

LHC Computing – the First 3 Decades

Jamie Sheirs^{1,*}

¹ CERN, 1 Esplanade des Particules, 1211 Geneva 23, Switzerland

Abstract. Computing for the Large Hadron Collider (LHC) at CERN arguably started shortly after the commencement of data taking at the previous machine – LEP – some would argue it was even before. Without specifying an exact date, it was certainly prior to when today’s large(st) collaborations, namely ATLAS and CMS, had formed and been approved and before the LHC itself was given the official go-ahead at the 100th meeting of the CERN Council in 1995. Approximately the first decade was spent doing research and development; the second – from the beginning of the new millennium – on grid exploration and hardening; and the third providing support to LHC data taking, production, analysis and most importantly obtaining results.

1 Introduction

Computing for the Large Hadron Collider (LHC) at CERN arguably started shortly after the commencement of data taking at the previous machine – LEP [1] – some would argue it was even before. Without specifying an exact date, it was certainly prior to when today’s large(st) collaborations, namely ATLAS and CMS, had formed and been approved and before the LHC itself was given the official go-ahead at the 100th meeting of the CERN Council in 1995.

Approximately the first decade was spent doing research and development; the second – from the beginning of the new millennium – on grid exploration and hardening; and the third providing support to LHC data taking, production, analysis and most importantly obtaining results.

Just as with LEP computing, the challenges looked at some (many) points both insurmountable and the costs unaffordable. At the time of writing, we have perhaps only reached a false summit but we are now joined by other disciplines facing similar issues. Thus, we can look forward to the future – both of the LHC and any eventual “successor” machines, with cautious optimism.

Whereas some of this story is told in the book “From the Web to the Grid” [2], the purpose of this document is not to repeat what is written there. Rather, it attempts to capture a few additional elements of the story and to provide pointers to documents that still exist recording the path, on occasions painful, to what can only be considered (great) success – not just in terms of physics but also computing. Thus, it is far from being a complete record of LHC computing but at least part of what entered the consciousness of this particular being.

Whilst many may consider that it is (far) too early to even contemplate a “history” of LHC computing, a significant amount of information has already been lost; more still is hard to find and the non-silicon based memories are fading even faster. Thus, this document maybe considered a “stub”, to eventually be completed with material from relevant CHEP presentations, CERN School of Computing lectures, as well as that from documents referenced below.

Just don’t leave it too late...

2 Mass Storage

Whereas the issue of storage was always a major concern for LEP, it was close to a nightmare when considering the data volumes and rates foreseen at the Superconducting Super Collider (aka “Clyde”) and / or the LHC. Recall that it was originally foreseen that LEP data taking would at least start with 6250 bpi round tapes, with the move to “3480” cartridges with a capacity of 200MB – only a very small increase over that of the open reel volumes – being a major step forward. Given that of the order of 1 million such cartridges would have been needed per LEP experiment, the LHC would need something like 1 billion! Therefore, it was not surprising that other technologies, such as high capacity video devices, were investigated in the early 1990s. Whilst these did good service for some high data volume experiments at that time, such as the NA49 experiment at CERN, fortunately technology moved on significantly by the end of the decade and beyond. Of course, the startup date of either of the

* Jamie.Shears@cern.ch

then-future hadron colliders was unknown at the time but predicting the hardware of choice 1 – 2 decades into the future remains essentially impossible.

From about the same time, i.e. the early 1990s, HEP began to look at the IEEE Reference Model for Mass Storage Systems that had been developed through and around a series of conferences and workshops that continue to this day: <https://www.storageconference.us/history.html>.

The model as it stood at the time, together with some systems that claimed to be based on it, can be found in the lecture notes from the 1993 CERN School of Computing. These have been extracted and converted back to machine-readable form at <https://zenodo.org/record/4139255>.

Some of today's WLCG sites still run systems based on those described nearly three decades ago and / or claim to have been designed according to the reference model.

3 The Rise and Fall of Object Databases

Whereas, at the time of LEP, the programming language of choice was expected to remain “Fortran[†] forever”, the possibility of different languages, and in particular those supporting the object-oriented paradigm, were increasingly explored (and the rest, as they say, is history).

A series of R&D projects were kicked off in the mid-1990s, including MOOSE (RD41), RD44 (aka GEANT-4) and RD45 – the latter to study object persistency, that many took to imply databases, and their potential applicability to the storage of event data in HEP[‡].

The proposal that led to the RD45 project, as well as the regular status reports and other documents, can be readily found via the CERN Document Server (using search string “RD45”).

4 The Experiments' Computing Technical Proposals

Towards the end of 1996, ATLAS [3] and CMS [4] published Computing Technical Proposals[§]. Both of these are relatively slim documents of just under 100 pages each. Written at a relatively early stage of the above-mentioned R&D project, they capture very much thinking of the time, as well as a possible model of regional centres that was carried over into later projects.

Whilst reference is made to technology tracking working groups that went under the collective name of PASTA (the acronym did not cover the networking WG), the reports (1996, 1999) cannot be found, although a reference to them, from 2003, is so far still available [5].

It is surely not surprising that many of the explicitly named candidate “solutions” did not make it to the data-taking phase of the LHC – yet another reminder about the challenges of predicting the future.

Can this really be different for any future fcc-xx machine?

