Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

LEP to an fcc-ee – A Bridge Too Far?

Introduction

This paper attempts to précis key documents, presentations and papers that are still available concerning the preparation for, and execution of, (mainly offline) computing for the Large Electron Positron collider (LEP) at CERN.

The motivation is not only to capture this information before it is too late – which in some cases it already is – but also should it be of interest, even if only anecdotally, in the preparation for future, "similar" machines. Specifically, as a result, the interested reader should be aware of what electronic documents are available together with direct pointers to be able to find them.

As further motivation, the 2020 update of the European Strategy for Particle Physics contains the following statements:

"The vision is to prepare a Higgs factory, followed by a future hadron collider".

"Given the unique nature of the Higgs boson, there are compelling scientific arguments for a new electron-positron collider operating as a Higgs factory".

[Pre-amble]

"An electron-positron Higgs factory is the highest priority next collider".

[High-priority future initiatives]

The document also states:

"Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation."

Therefore, a summary of the current status in these areas, particularly with regard to the former electron-positron collider (LEP) is called for.

In 2020, we are (were) at approximately mid-point between the end of data taking at LEP and the possible start of data taking at a future *fcc-ee* (2039?). Aside from a small level of data loss, the vast majority of the data from the LEP experiments is not only still available and usable but also analyses continue at a

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

low level from all 4 experiments (ALEPH[1], DELPHI, L3 and OPAL for those recently arrived from another planet. Welcome!). However, a significant amount of information from this era has already been lost, sometimes through in-action but also through policy, as deplored in the document "Software Preservation and Legacy Issues at LEP" (https://zenodo.cern.ch/record/2653526).

In this document we summarise key elements from pre-LEP planning documents (Computing at CERN in the LEP era, Computing at CERN in the 1990s and the MUSCLE report) – all of which have been scanned and stored electronically in the CERN document server as part of an exercise circa 2015 to capture as much of the documentation related to the CERN Program Library (CERNLIB) that was no longer available online (see the "Software Documentation" collection in the CERN Document Server).

We then turn to papers presented at the Computing in High Energy (and Nuclear) Physics conference series (CHEP[2]) to summarize "tip of the ice-berg" status reports regarding computing for, and by, the LEP experiments.

Finally, we summarize some of the work done in the context of the DPHEP[3] Study Group and (later) Collaboration regarding the preservation of LEP (and much other) data and raise some possible issues for the future.

PREPARING FOR LEP

Discussions on the long-term future of European High Energy Physics took place in the late 1970s, although not part of a formal "strategy update" as we know it today. Despite a break in tradition with proton accelerators and colliders at CERN, an electron-positron collider was considered the most appropriate machine at that time with which to advance our understanding of fundamental interactions. The Large Electron Positron, for some reason always written e^+e^-, collider project was approved in 1981 and entered into operation already in 1989 – a similar time period that the Future Circular Collider strategy update documents estimate for a possible *fcc-ee* machine. Much has been written about

---

[1] At the time of writing, it is relatively easy to find someone from the former ALEPH collaboration, as many still proudly display an ALEPH poster on their office door. To get a flavour of what life was like inside such a collaboration, try reading "The ALEPH Experience": https://cds.cern.ch/record/897514/files/cer-002567958.pdf.

[2] The proceedings from these conferences, widely distributed in hard copy and later also digitally, are already scarce, particularly concerning the LEP data taking period. More detailed reports, such as notes internal to the individual experiments, may be found in personal archives but are more likely, in the general case, to be lost or destroyed.

[3] DPHEP is often understood to be simply "data preservation in HEP", whereas the full title emphasized data preservation **FOR LONG-TERM ANALYSIS**.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

the engineering, detector and physics challenges of the LEP machine – here we focus on the computing (and primary offline) aspects.

A series of significant studies into LEP-era computing were undertaken, specifically:

1. Computing at CERN in the LEP era (May 1983 – available at https://cds.cern.ch/record/98880/files/)
2. The MUSCLE report: The Computing Needs of the LEP experiments (CERN Data Handling Division DD/88/1 https://cds.cern.ch/record/184457/files/CM-P00059881.pdf)
3. Computing at CERN in the 1990s (July 1989 – https://cds.cern.ch/record/206085/files/Computing%20at%20CERN%20in%20the%201990s.pdf)

As a historical note, the first of these reports was scanned semi-destructively, i.e. a hard copy was taken apart (it was stapled by some sort of industrial stapler), scanned by section using the photocopier in B513 at times that did not annoy other users of the machine too much, and then merged into a single PDF file. The second, relatively slim, document was disassembled and scanned via the hopper whereas the final, being bound, was scanned page by page. Errors (mis-scanned or missing pages) may have been introduced as a result. Whether the source files of any of these documents remain is doubtful: the MUSCLE report being marked up, if memory serves correctly, using VAX Document – an SGML-like markup language – that may therefore exist on some archive tapes but without the necessary software to process it, even if found.

As a general remark, it is interesting to note that these documents were prepared by relatively small groups of people, particularly when compared to LHC-era studies[4], such as those undertaken by RD44, RD45, RD47, MONARC etc. This is no doubt partly due to the dramatic increase in the capabilities of collaborative tools, even though a significant amount of international travel was also involved for LHC and HL-LHC era studies and related reports.

Computing at CERN in the LEP era

This document – also known as "the green book" – was published in May 1983. It is divided into 6 parts:

1. A cover letter for the Director responsible for Computing (Erwin Gabathuler);
2. Report of WG1 – data acquisition and monitoring;
3. Report of WG2 – networks and telecommunications;

[4] Indeed a future "LHC historian" may wish to unearth some of the work from these projects, their websites already being inaccessible since many years but some status reports existing in CDS, with more information, amusingly, at mirror sites outside CERN.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

4. Report of WG3 – general computing seervices (sic) and facilities;
5. Summary of the final report of the steering group for microprocessor standardization.

The cover letter proposed a possibly biennial workshop of those involved so that the recommendations of the document could be regularly updated – a sentiment that one can be sure was most welcome to those who had contributed to this work. Whilst the other sections are (highly) unlikely to be of direct relevance to any future machine, they support the well-known adage that predicting is difficult, especially when it involves the future. Furthermore, they paint the background against which the initial papers regarding computing for LEP can be understood.

It is worth noting how the LEP Computing Planning Group (LCPG) worked, with a total of 17 meetings and 3 working groups (listed above). WG3 on general computing services, held an open meeting to solicit user input some 6 months after it had been setup, with a joint meeting between WG2 and WG3 some months later to invite comment on their draft reports. It is unlikely that such a seemingly closed approach would be acceptable today, let alone well into the future, but it was perhaps unavoidable given the somewhat primitive collaborative tools of the time and being well before the arrival (and departure) of cheap travel.

The LCPG made a total of 22 recommendations, the full text of which can be found in the scanned document. They covered the following areas:

1. General network architecture;
2. Basic communications infrastructure;
3. Standardisation of high level protocols;
4. External Particle Physics Network;
5. CERN Backbone Network;
6. Local Area Networks; (by now one should be getting a feeling of just how much of a concern the whole issue of networks and connectivity was)
7. Terminal connections;
8. Data acquisition;
9. Support Services (CERNET, Electronics Pool, Online Computer Pool teams);
10. "Private" computer centres;
11. Central batch capacity;
12. Remote print stations;
13. Peripherals for Central Computers (mass storage system, tapes and disks);
14. File server;
15. Computing for engineers;
16. Interactive computing (!);
17. Personal Workstations;
18. VAX computers (my entry into CERN);
19. IBM interactive computing – evaluation of VM/CMS;

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

20. Interdivisional Cooperation;
21. Collaboration with Outside Institutes;
22. Data Acquisition – Detailed Planning.

To simplify into just three domains – computing in general, storage and both on-site and off-site networking – it is striking how much uncertainty there was, how much heterogeneity and how daunting the challenge of LEP computing must have looked at that time. For example, it was believed that data recording would be made to open reel 6250 bpi tapes[5] – similar to those that one sees on old movies that rotate forwards and backwards making "*gloop gloop*" noises. Apart from the relatively unreliability and lack of practicality of these tape volumes, the sheer number that would have been required would have been of the order of one million – vastly exceeding the physical storage limits of the CERN computer centre and leading to some older tapes being stored in the former ISR tunnel (where they rapidly decayed and became unreadable).
Turning to computers, the role of mainframes versus emulators versus Personal Workstations was very unclear – and remained so for a decade or more. To quote directly from the document[6]:

"CERN lacks experience of personal workstations such as PERQ and APOLLO. … Support the purchase of such systems for evaluation during 1983".

Later on we read:

"… it is not just the top of the market (PERQ and APOLLO) which could be of interest, but also low end (DEC, IBM, and why not APPLE?)".

Why not, indeed.

Whilst the optimal (short-term?) strategy for interactive computing was not clear at that time (VM/CMS on IBM and compatible mainframes, VAX/VMS systems and / or personal workstations, the report, perhaps unwittingly, presaged something that be a feature of LEP-era computing, namely regular, and often painful, migration from one platform to others.

Offline Data Processing Facilities

An aspect that was covered only very briefly – so briefly that the entire section is reproduced below – is that of offline data processing.

---

[5] *"but the door should be left open for new products such as video disks".* It still is…
[6] Although this section talks about "personal computers", it seems doubtful that this referred to "PCs" of the time, such as the IBM PC, the DEC Rainbow and so forth. The glossary of this section describes a Personal Workstation (sufficiently cheap as to be dedicated to a single user but differentiated from a Personal Computer (PC – abbreviated not used in the text) by its greater power.

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

In the LEP context, support for offline data processing will be most meaningful in the area of well defined support packages, of which the GEM memory management (system?) *is an example, and in the area of well defined support facilities, of which the Patchy source code management system is an example. Other areas which are growing in importance for physicists, and where a concentrated support effort might be needed in order to provide the required facilities, include advanced interactive graphics* (what other kind is there?)*, database management, and user transparent network interfaces.*

Owing to the size of the LEP collaborations, standard practices will have to be adopted in many areas of the data processing. These standard practices will either involve the use of accepted tools, or the adoption of a set of clearly stated rules. It will be very important that the DD and EP divisions participate in LEP experiments and that they foster effective collaboration between the experiments and also among the whole European particle physics community.

This possibly explains the author's distaste for naming specific solutions, particularly when talking about the future. Whereas Patchy was indeed used throughout the LEP era, and even beyond, the whole question of "proprietary" source code management systems was an area of hot debate during the same period. GEM – a system so complicated that the well-known quote regarding general relativity[7] could equally well have applied to it – never saw production status. Instead, two former packages in the same domain, namely Hydra and ZBOOK, were superseded at CERN by ZEBRA, with other largely similar systems at various laboratories around the world. ZEBRA was used by many of the larger packages in the CERN Program Library of the time (CERNLIB), including HBOOK, PAW and GEANT.

Data Communications

Data communications of all sorts – both on and off site – was an area of considerable concern and uncertainty. No obvious "winner" was in sight at the time and the phrases "Cambridge Ring", "Token Ring", "Token Bus" and even the word Ethernet could often be heard during animated conversations in CERN's Restaurant 2. The report cites a number of high-level Use Cases, such as file transfer, e-mail, remote computer access, for which an agreed set of high-level protocols was required. Even as LEP computing went "live", there was still considerable divergence in this domain with multiple "proprietary" networks (DECnet, IBM's SNA, EARN and BitNET, Apollo Domain etc.) and both mail and file transfer gateways between them. X.25 is cited at various points in the report as (at least) a strong candidate for wide area networking in the LEP era. But a

[7] *Professor Eddington, you must be one of three persons in the world who understands general relativity'*... And Eddington's reply was, '... I am trying to think who the third person is!'
Source https://www.physicsforums.com/threads/only-three-people-understand-general-relativity.935624/

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

dark horse had entered the race, even though it took well into the LEP data-taking era for it to pull into the lead, as we shall see in more detail later.

