

THE FRONTIERS COLLECTION

THE FRONTIERS COLLECTION

Series Editors:

A.C. Elitzur L. Mersini-Houghton M. Schlosshauer M.P. Silverman R. Vaas H.D. Zeh
J. Tuszynski

The books in this collection are devoted to challenging and open problems at the forefront of modern science, including related philosophical debates. In contrast to typical research monographs, however, they strive to present their topics in a manner accessible also to scientifically literate non-specialists wishing to gain insight into the deeper implications and fascinating questions involved. Taken as a whole, the series reflects the need for a fundamental and interdisciplinary approach to modern science. Furthermore, it is intended to encourage active scientists in all areas to ponder over important and perhaps controversial issues beyond their own speciality. Extending from quantum physics and relativity to entropy, consciousness and complex systems – the Frontiers Collection will inspire readers to push back the frontiers of their own knowledge.

René Brun • Federico Carminati
Giuliana Galli Carminati

Editors

From the Web to the Grid and Beyond

Computing Paradigms Driven by
High Energy Physics

 Springer

Editors

René Brun
Federico Carminati
CERN
Geneva
Switzerland
Rene.Brun@cern.ch
Federico.Carminati@cern.ch

Giuliana Galli Carminati
Hôpitaux Universitaire de Genève
Unité de la Psychiatrie du
Développement Mental
ch. du Petit Bel Air 2
1225 Chêne-Bourg
Switzerland
Giuliana.GalliCarminati@hcuge.ch

Series Editors:

Avshalom C. Elitzur
Bar-Ilan University, Unit of Interdisciplinary Studies, 52900 Ramat-Gan, Israel
email: avshalom.elitzur@weizmann.ac.il

Laura Mersini-Houghton
Dept. Physics, University of North Carolina, Chapel Hill, NC 27599-3255, USA
email: mersini@physics.unc.edu

Maximilian A. Schlosshauer
Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences,
Boltzmannngasse 3, A-1090 Vienna, Austria
email: schlosshauer@nbi.dk

Mark P. Silverman
Trinity College, Dept. Physics, Hartford CT 06106, USA
email: mark.silverman@trincoll.edu

Jack A. Tuszynski
University of Alberta, Dept. Physics, Edmonton AB T6G 1Z2, Canada
email: jtus@phys.ualberta.ca

Rüdiger Vaas
University of Giessen, Center for Philosophy and Foundations of Science, 35394 Giessen,
Germany
email: ruediger.vaas@t-online.deH.

Dieter Zeh
Gaiberger Straße 38, 69151 Waldhilsbach, Germany
email: zeh@uni-heidelberg.de

ISSN 1612-3018

ISBN 978-3-642-23156-8

e-ISBN 978-3-642-23157-5

DOI 10.1007/978-3-642-23157-5

Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2011941933

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

Modern High Energy Physics (HEP), as a science studying the results of accelerator-driven particle collisions, was born after the Second World War, at the same time as computers. HEP research has been constantly limited by technology, both in the accelerator and detector domains as well as that of computing. At the same time High Energy physicists have greatly contributed to the development of Information Technology.

During all these years, the Conseil Européen pour la Recherche Nucléaire¹ (CERN [1], located in Geneva, Switzerland) has been a privileged place for the evolution of HEP computing. Several applications conceived for HEP have found applications well beyond it, the World Wide Web (see Chap. 2) being the most notable example. During all these years HEP computing has faced the challenge of software development within distributed communities, and of exploiting geographically distributed computing resources. HEP computing has been very successful in fulfilling its mandate, and the quantity and quality of scientific results produced by High Energy Physics are a direct consequence of this. In this activity, HEP computing had to face problems and situations that anticipated those that Information Technology at large would meet years later. In many respects, HEP computing has been today where computing was going to be tomorrow.

This book describes the evolution of HEP computing, and in particular those aspects that have been most innovative. These aspects are described by contributions from different authors. The reader is presumed to be familiar with the field of computing, from software engineering to programming languages. No specific physics knowledge is required. The subject has been treated in the past mostly in conference proceedings or in specialised books on physics computing, but, to our knowledge, no other book has given a complete account of the evolution of HEP computing.

