Scientific Computing at LEP

LEPfest 2000 Science Symposium
CERN, Wednesday 11 October 2000

David Williams

Slides at http://cern.ch/davidw/public/LEPfest.ppt

Caveats

- I found this harder than I imagined it would be
- No single person knows all parts of the story
- If you want to do a proper history and review it would take many months of research
- So I try to paint a picture
- Surely a personal picture, with all the biases that implies
- And in no way am I trying to assign credit to individuals

- I have almost surely forgotten to mention some important aspects of the puzzle

Acknowledgements

- Experiments plus early MEDDLE members
- Especially
  - Carpenter, Cass, Cattaneo, Charpentier, Couet, Giani, Harvey, Hoejenbos, Gokieli, Jacobs, Jones, Knobloch, Newman, Robertson, von Rüden
- Apologies to the people I forgot!

- But the errors and omissions, and there are surely some, are all mine...

More than fifteen years

- Half a human generation, or five Internet generations

- Today's really smart searchers after the Higgs were born in 1975-77
- They are as old as Microsoft [founded in 1975; released Windows 1.0 in 1985]
- They were 6-8 years old when we published "Computing at CERN in the LEP Era" in May 1983
- 8-10 years old when MEDDLE first met in May 1985
- If these whizz-kids' parents had a computer then it was likely to be a Sinclair Z80, a Commodore C64 or a BBC Micro
- Cisco didn't exist. [It shipped its first product in 1986 and went public in 1990.]

Commodore C64

The nature of the problem - People and their interactions

- I suspect that the LHC experiments are running at the limit of what is feasible…
- Not the amount of funds that can be assembled
- Or the complexity of the detectors
- But the possibility of keeping such a large number of very smart people working enthusiastically and actively towards a common scientific goal

- Until the mid-1980s HEP's "computing problem" was often thought to be about obtaining enough processor power
- Then we worried about storage capacity
- The real problem has always been, in my opinion, getting people to collaborate on a solution

PLANNING

A necessary or an unnecessary evil?

Brecht

- Original

  Ja, mach nur einen Plan
  Sei nur ein grosses Licht
  Und mach dann noch ‘nen zweiten Plan
  Gehn tun sie beide nicht

  (Dreigroschenoper - Das Lied von der Unzulänglichkeit menschlichen Strebens)

- Translated

  Now let’s make a plan
  Show forth a shining light
  And then let’s make a second plan
  For neither will be right

Pushing to get computing on the agenda

- In November 1979 ECFA set up a "committee on data handling and processing" (co-chaired by Lillestol and Rimmer)
- That committee spawned ten working groups in Spring 1980 (SG5 on Networking, run by Sendall, was very influential)
- Although the "General Meeting on LEP" held at Villars in June 1981 had no plenary session on computing issues, there was a lot of work by the ECFA WGs and there is a rather extensive discussion, showing much wisdom, by John Mulvey in his closing remarks…including
  - There have been passionate arguments this week over the choice of a standard among the different schemes available for data management.

The first plan

- The 1983 Green Book - Computing at CERN in the LEP Era
- Commissioned by Erwin Gabathuler (then Research Director)
- Three sections - DAQ, Networking, Offline (and a special annexe on Microprocessor Support)
- Plus a Steering Committee Report which was basically an argument for serious funding of LEP computing

- Deliberately gave no funding total. Work over the following year concluded that CERN would need to invest some 50 MCHF over the following 5 years to implement the recommendations of the plan.

Green Book networking

- Good stuff on architecture
  - Move from CERNET to commercial backbone plus special purpose nets (controls, pits, buildings, inside computer centre, phones, …..)
- Choice of "standard" LAN
  - CERN should choose and support a standard LAN
  - Ethernet was stated to be the obvious choice (but we stopped short of recommending it exclusively)
- Digging deeper
  - discussed Cambridge Ring and CATV Metronets
- Estimates of "users"
  - 10-50 per experiment at CERN
  - 19-100 per experiment outside CERN (Europe, NA, ROW added)
  - I looked for 64 kbps per student but couldn't find it

Green Book DAQ

- Fastbus
- Built-in processor power in front-end systems
- VAX will be the preferred family of 32-bit computers
- Adequate staffing of support services

Green Book Offline (i)

- Be careful of "private" computers
- Double central batch capacity by 1987
  - In retrospect we all had a poor understanding of Moore's Law and its vital long-term impact
- Emulators to be driven by individual experiments
- Need for general (networked) file system
  - Conceptual conversion of MSS
  - But no concept of worldwide system (like NFS or AFS)
- Need for Microprocessor Support
- Computing for Engineers
  - "In general the level of computer support for engineers at CERN is inadequate and needs to be improved"
  - Modern system needed for electronics design and VLSI

Green Book Offline (ii)

- Interactive
  - This was the new idea - Did we need it? Was it a toy?
  - Stay flexible - We were afraid to choose at this time - So do everything (CERN compromise!)
  - Follow workstations
  - Install a central VAX, which was an 8600
  - VM/CMS pilot service (we were running MVS/Wylbur)

Green Book - Microprocessor Support

- A final Appendix to the Green Book covered a topic which had been addressed separately, Microprocessor Support
- Much wisdom about the subject in general
- And a recommendation to install a VAX-11/780 running Unix as the central support system
- With connections to CERNET and Index!

- CERN's first mainline Unix service

- Jobs for two DWs

Green Book general

- Recommendation 20

- We recommend that the Director responsible for computing should appoint a Deputy for Computing and Communications. We envisage this person being responsible for ensuring inter-divisional cooperation on a range of topics, including, for example, office automation, various special-purpose computing projects, and networks. In addition s/he would be the natural person to organise the inter-laboratory collaboration required for a successful European particle physics network

- Not accepted!

An interlude to reflect on money

- Computing comes last (in time, I mean)
- And tends to get funded last
  - Cannot make very sensible estimates 8-10 years before serious data
- But normally sufficiently well, in the end

- LEP construction was not an easy time for CERN (tunneling)
- Led us into complicated stories of what CERN could provide as opposed to what experiments would have to provide for themselves
- Probably should teach us that when we do develop sensible estimates for computing they should be respected
- Very important to start well, since expectations are now high at machine start-up (of any machine..)

Ongoing planning - MEDDLE

- MEDDLE was created in May 1985
- Meeting between EP-Electronics, DD and the LEP Experiments
- Roughly four meetings per year
- Brought together people responsible for CERN services (Director(s), DL(s), GLs) with the representatives of the experiments.
- Was it for planning, or more for discussions?

