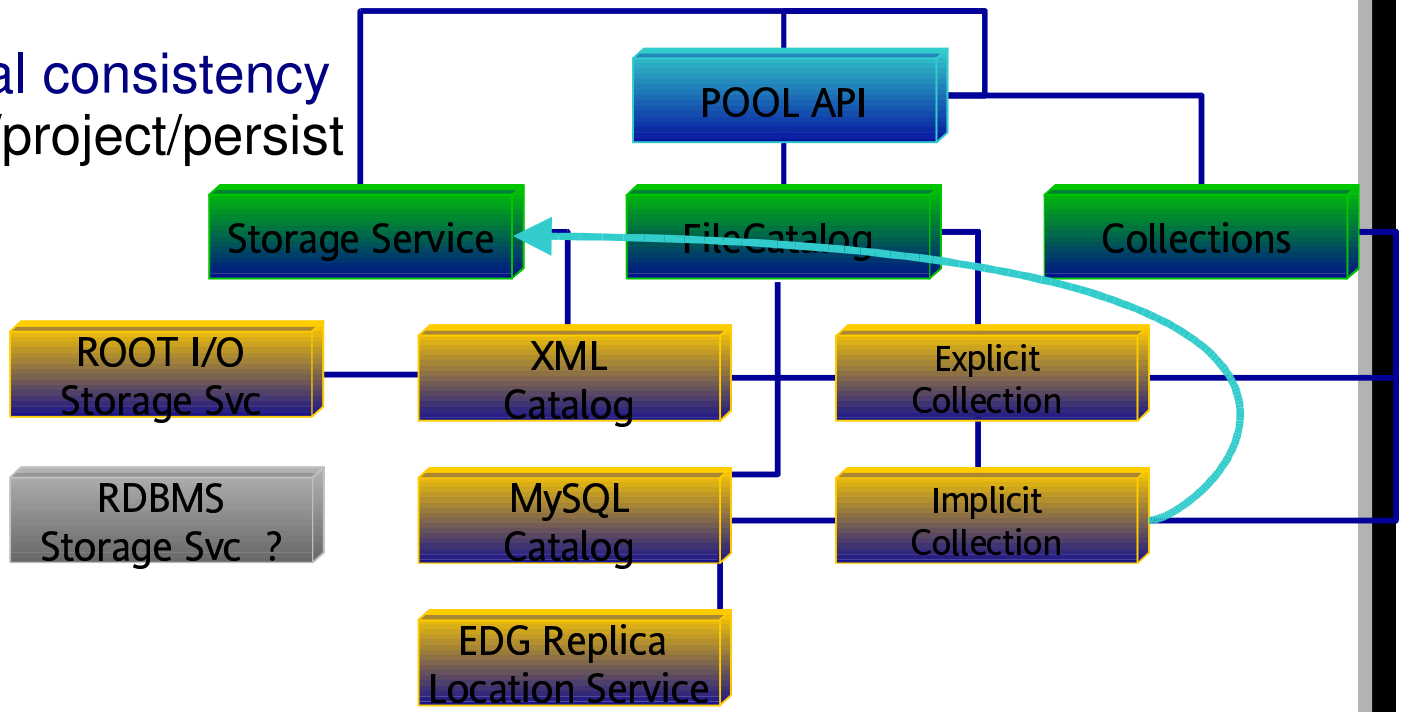

Underrepresented

Collaborative Tools (only lightly represented)

- vital for distributed collaboration on software development and physics analysis
- HEP needs "*culture of collaboration*" (H.Newman)
  - distributed and remote collaboration should be the norm
- not solely, or even predominantly, a matter of tool development in the community
  - how is the exponential commercial side evolving and how can HEP leverage it (take the Web as an example)?
  - what is the evolutionaly path, strategy, role for community-developed tools (e.g. VRVS)?
  - why is the user experience often poor?
    - poor physical facilities/configurations, instabilities, heterogeneous tools/protocols, support issues, ...?
    - current experience sometime competes unsuccessfully with the telephone, despite all the shortcomings (... I can confirm this)

major strides in networking

- *"HENP a leading application driver and a co-developer of global networks"* (H.Newman)
  - require rapid global access to event samples and analyzed physics results drawn from massive data stores
    - PB by 2002, ~100PB by 2007, ~1ExaByte by ~2012
  - rate of progress >> Moore's law
    - factor ~1M in 1985-2005 (~5k during 1995-2005) in global HENP network bandwidth
    - factor 25-100 gain in max. sustained throughput in 15 months on some US plus transatlantic routes
  - network providers see HENP as an opportunity because HENP pushes real production applications

Networking

- mostly concerned with the question “how to get rid off white spots on the globe?”
  - *“as the pace of network advances continues to accelerate, the gap between the economically favored regions and the rest of the world is in danger of widening”* (H.Newman)
  - *“Addressing the digital divide in networking”* (H.Newman)

DESY

- ... I will come back to this a little bit later ...