¹The Conseil Européen pour la Recherche Nucléaire, or European Council for Nuclear Research is the name of a provisional body founded in 1952 with the mandate of establishing a world-class fundamental physics research organisation in Europe. At that time, pure physics research concentrated on understanding the inside of the atom, hence the word “nuclear”.

High Energy Physics in a Nutshell

The discovery in 1896 of the natural transmutation of the elements via radioactive decay by the French scientist Henri Becquerel, while working on phosphorescent materials, coincided with the discovery of “rays” composed by sub-nuclear particles of varying energy. These rays were soon used to probe the structure of matter, and using them the New Zealand scientist Ernest Rutherford discovered in 1909 that matter is composed by atoms with a very dense and small positively charged nucleus and negatively charged electrons floating in virtual emptiness. In 1912 the Austrian-American physicist Victor Hess during a hot-air balloon flight discovered the existence of highly energetic rays coming from outer space and hitting the earth after having traversed the atmosphere. While the origin and nature of cosmic rays still poses some of the most formidable puzzles in physics, this discovery was instrumental in many ways. The history of the universe and the nature of matter were since then intimately linked; the study of our origin and of the infinitely big now relies on advances in our understanding of the microscopic nature of matter, and vice versa.

Cosmic rays and the discovery of antimatter.

In 1936 the American physicist Carl David Anderson, studying cosmic rays, discovered antimatter, earlier postulated by the British physicist Paul Dirac. Cosmic rays also provided physicists with probes of energies much higher than those coming from nuclear transmutation. This was an important step in the investigation of the nature of matter. In his 1924 doctoral thesis, the French physicist Louis-Victor-Pierre-Raymond, 7th Duc de Broglie, introduced the successful hypothesis that with each particle is associated a wave whose length varies as the inverse of the particle energy. The “resolution power” of an instrument, that is its ability to distinguish small details, varies with the inverse of the wavelength of the radiation employed. Highly energetic particles are excellent probes into the sub-nuclear world, as the associated waves have very short length.

While cosmic rays provided very short length waves, their energy, as well as their arrival angle, was subject to a wide statistical distribution, making systematic studies very difficult. Physicists then decided to produce high energy rays in the laboratory by accelerating charged sub-nuclear particles and nuclei with electromagnetic fields. In the 1930s the British physicists John D. Cockcroft and Ernest Walton, the American engineer Robert Van de Graaf and a Soviet team in Kharkov managed to accelerate nuclei via the creation of an intense electric field. Vacuum tubes were used to accelerate electrons, and this led to the discovery of the photoelectric effect in 1887 by the German physicist Heinrich Rudolf Hertz. The explanation of this effect by Albert Einstein in 1905 led to the formulation of a corpuscular theory of light and was one of the founding steps towards modern Quantum Mechanics.

The birth of High Energy Physics

The use of very intense electric fields was however impractical due to the difficulty to generate and control them. Particles could be accelerated by guiding them repeatedly through lower intensity electric fields, but this required very long “linear” machines to achieve high energies. A major breakthrough was achieved in 1929 by Ernest O. Lawrence at the University of California, Berkeley. He managed to accelerate electrons with a relatively low voltage by passing them several times through a potential difference. Lawrence’s machine used a magnetic field to keep the electrons on a circular path, and an alternating voltage gradient to accelerate them. Lawrence’s first machine, called a cyclotron, was made out of brass, sealed with wax and had a diameter of ten centimetres; it is said to have cost \$25 in all. It was the first particle accelerator and it opened the way to the creation of sub-nuclear and nuclear rays of ever-increasing energy (hence the name High Energy Physics). Almost all high energy particle accelerators today derive from Lawrence’s machine. Due to fundamental physics laws, the diameter of the accelerator has to increase with energy, so modern machines have diameters of several kilometres and are hosted in underground tunnels. Several particle accelerators have been built since the end of the Second World War. Major laboratories exist in the United States, Japan, Europe and Russia. The largest of them all is at the European Conseil Européen pour la Recherche Nucléaire (CERN) in Geneva, Switzerland.