5 Towards a (Network) Model of Regional Centres

A project that influenced significantly thinking regarding LHC computing models was MONARC – Models of Networked Analysis at Regional Centres for LHC Experiments [6]. Although there was clearly a preference for a network-centric model – it some ways not too dissimilar to that deployed for LEP and other experiments – it was far from clear whether the necessary bandwidth would be either available or affordable. Hence, a “fallback” solution with a very limited number of such (“continental?”) sites was considered with “tapes in a suitcase” as distributed mechanism.

But you have to get the data in and out too!

6 Enter The Grid

Around the time that the MONARC project was wrapping up and coinciding with the dawn of the new millennium, the potential of grid computing greatly enthused many in HEP. This led to a series of projects, such as the EU DataGrid in Europe and others elsewhere, complemented by an LHC-oriented

[†] Fortran continues to be developed, with the most recent version of the standard dating from 2018.

Apart from support for object orientation that came many years ago, additional features include “co-arrays” that could prove useful for exploiting parallel architectures.

[‡] The issue of databases in HEP has featured regular in Conferences on Computing in HEP, is covered in the book “From the Web to the Grid”, and can also be found in a slightly shorter version in <https://cds.cern.ch/record/1056959?ln=en>.

[§] The ALICE experiment was approved in early 1997 and LHCb in the 2nd half of 1998. All experiments produced more detailed Computing Technical Design Reports “in the 3rd millennium”.

one, initially the LHC Computing Grid that is now known as WLCG – <https://wlcg.web.cern.ch/>. For this project, a wealth of documents [7] can be found, from design reports, through memoranda of understanding to on-going service reports – as well as plans for future LHC runs and upgrades [8].

Given how well the grid has performed, it is easy to overlook the considerable effort that was required to get to this point. So called “Service Challenges” were first proposed in 2004 [9], and these were followed by Common Computing Readiness Challenges as well as Scale Tests of Experiment production – the latter taking place in two phases in 2009 [10] – “just in time”.

7 “I think we have it”

... was how the Director General “summarized” the presentations by ATLAS and CMS on their searches for “the” Higgs boson. This continues to be studied in greater detail with ever increasing amounts of data, requiring more and more computational resources. However, the task of “getting the Grid to work”, as opposed to keeping it running and extending it, could be considered to be done. This was the conclusion of the referees of the EGI InSPIRE project that had a work-package [11] devoted to supporting HEP and other “heavy users” of the grid, as well as senior physicists (well, one at least) from the LHC experiments.

Job done – time to move on!

8 Whatever Next(?)

Future colliders – both lepton and hadron – are under consideration that will likely be natural successors to LEP and the LHC respectively. Starting possibly half a century after each of their forebears, technology for computing, networking and storage will have all advanced so significantly as to “seem like magic” to those working on the former machines.

The challenges that will be faced will surely be different but if there are lessons that can be learned from the past you need to start by recording – and preserving – as much information and “knowledge” as possible. Sounds easy, but it isn’t, as the LEP “experience” tells us.

If the past is indeed another country, what is the future?

9 Some Suggestions

The following suggestions, made during the presentation associated with this paper at vCHEP2021, focus on the record of HEP Computing that is associated with the CHEP conferences.

1. Ideally, there should be a clearly and unambiguously named “collection” covering all CHEP conferences – past, present and future;
2. This collection should be maintained in a “Trustworthy Digital Repository” (TDR) according to widely agreed definitions of this term;
3. For current and future conferences, where the material can be considered to be “born digital”, it should be stored as such, along with all appropriate meta-data (templates, mark-up instructions etc.);
4. For past conferences, as many missing papers as possible should be recovered from hard-copy or other sources;
5. In all cases, the material should be complemented by machine-readable versions, e.g. by passing through a (named) OCR program;
6. “Best practices” should be followed, such as the assignment of Digital Object Identifiers (DOIs) and the use of Open Researcher and Contributor ID (ORCID) – or the equivalent standards in the future;
7. In all cases, the route via which material is entered into the TDR should be clarified. Currently, is it better than “anyone’s guess”?

Ideally, the above would be part of a CERN / HEP-wide initiative regarding data / information capture and preservation but that is way beyond the scope of this paper and even this author!

Bon chance!

References

- [1] The story of computing for LEP is at least partially told in <https://zenodo.org/record/4139583>.

- [2] “From the Web to the Grid” – <https://cds.cern.ch/record/1141328?ln=en>.
- [3] The ATLAS Computing Technical Proposal – <https://cds.cern.ch/record/322322?ln=en>.
- [4] The CMS Computing Technical Proposal – <https://cds.cern.ch/record/322321?ln=en>.
- [5] Reference to PASTA at https://it-dep-fio-ds.web.cern.ch/Documentation/tapedrive/pasta_tape02.html.
- [6] MONARC reports: see <https://cds.cern.ch/record/510694?ln=en>, for example.
- [7] The WLCG “document repository”: <https://wlcg-docs.web.cern.ch/>. A safe place?
- [8] The TDRs, including those of the LHC experiments and their upgrades, can be found via <https://cds.cern.ch/collection/LHCC%20Public%20Documents>.
- [9] See, for example, <https://cds.cern.ch/record/853112/files/cer-002536189.pdf>.
- [10] See, for example, https://cds.cern.ch/record/1322146/files/CR2011_015.pdf.
- [11] See <https://documents.egi.eu/public/ShowDocument?docid=742>.