Open Questions
Concerns
Challenges

Distributed Analysis

- what will it look like? What development line(s) are leading there? Still very much R&D pursued in multiple directions
  - several models with varying degree of Grid exploitations and distributed character
- some efforts reported which are incrementally extending established analysis tools into Grid-based analysis
  - PROOF, JAS, ...
- others working from various starting points
  - Genius (genius.ct.infn.it), Ganga (www.cern.ch/ganga), Clarens, ...
- not a lot of sharing/collaboration above the middleware level
- what analysis improvements will the Grid really provide? (this question was the organizers worth a “Panel discussion” ...)

Distributed Analysis

- *" for a long time data analysis has been the last wheel of the car"* (R.Brun)
  - clear message from current generation: "don't leave data analysis systems and infrastructure too late" – it will lead to problems
  - vastly more true when talking about doing globally distributed analysis, for the first time
    - unprecedented volume and complexity make distributed analysis both very difficult and mandatory
- distributed analysis is the *"crucial test for grid technologies"* (R.Brun)

Distributed Analysis

- How to ensure that senior people can contribute directly to physics analysis?
  - how to interpret the fact that PAW usage is still rising?
  - has everyone bought the C++/OO paradigm shift?
  - Is HEP developing and/or providing the right tools?
- Is there enough engagement of senior physicists in the exploratory work being done on future physics analysis environments?

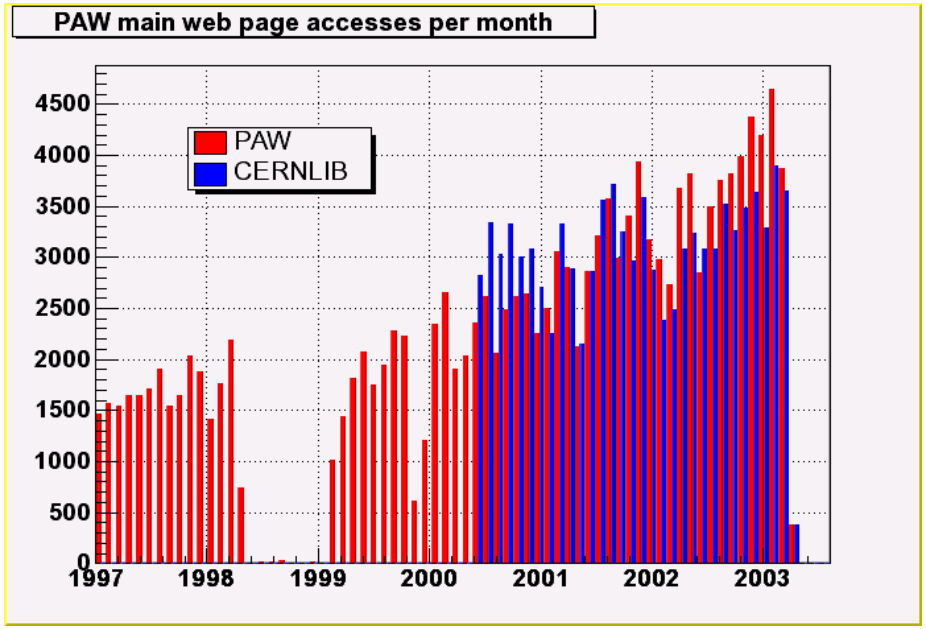

... my favourites ...

IGUANA

- Interactive Graphics for User ANAlysis
  - www.cern.ch/iguana
  - focus on
    - interactive detector and event visualization
    - access to services: data access, application execution, ...
    - interfaces to other tools components (e.g. RooT, JAS, Geant 4, ...)
  - provide an easy-to-use coherent interactive graphical interface for the physicist
  - applications
    - ORCA visualization (CMS reconstruction framework)
    - Geant 4 visualization. e.g. OSCAR (CMS simulation framework), CMS XML geometry description, ATLAS simulation, ...