The role of CERN

Built on the Franco-Swiss border, CERN was founded in 1954 and it is currently funded by 20 European nations. CERN is now operating the Large Hadron Collider (LHC) which is colliding proton beams at energies up to 7 Tera (10^{12}) electronVolts (to be increased to 14 Tera electronVolts after 2013), and lead nuclei beams at energies up to 2.76 Tera electronVolts per nucleon (to be increased to 5.5 Tera electronVolts after 2013). The LHC is the largest scientific machine ever built. CERN has financed and built most of the accelerator complex, with important contributions from the United States, Russia and Japan. Four experiments have started to analyse the results of particle collisions around the 28 km LHC ring. Experiments are designed, financed and built by large international collaborations of hundreds of physics institutes, university departments and laboratories – including CERN’s own Experimental Physics Division – comprising thousands of physicists from four continents.

A modern High Energy experiment is typically built at the “intersection point” of a particle accelerator. It is here that counter-rotating beams of particles cross each other generating particle collisions. The energy of each collision transforms

itself into mass, according to Einstein's famous formula $E = mc^2$, generating a host of particles not observed on Earth under normal conditions. The "detectors" are very large scientific instruments that surround the interaction points and detect the particles. Several million collisions occur each second, and, after discarding the uninteresting ones, the surviving ones are registered for further processing at data rates that can reach few Gigabytes per second.

It is perhaps important to note that particle collisions create a very high temperature and energy density, similar to what is believed to have existed at the beginning of our universe immediately after the Big Bang. In this sense, a particle accelerator is also a "time machine" which reproduces the condition of the universe soon after its birth, even if only for a very tiny amount of matter. Again, the understanding of our origins is linked with the microscopic composition and behaviour of matter.

Computing in HEP

Since its beginnings, High Energy Physics has relied heavily upon computers to extract meaningful physics results from observations. This is due to several reasons, which have to do with the statistical or probabilistic nature of the underlying theory and with the large amount of collected data that has to be treated electronically.

Very large computing needs and complexity.

This has taken on very large proportions with the latest generation of experiments, which have produced several hundred TeraBytes of data. The new experiments which came on-line in 2009 at CERN are producing a few PetaBytes (10^{15}) of data per year. These data need to be pre-processed several times before being properly analysed. The design and understanding of the experimental apparatus require the computer simulation of its response, and this is a very demanding task in terms of computing resources, as it generates an amount of data comparable with that of the real experiments. However the complexity of computing for High Energy Physics is not only due to the sheer size of the resources needed, but also to the way in which the software is developed and maintained. Laboratories that have an accelerator, such as CERN in Geneva, Fermi National Accelerator Laboratory in Chicago, the Brookhaven National Laboratory, or Stanford Linear Accelerator Centre host high energy experiments. As we have mentioned, an experiment is built by a large international collaboration comprising thousands of scientists coming from hundreds of institutes in a few tens of countries. These researchers work rather independently on the same code, meeting only rarely and with little hierarchical structure between them. Requirements change very frequently and the problems to be solved often push the boundaries of scientific knowledge, both in physics,

and also in computer science. It can be said that scientific research in fundamental physics has its limit not in the creativity of the researchers, but in the current state-of-the-art of engineering and computer science.

It has happened several times in the past that physicists have been at the origin of important developments in the field of computing. Notable examples are the invention of the Web at CERN and, more recently, the CERN-led development and deployment of the largest Grid (see Chap. 3 for the description of the concept) in operation in the world for the storage and processing of the data coming from the LHC experiments. And while the Grid is still developing its full potential, the attention of the HEP computing world is already moving to computing Clouds and virtualisation (see Chap. 6) at the boundary of knowledge in distributed computing. These results have been given considerable attention by the media. However there is relatively little published material on the general history and development of computing in High Energy Physics, which presents many other interesting and innovative aspects. For instance, programme code has always been shared between physicists, even beyond the boundaries of a single experiment, in a very similar way to what has now become known as “open source” software development. The management of the development of a large piece of software driven by very dynamic requirements has given rise to techniques very similar to those that have been independently developed by “Agile Software Engineering” technologies.

Most scientific communication between fundamental physicists is centred on their scientific results, with little attention paid to the discussion of their activity as Information Technology professionals (see Chap. 5).