- Certainly fulfilled a useful role in the early years
- Metamorphosed when other experiments grew in (computing) importance

Early Meddling

- Among other things the first meeting discussed how to handle the list of 27 topics which had been requested for discussion!
- Four follow-on meetings were scheduled:-
  - Networks
  - Support for PCs, workstations, and terminals
  - Offline
  - Online
- And the presentation on networking says:-

  Money is the only weapon!

  Therefore do not buy from non-OSI suppliers

A different sort of plan - MUSCLE

- In January 1988 (18 months before first data) a much more detailed scenario was developed, with estimates for processor requirements at CERN, data volumes, tape mounting rates and external network needs.

- It concluded that in the first year of running each experiment would need some 13.5 CERN units of processor power at CERN, and a similar amount outside
- It developed a model (extremely simplified precursor of MONARC) about how much I/O would be needed to (re) process and access the data

More MUSCLE

- ….."we insist that a dramatic increase in communications bandwidth will be essential if the hopes of distributing a significant fraction of the LEP computing load are to be realised."

- ……."Although the bandwidth of a lorry-full of recorded cartridges is very high, such transport has to be planned days in advance, and it is certainly not flexible."

Regional Centres

- I am fairly sure that this report was the first use of the term Regional Centre
- To "qualify" a centre needed:-
  - 2 CERN units at start-up, doubling two years later
  - 10 Gbytes!! of disk space at start-up, doubling two years later
  - access to 100s of 3480 cassettes at ~10 mounts/hour
  - 64 kbps link to CERN at start-up, upgraded to 2 Mbps two years later
  - OS and software environment as close to that at CERN as possible
  - sensible support services
- The recommendation on bandwidth had an important and beneficial impact

The second plan

- Computing at CERN in the 1990s
- Started towards the end of Herwig Schopper's period as DG, in 1988, by John Thresher
- A bid for greater realism, in the sense that the computing problem was better understood, and the end of paying off the "LEP debts" was in sight
- So perhaps more money would be available, and a rather more credible scenario developed

- But by the time it was ready there was a new DG, with his own ideas. The new plan didn't fit well with those ideas

- It contained lots of ideas, many of them good. But much of the effort was wasted. Worth more study!

Ongoing planning - the COCOTIME annual review

- MEDDLE was created in May 1985
- COCOTIME had existed since ~1970, and was responsible at that time for allocating batch time among the experiments
- Starting in 1994-95 it was morphed into an annual review with the experiments of the status of their computing and their technical and financial needs for the coming year
- Director, CN/IT management, and experiments (plus CCCR) then converge on the following year budget allocations and on the technical plans for implementation during the shut-down
- Rather effective

THE PROCESSOR CAPACITY SAGA

What was a CERN unit?

IBM 370/168 - the CERN Unit

Two CERN units per experiment

- In the eighth meeting of the LEPC, held in June 1983, in the Open Session…
- "With respect to computing the Director-General stated that CERN cannot provide more than the equivalent of two 370/168 per LEP experiment."

- Today's fast PCs are about 100 CERN units (370/168 power)
- In February 1983 ALEPH had estimated that they needed 12 CERN units for start-up (1988)
- The announced limitation was a very important political statement to the experiments
- And introduced a lot of "private" computing

Three CERN units per experiment

- By the time of the tenth LEPC in October 1985 we can read that…
- I[an] Butterworth presented the policy of the Directorate viz-a-viz data handling in the LEP era….

  [CERN] will, however, provide the equivalent of three IBM 168 per LEP experiment. (50% increase in 2+ years).

- We should note that this commitment was actually preceded by a funding discussion (upping the computing investment line from 6 to 8 MCHF per year) which indicated an understanding that CPU power was not the only issue.

Siemens 7880 in 1985

Twelve CERN units in 1989

- In Jan 1988 the MUSCLE report had suggested that each experiment needed access to 13.5 CERN units at CERN in 1989
- Following the installation of the Cray X-MP/48 in 1987 it became possible to allocate twelve CERN units per experiment in 1989
- Even if not universally loved (vectorisation, access) the Cray was for some time very cost-effective and provided an extremely effective batch service during its working life at CERN
- It was also our first mainframe running Unix, and CERN had a major role in pushing the rapid spread of that change among Cray users

How much CPU did we really use?

- In the year 2000 each LEP experiment has had access to roughly 2000 CERN units of batch capacity at CERN
- Or roughly 1000 times what they were warned in 1983 would be the limit of CERN's possibilities

- Moore's law (which predicts the doubling of the processor power every 18 months at roughly constant cost) predicts an increase of 1000 times over 15 years. QED.

BATCH
CERNVM to SHIFT

Types of batch

- In 1983, batch was batch, and you ran it on mainframes
- But times were changing
  - The PC-XT was the state of the art (running with a 4.77 MHz 16-bit Intel 8088)
  - The Macintosh, and the 68020, would be introduced in 1984
  - "Graphics workstations" were starting to appear, with the pioneering, but ultimately unsuccessful PERQ from Three Rivers/ICL leading the way
- We started to distinguish simulation from data intensive batch

CSF

- In December 1991 the Central Simulation facility was assembled, based on 45 HP 9000/720 series workstations (Snakes) running Unix (HP-UX).
- This was CERN's first computer "farm" based on commodity equipment. Previous farms had been based on various generations of emulators.
- It offered a very significant price-performance improvement over conventional mainframes, with a software environment which was rather reasonable.

CSF in 1992

SHIFT

- Another step was to explore the possibility of carrying out much more compute-intensive tasks than CSF could support
- Work (HOPE) based on Apollo had not really generated enough I/O power
- Then a joint project (SHIFT) between IT, Opal (Indiana), and Silicon Graphics, and using UltraNet to provide localised high-throughput I/O, was started in mid-1990, and became operational for the 1991 run.
- Over a few years this architecture matured and became CERN's default batch processing architecture. It then naturally developed into the PC farms that we can see today..