To be fair, the report did acknowledge these issues:

"Since, even in the medium term (5-7 years), we will have to live with the present multiplicity of protocols, we believe that the only way of implementing successfully access to services across networks (interworking) will depend on translation between protocols. We must point out that these problems of protocols and interworking are probably even more severe for the external connections than for those on the CERN site... As far as possible, common solutions should be implemented for off-site and on-site connection problems."

Given these concerns, it is perhaps surprising that the report states earlier on:

"The outside networking will be based on the X.25 public packet switching services."

Human Communications

Although this was not studied as a specific topic, we see more and more hints of the need for collaborative work across CERN and outside sites. The famous 1/3 – 2/3 rule is mentioned, which is still in effect in 2020 (with the sum of resources at each Tier in WLCG being approximately the same, thus 1/3 Tier0, 1/3 the sum of Tier1s and 1/3 the sum of Tier2s). Not only is electronic mail mentioned at several points in the report, but with no clear solution (a dedicated e-mail machine perhaps?) but also the need for outside collaborators to be able to access / download the latest source code "correction sets" and / or calibrations in a "reasonable" (minutes?) period of time. At a time when many of us are "always online" it is hard to imagine / recall a period when the opposite was true. Indeed, communications between sites were still often carried out using Telex machines (you wrote your message and gave it to a Telex operator, who then mis-typed it and sent it off) that were only in the 1980s being replaced by (tele)Fax machines. Great excitement!

Conclusions on the "Green Book"

The report clearly contains much more detail than can be summarized above and reflects even more work and thought that went on "behind the scenes". Unfortunately, much of this information has been lost – some to be obtained in hard copy from secretariats that have not existed for decades, others stored on the IBM MVS system (possibly to be found on archive tapes for anyone with sufficient perseverance). Even that information that is still available – namely in the report itself – assumes a "knowledge base" that is rapidly decaying, such as the structure of the CERN Program Library at the time (never explicitly laid out), the debates on programming languages (not everyone believed that it would always be Fortran, and so it turned out) and many other aspects of CERN's culture at the time.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Whilst the report made a number of technology predictions and recommendations that turned out to be incorrect, the broader strokes, such as the need for collaborative working, are still equally valid today.

Some of the explicit cost estimates included in the report – such as 320KCHF for 8MB of mainframe memory – maybe of interest to those preparing today and tomorrow's budgets. Every aspect seemed to present an almost insurmountable challenge at the time.

In summary, it was certainly the intent of the authors and contributors to do absolutely the best job that they possibly could, under the circumstances. And so they did.

The MUSCLE Report

The MUSCLE report, based on the Meeting(s) to Understand the Specific Computing needs of the LEP Experiments, is a slim document, bound in red for a change, of just 42 pages. Before diving into details, it is interesting to contrast this with, for example, the (W)LCG[8] Memorandum of Understanding and / or Technical Design Report. Its scope is much less – it does not talk about detailed services and responsibilities, associated service levels, availability and so forth. However, the first chapter of the LCG TDR is even shorter – a mere 30 pages – but somewhat narrower in scope. Perhaps, if nothing else, a MUSCLE-like structure for a future *fcc-ee* would be sufficient?

The document structure is repeated below – the full report should be consulted for details.

Chapter 1 – Introduction
Chapter 2 – General elements of LEP computing
Chapter 3 – Data storage
Chapter 4 – On-site links
Chapter 5 – Estimating the requirements for LEP computing
Chapter 6 – Computing outside CERN
Chapter 7 – Distribution of the computing activities
Chapter 8 – The impact of "private" mainframes
Chapter 9 – Resources
Chapter 10 – Recommendations
Appendix – Comparisons with other experiments (e.g. UAn, HERA)

Some details from the chapter on recommendations are worth repeating as they may appear familiar to anyone involved in LHC Computing.

---

[8] If we have to talk about acronyms, I have tried to explain the confusion between LCG and EGEE (look it up) as follows. If you write "LCG" as a name, it might look like "Elsie Gee". So, E.Gee. *Quod erat demonstrandum*.

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

"Many of us… are only now starting to understand the enormity of LEP computing".

"Physicists should not find themselves in a situation where they have to come to CERN to obtain adequate computing resources".

The latter led into recommendations that could be read as recommending analysis and simulation centres. Also now known as Tier1s and Tier2s?

There were also specific recommendations regarding on and off site network links, as well as appropriate mass storage systems available to all of the central systems (Cray, IBM and compatible, VAXcluster).

This was left as an open question, as was the management of the many IBM 3480-style cartridges. Following (I believe) a vendor presentation[9] to a MEDDLE meeting a new task force was setup to look into these problems. Inspired by the naming of the MUSCLE report, this task force, that included many members from the MUSCLE team, is named the FATMEN report, bound again in green (See https://cds.cern.ch/record/196909/files/CM-P00059915.pdf). Further details of this work and other related activities can be found in "Databases in High Energy Physics – A critical review[10]": https://cds.cern.ch/record/1056959?ln=en.

The emergence of the IBM 3480 and subsequent cartridges with the (approximately) same form factor is one of a number of historical accidents that contributed to the continued availability of most of the LEP data until this date. Not only were these cartridges dramatically more reliable than the "open reel" tapes that they replaced but they were also much better suited to robotics. Increases in cartridge capacity were initially rather slow, starting at just 200MB, recent versions can store a few tens of TB. A cartridge capacity of 100TB – the approximate data volume of an individual LEP experiment for a cost of around $100 (not including the robotic and other infrastructure costs – largely born today by the needs of the LHC experiments) may even be reality (or old news) by the time you read this document!

---

[9] Faced with the challenge of making data available to the 3 main systems in the CERN Computer Centre at the time, one vendor proposed giving access on Mondays, Wednesdays and Fridays to one system, Tuesdays, Thursdays and Saturdays to another and quite possibly Sundays and public holidays to the 3rd. To quote the great Swiss musician and critic, A. Cortot from a different context, *"the utter childishness of this notion merits no further comment".*

[10] An earlier version of this work appeared in the proceedings of CHEP 2006 for which it was originally prepared. A very slightly extended version can be found in "*From the Web to the Grid and beyond : computing paradigms driven by high-energy physics*": https://cds.cern.ch/record/1141328?ln=en.

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

To this author, the MUSCLE report remains clear, comprehensive whilst also succinct. Although its detailed structure may not be appropriate in close to two decades time, it might well serve as guidance for a "Meeting to understand the specific computing requirements of the *fcc-ee* experiments".

Computing at CERN in the 1990s

Published only a short time after the previous report (on the auspicious date of 14 July 1989), it is much more comprehensive in scope. It also seems to have stood the test of time better than the previous "green book" that looks positively antiquated in comparison. Not only does it contain more concrete recommendations but also one can see its impact in the CERN of today (2020). Not bad for a 30 year old report[11]!

Given that the document was published at the time of LEP startup, we should have had a pretty good idea of what we were doing. However, there were still a number of statements or proposals that did not stand the test of (much) time, or else that developed in a way that was probably not foreseen when the report was being written. These will be covered in papers from CHEP conferences that we will cover next.

In addition to an Executive Summary, there were the following chapters:

- Computing for Experiments;
- CERN Data Networking Requirements in the 90s;
- Accelerator Computing Requirements in the 90s;
- Computing for Engineering at CERN ditto;
- MIS Computing ditto;
- Theory Computing ditto.

The report made a number of recommendations to "the CERN management". With the exception of then CERN Director General (Carlo Rubbia), the steering committee that included division leaders (aka "department heads"), senior physicists from inside and outside CERN as well as the CERN director responsible for computing (there were in fact 3 research directors at the time), could easily have been confused for the CERN management.

The 10 recommendations listed in the Executive Summary were paraphrased where felt necessary / appropriate:

[11] In addition to the version in the CERN Document Server, an electronic, and hence searchable version, was discovered at http://pcbunn.cithep.caltech.edu/jjb/ngb.html (now also in CDS).

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

1. Increase computing resources and review in the context of the CERN Scientific Programme;
2. HEP-CCC (more Cs would have easily been possible: Computer Centre Coordination Committee) should evolve … with a strong mandate to coordinate European HEP Computing;
3. Coordination of mainframe production resources across European sites;
4. "Private" computing facilities should be carefully controlled (but how?);
5. Major Components of the Computing Plan (10 sub-bullets itemized below – the fortuitous discovery of an online version of the NGB changes things!);
6. 5-year staffing plan (more, of course);
7. Alternative Staffing methods (ditto);
8. Training (see 2020 EPPSU!);
9. Advanced development projects (ditto);
10. The next steps that "we", the CERN management should define for the next 5 years.

The Key Elements of the Computing Plan

We recommend that the plan for computing in the next five years must provide for the following key elements:--

Long-distance networking

1. *Communications facilities* between CERN and the collaborating institutes *must be improved in order to facilitate decentralized physics data processing and analysis*. The bandwidth proposed for the major paths has been expressed as *one glass-fibre equivalent*, currently meaning 2 Mbits/s. (SRR 17, section 3.3.3)

An evolution to more Distributed Computing

2. It is widely believed that physics analysis will profit from *cooperative processing* between powerful graphics workstations and the central computing facilities. (section 3.2.5)
3. This alone will require a *substantial upgrade of the on-site networking*, especially in achieved bandwidth. FDDI products have an essential rôle here. (SRR 16, section 3.3.2)
4. Additional, improved *centralized support for workstations* will be necessary. (SRR 9, section 3.2.8)

Data Storage and Handling

5. The central computer services will need to put particular emphasis on *high data volumes*. This will have major consequences on the *handling, storage and copying of data*, which will also need to be reflected in the larger processing centres outside CERN. The estimates for the overall requirements of the LEP experiments in terms of storage

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

capacity and processing power contained in the MUSCLE report, have been largely confirmed by the working groups. From those estimates we make the following recommendation intended to cover *all CERN requirements in 1991/1992.*

- Make provision for sufficient *cartridges and storage space to handle at least 50 Terabytes of data* (250,000 cartridges at current densities). (section 3.2.6)
- Provide at CERN *an automated cartridge storage facility* capable of handling 8 Terabytes (40,000 cartridges). Similar facilities will also be necessary at the major outside institutes. (section 3.2.6)
- Provide a new *cartridge/tape copying and distribution service at CERN.* (SRR 5, section 3.2.6)
- Provide about *1 Terabyte of disk space at CERN* to allow fast access to the most active percentage of the data in a *central data repository*. (SRR 4, section 3.2.6)