IT innovation originated in HEP went often unnoticed outside the field.

The result is that several computing techniques and concepts, developed and applied in the context of High Energy Physics, have had broader success after being “reinvented” and developed independently, maybe years later. The main reason for writing this book is to describe the world of computing in High Energy Physics, not only through its results, but by describing its challenges and the way in which HEP is tackling them. The aim is to provide a needed “theoretical” justification for the techniques and “traditions” of HEP computing with the intention of making them better known, and possibly promoting dialogue and interdisciplinary collaboration with other disciplines that use computing either as a research field per se, or as an important tool in their activity. This seems to be a very appropriate moment to do it, because a generation of experiments has now come on-line after almost 20 years of development, and therefore we can now describe a mature and accomplished situation in computing for HEP. The next generation of experiments will almost certainly introduce radical changes to this field, or at least we hope so.

A large part of this book talks about the development of HEP computing at CERN and therefore the reader may question the legitimacy of the claim that it deals

with the history of HEP computing worldwide. This is indeed a delicate subject, and CERN has been at times accused of “software imperialism” and of trying to impose its software choices on the worldwide HEP community, while hindering or undermining the development of independent software projects elsewhere in institutions collaborating with its experimental programme.

The special role of CERN in the development of HEP computing.

These criticisms may not be totally unfounded, but we also believe that the software development process in HEP is essentially a Darwinian process, where the ideas that eventually prevail have some degree of fitness to the environment. This does not make them necessarily the best in all possible respects, but it definitely tells us something about their adaptability to different experimental conditions, size of collaborations, and computer hardware and operating system diversity.

HEP is a very distributed and collaborative activity, and so is its computing. It is therefore an undeniable fact that only a fraction of HEP code has been written at CERN, but it is also true that the fundamental software infrastructure adopted by all HEP experiments in the 1980s and through the 1990s was the CERN Program Library, and from the late 1990s the ROOT (see Chap. 1 for a description of the ROOT genesis and evolution) package, developed at CERN by two of the authors of this book and their team. While the whole HEP community has largely contributed to both projects, their design, development, maintenance and distribution has been centred at CERN.

From the point of view of the subject, this book is largely complementary to existing literature. However the information contained in this book is original in the way it is treated and presented. Our objective is not to describe algorithms and computing techniques, nor do we concentrate only on a few specific aspects such as programming language or development methods. Our aim is to give a comprehensive description of the evolution of computing in HEP, with examples. The closest analogy will be with the material presented at the Computing in High Energy Physics (CHEP) conference series, but presented in a set of chapters providing a methodological and historical account and covering the major aspects and innovations of HEP computing.

A book for education and research.

This book is intended both for education and for research. For students of physics and other scientific disciplines it provides a description of the environment and the modalities of the development of HEP software. It often takes young researchers a relatively long time to become active in HEP computing, due to the

relative uniqueness of its environment and the fact that it is not documented or even less taught at university. For physics researchers this book will be a useful documentation of the history of one of their most important tools, i.e. computing. For researchers in computing and its history, the book will be an important and original source of information on the evolution of HEP computing in the last 20 years from a very privileged point of view.

The Structure of the Book

As mentioned above, this book is written by different authors, each one of them having his or her own style and point of view. As we said, this is yet an untold story, which spans more than 20 years of intense development, and some of the main actors often held widely diverging point of views on the course to take. The subject matter itself is very diverse, ranging from computer hardware to programming languages, software engineering, large databases, Web, Grid and Clouds (see Chaps. 3 and 4). It is not surprising that this is reflected in the material of this book. While we believe that this is one of the main assets of the book, we are also aware that it may introduce some disuniformity and lack of coherence. To help in alleviating this danger, we have tried to give some formal coherence to the appearance and structure of the different chapters. Each chapter starts with a short statement, followed by an introductory paragraph. It ends with a conclusion section on the lessons learned and, where appropriate, a look at the future. Grey filled boxes indicate important concepts developed in the text. We have tried, where possible, to cross-reference the various chapters, in order to allow the reader to easily recognise related concepts in material with which he or she may not be familiar.