IGUANA

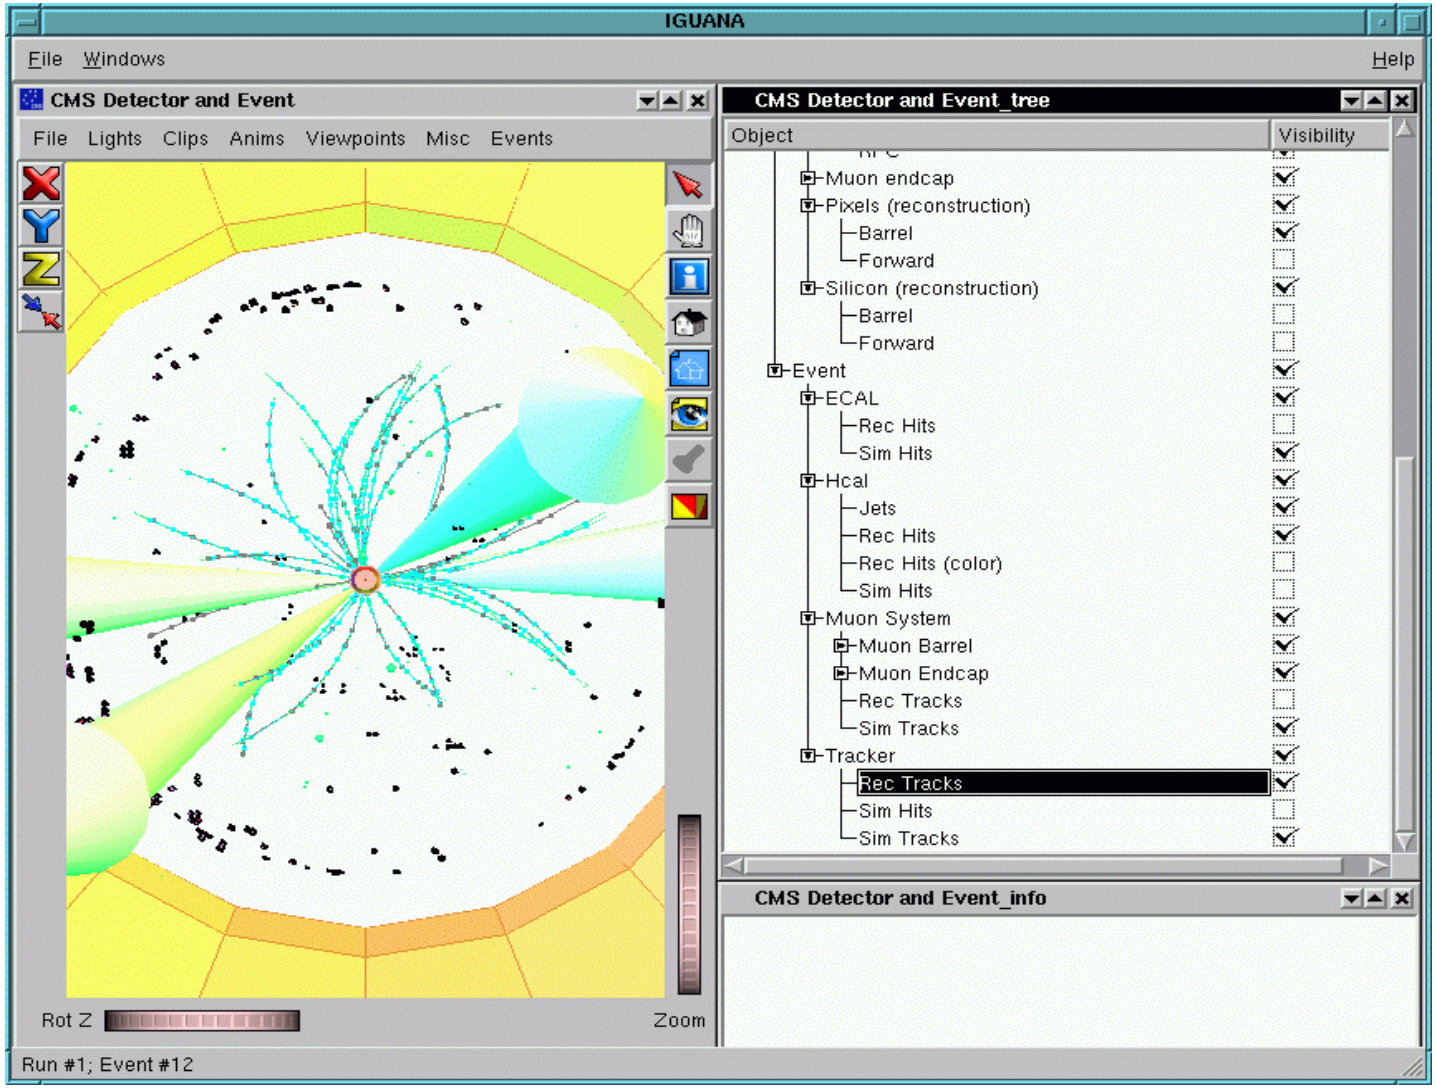
IGUANA
File Windows Help
CMS Detector and Event
File Lights Clips Anims Viewpoints Misc Events
Rot Z
Zoom
CMS Detector and Event_tree
Object
Visibility
Muon endcap
Pixels (reconstruction)
Barrel
Forward
Silicon (reconstruction)
Barrel
Forward
Event
ECAL
Rec Hits
Sim Hits
Hcal
Jets
Rec Hits
Rec Hits (color)
Sim Hits
Muon System
Muon Barrel
Muon Endcap
Rec Tracks
Sim Tracks
Tracker
Rec Tracks
Sim Hits
Sim Tracks
CMS Detector and Event_info
Run #1; Event #12