Processor Power

6. The CERN physics programme will require at least 600 units [Footnote: The CERN unit of computing power is historically the power of an IBM 370/168 or a VAX 8600. As a rough indication, this may be taken as 3 IBM MIPS or 4 DEC VUPS.] of CPU power by 1991--1992.
    - At least *150 units* of batch processor power should be provided on general--purpose computers *in the CERN Computer Centre.* This is to be compared with the existing 70/75 units available for batch. This power is necessary not only for conventional batch work, but *in order that the centre can fulfil its increasingly demanding rôle of supporting distributed computing.* (SRR 6, section 3.2.7)
    - This implies the availability of *twice this capacity, namely 300 units, in the outside institutes.* Taking into account centres equipped for production, this capacity is currently around 150 units. (Appendix C)
    - Additionally, we recommend that methods be investigated for providing substantial processing power in the form of *advanced parallel computing facilities* that offer potentially cheaper computing for relatively stable programs. Integration into normal Computer Centre operation is essential. (SRR 7, section 3.2.7)

      Parallel computing systems have in addition considerable interest for on--line systems and applications for theoretical physics.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Computing at the Experiments

7. *Centralized support for hardware and software components* to be used in the construction of on-line systems at the experiments, as well as *design support* to aid their effective integration, will still be necessary, both for smaller experiments, and to provide components for the larger experiments. (SRR 12, section 3.2.9)
8. The extreme environment of future hadron colliders requires major advances, and forces a tight coupling of the issues of detector design, detector digitization, data compression and triggering. We recommend setting up *coordinated pilot projects* for introducing new high-level design methodologies for both software and hardware, and for acquiring familiarity with the application of modern techniques to future experiments' real--time problems. (SRR 11, section 3.2.9)

Software diversity

9. The desire to reduce diversity by moving to *common software environments* is a strong theme from the working group reports, in areas ranging from operating systems to end--user packages, (SRR 14, section 3.2.10). In particular:
    - *Unix* should be investigated as a common operating system, given that new initiatives (e.g. OSF, AIX) are encouraging. (SRR 13, section 3.2.10)
    - The use of the *ORACLE* DBMS should be consolidated at CERN and generalized within the community. (sections 3.2.10, 3.4.1, 3.5.6, 3.6.3)
    - A set of *CASE* tools should be selected, and a suitable license negotiated so that all participating institutes may benefit. (section 3.2.10)

MIS

10. The *MIS* environment at CERN should be *modernized and made coherent*. It should include a *CERN corporate data model* using a *unique DBMS* (ORACLE). This implies the replacement of the currently used major corporate applications. A comprehensive *electronic forms handling system* should be implemented and *EDI* should be used, where possible, for corporate applications. (SRR 22 and 23, section 3.6.3)

Computing for Engineers

11. Experiments and accelerators depend on *front-line technology* that can only be achieved through the use of *CAE tools.* CERN is late investing in this area. (SRR 20, section 3.5.1)

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

12. A large percentage of CERN staff were educated before computer-aided tools became part of the syllabus. If CERN is to gain full benefit from CAE techniques, *staff at supervisory as well as technical levels will need training in modern design methodologies.* (SRR 21, section 3.5.1)

Accelerator Control

13. *Accelerator control systems* should be increasingly based on *commercial hardware and software*, using a layered approach with common design elements and a *three tier LAN*. (sections 3.4.1, 3.4.2)
14. *Strong central support* for the computing needs of the accelerator divisions must be consolidated. (SRR 19, section 3.4.7)

Computing for Theorists

15. A policy decision is required as to the fraction of CERN's vector processing power that should be devoted to theoreticians working at CERN. It is recommended that this fraction should not exceed 10% of the total available. (SRR 24, section 3.7.3)

Selected Highlights from the WG Reports

In 1989, the penetration of Unix in the Physics Community, at least at CERN, was very low (the Apollo Aegis operating system being definitely not Unix). However, this was soon to change in a big way.

The use of Oracle in the Physics Community was rather patchy in the LEP era, as we will see from papers presented to CHEP.

The use of Software Engineering tools and methodologies was hotly debated with very different results according to experiment / community.

Both local and wide area networking was still an area of uncertainty and change: FDDI in particular never took off as predicted, nor did OSI conformant solutions.

A need to reduce Software Diversity

*The general desire to attack the diversity of software can be found in many of the working group reports. In the report for experiments the strongest recommendation in this area proposes "**to start a Unix service on the central IBM service at CERN as soon as adequate software is available from IBM.**" An initial small--scale service should receive further resources according to user demand. In addition there are a number of recommendations aimed at improved organization and coordination of the HEP software community.*

This is not the first time that a pilot Unix service on the IBM has been proposed. What is new is that there are now sufficient good reasons for such a pilot

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

investigation to make sense. Apart from the penetration of Unix into the HEP world though its appearance on workstations, the Cray and various other machines, several of the traditional arguments against Unix in this rôle are decreasing in validity. One may note in this context the previously bad reputation of FORTRAN compilers available with Unix, which gave inefficient run time performance compared with those available in proprietary systems (e.g. MVS, SCOPE, VMS, VM), There has been no seriously supported version of Unix on IBM systems. The recent initiatives of IBM in the direction of the AIX operating system, (AIX will run the same FORTRAN compiler as VM/CMS and MVS), and the creation of the Open Systems Foundation [Footnote: There are currently nine sponsoring manufacturers and more than 90 full members.] by many manufacturers, show an increased interest in Unix from the computing world in general.

Nonetheless, one is obliged to take into account the special needs of the HEP community, particularly in the areas of very substantial tape handling and of batch schedulers. These are not yet guaranteed to be solved by AIX on the IBM, even if they have indeed been solved on the Cray. Indeed the requirements for Unix running on such large systems have been tackled in relatively few places as yet. However, referring back to [the previous report?], *it is instructive to note that the VM/CMS service at CERN was a user requirement during 1983, a Directorate decision during 1984, and an embryonic user service during 1985.*

It is recommended that a pilot Unix service is started on the central IBM computers at CERN as soon as practicable.

The recommendations concerning the organization and coordination of software issues in the HEP software community include the proposal for a Software Support Committee. Such a body could help to reduce the diversity of software by making successful commercial and HEP--specific products more widely available. It could also evaluate the benefits of writing HEP--specific tools against the cost of manpower, were a suitable commercial package not available.

HEP-CCC should establish a Software Support Committee as one of its technical working groups.

*The crucial rôle that data base techniques and access methods have to play in the management of the huge volumes of data expected in the next decade of computing at CERN is recognized. The working group recommends "**the consolidation of the use of ORACLE in the HEP community for this purpose.**" This recommendation is strengthened by the request for HEP--wide coordination of the definition, writing and installation of data base orientated software for large physics experiments. It is felt that this coordination could also come under the umbrella of the Software Support Committee.*

Additionally, the working group report highlights the success of some larger experiments in using software engineering tools to design and implement their physics codes. They recommend the support of software tools and methodologies

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

[Footnote: A methodology is a set of "methods", typically purchased as a collection of CASE tools.] , including their selection and evaluation. ***It is believed that, in view of the distributed nature of the software development effort amongst collaborating institutes, such tools will become essential for coherent and correct code.***

CERN Data Networking Requirements in the Nineties

- *The networking between the outside institutes and CERN must make a substantial jump in bandwidth (to at least 2 Mbits/s) in order to provide new possibilities and to overcome barriers to distributed working.*
- *One of the world's largest Ethernets exists at CERN; this needs continuous upgrading in bandwidth. Presently, FDDI [Footnote: Fibre Distributed Data Interface, ANSI X3T9.5, ISO DIS 9314] products offer the best way by which to proceed.*

The Committee considers that CERN is well placed to exploit the FDDI products that will appear imminently on the market in order to provide this general--purpose high--speed backbone, and recommends active continuation of the FDDI programme.

*The working group recommends that "****access to CERN at speeds of at least 2 Mbits/s is required from major HEP sites in Europe, to allow decentralized analysis of the LEP data, as soon as this analysis starts in earnest. Higher speeds (equivalent to those attainable on a LAN) will be required as soon as practicable, to allow effective use of workstations over geographical distances.****"*

The provision of these high speed links is of the highest priority*, and initiatives that are being actively pursued are strongly supported. It is recognized that the implementation of 2 Mbit/s links on CERN site, namely CERNET, enabled a significant change in the number of users able to work away from the Computer Centre. The recent success of T1 (1.5 Mbits/s) long distance links within NSFNET [Footnote: The network of the American National Science Foundation, currently upgrading some links to DS3 (45 Mbits/s).] in the United States is additional evidence of the quantum jump in user productivity that can be achieved through such bandwidth. Progress towards the desired goal of the HEP community to overcome geographic barriers, and to facilitate decentralized physics data processing/analysis, is dependent on the existence of such links. The required bandwidth has been expressed as "****one glass--fibre equivalent****".*

Improved communication facilities between CERN and the collaborating institutes must be provided. The major paths should operate at a bandwidth of at least 2 Mbits/s.

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Summary of "the New Green Book"

We have now set the scene for the start of LEP data taking. In the few brief years since the original green book, much had already changed. Distributed computing is to the fore, although there is still the feeling that mainframes will continue to contribute significantly. We see a significant change in storage media with the arrival of the 3480 cartridge, and with it robotics. But FDDI? Unix on the IBM mainframe? Well, you can't win them all.

It will not be surprising to see the above issues discussed at Computing in HEP conferences. What was perhaps less obvious was how rapidly things would change.

CHEP FOR LEP

Conferences on Computing for High Energy and Nuclear Physics (CHEP) have been held at roughly 18 monthly intervals since the mid-1980s. By tradition, they rotate around the world: Europe, the Americas and elsewhere. Typically attended by several hundred people, the proceedings – taking those from CHEP 92, held in Annecy, France, as an example – run to approximately 1000 pages. The format of CHEP conferences includes (invited) plenary talks, parallel sessions, posters and a final morning of track and conference summaries.

Although some images from CHEP conferences still exist at the time of writing[12], the agendas and conference proceedings are harder to find. Fortunately, scanned versions exist in the CERN Document Server, occasionally with illegible pages and typically low resolution, and hard copies (probably) exist in the libraries of CERN and other institutes.

These proceedings, also known as "the annals of HEP computing", provide a high-level snapshot of both achievements and plans.

The proceedings of most of these conferences can be found with the search string "International Conference on Computing in High Energy Physics", in the collection "Conferences", although in a few cases the "International" needs to be omitted. For those conferences that took place from the year 2000 on, "Nuclear" needs to be included, i.e. "(International) Conference on Computing in High Energy and Nuclear Physics", or else "Physics" omitted.

You can also try InspireHEP for some "alternative facts" perhaps:

https://inspirehep.net/conferences?sort=dateasc&size=25&page=1&start_date=all&q=series.name%3A%22CHEP%22

[12] Possibly still available at http://event.twgrid.org/chep2010/previouschep.html.

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

For CHEP 2012 held in New York and CHEP 2019 held in Adelaide the proceedings seem to be missing from CDS. In the latter case, presentations can be found in Zenodo.cern.ch using the search string "CHEP Adelaide" (and perhaps others) and the full proceedings are yet to be published.

Since CHEP 2006, these events have been typically preceded by what were initially WLCG "Service Challenge" workshops, that later became WLCG Collaboration Workshops – see https://indico.cern.ch/category/890/.

The (expired) CHEP 2012 website says, however, *"The **CHEP** 2012 organizers will publish all accepted contributions (both oral and poster presentations) in the conference **proceedings**. The organizers have arranged to publish the **proceedings** electronically in the Journal of Physics through IOP Publishing, about 2 months after the conference."* And indeed this seems to be the case (see table below).

These conferences are listed below in chronological order – one of the few cases in which it makes sense (as opposed to reverse order).