It is important to mention at this point that the opinions expressed in this book are those of the authors and may not reflect the official position of their institutes.

Chapter 1 is a broad overview of the evolution of the programmes, the technology and the languages of HEP computing over the last 20 years. The sheer evolution of the computing hardware in these years has seen a 10,000 increase in processing speed and a 100,000 increase in live (RAM) memory. This is probably the largest technological evolution of any field of technology. The chapter describes how HEP computing coped with this breathtaking evolution to provide the needed resources for the scientific endeavours of HEP.

Chapter 2 recalls the story of the invention of the World Wide Web inside the HEP community. This is one of the best-known, but also one of the most striking examples of “spin-off” from fundamental research. It was a paradigmatic event that deserves to be studied in detail. Few inventions have changed the face of our world so quickly and so deeply as the Web has done, and yet it was invented with the immediate purpose of helping to share images and information related to HEP. In an era where there is so much debate on the relations between applied and fundamental science and on their respective roles and impact on society, this extreme example of serendipity in fundamental science is certainly worth considering in detail.

Chapter 3 describes the genesis and evolution of the worldwide LHC Computing Grid, the largest Grid in operation in the world to date. Extending the Web paradigm of sharing distributed information to include the sharing of distributed computational resources, the Grid aims at homogenising the computing infrastructure, making it ubiquitous and pervasive, and freeing it from the space and time availability constraints of any single physical instance of computing and storage. The consequences of this process are yet to be fully understood and appreciated, but their impact on society at large can be potentially larger than that of the World Wide Web. While the idea was born in the U.S. academic world, the first large scale realisation of the Grid paradigm was realised for the needs of the LHC experiments via a CERN-led international effort.

Chapter 4 describes the Grid in operation. While the middle-ware technology is extremely complex and still in a state of “frantic” evolution, the deployment of the Grid in over 50 countries and 300 institutes around the world has opened a whole new set of opportunities and challenges for international collaborations. Scientists and system administrators from widely different cultures and time zones are learning to work together in a decentralised and yet tightly coupled system, where the global efficiency depends on the efficiency of every single component. Researchers in economically and digitally challenged countries suddenly find themselves with the opportunity to access the largest computing infrastructure in the world and to perform frontier research with it. With the access to data and computing resources becoming completely democratic and ubiquitous, a new set of rules has to be invented for this new virtual community to work together.

Chapter 5 describes the software development process in HEP. Software Engineering has been invented, or rather postulated in the 1950s to apply “solid engineering methods to software development”. Although a staggering amount of software has been developed, some of which very successfully, the varying degree of success of software projects is still surprisingly wide. Software Engineering has offered some very interesting insights in the process of software development, but has failed to develop “solid engineering methods” that can transform software development into a predictable process. HEP has developed software for many years, often very successfully, without ever adopting a specific software engineering method. Nevertheless it had its own method, although this was never formalised or described. This chapter describes the software development methodology which is proper to HEP and discusses its relations with the recent advances in software engineering.

Chapter 6 describes the impact of virtualisation on Grid computing. The Grid computing paradigm proposes a conceptual virtualisation of the physical location and other features of a set of computational resources. It is therefore quite natural to push virtualisation one step further and virtualise the hardware within computing centres, in order to increase the uniformity of the Grid components and to help application developers and end users to cope with the different physical platforms. Moreover this makes reconfiguration and upgrade of the underlying systems more transparent for the end users, increasing the overall robustness and resilience of the Grid infrastructure. Finally, virtualisation techniques offer opportunities to increase

security of the applications running on the Grid, which is one of the major concerns for such a widely distributed system. The newly born Cloud concept adds another level of abstraction, but also of opportunity, to this virtual infrastructure.

Chapter 7 deals with the exploitation of parallel computing. HEP computing features a kind of parallelism which has been defined as “embarrassing”, as the workload comes in units of “events” (the result of particle collisions) which are independent from each other. The result of the processing of each event can be calculated independently and statistically combined with the other results in any order. As a consequence of this, parallelisation has been used since a long time in HEP to reduce the time-to-solution and to allow the interactive optimisation of the analysis algorithms. However, parallelisation had stopped at the “boundary” of the code that treats each event. The advent of multi and many-core machines and of powerful Graphics Processing Units (GPUs) and the large memory requirements of modern HEP codes make it less efficient to continue assigning an event to every core, and therefore parallelisation has to happen also within the code that treats each event.

