A Roadmap for HEP Software and Computing R&D for the 2020s

Eckhard Elsen¹

Published online: 29 November 2019
© The Author(s) 2019

For a while it seemed as if Scientific Computing was doomed; the fact that technological advancement no longer obeyed Moore's law, which had so conveniently allowed an ever-increasing number of computing tasks to be performed, had become obvious. The electrical power density in silicon cannot easily be increased further. And storage technologies are hitting similar limits; storage on disk is limited by density and mechanical positioning, which affect readout speed—yet the need for more computing power and storage space is undeniable.

The Large Hadron Collider at CERN in Geneva is one of the world's largest producers of scientific data. In the course of various experiments, it has produced roughly an exabyte of raw data. This treasure trove of data forms the basis of all physics analyses from the LHC. Further, given the nature of the research conducted at the LHC, the data from the first runs in 2010 is as valuable as that from today, or that from 2035. The experiments are statistically limited. An upgrade of the LHC, the High-Luminosity LHC, will go online in the mid-2020s and feature upgraded detectors with much finer readout granularity. By that time, the computing demands will have increased tenfold if the same computing paradigms are used as today.

Other large-scale research facilities also have a growing need for scientific computing resources; sensors are getting smaller and are sampling faster. As a result, there is more and more data to be investigated and analysed. Examples include the large telescopes (fields) that are now being designed, the plans for gravitational wave detectors, and fast sampled scattering data from light sources—to name but a few.

The situation is exacerbated by the fact that the complex readings taken by these instruments can only be understood using advanced simulation. In many cases, the computing

demands for simulation exceed those for reconstructing the data.

Quo vadis scientific computing?

The roadmap [1] represents a bottom-up discussion in the particle physics community, which ultimately produced a Community White Paper, the basis of this report. The community took stock of current needs and attempted to identify target areas for increasing efficiency, e.g. by adapting computing models, using new software algorithms, exploring new hardware architectures, and using modern software tools to effectively enhance detector resolution. They also recognised the need for data preservation, i.e., making scientific data accessible in a form that can also be understood by laymen in the future.

The result: a Programme of Work for the field as a whole, a multifaceted approach to addressing growing computing needs on the basis of existing or emerging hardware. And following in the fine tradition of open source software, it has been organised as a community effort, guided by the particular interests of individuals. Having seen these experts' enthusiasm, I am convinced this will yield excellent and unexpected results.

While it is heartening to see the field almost naturally self-organise, it also comes with a tremendous down side. All too often, the experts engaged in the effort will not be recognised as the key drivers of breakthroughs in analysis, even though they are who made the analysis technically feasible in the first place. An academic career is not readily open to them. Universities reward breakthroughs in physics, instrumentation or theoretical information technology, but rarely do the same for the handling of (truly) big data. In the current era of Artificial Intelligence, perhaps those persons should be recognised who apply this extra intelligence to help understand scientific results and improve detector resolution. Where are the faculty positions that we need in order to acknowledge the efforts of at least the very best of these experts, before they are lured away to work in industry?

Will the community effort be sufficient? Yes and no. After a few years, the effort will indicate potential directions for investment. Public funding authorities, who provide for

✉ Eckhard Elsen
eckhard.elsen@cern.ch

¹ Research and Computing at CERN, Geneva, Switzerland

the operational funds, have generally adopted the philosophy that the later an investment has to be made, the better. I am sure they will read the outcomes of the Programme of Work described in this text with great interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Reference

1. The HEP Software Foundation, Albrecht J, Alves AA et al (2019) A roadmap for HEP software and computing R&D for the 2020s. *Comput Softw Big Sci* 3:7. <https://doi.org/10.1007/s41781-018-0018-8>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.