Date	Location	Proceedings
25 - 28 Jun 1985	Amsterdam[13]	https://cds.cern.ch/record/108031?ln=en
2 - 6 Feb 1987	Asilomar, CA	https://cds.cern.ch/record/110823?ln=en
10 - 14 Apr 1989	Oxford	https://cds.cern.ch/record/205278?ln=en
9 - 13 Apr 1990	Santa Fe, NM	https://aip.scitation.org/toc/apc/209/1
11 - 15 Mar 1991	Tsukuba, Japan	https://cds.cern.ch/record/275229?ln=en
21 - 25 Sep 1992	Annecy	https://cds.cern.ch/record/223307?ln=en (Complete proceedings exist as a CERN Yellow Report)
21 - 27 Apr 1994	San Francisco	https://cds.cern.ch/record/277254?ln=en https://inspirehep.net/conferences/968795
18 - 22 Sep 1995	Rio de Janeiro	https://cds.cern.ch/record/289231?ln=en https://inspirehep.net/conferences/969392
7 - 11 Apr 1997	Berlin	https://cds.cern.ch/record/305457?ln=en
31 Aug - 4 Sep	Chicago	https://inspirehep.net/conferences/9712

[13] This is often referred to as the first CHEP conference. However, the corresponding CDS entry says it was *"the 3rd in a series of conferences, previously held in Padua (1983) and in Guanajuato (1984)"*. Elsewhere we find: In addition there are two events which may not be considered CHEP proper but seeded the path to them:
1983 - Padova, Italy - March 23-25, 1983 - Three day in-depth review on the impact of specialized processors in elementary particle physics
1984 - Guanajuato, Mexico - May 8-11, 1984 - Symp. on Recent Developments in Computing, Processor and Software Research for High Energy Physics

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

1998		33
7 - 11 Feb 2000	Padua	https://cds.cern.ch/record/406721?ln=en
3–7 Sep 2001	Beijing	http://www.ihep.ac.cn/~chep01/ (dead link)
24–28 Mar 2003	San Diego	http://www-conf.slac.stanford.edu/chep03/ (dead link)
27 Sep - 1 Oct 2004	Interlaken	https://chep2004.web.cern.ch/ and https://chep2004.web.cern.ch/Programme/contributionlist_full.html
13 - 17 Feb 2006	Mumbai	https://cds.cern.ch/record/824920?ln=en
2 - 7 Sep 2007	Victoria	https://cds.cern.ch/record/947732?ln=en
21 - 27 Mar 2009	Prague	https://cds.cern.ch/record/993686?ln=en
18 - 22 Oct 2010	Taipei	https://iopscience.iop.org/volume/1742-6596/331
(May 2012)	New York	https://iopscience.iop.org/volume/1742-6596/396
14 - 18 Oct 2013	Amsterdam	https://iopscience.iop.org/volume/1742-6596/513
13 - 17 Apr 2015	Okinawa	https://iopscience.iop.org/volume/1742-6596/664
10 - 14 Oct 2016	San Francisco	https://iopscience.iop.org/volume/1742-6596/898
9 - 13 Jul 2018	Sofia	https://indico.cern.ch/event/587955/
(4 – 8 Nov 2019)	Adelaide	http://chep2019.org/

CHEP 1985

A fairly lengthy paper (17 pages) from the leader of the CERN Data Handling (DD) division at the time – Paolo Zanella – discusses Trends in Computing for HEP (https://cds.cern.ch/record/161215/files/cer-000072548.pdf).

It states that ***"To satisfy the computing needs of … the LEP era, a straightforward increase in overall capacity is not enough. New facilities for data acquisition, storage, processing and communications are also required".***

Clearly drawing on the work in the original "green book" (Computing at CERN in the LEP era), the paper examines each of these areas in detail.

The reader will notice that explicit capacity figures look almost unbelievably modest by today's standards but we should not underestimate how much of a challenge they were at the time. The revolutions in computing that took place during the LEP era will be covered as we walk forward in time – this was a period when some people (working for Digital Equipment Corporation – look what happened to them!) claimed that no-one ever needed a computer faster

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

than a VAX[14] 11/780 (with its 5MHz processor). Papers still exist hinting at the purchase price of computers, storage and memory at the time and they were as high as the performance and capacity was low.

Zanella's paper estimated that 10 – 15 "CERN units" – equivalent to a VAX 8600, perhaps 4 times more powerful than an 11/780, or else an IBM 370 / 168, would be required per LEP experiment. This number appeared regularly in discussions at that time with big question marks regarding how many CERN units CERN itself could afford to provide – typically much less than the above estimate!

A number of statements provoke at least a chuckle, such as ***"FORTRAN is probably the only perennial standard that will never be questioned"***, although I also like the opening phrase ***"Even the most reluctant physicists are now convinced that computers are critical for the success of their experiments".***

In the former case, it is worth pointing out that not only did CERN and other HEP labs invest heavily in the Fortran[15] standardization process but that the language, in its various forms, did serve HEP well for many decades. Whether the latter phrase is now widely accepted or not, the then-current Director General of CERN, Rolf-Dieter Heuer, did point out, during his summary on the "Higgs discovery day" that the machine, the detectors and the associated collaborations, as well as the Grid had all been required to achieve that memorable result.

Covered in more detail in other papers, Zanella's report discussed also the importance of networking and the various applications that required it, as well as alternatives to mainframes, including (small) farms[16], emulators and also powerful workstations. All aspects of computing were already highly heterogeneous, necessitating, for example, mail and file transfer gateways.

Not available online but an idea that will come back later was the possibility of using satellite communications – there had already been such a project "STELLA" in the early 1980s but have a look at the roof of B513 at CERN if you ever have the opportunity.

---

[14] VAX (architecture) / VMS (operating system) computers were incredibly popular with users – probably because they were "user friendly". A faster machine – a VAX 8600 – was installed around that time for the LEP experiments with disks of less than 500MB capacity per experiment (plus a "system" disk and a "scratch" area). See https://cds.cern.ch/record/1721528. Whether the author features in that photo is hard, even for him, to tell. But no, not with that shirt.

[15] Prior to the Fortran 90 standard, FORTRAN was written in uppercase.

[16] To give an idea of the scale, a set of IBM RS6000 machines at CERN in the early 1990s used the names of some of the Hawaiian islands as node names – Oahu, Maui, etc. Even though the complete archipelago consists of 137 islands, few can name beyond the 7 inhabited ones.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

CHEP 1987

Most papers are available[17] in a special edition of Computer Physics Communications, available to CERN account holders through Science Direct (true also for some other CHEP conferences):

"We are very pleased to include in this issue of Computer Physics Communications nearly all the papers presented at the conference on Computing in High Energy Physics held at Asilomar State Beach, California, 2—6 February 1987. This conference is one in a series that started at Bologna in 1980, followed by CERN in 1981, Padova in 1983, Guanajuato in 1984, and Amsterdam in 1985. Those conferences were a major forum for papers on all aspects of computing in high energy physics."

Apart from the location, this must have been a groundbreaking and exciting conference. Some selected papers:

- DEVELOPMENT OF SOFTWARE FOR ALEPH USING STRUCTURED TECHNIQUES
- USE OF THE ADAMO DATA MANAGEMENT SYSTEM WITHIN ALEPH
- PAW - TOWARDS A PHYSICS ANALYSIS WORKSTATION
- SOFTWARE ISSUES FOR LARGE DETECTORS (not directly related to LEP)
- TWO YEARS OF REAL PROGRESS IN EUROPEAN HEP NETWORKING: A CERN PERSPECTIVE
- DATABASE SYSTEMS FOR HEP EXPERIMENTS
- THE DELPHI INTERACTIVE ANALYSIS AND TANAGRA (not summarized)
- PERSONAL COMPUTERS IN HIGH ENERGY PHYSICS

In Software for ALEPH using structured techniques by Gottfried Kellner, he states:

"Software systems are very often badly designed, badly or not documented, delivered too late, very costly to modify and unstable against upgrades. Our own environment, i.e. High Energy Physics, is no exception.

Why is this so? It is certainly true that concepts of software engineering were only introduced some 15 years ago. The tendency is to write code immediately — because one knows what one wants —and think afterwards how to modify it to include other options. Software writing is still an art — some people are very good at it but they don't (or can't) explain to other people how they do it."

Whilst the use of Structured Analysis / Structured Design was rather successful within the ALEPH collaboration, there were certainly differing opinions and

[17] Some are also available in the CERN document server but these appear to be scanned internal notes, e.g. DD notes, rather than the published proceedings. The Science Direct copies are typically higher quality as a result.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

experiences across the other LEP experiments. However, can one disagree with Gottfried's introductory remarks, even today?

The paper on the use of ADAMO within ALEPH explains:

ADAMO is a data management system for defining data and manipulating them from FORTRAN programs. A form of the Entity—Relationship (ER) model and the data flow diagrams of structured analysis have been combined to provide a system suited to algorithmic work. The FORTRAN interface consists of calls to a subroutine package with a few data manipulation primitives. Data items are visible in variables carrying meaningful names; the code is compact and readable. The ER data structures, mapped onto tables, are compatible with a relational data base management system. Tools have been written to copy tables to and from a relational data base. The system has been used for several applications in the fields of event simulation and reconstruction, online systems, detector description, support tools and packages.

The paper on PAW – the well-known Physics Analysis Workstation – explains:

In recent years interactive data analysis and presentation packages have become more and more attractive on conventional timesharing systems. The advent of personal workstations with their good response time and powerful graphics capabilities is making these systems even more popular. PAW has been designed in that spirit. It combines the best features and the experience acquired with the existing systems together with the user friendliness now possible on modern workstations. Rather than being a monolithic system, which may be well suited to one operating system, PAW is built with modules which can also be used in the batch environment or by other interactive programs. PAW has been designed to work on many different configurations.

The paper on progress with European HEP networking (worth reading for all the details) starts:

The last two years have been marked by real progress in networking in HEP. Home-made developments, studies and plans have given way to the use of real networks involving hundreds of HEP and other computers and based on externally produced software and hardware. Within the last year, the first generation of industrial software products following some of the international standards for networking have become available. Related developments are taking place in networking for on-line systems and indeed the LEP experiments are distinguished by their heavy and crucial reliance on both local and wide-area networks. This paper describes the progress made at CERN since the last two years and looks at perspectives for the future.

The paper concludes by commenting on the:

- emergence of Ethernet as the favoured LAN;
- emergence of DECnet and TCP/IP as the de facto LAN protocols;

- precursors of truly distributed computing systems (RPC) and file systems;
- dramatic growth in WAN infrastructure and traffic;
- unfortunate degree of protocol diversity in the WAN;
- early usage of, and commitment to, OSI protocols;
- use of gateways;
- **heavy dependency of LEP collaborations on networking.**

The paper on Database Systems for HEP – a topic that comes back regularly:

"… reviews the role of commercial and home-made database systems in HEP experiments. … draws heavily on the qualitative and quantitative evaluations performed by the L3 offline computing group, but also reviews some of the activities of other experiments. … considers to what extent features, such as query language, Fortran interface, efficiency, portability, robustness, security and cheapness, are needed for various applications and are offered by existing systems. The commercial systems ORACLE and SQL/DS, and the HEP-made 'access methods' KAPACK and ZEBRA-RZ will be discussed in some detail. ***Finally, … presents a brief overview of the L3 database system."***

The paper on Personal Computers in HEP, whilst describing many of the problems that existed at the time, is interesting in that it pre-dates by as much as a decade other proposals to investigate the use of PCs for HEP – by which time that had of course become significantly more powerful and many of these early problems had been overcome.