Chapter 8 deals with the legal aspects of intellectual property law for HEP software developers. The question of the intellectual property of the software developed by large HEP collaborations is technically very complicated. The experimental collaborations are not legal entities, and the software is produced by tens, at times hundreds of authors from different institutions who spend some time in the collaboration and then move on either to other experiments or to a different career. Questions such as who is the legal copyright holder or what is the legal status of the software are in fact complex problem of international law that must be understood and correctly addressed in collaborations that can last several years and involve thousands of researchers.

Chapter 9 deals with databases. HEP computing is essentially a data-centric problem. The principal activity is data transformation and data mining. In this sense HEP computing is similar to other forms of computing that deal with large data-sets, which are becoming more and more commonplace as sensors and measurement instruments become inexpensive and ubiquitous. It is therefore natural that databases occupy a central place in HEP computing. The efficiency of data and meta-data access is perhaps the most important feature in determining the success of an experiment’s computing infrastructure. As the focus is put on ordered data transformation and data mining, the transactional aspect of relational databases is not particularly critical for HEP computing, and therefore classical or object-oriented databases occupy a “niche” in HEP computing, while the bulk of the data is held in custom designed data stores.

Chapter 10 deals with obtaining fast and reliable access to distributed data. Grid computing proposes a paradigm for transparent and ubiquitous access to distributed high-end computing resources, but globalised data access is still a very hard problem to solve. For truly transparent data access, local and remote operation should be transparent for the applications. Moreover the different features of data access over Wide Area Networks (WAN’s) and Local Area Networks (LAN’s) should be optimised and, in some sense, “masked” by the protocol. The systems providing

data access should also hide the diversity of the data storage technology from applications and virtualise the data repository, thus eliminating the need for an application to worry about which storage element holds a given file or whether it has been moved or not. These are very difficult, and in some sense contradictory requirements to satisfy, and HEP computing has accumulated very substantial experience in dealing with these problems over the years.

Chapter 11 attempts to go beyond the technical description of HEP computing to consider the larger picture of the development of an increasingly interconnected system of machines and people sharing computing resources and information in a virtual universe, ubiquitous and largely independent of actual space and time. This virtual interconnected world resembles more and more the realisation of a human nervous system on a planetary scale. As already said, HEP computing is today where everyday computing will be tomorrow, and it is interesting to draw a parallel between the capacity of the human brain to represent and understand the universe, and the construction of a world-wide network of connections and nodes where information is stored, exchanged and processed as in the neurons of our own brains. This introjection of the universe and projection in the physical (and virtual) reality of our own brain is intriguing and worth considering at the end of our book.

The Authors

René Brun is a physicist working at CERN since 1973. While working on the detector simulation, reconstruction or analysis phases of the NA4, OPAL, NA49 or ALICE experiments, he has created and lead several large software projects like GEANT, a general detector simulation system, the data analysis system PAW, and ROOT a general experiment framework for data storage, analysis and visualisation. All these systems have been or are used today by thousands of people in the scientific community.

Predrag Buncic. Following the studies at Zagreb and Belgrade Universities, Predrag Buncic stated his carrier as a physicist in NA35 experiment at CERN where he worked on a streamer chamber event reconstruction and where he quickly discovered his passion for scientific computing. In 1994 he joined the NA49 experiment to work on challenging problems of data management as well as reconstruction, visualisation and data processing. In 2001 he moved to ALICE experiment at LHC where he initiated AliEn project, a lightweight Grid framework that later served as an inspiration for the first gLite prototype. For this reason he joined EGEE project in 2004 and worked for two years in CERN/IT on Grid middleware architecture. Since 2006 he is working in CERN/PH Department and currently leading Virtualisation R&D project in PH/SFT group.