HOLDING THE DESKTOPS TOGETHER
From CERNVM to NICE and PLUS

The dual role of the mainframe

- In 1965 a computer was mainly a powerful calculator
- By 1975 a mainframe processed data
- And by 1985, in addition to that role, it had started to provide the "desktop environment"
- Initially mainly for dumb terminals connected via Index

- But the number of users, and their needs, would grow fast:-
  - Offsite access
  - Spread of LANs onsite
  - Spread of computer use
  - Experiments wanting to work collaboratively (email, sharing programs and data)
  - Hence, wish to see "uniform logon environment"

A growing number of users

- This "uniform environment" was provided *de facto* by CERNVM, and the number of users grew rapidly, even as the use of PCs and Macs took off
- It reached a peak of ~4200 users/week in 1992-93
- Eventually the CERNVM service was turned off, after sterling efforts by many people to make a migration possible, in the middle of 1996

Weekly interactive users 1987-2000

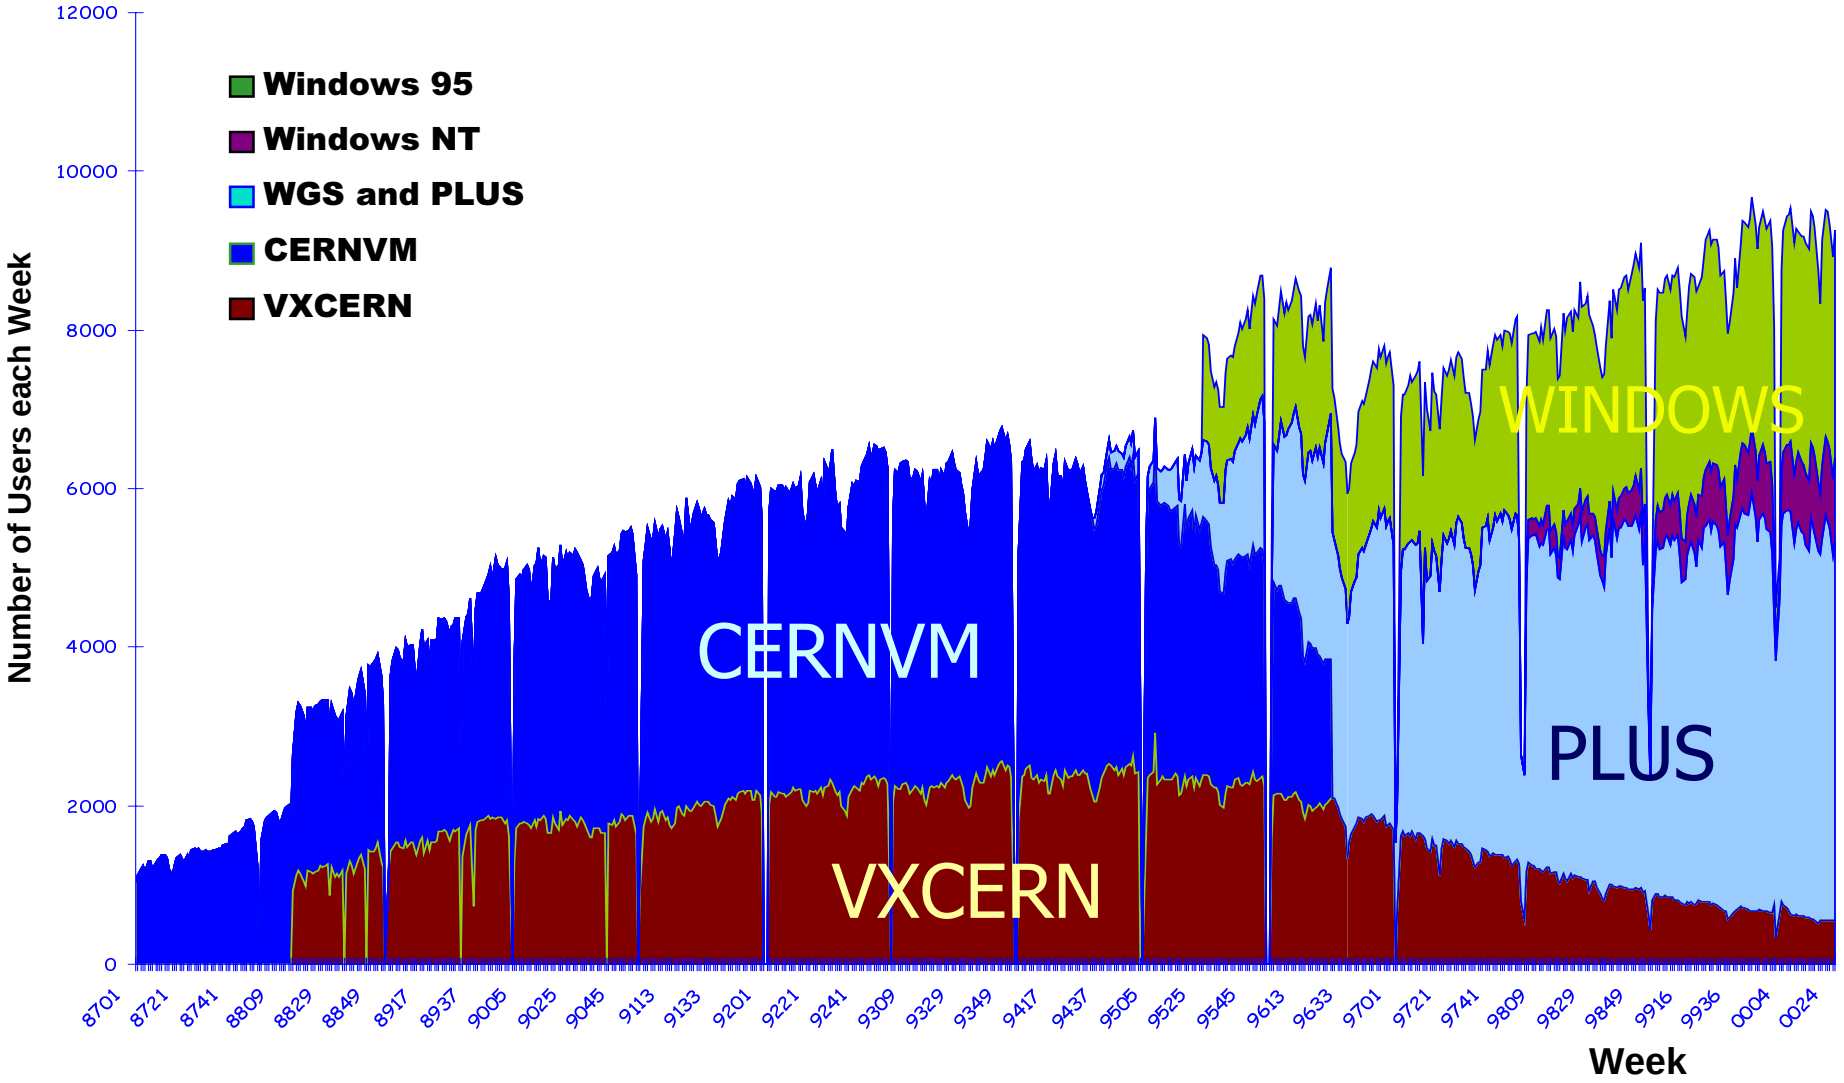

NICE

- For people working on PCs, the NICE service, based on Novell, had provided many of the facilities needed by a "standard desktop"
  - "standard" applications
  - shared file system, able to reflect division/group/service structure
  - mail and messaging
- NICE was a very advanced concept for its time and we probably missed a chance for a significant piece of technology transfer here