CHEP 1989

ARE (WERE) WE THERE YET?

This conference that was held just months before the beginning of data taking at LEP asked the obvious question: ***Computing on the eve of LEP data taking: Are we ready?*** This paper, by David Williams, the leader of the Data Handling division at the time stated:

> *"It is rather amazing that the computing for LEP appears to be in such good shape. But much work remains ahead of us. The data volume that we are faced with is one* (to) *two orders of magnitude more than we have faced before."*

The paper talks of the data volumes that would come from the accumulation of 10 million events (per experiment):

- 4 terabytes of raw data (20,000 3480 cartridges)
- 2 terabytes of simulated data (10K cartridges)
- 1-4 terabytes of summary data. (etc.)

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Whereas the final data volumes were significantly larger – around 100TB per experiment – it is hard to put oneself in the mindset where these volumes were "frightening".

The paper covered much more than just the computing aspects, as the abstract shows:

"The physics objectives of LEP, the status of the LEP machine, and the status of the LEP experiments are briefly reviewed. An analysis is made of the state of readiness of the computing resources, both hardware and software, needed to record, transport and process the data coming from the LEP experiments, and to extract the physics from this huge volume of data. A tentative conclusion is proffered that we are "ready" to face the challenge."

Despite the uncertainty regarding local area networking in the Green books, the LEP experiments made extensive use of Ethernet, that was bridged (due to its limitations at the time) to cover distances up to 15km. Given the growth in personal workstations and PCs, reflected in the increased traffic, led to the expectation that saturation would be reached within 2 years. There was hope that the FDDI standard, with 100 megabits / second per fibre, could alleviate the situation.

In addition, there was great pressure to make available at least 2Mb/s to outside institutes – not to move data but to keep code and calibrations in sync, together with some key events and results. However, there was significant concern about not only the availability but also the costs associated with these connections – perhaps 5 times higher than those seen in the US!

Whilst offline computing at CERN was still dominated by "mainframes" (IBM and compatible, Cray, VAX systems), there was increasing use of workstations that offered not only excellent graphics but also significantly cheaper CPU than the mainframes, but the recent emergence of RISC processors that held great promise.

In summary, he highlighted reasons to be cautiously optimistic (provided that sufficient financial resources were made available) – a similar picture that I also painted regarding our readiness for LHC computing, following the WLCG Service Challenges, at CHEP 2006 in Mumbai.

BEYOND ETHERNET - FUTURE LANS

This paper concluded that:

"Ethernet is still a great success and will stay with us for a long time. Rather than be replaced, Ethernet will be used together with new high bandwidth LAN technologies to give hybrid "super LANs" covering sites as large as CERN's.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

General workstations and PCs will remain Ethernet-based, but emergent super-workstations and associated file-servers, imaging applications, and high throughput connections to central resources be based on FDDI."

This turned out to be fairly accurate prediction, holding true until well into the 1990s, before Gigabit Ethernet was rolled out as a replacement. As we will see, some special purpose networks were used in parallel in certain cases.

OVERVIEW OF HEP WIDE AREA NETWORKING: PRODUCER PERSPECTIVE

Complementary to the above was a paper on wide area networking. Despite significant advances since the report at the previous CHEP, there were still major difficulties. The situation remained highly heterogeneous, with relatively low bandwidth but expensive connections. The situation regarding OSI networking remained unclear whereas TCP/IP was forecast to be in use "everywhere" within 18 months.

1. The HEP leased line network consisted of 20 connections over 3 continents, ranging from 64 kbps down to 9.6 kbps;
2. The HEP X.25 network ran on top of this leased line infrastructure;
3. DECnet consisted of over 10,000 nodes in some cases over different X.25 infrastructures;
4. The HEP SNA / RSCS network linked a number of IBM based computer centres;
5. Mail and file transfer gateways existed, such as GIFT (file transfer) and MINT (mail at CERN).

Despite these "successes", a long list of reasons for user unhappiness was given, covering not only insufficient bandwidth but also many "usability" issues. Well, things could only get better(?)

UNIX FOR HEP

In the context of the latter (the Cray ran a version of Unix called Unicos), a paper from Dietrich Wiegandt discussed the possible use of Unix for HEP (recall that the New Green Book recommended starting a Unix service on IBM as soon as possible):

"Until a few years ago, UNIX was hardly known outside of computer science institutions. The spectacular growth of the number of installed UNIX systems since 1982, now approaching one million, has surprised even experts. Looking back to the very beginning in 1969 and following the development of different UNIX versions, with a quick look at some aspects of the UNIX philosophy may help to understand this phenomenon.
Standards for the UNIX system are emerging; UNIX is the preferred operating system for innovative computer architectures, yet the HEP community seems to remain skeptical. The success of UNICOS, the UNIX-derived supercomputer

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493
operating system of CRAY Research, and the use of UNIX-derivatives on workstations like Apollo and Sun, may help to change this attitude."

Dietrich concluded with a relatively long list of requirements for "Unix in a computer centre", that included: batch environment, high quality Fortran support, magnetic tape support, resource limitations (disk, CPU, etc) on a per-user basis, good networking including distributed filesystem support, access protection and accounting facilities. It is interesting to note the explanation that he provides for the very terse commands[18] that are typical of Unix that came from the environment in which it was originally developed. The fewer key strokes the better.

Finally, in the area of systems and infrastructure, whilst there were a number of talks concerning transputers, we can probably safely assume that a come back on the timescale of an *fcc* is unlikely.

OVERVIEW OF THE ESSENTIAL TOOLS

"The essential tools needed for the analysis of HEP experiments are software, networking, data handling, and CPU power. Leaving most software topics to other papers, this paper describes how we can use today's hardware effectively for HEP data analysis. It describes a new but viable model for HEP computing and data handling on a single site, and outlines the challenge of extending this environment to the wider university-based community. Finally it touches on the revolution in software design that may be triggered by the revolution in computing technology."

In this model, the mainframe assumes the role of data server, feeding and receiving data from multiple workstations. Whilst this might work well as a centralized solution, it leaves remote physicists somewhat ostracized. By coupling similar systems at a small number of remote sites (file server, DB server, workstations) and ensuring synchronization between the sites, for which only modest networking is required, this problem is solved. (Bulk data transfer still required the "suitcase" model – something that was considered up to the design stage of LHC computing, at least as a backup should networking be either too expensive or too poor).

PAW – THE PHYSICS ANALYSIS WORKSTATION

An update on the well-known and highly successful PAW package stated that it had been implemented at the time on IBM (VM/CMS and MVS/TSO), VAX/VMS, Apollo (Aegis and Unix) as well as Sun (Unix). A partial implementation for the MAC existed and a version for the Cray was in preparation! (Did this ever exist? If so, why?)

---

[18] For some reason, there were people who preferred to use a Unix "shell" on top of the more intuitive VMS. A senior programmer once explained that he couldn't remember how to create a directory (create/directory) but luckily was able to use the Unix *mkdir* command instead. Obvious, really.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

THE ALEPH EVENT RECONSTRUCTION FACILITY: PARALLEL PROCESSING USING WORKSTATIONS

This paper described one of an increasing number of the use of multiple workstations where previously a mainframe would have been used. In this case the features of the VMS operating system were used to advantage and similar solutions, based either on VMS or increasingly Unix, were a feature of the early days of LEP.

"Event reconstruction of high energy electron—positron interactions, which will be observed in the ALEPH detector at LEP, represents a computational task which requires very large compute power, large program memory and only moderate input/output capacity. The intrinsic event-level parallelism of this task allows a solution based on an array of loosely coupled computers. A cost-effective solution is presented using commercial hardware, namely an Ethernet-based local area cluster of VAXstation processors. The parallel execution of the reconstruction program is achieved with minimal code modifications, by exploiting the fact that each processor runs a full operating system (VMS) and is able to do autonomous random access input/output. Results are shown of input/output performance and CPU utilization efficiency. These results were obtained on a prototype consisting of 6 VAXstation 3200 nodes and a microVAX 3600 file server. The final system will start with 12 such nodes, giving a total capacity of 8 Mflops."

THE OPAL EVENT SERVER

This paper presented a hybrid approach, using a CERN mainframe and workstations and other computers both in the local and wide area. Notwithstanding the relatively low bandwidth of the network connections available at the time, they were still sufficient for messaging and the retrieval / distribution of small event samples.

'The OPAL experiment at LEP may record and simulate in detail on the order of five million events per year. Many analyses may be performed on the distributed summary information of subsets of these events and on rare occasions a physicist may require access to the full data to verify the detailed interpretation of the event. The OPAL event server was conceived as a service machine on the CERN central IBM system CERNVM. The server receives mail and messages from workstations and computers on the CERN site and in the OPAL institutes world wide requesting information for specified events. The information is extracted from a library of *magnetic cartridges and transmitted to the requestor in an ASCII format."*

DATA STRUCTURES AND DATA-BASE SYSTEMS USED IN HEP: MODELLING AND IMPLEMENTATION

"The importance of data modelling and data-base design based on software engineering techniques is discussed with particular consideration of applications in high energy physics experiments. As an example, the approach chosen by the ALEPH collaboration is presented in some detail. A survey of data-base applications

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

used in various experiments showed that the requirements are almost identical. Therefore, a proposal is made to design general-purpose packages including tools for data exchange and documentation."

This paper argued that Structured Analysis / Structured Design and data modelling techniques were considered vital in designing software for large experiments. This was certainly felt to be the case within ALEPH but was a topic that was hotly discussed at both this and future CHEP conferences. The paper continued with another survey of "database" applications in HEP and concluded that *"there was no doubt that for larger experiments and experiment-wide database was needed. However, the choice between commercial and home-grown is still far from obvious".*

USING A DATA MODEL FROM SOFTWARE DESIGN TO DATA ANALYSIS: WHAT HAVE WE LEARNED?

This almost companion paper discussed:

"The ADAMO data system (based on the entity-relationship model) *is being used in a number of particle physics experiments. Experience with it indicates that data modelling is a powerful program design method that extends across the whole software life-cycle, although existing support tools are not yet satisfactory. The entity relationship data of ADAMO can be handled by various programming languages* [including FORTRAN] *, and can be used effectively in interactive data analysis".*

SUMMARY OF CHEP 1989

Even if Williams' paper was perhaps a little optimistic, and with shades of some of the disruptive changes about to sweep over us (we already saw innovative use of workstations working with or replacing mainframes), the outlook for first data taking and analysis did look rather promising. There were clearly some pieces of software that were either not really ready or else not mature and "bug-free" but then this statement is probably always true. More formal methods for program design had both enthusiastic supporters as well as detractors – something that is probably still the case today.

CHEP 1990

e-proceedings at https://aip.scitation.org/toc/apc/209/1.

Of the papers that can be (easily) found online, that on Databases for Large Detector Systems – Experiences at LEP – is perhaps the one most directly related (to LEP): https://cds.cern.ch/record/207117/files/CM-P00065564.pdf. This paper describes the various "databases" used by the different LEP experiments.