Federico Carminati is presently Computing Coordinator of the ALICE experiment at LHC. After getting his Master in Physics at the University of Pavia, Italy in 1981 he worked at Los Alamos and Caltech as particle physicist before being hired by CERN in the Data Handling Division. He has been responsible for the CERN

Program Library and the GEANT detector simulation programme, the world-wide standard High Energy Physics code suite in the 1980s and 1990s. From 1994 to 1998 he worked with Nobel Prize winner Prof. Carlo Rubbia at the design of a novel concept of accelerator-driven nuclear power device.

Fabrizio Furano is a Staff member at CERN, working for the Grid Technology group, in the IT department. He has an extensive experience related to software engineering, performance, simulation and distributed systems, acquired in his activity in the telecommunications industry, university and High Energy Physics computing. From a position related to software development and contact center technology, he moved to work with INFN (Italian Institute for Nuclear Research) for his Ph.D. studies in Computer Science and for a post-doctoral grant, during which he was contributing to the Computing Model of the BaBar experiment at the Stanford Linear Accelerator Center. He taught C++ Programming as Assistant Professor at the University of Padua (Italy) and Software Engineering as Adjunct Professor at the University of Ferrara. He moved to CERN in 2007, to work on Data Management technologies for the LHC experiments.

Giuliana Galli Carminati is presently senior psychiatrist responsible for the Unit of Mental Development Psychiatry (UPDM), University Hospitals of Geneva (HUG), Switzerland. After getting her degree in Medicine at the University of Pavia (Italy) in 1979, she obtained specialisations in Laboratory Medicine and in Psychiatry and Psychotherapy, as well as a Master in Group Therapy and a Doctorate in Psychiatry (Geneva University) in 1996. In 1998 she also got a doctorate in physics (Laurea) at the Tor Vergata University (Rome). In 2008 she obtained the title of Privat Docent at the University of Geneva. Her research activities deal with the treatment of Intellectual Disability and autism and the Quality of Life of intellectually disabled patients. She is particularly interested in the relations between matter and mind, and in particular in the application of Quantum Information theory to the modelling of human psyche, subject on which she has authored some papers.

Andy Hanushevsky obtained his B.A. in Geology at the University of Rochester and his M.S. in Computer Science at the Cornell University. After having worked at the Xerox Corporation, the NPD group and the Cornell University as system programmer, he joined the Stanford Linear Accelerator Laboratory in 1996 as Information System specialist. There he worked on extremely scalable data access systems, high performance protocols for distributed systems and parallel and multi-threaded algorithms. Together with F. Furano he developed the Xrootd – Scalla system which has become a standard tool for high performance local and distributed data access for High Energy Physics experiments worldwide. He was awarded a U.S. patent in 1996 for a method to read hierarchical distributed directories without obtaining locks and the 1993 IEEE Certificate of Appreciation for MSS Standards Work.

Patricia Méndez Lorenzo was born in Salamanca (Spain) where she studied Physics at the University of Salamanca. She completed her Diploma thesis and her Ph.D. in Particle Physics in Munich (Germany) at the Ludwig Maximilians University. Since 2002 she is working at CERN. She has been part of the WLCG

Project until January 2010. Inside the WLCG Project she has been the Grid Support responsible of the ALICE experiment and the responsible of the NA4 effort for Particle Physics inside the EGEE-III project. She has also collaborated with a large range of applications ported to the Grid.

Lawrence Pinsky is the current chairperson of the Physics Department at the University of Houston, an experimental particle physicist who recently served as the Computing Coordinator for ALICE-USA as a member of the ALICE Collaboration at CERN and the ALICE-USA WLCG representative. He is also a licensed Attorney at Law in the State of Texas as well as being a licensed U.S. Patent Attorney with both a J.D. degree and an LL.M. in Intellectual Property and Information Law. He also has regularly taught courses in Intellectual Property Law at the University of Houston Law Center.

Fons Rademakers received his Ph.D. in particle physics from the Univ. of Amsterdam in 1991 for his work on event displays and data analysis for the DELPHI experiment at LEP. Since then he has worked at CERN and been involved in designing and developing data analysis programs. In 1991 he joined the PAW team of René Brun where he developed amongst others the column wise-ntuples, the PAW GUI and the PIAF system. In 1995 he started with René Brun the ROOT project and has been involved in all aspects of the system since then. In 2001 Fons joined the ALICE collaboration and has worked as software architect on the initial version of AliRoot. In recent years his special attention has gone to high performance parallel computing using PROOF.