PLUS

- Public Logon Unix Servers
- The provision of similar (to NICE) services for the Unix world was harder
- Harder because of the need for heterogeneity, and because there seemed to be even less commercial interest
- But it was successfully done, and enabled CERNVM to be switched off, and bring the era of the mainframe to an end at CERN
- Though the change was inevitable, and beneficial, it has proved to be hard work under the covers to retain the degree of coherence that the mainframe's "single system image" provided

SOME TECHNICAL ASPECTS

Networking

- Before we started working on LEP, we were not networked
- Now that we reach the end of working on LEP, we cannot imagine what HEP computing without networking could mean
- CERN's main network connectivity in 1984/85 was a 64 kbps X.25 connection to the Swiss PTT's Telepac service
- Mainframe to mainframe connectivity (IBM to IBM, and later to VAXes) also came via EARN

- Now we have a total of hundreds of Mbps of connectivity available to France, to Switzerland, to Europe, to the USA, and to many other destinations.

Network protocols (i)

- In the early days of LEP computing (83-87) there were big battles about protocols
- To encourage competition many people wanted to avoid "proprietary" protocols
- And they pushed the concept of ISO's Open System Interconnection model

- In the real world of fast moving IT, OSI lost out to TCP/IP, which formed the basis of the Internet

- TCP/IP was very "non-proprietary"
- And it is a pity that so much time and effort was wasted painting it as "American" rather than "European"

Network protocols (ii)

- And I will not detain you with the Cheapernet/UTInet saga

e-mail

- Internal (onsite) e-mail came into use in the late 70s
- As far as I can ascertain the reasonably general use at CERN of offsite e-mail dates from 1983, following the connection of Wylbur via the PRIAM Vax to the world of uucp mail
- I have stated, and I do not think that it is an exaggeration, that we could not have built the LEP experiments without e-mail
- I must also say that I am full of admiration for the people who have worked on CERN's e-mail services over the years. Since we all use e-mail, we all think that we understand it. It is, in fact, a fearful technical challenge at this scale.

3480 cartridges

- The MUSCLE report (January 1988) recommended that 3480 cartridges should be used for the storage of LEP data
- There were times in late 1988 and early 1989 when this looked to be a rash proposal, since it required Digital to deliver 3480 drives connected to Vaxes. And that was a project with delays which increased by one week per week
- But everything was finally "all right on the night"

- Data volumes
- ALEPH wrote 1.5 Tbytes of raw data at LEP1 and 2.1 Tbytes of raw data at LEP2.
- Or roughly 20,000 cartridges (of 200 Mbytes capacity)
- 20 Gbytes is now a standard cartridge capacity

3480 units

3480 cartridge (200 MB)

STK silos in 2000

VAX-VMS

- "The" key workhorse computer of LEP data handling.
- All experiments used them at the heart of their online system, and farms, and often built them into crates
- And VXCERN was at the basis of much off-site connectivity, especially e-mail and file transfer
- And many outside institutes and universities had VAXes
- Clustering, tight and loose, DECnet, good support for Fortran-77 and other languages - including excellent debugging were key factors. Also the whole VAXstation line.

- In the end another convincing proof that proprietary systems, however excellent, are doomed.

VAX Area in 1993

Code management systems

- PATCHY
- Historian
- CMZ
- CVS

- Names to reminisce about
- I decline to try to draw any lesson

Distributed file systems

- CERN started its AFS service in May 1992
- AFS = Andrew File System (from Carnegie-Mellon Uni)
- Heterogeneous, wide-area shared files
- Never really accepted into the main stream of manufacturer's OSs
- But, as the effective bandwidth and reliability of international networks improved, AFS proved extremely valuable for software distribution and development among the LEP collaborations
  - "from the perspective of the person involved in the development, debugging and maintenance of our huge code the most important technical breakthrough which happened during this time was the advent of AFS"

Disk space

- The move from mainframes to workstations was not the only sea-change which affected LEP computing
- Disk technology changed just as dramatically
- The MUSCLE recommendation (January 1988) was that regional centres needed *at least 10 Gbytes of disk*
- MUSCLE also discussed the desirability (for faster analysis) of each experiment having access *100 Gbytes of disk* at CERN
- In the year 2000, each LEP experiment has access to *~2.5 Tbytes of disk*

SOME COLLABORATIVE SUCCESSES

For the HEP community

Starting with some CERNlib “products”

Delivering software

- No longer duplicating a card deck or magnetic tape!
- Over the network
- To the workstation
- Configured to some agreed specification
- The world of CERNlib, ASIS, AFS, CVS, SRT, ….
- The steady improvements here helped to deliver two really major packages which enabled LEP physics...

PAW

- Discussed in MEDDLE at the September 1986 meeting (first meeting with John Thresher present as responsible Director)
- Towards a Physics Analysis Workstation (December 1985)
  - Bock, Brun, Pape and Revol
- As the world moved to graphics and interaction, trying to bring "everything" together to create a convenient interactive interface for the physicist doing analysis
- A real *de facto* standard and huge success