This document starts:

The physical size of the LEP detectors at CERN, the quantity of information which they generate, and even the number of collaborators needed to build and run them, has shifted a large fraction of the associated software problems into the area of data management. Managing large volumes of data of all kinds for a variety of applications, from transport and storage to retrieval and analysis, has become, and will remain, the primary concern of many HEP programmers. All four LEP experiments have used a mixture of commercial and home-grown database management systems to exploit their detectors. Despite some efforts to achieve coherence the same data appear in different guises in different systems, and essentially equivalent code has been written several times over to access and manipulate these data. What can we learn from the LEP experience to ensure that, for detectors at the LHC and SSC, we overcome the difficulties inherent in the information management challenge which they present?

It then describes the use of various database systems – homegrown or commercial – by the different experiments. These included:

- The ALEPH collaboration made a concerted effort to adopt a coherent data view by using the Entity-Relationship Data Model as the basis for all data descriptions, from detector geometry to collaborators' home and e-mail addresses. The modelled data are stored in tables which are created, filled and accessed using ADAMO software;
- The ALEPH online database is also held in ORACLE and is used for run control and status recording;
- The ALEPH personnel inventory data are also kept in a central ORACLE database;
- The DELPHI detector description is held in CARGO, a home-grown database system which has its own internal hierarchical structure and which uses a modified version of the Keyed Access Package KAPACK to implement the access routines;
- The OPAL and L3 detector descriptions and event data are held in ZEBRA formats;
- Both experiments have implemented a data management system on top of the ZEBRA direct file access routines RZ, for use in the analysis phase, whether on- or ofl-line. The L3 package is DBL3; the OPAL package is OPCAL.

In addition, under the heading of Tape Management, it says:

Real progress was made in 1989 with the introduction of the Distributed File and Tape Management System FATMEN. The requirements for FATMEN were compiled by the LEP experiments plus UAl and UA2, in collaboration with the CERN Data Division. The experiments[19] *are in the process of converting to the new facility. FATMEN is a networked system which supports access to stored data in a*

[19] ALEPH never converted but several outside experiments, including at FNAL, did. See, for example, https://cds.cern.ch/record/401013/files/p759.pdf.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

consistent manner, regardless of the host operating system or the location or type of storage medium. FATMEN offers bookkeeping facilities for file storage and retrieval for physics data analysis e.g. the management of cartridge robots and dim farms. It has a command line interface and a FORTRAN callable interface, and has been installed on almost all of the HEP mainframes and workstations.

FATMEN uses the ZEBRA package to provide memory and data structure management and file exchange. The ZEBRA RZ direct file access routines are used to manage catalogue information stored in a hierarchical direct access directory structure. File identifiers are UNIX look-alikes, for example, //CERN/DELPHI/RAWDATA/RUN20. A ZEBRA server, resident at the production centre, provides remote clients with access to the FATMEN File Catalogue, maintained under ORACLE; at CERN the IBM CERNVM ORACLE RDBMs is used. A combination of ZEBRA file access facilities, local journal files, ORACLE journaling and roll-back is used to assure reliability of the catalogue, which is eventually expected to contain references to hundreds of thousands of datasets.

On a personal note, having been directly involved with this package and its support, it was surprising how reluctant some people were to move away from referring to data by a tape volume and file sequence number to something more "meaningful", as well as operating system and location independent.

CHEP 1991

This conference was the first after LEP data taking had started. It is interesting to see the continued role of the mainframe already being questioned, the evolution of networking, as well as "new" techniques, such as direct access to events (as opposed to the more traditional sequential processing of files).

Unfortunately, not all of the papers are available online, with a small sub-set in CDS as scanned PDF files. In addition, 70 titles can be found via https://inspirehep.net/conferences/967381, a few of which also have associated PDF files.

HEP Computing – Where are we now? Open questions for CHEP 91

The paper that accompanied this talk – seemingly written after the conference – is much more a summary of the then-current situation than open questions. One of these was clearly the relationship between the SSC and the LHC where there was a certain understandable amount of rivalry. At the conference, there was much debate on the role of, as well as the difficulties in moving to, Unix. The use and value of formal methods in program design was also heavily debated, as was (the inevitable?) move to some form of distributed computing.

Some of the topics raised included:

- The RISC revolution, where some of the latest processors where being bench-marked at speeds similar to those of single CPUs in CERN's

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

mainframes but there were questions regarding management as well as FORTRAN support;

- Operating systems: VM/CMS and VAX/VMS versus Unix? *"Any large scale conversion to Unix is likely to take 5 years or more"* – things moved much quicker than this;
- Distributed computing and networking (two separate but clearly interlinked topics);
- Storage media – increases in capacity for IBM 3480 style cartridges versus cassettes (dual reel – from DAT and Exabyte through to Ampex and Sony) – optical disks were already failing on their promises regarding both capacity and speed;
- Finally two less obvious but distinct topics, namely vectorization and neural networks. Are there clear answers regarding these even today?

Is the rôle of the mainframe terminated? : mainframes versus workstations

This detailed paper, from David Williams again, asked an important question that divided people, at least in CERN IT[20], into two distinct camps. Detailing the characteristics of the two options, David announced one more round of IBM upgrade at CERN, in parallel with the move into production of the Scalable Heterogeneous Integrated Facility Testbed – SHIFT[21]. (SHIFT could be considered a 2nd generation attempt to use (powerful) workstations, following on from a joint project between the OPAL experiment and HP – HOPE – that was a Monte Carlo simulation farm. There were several other experiment-led projects using workstations, some of which have already been described here.

HEPnet in Europe : status and trends

In this talk Francois Fluckiger reported on the progress since the previous CHEP. He explained how:

"the three major objectives or predictions have been achieved (enter the megabit era, deploy TCP/IP, and install a TDM infrastructure) and what is still needed to complete this phase. It also sketches what the next generation of HEPnet may look like: a set of new multi-media services running on hundred megabit bandwidth probably using SDH and ATM technology".

His conclusions, in the form of 9 points, were as follows:

1. *HEPnet in Europe has entered the megabit era;*
2. *Eastern and central Europe is now partially connected;*
3. *TDM (time division multiplexing) and TCP/IP technologies are and will continue to be heavily used over the coming 4 years;*
4. *OSI usage has not yet really taken off;*

---

[20] At the time it was the Computing and Networks Division – the "Data" had been dropped in the process.
[21] The Testbed word was later dropped and the T from "facility" used in its place.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

5. *HEPnet is the fastest network in Europe;*
6. *CHEOPS will investigate spooled data transfers;* (see CHEP 92)
7. *Video and voice services are being considered;*
8. *Prospects for 34/140 Mbps wide area experiments exist;*
9. *By 95/96, technology of choice seems to be SDH transmission and ATM switching.*

Computing at LEP

This paper, by Manuel Delfino, gave:

"A general review of computing at LEP is given, with emphasis on event reconstruction, data flow and storage, and physics analysis. The description is from the point of view of the strategies used to solve the problems; the reader is referred to the parallel session talks in CHEP91 for technical details and information which is specific to each of the four LEP detectors."

To some extent, if the MUSCLE report represented the theory, this paper described the practice – as well as future directions. Manuel highlighted the speed with which the LEP collaborations had turned data into papers, with about a dozen submitted for publication by the end of 1989 – the first year (months) of data taking. This he attributed to not only the careful planning carried out over a number of years between the host laboratory and the collaborating institutes, but also due to the involvement of industrial partners. He also highlighted the use of "emerging technology", namely workstations, where financial constraints were a contributing factor.

He described the data flow for the LEP experiments, noting that they produced their DSTs in quasi-real time. Whereas these and the raw data were recorded onto 3480 cartridges, analyses proceeded from "staged" disk copies that were also distributed to outside institutes. He noted the use of private workstation clusters – at that stage mainly VAX and Apollo but with Unix systems of various brands being installed. Physics analysis typically took place firstly in batch, to produce n-tuples that were then further analysed using PAW. In parallel, there was the simulation chain where the increased use of workstations was expected.

He then discussed the use of dedicated workstation farms in the data chain and noted how the LEP experiments had preferred the use of standard commercial components as opposed to emulators used at FNAL and SLAC.

Following a discussion on the importance of networking, both in the local and wide area, he noted the pragmatic choice of TCP/IP, with support for OSI largely unavailable.

He concludes by noting that LEP computing is *"just beginning"* but that *"planning for computing continues in the same spirit as before".*

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Quasi-online data processing with the ALEPH event reconstruction facility: one year of total success

This paper described:

"The ALEPH Event Reconstruction Facility (sometimes known as FALCON), delivers fully reconstructed event files for physics analysis within a few hours of data acquisition. The facility is fully automated and operates with minimal supervision. The immediate availability of data-summary files greatly facilitated timely publication of physics results and made sophisticated monitoring of the detector possible. The system is composed entirely of hardware which is commercially available and maintained, namely a VAXcluster using 12 VAXstation 3100 as batch processors. Events from a single data-acquisition run are processed in parallel, using modern disk random access techniques for work distribution to the processors, with the original time ordering of the event stream preserved on output."

What was particularly ingenious about this system was:

"The only unusual feature of the hardware is the use of dual-ported disks to transport the raw data between the Data Acquisition (DAQ) VAXcluster and the reconstruction facility. This is accomplished by exploiting the dual data ports on DEC RAseries disks (RA90 in our case), usually used for automatic failover in disks hosted by HSC disk controllers in a single cluster. In our case, one of the ports on the disk is connected to an HSC on the DAQ system, the other to a KDA5O0controller on the reconstruction cluster. The disk hardware automatically protects against simultaneous access by both clusters, and synchronization of access is controlled by simple DECnet messages".

Rapid access to event subsamples in large disk files through random-access techniques

Historically, analysis was often performed from data on tape. Where disks were used, they were frequently treated as "fast tapes" – namely accessed sequentially. Increasingly, and not just at LEP, the use of event directories was made, as described in this paper from ALEPH:

"Although large disks have been used to hold data for physics analysis for many years, their random-access capabilities have rarely been exploited. We present a description of the use of random access techniques, as implemented in the Event Directory concept of the BOS package, for physics analysis of data from the ALEPH detector at CERN. Only one copy per computer system of the full data sample is kept on large disks. Events are assigned to classes according to their topology, and users can access only events fitting a certain logical combination of these classes transparently. In addition, each user can save a private directory, containing pointers to events selected by any given algorithm, for direct access in later analysis steps. The impact of this technique on a variety of issues, from analysis job turn around to data management, is explored".

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

And in conclusion:

"The implementation of true random access to disk files in the BOS package has been used as a basis to organize a large fraction of the data analysis for the ALEPH detector, from event reconstruction to physics analysis. Management of data files and disk resources is greatly simplified, as only one master copy per computer installation is necessary and users do not generate large data files. The I/O load on the computers is reduced proportionally to the density of data being selected, resulting in substantial gains in efficiency. The ease with which event subsets can be shared results in more effective cooperation between members of the collaboration".

Missing Papers

The following papers have unfortunately not been found online. Whilst I believe that I remember each and every one of these talks, I have refrained from providing a personal summary – likely to be inaccurate and probably not representative after such a long period of time.

As a last resort, you only have to beetle off to the central library and request a copy of the full proceedings. If you do so, please consider scanning the document(s), OCR-ing them and adding them back to the CERN Document Server.