.. a title that made me curious ..

- *"The OptIPuter – Implications of Network Bandwidth on System Design"* (P.Papadopoulos)
  - data intensive scientific applications require exerimental optical networks, e.g. large data challenges in neuro and earth sciences ... and HEP
  - OptIPuter Research Project
    - large NSF project funded at $13.5M from 2002 – 2007
    - build two high-capacity networks with associated modest size endpoints
  - *"From SuperComputers to SuperNetworks"*
    - TeraGrid optimized for Computing
    - OptIPuter optimized for Network Bandwidth
  - *"key driving applications keep the IT research focussed"*
  - the ultimate answer to any network bandwidth limitations?

.. a talk on software engineering ..

- *"Software Engineering is catching up to us"* (F.Carminati)
  - *"High ceremony processes"* (formal paper documents, very detailed design models, distinct developer roles, formal processes, ...) are not an obvious success
  - *"Agile Methodologies"*, Extreme Programming (XP) is Software Engineers' response
    - extremely close to successful HENP working model
    - adaptive, simple, incremental, tight iterations, plan for change, adjust the methodology for your environment
  - means of responsibility formalizing and addressing – in a useful way – software engineering in HEP, and software management
  - I don't share all of Frederico's opinions, but the talk is fun to look at ...

.. Fortran is back ..

- ... Fortran 90 is used ....
  - " The Persint (Perspectively Interacting) visualization program for the ATLAS experiment" (D.Pomarède)
    - tool for debugging and optimizing the ATLAS muon pattern recognition and track reconstruction algorithm
    - www.cern.ch/atlas/groups/muon/persint.html
- why Fortran 90? (my interpretation of the speaker's answer)
  - because it's not forbidden
  - it's something fancy
- it was the exception !

.. Moore's Law ..

- S.Wolbers: it is not a substitute for
  - more efficient and faster code
  - smaller data sizes
  - thinking before doing
- increased computing drives the science and vice versa
- everyone would benefit from optimizing wherever possible