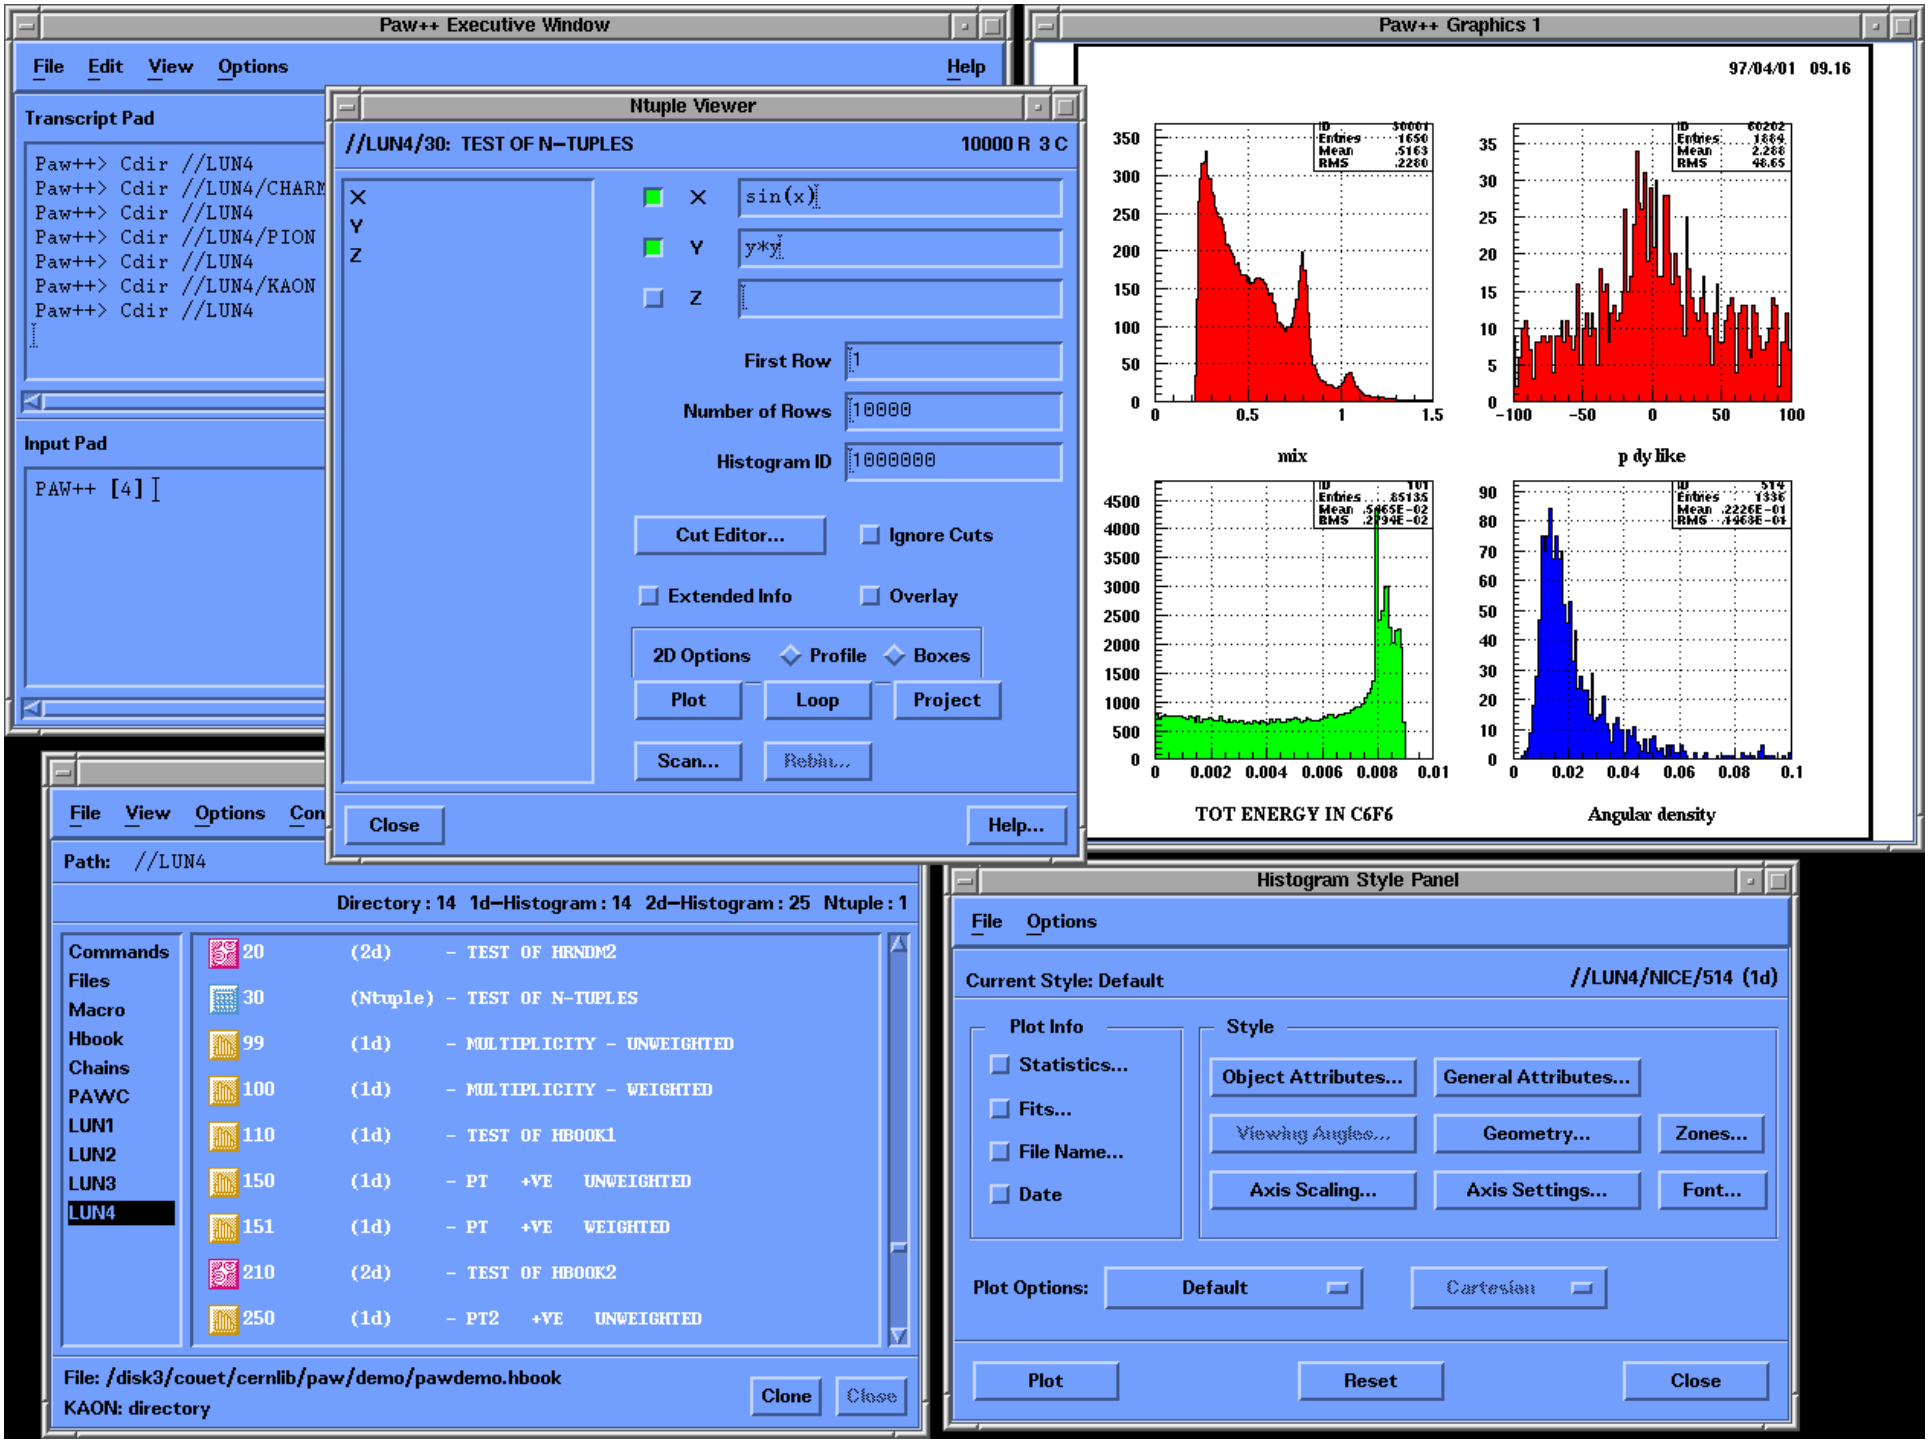
Paw++ Executive Window
File Edit View Options Help
Transcript Pad
Paw++> Cdir //LUN4
Paw++> Cdir //LUN4/CHARM
Paw++> Cdir //LUN4
Paw++> Cdir //LUN4/PION
Paw++> Cdir //LUN4
Paw++> Cdir //LUN4/KAON
Paw++> Cdir //LUN4
Input Pad
PAW++ [4]
Ntuple Viewer
//LUN4/30: TEST OF N-TUPLES
10000 R 3 C
X
Y
Z
X sin(x)
Y y*y
Z
First Row 1
Number of Rows 10000
Histogram ID 1000000
Cut Editor...
Ignore Cuts
Extended Info
Overlay
2D Options Profile Boxes
Plot Loop Project
Scan... Rebin...
Close
Help...
Paw++ Graphics 1
97/04/01 09.16
mix
p dy like
TOT ENERGY IN C6F6
Angular density
File View Options Con
Path: //LUN4
Directory : 14 1d-Histogram : 14 2d-Histogram : 25 Ntuple : 1
Commands
Files
Macro
Hbook
Chains
PAWC
LUN1
LUN2
LUN3
LUN4
20 (2d) - TEST OF HRNDM2
30 (Ntuple) - TEST OF N-TUPLES
99 (1d) - MULTIPLICITY - UNWEIGHTED
100 (1d) - MULTIPLICITY - WEIGHTED
110 (1d) - TEST OF HBOOK1
150 (1d) - PT +VE UNWEIGHTED
151 (1d) - PT +VE WEIGHTED
210 (2d) - TEST OF HBOOK2
250 (1d) - PT2 +VE UNWEIGHTED
File: /disk3/couet/cernlib/paw/demo/pawdemo.hbook
KAON: directory
Clone
Close
Histogram Style Panel
File Options
Current Style: Default
//LUN4/NICE/514 (1d)
Plot Info
Statistics...
Fits...
File Name...
Date
Style
Object Attributes...
General Attributes...
Viewing Angles...
Geometry...
Zones...
Axis Scaling...
Axis Settings...
Font...
Plot Options:
Default
Cartesian
Plot
Reset
Close