1. SHIFT: The Scalable Heterogeneous Integrated Facility for HEP computing, J.P. Baud(CERN), J. Bunn(CERN), F. Cane(CERN), D. Foster(CERN), F. Hemmer(CERN) et al.
2. The network and communications environment for a large collaboration by Mount, R P
3. Reality of software engineering in high energy physics, J. Knobloch
4. The L3 database system, B. Adeva, P. Bagnaia, S. Banerjee, L. Barone, D. Boutigny et al.
    - (A detailed paper on DBL3 from the same year can be found via https://doi.org/10.1016/0168-9002(91)90117-9)
5. Database management and distributed data in high energy physics: Present and future, L.M. Barone(Rome U. and INFN, Rome)
6. The DELPHI off-line packages: Difficulties and trends, G. Grosdidier(Orsay, LAL) for the DELPHI collaboration.

Summary of CHEP 1991

At this stage, everything still seemed to be going smoothly although further changes were expected. Increased "competition" between the SSC and LHC was visible and numerous activities related primarily to the former were presented. Of course, at that time, the foreseen starting dates of both of these machines were unrealistically close –later in the same decade. And the rest, as they say, is history.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

CHEP 1992

By CHEP 1992, much of the focus had moved on to the LHC, SCC, Hera, the Tevatron and other machines. There were numerous talks on Object Orientation, plus the use of more standard tools than those developed within HEP. The papers that are summarized below are those specifically oriented to LEP offline computing.

(On the opening morning, a status report of the recent "HEPLIB[22]" initiative was scheduled. Unfortunately, the speaker had not been permitted to travel and so I was asked, at around 22:00 on the previous evening, to stand in for him – but without any material. You can read the report in the CERN Yellow Report that covered this conference. This conference was also where the WorldWideWeb was first widely presented to the HEP community and where it won the award for the most influential paper).

The Offline Analysis System of the OPAL experiment

"We present the offline analysis system of the OPAL experiment at LEP. The software design, development and distribution system are described. The use of cpu power, and flow and management of data are tracked from the event builder through reconstruction and analysis at CERN and export to institutes. Particular attention will be given to our workstation-based projects to do as much mechanical data handling as possible centrally (at CERN), liberating physicists throughout the collaboration to develop better algorithms and physics analyses".

This paper, by Steve O'Neale, has many high-level similarities of those from other experiments – modulo some details regarding the computing hardware and architecture. Whilst OPAL physicists also followed courses on SA/SD, they were less enthusiastic about its practicality:

"Software engineering courses in SA/SD were taken by some OPAL physicists. Although SA data flow diagrams, maintained by the PROSA package, proved to be a valuable aid in design, the general application of the methodology was inhibited by the poor support tools".

More details were given regarding their offline software, although again there were similarities to be found with other experiments:

"The ZEBRA memory management package and Fortran were adopted as soon as they were available. Coding conventions emphasized strict adherence to the

---

[22] There had been a HEPLIB meeting at KEK between the two CHEPs that I attended on behalf of CERN. Arriving in Tsukuba by bus, I got into a taxi and asked him to take my to my hotel. After much heated discussion, with neither of us understanding the other, he gave in, and drove me to the hotel that was on the other side of the road. For some reason I did not include this in my report.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Fortran standard so that programs could, if necessary, be run on online and offline systems throughout the collaboration.
The simulation program GOPAL, based on GEANT, defines the detector geometry and tracks particles through this geometry to simulate their interactions and the response of each sub-detector of OPAL. The event reconstruction (or data reduction) program, ROPE, processes raw data and calibration constants to write a DST (data summary tape). The ROPE kernel is also used to read DSTs and write histograms, n-tuples and selected events. The event-viewing program, GROPE, may run the ROPE processors interactively. GROPE is based on GKS-3D and may also be run with PHIGS via GOSIP. In general, OPAL has made great use of and many contributions to the packages in the CERN Program Library".

He concluded:

"The success of OPAL's offline system is evident from our published physics results, our contributions to the common software of the LEP era, and the technical work cited in this paper. The workstation systems developed in joint projects with OPAL may also be of lasting value to the HEP community. We recommend central DST analysis facility to other experiments as, with current technologies and HEP budgets, it is much easier to install cpu power close to the data than it is to disperse bulk data throughout a collaboration. The planning, funding, and manpower for the analysis system should be given proper consideration at the proposal stage of an experiment".

The OPAL Monte Carlo production system

Here Steve O'Neale described how resources spread across the collaboration, i.e. not just in the local area, were used to build:

"The OPAL Monte Carlo Production System receives files of our vectors and yields simulated raw data and analysed data summary files. The management, monitoring and documentation is based on run and event numbers. Processing is performed on several workstation systems throughout the collaboration. Data collection and distribution is primarily on Exabyte 8mm cassettes with some use of IBM mainframes as servers for 3480 cartridges".

Specifically, he mentions resources (and their contributions) at CERN, Montreal, RAL, Saclay and Cambridge. He concluded by saying:

"The OPAL Monte Carlo production system runs successfully on multimillion-dollar mainframes and on many small clusters of workstations. The Monte Carlo physics analysis needs for 1993 can be met with an affordable number of workstations".

Data Storage and Network in Offline Analysis of ALEPH

In this paper Ronald Hagelberg discussed the use of mainframes together with VMS and Unix clusters for data analysis.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

"In large HEP Experiments, CPU is more and more distributed with the increasing use of workstations. A typical analysis accesses several Gigabytes of data. For computer platforms linked by fast networks, redundancy of data should be avoided to simplify management.

Taking the ALEPH experiment as an example, we show the approach taken for networking and data storage. ALEPH at CERN uses a CRAY, an IBM mainframe, a centrally-operated UNIX cluster "SHIFT", and a locally managed UNIX and VAX cluster for data analysis.

Efficient and transparent access to the data is ensured for workstation users by a fiber-optics based FDDI link between several computer platforms".

Distributed physics analysis in the OPAL experiment

Similar to the above in using a combination of mainframes, a VAXcluster and Unix stems, OPAL described:

"A software system for analysis of OPAL data was developed in which the physicist submits an analysis job in a generic language, the job is split over the CPUs which have easy access to the data and the results are merged and sent back to the user. Data access is local in this system, thereby avoiding network congestion. Direct access to events is provided through the use of event lists. Heavy use is made of CERN packages ZEBRA, HBOOK and FATMEN".

The system – Management of Analysis by Workstations (MAW – to rhyme with PAW) provided:

"a system in which low-cost workstations are used efficiently for physics analysis. A generic job syntax shields the user from the intricacies of the various operating systems. Event level parallelism and direct access to data are used to increase job turn-around. A standard programming language (FORTRAN) and standard output formats (histograms and Ntuples) are essential for achieving the parallelism. The system has low network requirements and can therefore be scaled up easily to the larger needs we expect for future HEP computing".

Experience with a UNIX workstation cluster at ALEPH

This paper, by J. Hilgart, discusses:

"A UNIX-based workstation cluster is being used by a growing community of physicists at ALEPH. These stations serve as both personal scientific workstations and batch processors in an offline environment. We describe the experience gained in managing such a cluster and offer advice on some useful system management utilities. Since sharing of a large disk pool (80 GBytes) resident on a VAX cluster is required, good network connections between the VAX and UNIX clusters were implemented, including use of remote file sharing software (NFS), and fiber-optic based computer links (FDDI). Efficient UNIX cluster to UNIX cluster networking is also realized utilizing software written at CERN".

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

In his conclusions he notes:

"The trend of CERN's computing services to offer consulting is extremely useful. To our dismay, we needed lots of support in areas where the software standards of Ultrix fall short of those for VMS. Specifically, we have been hampered by lack of effective software in cluster management, backups, compilers, debugger, and batch queues".

DELFARM, the DELPHI offline production farm

A hybrid farm, described by Michele Michelotto, addressed a common problem:

"The Delphi experiment at LEP, CERN requires a large processing power. The raw data coming from the online computers must be processed within few hours and several reprocessing steps are needed to reach final resolution. A farm of fast and cheap Unix workstations performs this task and produces the DST files for the whole collaboration".

But with a different hardware approach:

"The new farm makes use of seven Unix workstation (Decstation 5000/200, with the MIPS 3000 RISC processor rated 19 Specmarks, running Ultrix operating system and diskless) and of four Vaxstation 3100/76, each equipped with SCSI disks and a 3480 cassette drive. All were connected on a single Ethernet segment".

This "production" farm was complemented by a Unix-only "analysis" farm, where already-debugged code could execute in parallel, producing n-tuples and / or histograms as desired.

SHIFT from a physics point of view

This paper described the SHIFT system from an architectural, hardware and functional point of view. Whereas the hardware setup maybe of historical interest, the system is also described in this CERN Courier article, as well as elsewhere: https://cerncourier.com/a/a-major-shift-in-outlook/

Data management using the FATMEN package

The main features of the FATMEN package have already been described above. It inherited many ideas from DBL3, including the method of communicating updates between clients and servers as ZEBRA FZ files in ASCII format. Servers existed for VM/CMS, VAX/VMS and Unix. It allowed users to reference data, independent of media, operating system or location, using a "logical name" in pseudo Unix format, again was described above. Several hundred thousand names per experiment existed, stretching to the limit the capabilities of ZEBRA RZ that originally used 16 bit fields as record pointers.

HEPDB: a database management package for high energy physics

HEPDB was package that was very largely based on DBL3 but included an OPCAL-like interface. It was included in PACKLIB and (subsequent to this conference) adopted by a number of non-LEP as well as non-CERN experiments. Its scope was basically that of its parent packages, namely DBL3 and OPCAL.

CHEOPS: an efficient data dissemination system

This paper described an attempt to use satellite communications to distribute data and was an initiative strongly supported by the DG of the time. Unfortunately, it never made it to production status but we were not to know that at the time.

"The CHEOPS project is a collaboration between CERN, Demokritos (Greece), LIPAINESC (Portugal) and SEFT/HUT/CSC (Finland) helped by Technical University Graz (Austria). The main goal is to investigate bulk dataset dissemination between CERN and the High Energy Physics (HEP) laboratories of its member states. This requires high data network bandwidth which is an expensive resource in Europe. The key point for such a facility to become affordable is to take advantage of lower tariffs applied to communications lines during unsocial hours (mostly night time). This led to the definition of CHEOPS as an automat which queues data transfer request issued by physicists during working hours and performs the transfers with maximum efficiency when network capacity is available at low cost.
Another important aspect of a data dissemination facility is the fact that several CERN member-states are geographically far distant or do not yet have appropriate terrestrial infrastructure; this is becoming more serious as new states from Eastern Europe need access to CERN data. A solution is high speed point to point bit stream transmission by satellite, today a well understood technology".

Summary of CHEP 1992

From the presentations at this conference[23], the question on the role of the mainframe raised at the previous event seemed to have been largely answered. Yes, perhaps as a fileserver, but even this aspect was soon to disappear. Otherwise, farms were everywhere.

CHEP 1994

This conference featured a paper *"Data handling at CERN for LEP and future experiments"* (p. 245) by Mount, R P (but no online version).

However, 74 contributions can be found through InspireHEP, some of which have associated PDF files: https://inspirehep.net/conferences/971233.

[23] Who can forget the conference dinner and cruise on Lake Annecy with "Les Cernettes"?

Monday 19 October 2020
Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

CHEP 1995

This was the last CHEP conference during the LEP data taking era with multiple presentations from the LEP experiments, plus some covering CERNLIB developments. The conference home page no longer exists but the proceedings are available in the CERN library both on CD as well as printed copies. The following entry was found, with some difficulty, in InspireHEP: https://inspirehep.net/conferences/969392.