from www.intel.com

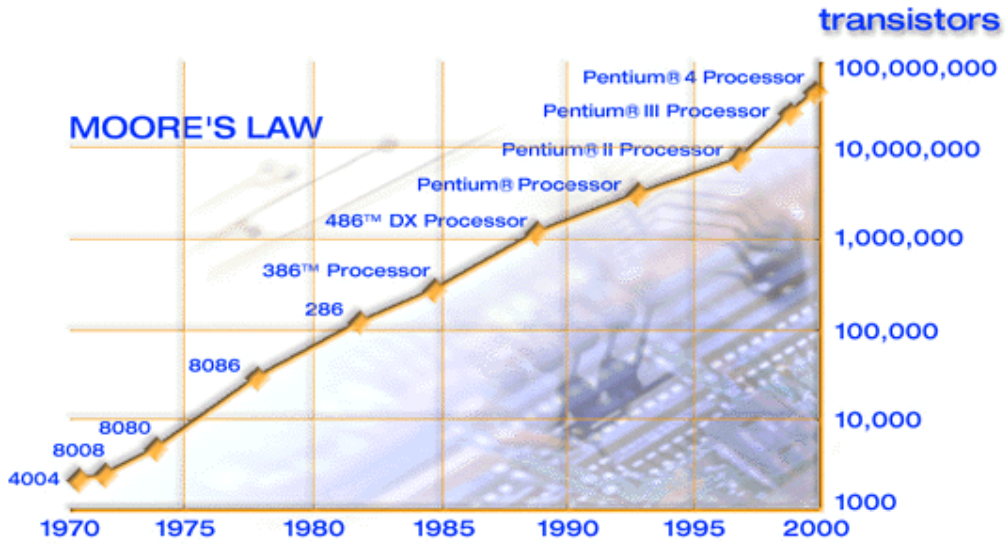

.. something to watch ..

- LHC (ATLAS, CMS, LHCb) data challenge (G.Poulard)
  - test prototype infrastructure before start of the real experiment, i.e.
    - physics event generation
    - event simulation (Geant 4)
    - background simulation
    - event reconstruction
    - data analysis
  - Data Challenge DC-2 (Q4/2003 to Q2/2004) will use LCG common software (e.g. POOL) and use (exploit) Grid technologies (LCG-1)

.. the change in style ..

- no handwritten transparencies anymore
- only few talks with foils
- the majority using electronic means, Powerpoint, PDF, quite a few StarOffice, in fact
- a lot of participants travelling with Laptops
  - my rough estimate: about 1/3 of the participants
  - quite a bit of noise (e.g. Windows beeps) in the auditorium
- ... how could we ever live without wireless ?

general remarks
(my personal view)

- no outstanding new topic (“highlight”)
- many developments reach maturity, e.g. “grid” technologies, OO/C++ based applications
- “ the Grid” is happening, it continues to be the “hot” topic, leaving its infancy, albeit (unfortunately) with a lot of diversity
- almost nothing on networking and collaborative tools
- main theme: consolidation, maturation
- about 400 participants (19 (+1) from DESY)

DESY's role

- DESY (again) underrepresented in plenary session speakers and parallel session convenors (apart from M.Ernst), despite of (usually) good talks in parallel sessions
  - we do (did?) not cover the “hot topics” in HENP Computing
    - OO/C++
    - Grid
    - Geant 4
    - ...
  - program comittee dominated by CERN and SLAC

DESY's role – my view

- for me this means
  - we need to more visibly promote our developments
    - we must more visibly defend our developments
      - disk cache is a DESY development
      - GAN was not represented, not even mentioned
    - we have our word to say
      - we are not doing fancy things
      - we have (made) the same experiences as everybody else (e.g. CDF, D0)
  - we should actively participate (i.e. provide visible contributions) to major developments (e.g. Grid, Geant4, ...)
  - concerning the Grid activities, we cannot stand away any longer
- we have to take care that we do not miss the train ...

Lessons

- *"Make it simple"* (R.Brun)
  - heavy frameworks often perceived as a serious obstacle
  - no more complexity than necessary
  - users want consolidation, ease of use, and stability
  - must consider also needs of the future (longer view of maintainability and evolution)
    - in the interests of long term stability
    - availability of personel
- *" Plan for change, don't be afraid to make major modifications" (S.Wolbers)*
- *" Coming experiments must learn from prior generations – give early attention to data analysis" (T.Wenaus)*

Conclusions

- C++ is a mature and accepted standard
  - several generations of C++ code in production experiments (BaBar, CDF, D0, ...)
  - maturation of tools into broad usage (Geant4, ROOT, ...)
  - no sign of major new language migration so far (Java, Python, ... more complementary than competitive)
- Grids are making great strides, but are still meandering
  - HENP is a successful and valuable partner with Computer Science
  - great for innovation but dismal for standardization
  - security is a bright spot
  - monitoring situation is disappointing

Conclusions

- Global Collaborative Computing must become a successful norm
  - but this is the area were HENP is strong
- HENP is an important partner for Computer Science (e.g. Grid, networking, ...)
  - HENP has the applications

... the next CHEP

- CHEP 2004
  - will be organized by CERN (on the occasion of their 50th anniversary)
  - will take place in fall 2004 in Switzerland (not at CERN)

... the end ...