GEANT3

- The basis of much (not all) of the main LEP simulation work
- A series of releases aimed at increasing functionality and reliability …
- Of which 3.14 (early days), 3.15 (end ‘91) and 3.21 (Spring ‘94) were some of the most important for LEP
- Effort then evolved/switched to GEANT4 (OO)

OPAL
GEANT 3.21

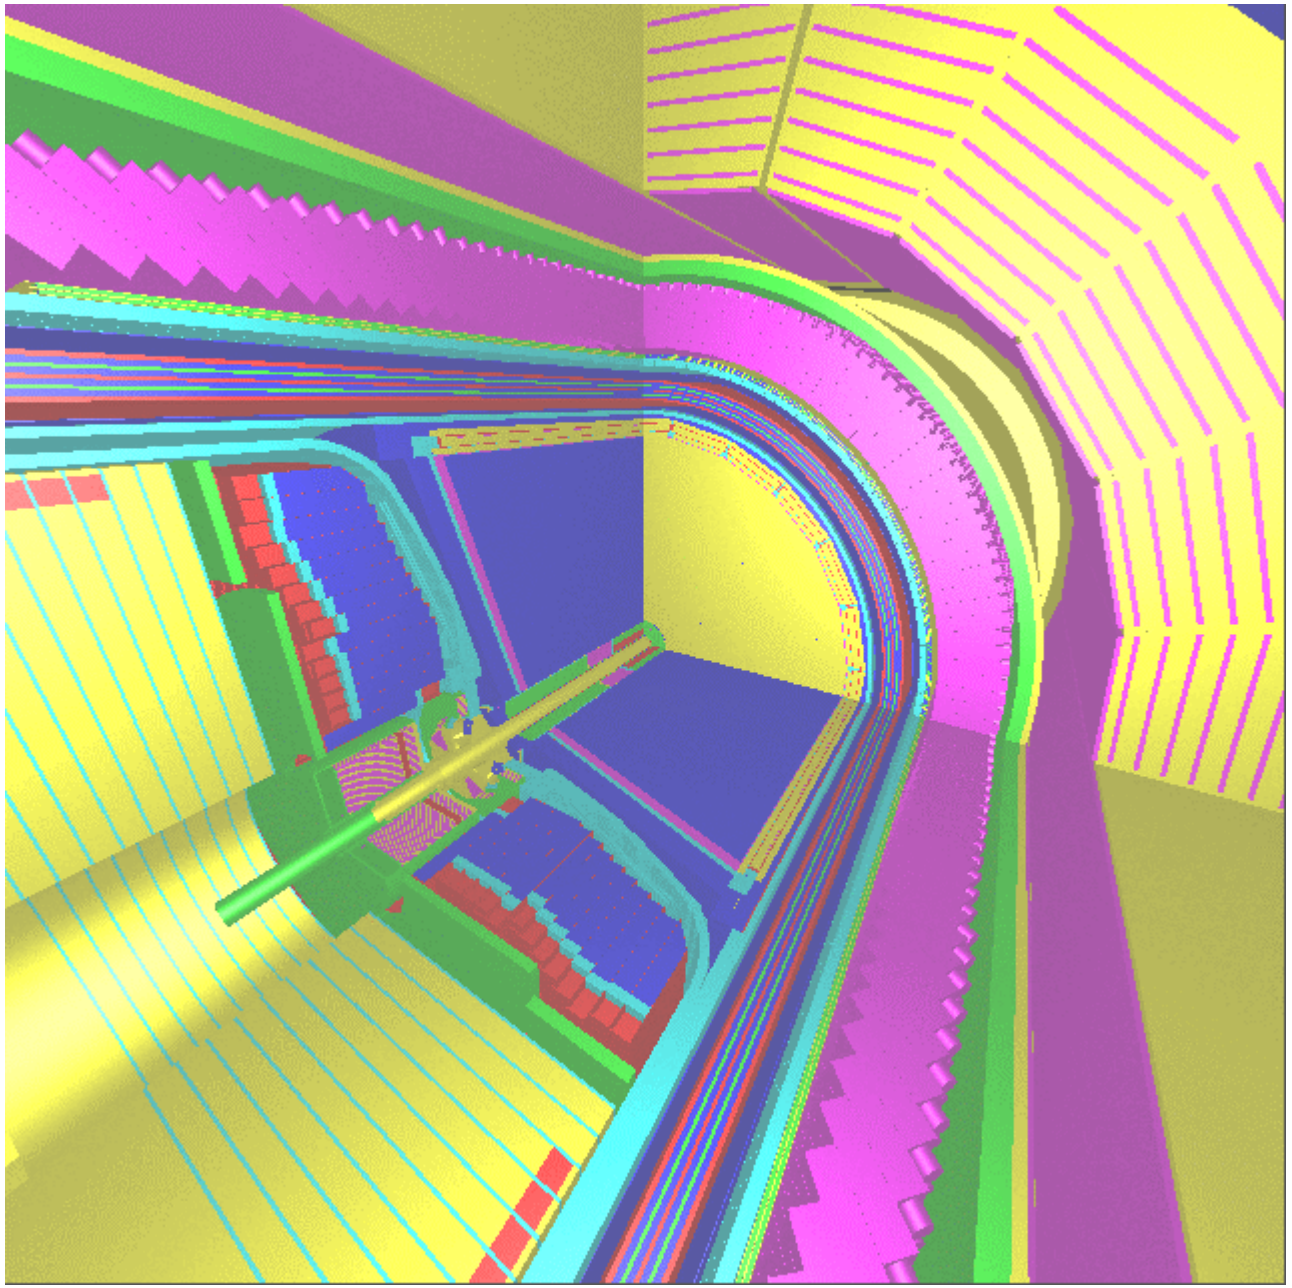

FATMEN

- Requested by the LEP experiments at MEDDLE at the end of 1988
- File And Tape Management - Experimental Needs
- Mapping a user name, such as Opal/Data/Run123/Muons, to a file on disk or on tape
- Was in rather extensive use by the end of 1989
- An early example of a "just in time" project!?
- Which sometimes needed the author to patch the server!

- A fundamental part of the LEP processing model
- And an area where we still need more work in the future (relation to grid model)

CHILDREN of CAMAC

LEP used two main hardware bus systems, FASTbus and VMEbus

HEP contributed a lot of technology transfer to industry in both cases

VMEbus

- Heavily used already by UA1
- Serious HEP contribution to the standard making process (and technology transfer in general)
- Industrial control systems
- LEP experiment control systems

- A bit small for front-end electronics
- So you either used Fastbus or some “non-standard” boards

FASTbus

- Form factor for front-end electronics
- Many/most of the original boards still in use today
- It was never the success that CAMAC was
- Most people feel that it was essential for LEP data acquisition

THE BIGGEST SUCCESS


© Sam Ogden
All Rights Reserved

WWW

- A great achievement
- A fantastic idea, at the right time

- I still find it interesting that neither industry nor academia got there first
- It proves something about the benefits of assembling together urgent needs, infrastructure and smart people, and letting them interact..

- And why it is exciting to work at CERN, and in computing
- And why we should not always listen to wise people who tell us that industry will always do better than we will….

OTHER PEOPLE'S VIEWPOINTS

As I said at the start, we all look at LEP computing from different viewpoints.

What about experiments and outside labs?

As seen by experiments (i)

- Online DAQ and monitoring
  - all experiments did a good job
  - often helped by central support, but less important than for several earlier generations of experiments
  - Digital sold a lot of systems - but had a good Joint Project team at CERN for many years
  - a lot of 680x0s and VAXes were built into intelligent front ends
  - all DAQ systems and all slow controls were significantly upgraded over time
  - and many modern software techniques (extensive use of C++ and CORBA) were adopted in that process. No time for OO kangaroos.
- Quasi-real-time reconstruction
  - FALCON proposed in February 1988
  - 6 CERN units of VAX 3100 (plus a big IBM disk)
  - Invaluable for running the experiment and quickly the *standard model*

As seen by experiments (ii)

- Writing the offline programmes
  - good but could still do better? And even better tools would help
  - too much duplication?
  - too little design??
  - I had better keep quiet
- Developing and operating the production systems
  - FATMEN and TMS were progress
  - But a lot of effort and ingenuity - little shared - went into this area
  - Where investment in grids will hopefully pay off for LHC

As seen by experiments (iii)

- Analysis from outside
  - By the time we got to LEP2 we can perhaps claim that it was feasible for a graduate student or post-doc based "at home" in much of Europe or the USA to take a serious part in an experiment's analysis while travelling rather little to CERN
  - A good step forward, even if still not perfect
  - And even if this does not work well yet for many other countries

As seen by outside labs (i)

- Did CERN dominate too much?
  - Is the glass half full or half empty?

- HEP-CCC
  - The High Energy Physics Computing Coordination Committee was created as long ago as 1982/83
  - With "director-level" representatives from several European countries, CERN and DESY. Plus representatives from ECFA, USA, and now perhaps other regions
  - Discussion but few decisions. Necessary but surely not sufficient.

As seen by outside labs (ii)

- Regional centres
  - France has a very strong centre at IN2P3, Lyon.
  - During the LEP period the HEP use of the Rutherford facilities has tended to decrease.
  - There has been very good network connectivity to Italy, and the internal INFN connectivity is also good. There is no single Italian regional centre, but it is conceptually a "distributed INFN centre".
  - Germany has it federal structure. My personal view is that it is a pity that DESY does not play a strong role as a regional centre for CERN computing
  - Several smaller centres have played strikingly important roles in LEP analysis
  - One that I know about personally was Florida State (ALEPH reprocessing runs) but I also know that there were many others

THE LESSONS

To remember at LHC?