Using WWW to improve software development and maintenance: application of the LIGHT system to ALEPH programs

Programmers who develop, use, maintain, modify software and faced with the problem of scanning and understanding large amounts of documents, ranging from source code to requirements, analysis and design diagrams, user and reference manuals etc. This task is non trivial and time consuming LIGHT, LIfe cycle Global HyperText is an attempt to solve the problem using WWW technology. The basic idea is to make all the software documents, including code, available and cross-connected on the WWW. The first application of this concept to go in production is JULIA?LIGHT, a system to convert and publish on WWW the software documentation of the JULIA reconstruction program of the ALEPH experiment at CERN.

CHEP 1997

No LEP-related papers found (but the proceedings can be found in https://www.sciencedirect.com/science/article/pii/S0010465598800040).

CHEP 1998

See https://inspirehep.net/conferences/971233 for PDF versions of 28 papers from this conference, none of which refer to LEP.

These do not include to a potentially relevant submission to this conference, namely N. Smirnov, F. Carena and Tz. Spassoff, *"Data Archiving in the DELPHI Experiment"*, referred to the paper below from CHEP 2K.

CHEP 2000

By the time of CHEP 2000, the focus had moved firmly to the newer and / or still future experiments. Only two submissions referred explicitly to a LEP experiment – both coincidentally DELPHI. One paper covered lessons learned over the previous decade (the DELPHI control system) whereas the other was more forward looking:

*"The paper (*http://chep2000.pd.infn.it/short_p/spa_f071.pdf*) presents the current status of the project for building object-oriented programming environment in the DELPHI experiment.* ***The main project goal is to make possible the analysis of the DELPHI data far beyond the end of the LEP data***

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

taking. *The special attention is paid to software reuse to provide a contribution to the CLHEP library".*

Whilst the goals of this project were indeed very noble, it underlines one of the key messages of the DPHEP Blueprint (https://arxiv.org/pdf/1512.02019.pdf) that came much later (May 2012), namely that constant effort is needed, rather than a one-off exercise (in this case migration from ZEBRA format to Objectivity/DB, long since dropped as a potential solution for the storage of LHC data).

This allows us to segue smoothly into the next topic, namely preservation and re-use of (LEP) data.

DPHEP 2009

In early 2009, at the first workshop of the newly formed Study Group on Data Preservation and Long Term Analysis in High Energy Physics, more commonly known as DPHEP, a presentation was made on Data Preservation at LEP. Despite the lack of any formal commitment from the start of the experiments, there was clearly a Data Preservation success story close to one decade after the end of data taking.

The abstract is reproduced below, with the full paper available at https://arxiv.org/pdf/0912.1803v1.pdf.

The four LEP experiments ALEPH, DELPHI, L3 and OPAL successfully recorded e^+e^- collision data during the years 1989 to 2000. As part of the ordinary evolution in High Energy Physics, these experiments cannot be repeated and their data is therefore unique. This article briefly reviews the data preservation efforts undertaken by the four experiments beyond the end of data taking. The current status of the preserved data and associated tools is summarised.

Some key points from the paper are repeated, should they offer inspiration to future data preservation efforts.

- *DELPHI for example has a software CD project which includes all DELPHI software in source form. It runs independently from AFS and includes everything to run the detector simulation and reprocess data, however not the physics generators.*
- *ALEPH went further and also conserved the computing environment. They distributed a 'mini-system' to each participating institute consisting of a laptop and an external disk containing the data and MC samples. It runs the Linux distribution used at the end of LEP as well the necessary ALEPH software to access it.*
- *OPAL attempts to keep its analysis environment alive by adapting the software to changes in the computing platform such as new OS versions and new staging software.*

…

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

- *There are also examples of parts of the software which stopped working: The event display programs of OPAL and DELPHI do not run on recent platforms. This is because they relied on commercial libraries and there is no person power available to re-write the complete event display at this stage.*

The software in question here is the PHIGS graphics package. At the time of LEP, a number of hardware vendors provided their own implementation whereas there were two commercial implementations – TGS in the US and G5G in France. These companies have since merged a new port of the libraries to a "recent" operating system was paid for by CERN. How long this port will remain usable for is yet to be seen.

The LEP experiments also relied heavily on software from the CERN Program Library (CERNLIB), including ZEBRA and also PAW. ALEPH used the Bos I/O system from DESY. CERNLIB has not been officially supported for many years now and some of its implicit assumptions, such as the ability to be able to store a relative address in a 32 bit integer, has already broken once, with a work-around at the link / load level.

CHEP 2012

CHEP 2012 marked a watershed moment in the domain of data preservation in HEP. Not only did it happen in the same month as the DPHEP Blueprint document (see https://arxiv.org/pdf/1512.02019.pdf) was published but also it included a half-day parallel workshop at which some significant announcements were made. Most importantly, CMS presented their Open Data policy – something similar to which was then or later adopted by the other LHC experiments. However, this is not intended to be a DPHEP paper: for more complete and / or recent news concerning the activities of DPHEP, initially adopted by ICFA as a "panel" and more recently transitioned from a Study Group to a Collaboration, please see the DPHEP Website: https://dphep.web.cern.ch/.

A not uncommon comment regarding CHEP is that the format, particularly the parallel sessions with time for just a small number of questions, does not permit the discussion that a workshop format would. Also, particularly regarding DPHEP, there has been a distinct lack of flexibility with the exception of CHEP 2012 in New York. (Gathering such a large fraction of the community together is not something that can be easily repeated and a few scattered talks throughout a conference in no way replaces e.g. a half-day, parallel, voluntary, walk-in walk-out workshop. A dedicated DPHEP workshop typically attracted 30-40 people. Even if all attended for an entire session, this could hardly impact on a CHEP conference overall).

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

It is also true that similar themes come up regularly, even on occasion during the same event, but without (for example) a panel session[24] to attempt to compare the various proposed solutions and / or experiences.

So do we actually learn long-term lessons from CHEP or just take point-in-time snapshots?

(One could also comment on the difficulty in finding the proceedings from past CHEPs, the discrepancies between different search routes, the poor quality of some of the scanned documents (which should only be used as a last resort and OCR-ed should the original published version not be available).

DPHEP "Full Costs of Curation" Workshop

Of the many DPHEP workshops, held either during the "Study Group" phase or else after DPHEP had evolved to become a Collaboration, one focused on estimating the "Full Costs of Curation" - https://indico.cern.ch/event/276820/.

Given that nothing is permanent, the overview of the workshop is repeated here.

The purpose of this workshop is to discuss the full "costs of curation" of HEP data, e.g. that from the LHC, over a period of several decades (say 1, 2, or 3).

To look that far into the future, the past can be used as a "guide" as to the scale of the changes that might occur.

25 years ago, the LEP collider was about to start operation - which was shortly followed by an era of significant change in offline computing that continues to this day.

30 years ago, the Z and W had only recently been observed, Wylbur and MVS were still the main central computing services at CERN and central interactive services (VM/CMS, VAX/VMS) were only just making their debut.

The "programming language of choice" was Fortran - albeit often with extensions at the language level (e.g. VAX Fortran) or through libraries (e.g. Hydra, ZBOOK, ZEBRA etc.)

(The default terminal settings - 24 x 80[25] - are a legacy from this era.)

---

[24] There was a so-called panel session on databases at CHEP 1992 but it was more a selection of prepared presentations with no real discussion about the different viewpoints. See CHEP 1992 above.

[25] Both of these are attributed to the use of punched cards (the former) and FORTRAN. Punch operators punched out 24 lines of code and then took a rest. FORTRAN statements were written in columns 7 to 72 with columns 1 – 5 being used for statement labels, 6 to indicate a continuation of a previous statement

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Storage capacity was minimal (the VXCERN service, originally based on a VAX 8600 or "Venus", had 6 disks of less than 500MB capacity each: one for each LEP experiment, one for the "system" and one for scratch space).

Today, mobile phones have internal storage of several GB, GHz multi-core processors and are pseudo-disposable devices.

What changes will take place in the future?

This is of course unknown, but we do know that there are clear scientific reasons for keeping the data from today's current experiments - as well as those from the recent past - fully usable until well into the future.

This workshop is about establishing the costs of such an exercise.

Having established these costs, according to a number of possible scenarios (e.g. "best case", "worst case"), this information will then be used for resource and budget planning.

Summary

The above represents an imperfect and incomplete review of computing for the LEP experiments based on the material that can still be found without too much pain – largely documents that can be found in, or referenced from, the CERN Document Server. Unfortunately, we have been rather poor at preserving the narrative, not to mention the software and much else from this magnificent part of CERN's history.

The LEP era saw multiple overlapping disruptive changes in the area of computing, networking and storage – no stone was left unturned. Starting firmly in the era of heterogeneous mainframe, batch oriented, FORTRAN-based computing with embryonic networks and collaborative tools, it saw the emergence of distributed computing, first in small clusters, later in large farms and finally grid computing. Not even the programming language was sacrosanct despite many predictions to the contrary.

Surely there are lessons to be learned from the LEP experience – the sometimes painful steps do not have to be repeated and the knowledge gained, also from (HL-)LHC and beyond should imply starting from a much higher base. (Networks may well have terabit / second capability or even more; storage at CERN may have exceeded the Exabyte barrier and CPU capacity – however it is delivered – will also have developed enormously since the era of the "CERN unit").

---

and "sequence numbers" in the trailing columns. Comment "cards" had a C in column 1.

Jamie.Shiers@cern.ch
https://orcid.org/0000-0001-8404-3750
DOI:10.5281/zenodo.4018493

Whereas LEP started with no "managed storage" (each experiment had a set of two letter prefixes for their 6 character tape volumes), there are now multiple copies of the remaining data (two on tape; one on disk at CERN alone) and "bit preservation" is performed using the same technologies as for LHC data.

Although adapting to these changes was often a challenge, consider where we would be if LEP data taking had ended in the VAXcluster / IBM / Cray era? Almost certainly, the data, software and much else would be unusable, as is the case for (most?) pre-LEP experiments[26].

The total data volume is now trivial but preserving the software, documentation, meta-information and knowledge certainly is not. There are those that argue:

- If / when a new *fcc-ee* (or other Higgs factory) is built, its precision and statistics will soon make the LEP data totally redundant;
- Even if not, can there be any information left that we failed to extract at the time, when we had fully trained active collaborations?

However, the cost[27] of keeping the data "alive" is so small; a minute fraction of what it cost to acquire and analyse initially.

Preservation is an activity that cannot be left to individuals but needs to be part of an overall policy of the (virtual) organisations involved. Much information, data and knowledge has already been lost regarding the once flag-ship LEP experiments. Consider also the many pre-LEP ones, where shadows remain in the Particle Data Book and in the HepData database.

Whilst in theory it is never too early to consider long-term data preservation, in practice it often turns out to be too late. ***Schade.***

Acknowledgements

The authors of the individual papers and reports summarized in this document deserve all the credit for the hard work and constant innovations that were a key characteristic of computing for LEP. The humble scribe deserves all the blame for any errors, omissions or misrepresentations.

[26] OK, not JADE, but what about the bulk of those experiments listed as "completed" in the CERN Grey Book?

[27] The author has argued that the cost of bit preservation tends to zero and that of software (and other) preservation tends to a fraction of an FTE. But in neither case exactly zero. See https://indico.cern.ch/event/276820/. Whilst the data is being used even for rare publications, this "value" is normally considered greater than the ever-diminishing cost.