Les Robertson was involved in the development and management of the computing services at CERN from 1974, playing a leading role in the evolution from super-computers through general purpose mainframes to PC-based computing fabrics and finally distributed computing Grids. He was active from the beginning in planning the data handling services for the experiments that use the Large Hadron Collider (LHC), and led the Worldwide LHC Computing Grid Project (WLCG), set up in 2001 to prepare the computing environment for LHC. The project included the development and support of the common tools, libraries and frameworks required by the physics applications, the preparation of the computing facility at CERN, and the development and operation of the LHC Grid, which integrates resources in more than 140 computing centres in 35 countries around the world. Prior to coming to CERN he worked in operating systems and network development in the computer industry in the UK. He has now retired.

Ben Segal got his BSc in Physics and Mathematics in 1958 at Imperial College, London and his Ph.D. in Mechanical and Nuclear Engineering in 1961 at Stanford University (U.S.). He started working at CERN in 1971 on high-speed computer networks as well as an early satellite data transmission system “STELLA” (1978–83). He also coordinated the introduction of the Internet Protocols at CERN (1985–89) and encouraged the introduction of Unix (then Linux) at CERN since the early 1980s. He participated in setting up the Internet Society (ISOC) Geneva Chapter in 1995. He was a mentor to Tim Berners-Lee who invented the World Wide Web at CERN (1989–91). From 1990 was a member of a small team which developed the “SHIFT” system, migrating CERN’s computing capacity from central mainframes

to distributed Unix clusters. He has worked since 2004 in the developing area of “volunteer computing”, with its great potential for public involvement in science and education. Projects include LHC@home, MalariaControl.net and Africa@home. He is also co-founder in 2009 of the CERN-based Citizen Cyberscience Centre.

Jamie Shiers currently leads the Experiment Support group in CERN’s IT department. This group plays a leading role in the overall Worldwide LHC Computing Grid (WLCG) project, with a strong focus on service and operations. He has worked on many aspects of LHC computing since the early 1990s, moving to the Grid service area in 2005 when he led two major “Service Challenges” designed to help bring the service up to the level required for LHC data taking and analysis. He is a member of the Management Board of WLCG and has authored numerous articles on Grid and Cloud computing. Dr. Shiers received a Ph.D. in physics from the University of Liverpool in 1981, following a degree in physics at the University of London (Imperial College) in 1978. He has worked in the IT department at CERN for the past 25 years in a wide variety of positions, including operations, application development and support, databases and data management, as well as various project leadership roles. Prior to this he worked as a research physicist at the Max Planck Institute for Physics in Munich, Germany and as a guest physicist at CERN.

Acknowledgements We gratefully acknowledge the help of the many colleagues with whom we have discussed the idea and the content of this book. Trying to mention them all would inevitably mean forgetting some. But a special thanks goes to Michelle Connor for having proofread the manuscript with intelligence and attention.

Geneva, Switzerland

Federico Carminati

Reference

1. <http://www.cern.ch>

Contents

1	Technologies, Collaborations and Languages: 20 Years of HEP Computing	1
	R. Brun	
2	Why HEP Invented the Web?	55
	B. Segal	
3	Computing Services for LHC: From Clusters to Grids	69
	L. Robertson	
4	The Realities of Grid Computing	91
	P.M. Lorenzo and J. Shiers	
5	Software Development in HEP	115
	F. Carminati	
6	A Discussion on Virtualisation in GRID Computing	155
	P. Buncic and F. Carminati	
7	Evolution of Parallel Computing in High Energy Physics	177
	F. Rademakers	
8	Aspects of Intellectual Property Law for HEP Software Developers	201
	L.S. Pinsky	
9	Databases in High Energy Physics: A Critical Review	225
	J. Shiers	
10	Towards a Globalised Data Access	267
	F. Furano and A. Hanushevsky	
11	The Planetary Brain	289
	G.G. Carminati	

Glossary 311

Index 337

Name Index 351