Fifteen years is a long time (i)

- Moore's Law predicts 1000x times over 15 years
- Everything will change
- Far more than you ever imagined
- You will use things that are not yet invented, or even thought about
- LHC computing will belong to the wireless age. Always on and on everywhere. We will look back on LEP computing as the stone age.

- You will replace most of your original software over time
- To help that process we must design in layers, and keep things as modular as possible
- Good architecture; Good design; Good implementation

Fifteen years is a long time (ii)

- Moore's Law helps computing to get cheaper per unit of processor power. But IT continues to become more pervasive in science and society, and that is likely to continue during the LHC era. This produces a need for increased investment, and we did not take enough account of that factor during the LEP period.

- The real challenge is getting people to collaborate!

- CERN accelerators, sooner or later, exceed their design luminosity!!

Particle physics is a fertile source of IT challenges (i)

- Sometimes we are too arrogant
  - Claiming that our requirements are so special, and that only our special HEP solutions are appropriate
- Sometimes we are too defensive
  - Claiming that industry has left us behind, and that our old in-house approaches to writing software (for example) are very dated
- The reality is still that we are one of the world's leading sources of scientific data
- With a truly distributed global user community
- And that our needs challenge today's reality
- And that our approaches in several areas, including the development of large software and data grids, are at the very leading edge

Particle physics is a fertile source of IT challenges (ii)

- We need to understand better the capabilities of industry and the computer science community
- Watch what is happening in crucial areas
  - Storage and access in the broadest sense
  - Much higher bandwidth
  - Different wide-area network architectures
  - Construction and evolution of large software packages
  - All leading to Data Grids
- And collaborate more with these people - we all have a lot to win
- We have been saying this for almost twenty years now, and not acting seriously enough!

There is no substitute for enough resources

- Especially people
- And especially if we are supposed to help transfer our know-how to other sciences, and even to industry

LEP analysis will not stop with the beams

- People will want access to LEP data for many years yet
- And I guess that for at least the next years there will be very intense scientific activity
- So the computing environment supporting that access had better not disappear too quickly either!

PC farms in 2000

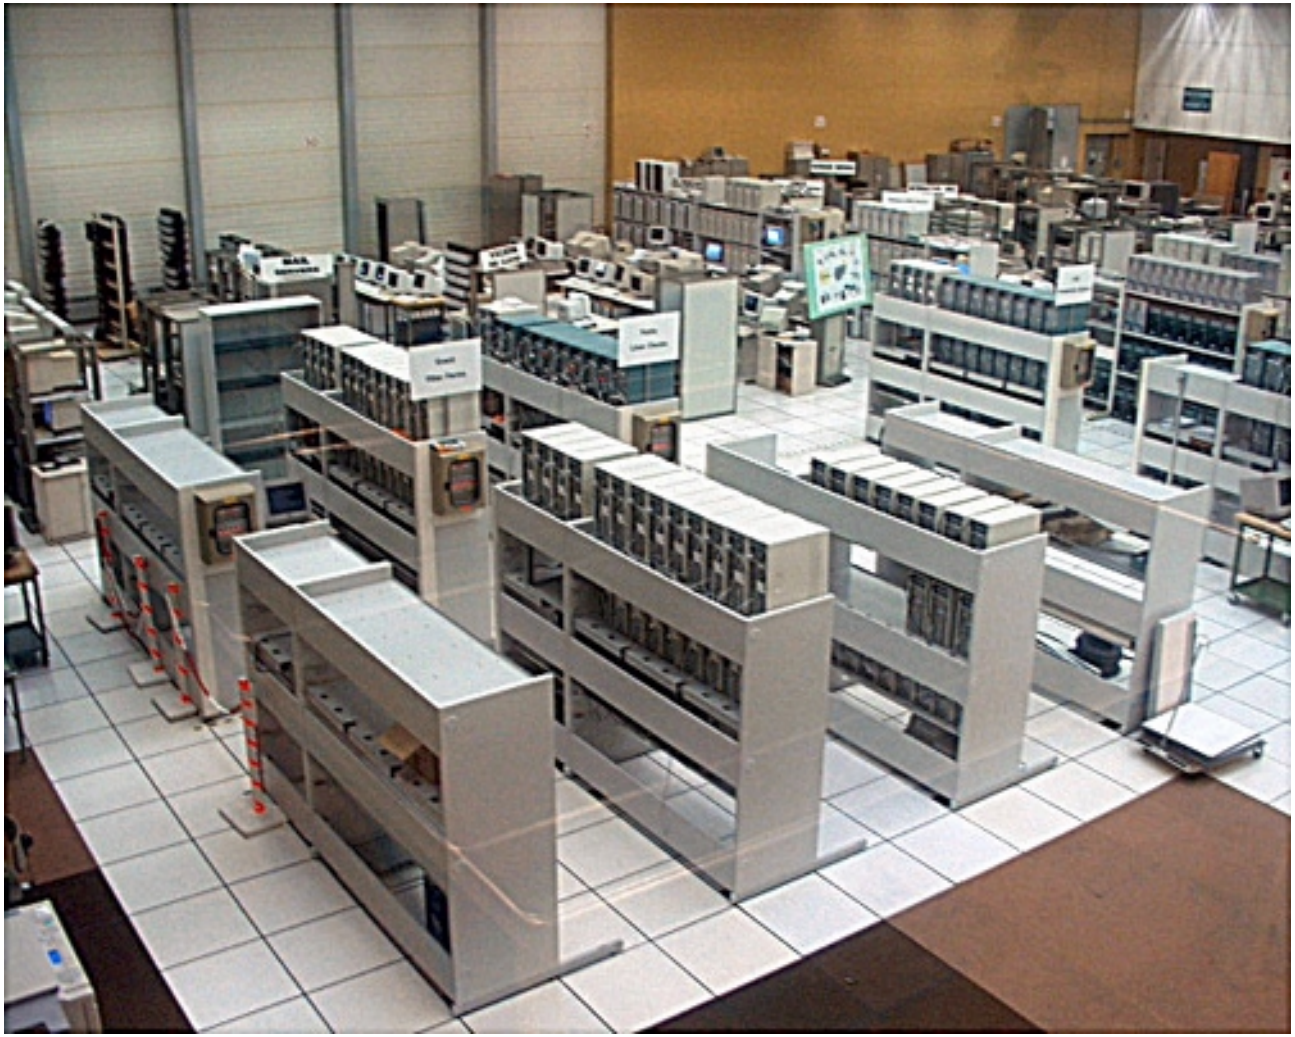

Computer Room in 1985

My final remark

- It needed many people to achieve all that has been achieved with LEP computing
- They are the people that we should be expressing our thanks to today