Paving the way to the Internet and Web

In the early 1980s, CERN's external connections amounted to two minuscule leased lines operated at 9.6 kilobits per second (kbit/s), one to the Rutherford Appleton Laboratory (RAL), the other to CEA, Saclay. Ten years later, CERN had become the centre of a large star-shaped network, by far the largest internet hub in Europe. Today the LHC Optical Private Network (LHCOPN), a network connecting CERN to the LHC T1 sites operates multiple 100 Gigabit per second (Gbit/s) links. How did CERN get there?

Facing technological challenges has frequently been at the root of developing leading-edge technologies for CERN engineers. In turn, mastering a technology provides the confidence necessary to undertake very ambitious projects. Such was the case for networking, where CERN developed as early as the 1970s outstanding expertise on network protocols.

Facing the need for an open (i.e. based on non-proprietary technology), high bit-rate communication system to interconnect on-site a variety of computer brands — mainly in order to transfer files — CERN launched CERNET [26], the most ambitious networking project of its time. Communication between computers requires the use of common rules. These rules are called communication protocols. They are building blocks stacked one on top of the other in a layered structure, and are in practice implemented as modules of hardware, software or a mix of the two. They broadly divide into two types. The first type ensures the transport of raw information data between two network points, independent of its significance or use. The second type of protocol offers the end user a tangible service, e.g. email. CERNET implemented a genuine and complete suite of layered protocols, ranging from the physical and electrical plug to the CERN File Transfer Protocol running on half a dozen different brands of computers. CERNET was put into operation in 1977. This is a unique example of a complete multi-brand networking stack developed by a single organization. It was also the fastest packet-switching network of its time [Highlight 9.3].

Probably more important for the future was the type of technology adopted by CERN for one of the layers, called the packet layer. All modern networks use the packet principle where the data stream is chopped into individual pieces called packets. Packet protocols are divided into two fundamentally opposite classes: the connection-oriented class where the two end-computers must establish a connection before an information packet can be sent — the equivalent of a telephone call, and the connectionless class, where packets are independent of each other, as letters posted in a letter-box, and can be sent at any time. Connection-oriented networks check first via the call set-up mechanism that a fixed end-to-end route is available. The main advantage lies in the fact that commercial network

operators know the duration and the amount of traffic exchanged over each connection and they can charge the user on a per-call basis. The drawbacks are the extra delay to setup the call before any useful data can be sent and the additional complexity. Connectionless networks are not aware that a series of packets may belong to a flow. As letters, individual packets carry the full destination address. The drawback is that no prior check is made that a non-congested route is available when new packets are admitted. The chief advantages are the no-delay admittance of packets, higher resilience as routes are not pre-assigned, and more generally, simplicity.

CERN opted in the early 1970s for the latter, the connectionless approach. A prescient decision, since connectionless is precisely the philosophy of the Internet Protocol, the famous IP. By creating, 15 years before the explosion of IP, a deep expertise in IP-like technology, CERN paved the way for its future role in the Internet infrastructure and also for its invention of the Web.

The emergence of HEPNet

In the early 1980s, only RAL and Saclay were connected to CERN, via slow analogue leased lines funded by these two organizations. They were mainly used for remote login to CERN mainframes. In 1983, a proposal was made by CERN to connect more sites and offer additional services (File Transfer, Job submission and email in addition to remote login). This initiative was coined HEPnet. As alternative to leased lines the use of the emerging public packet switched networks was proposed by the national PTTs, monopolies at that time. The service offered by the PTTs was based on X.25 technology, a connection-oriented protocol, which is the exact opposite of CERNET and IP. CERN connected to the Swiss PTT X.25 network, allowing computers to make calls — as over a telephone network — to and from any other HEP site connected to their domestic X.25 network. To facilitate file transfers, CERN developed a gateway between CERNET and X.25, which was released in 1983. This development was one of the very few operational gateways between connectionless and connection-oriented networks.

The main drawback of X.25 networks was their per-call charging principle modelled by the PTTs on the telephone, making difficult any planning of the telecommunication costs. The X.25 PTT service was progressively abandoned and new leased lines — initially analogue, then digital at 64 kbit/s — connected additional HEP sites to CERN. In parallel, in 1982 IBM proposed to several research computer centres in Europe, to establish and fund for four years a network based on the BITNET (Because It's Time NETwork) model in the USA. The initiative was called EARN (European Academic and Research Network) [35]. Additional analogue links were installed and CERN became in 1983 the largest hub of EARN. At that time, all of these links operated in a proprietary way.

CERN: Precursor in Internet Technology

In 1985, a TCP/IP coordinator was appointed at CERN and the LEP control system adopted this technology. CERN was one of the first sites in the world to receive for evaluation early versions of IP routers from the future market leader. It was at the end of 1988 that CERN started to operate leased lines with the TCP/IP protocols. By then, more leased lines had been set up, extending to Boston and Tel-Aviv. In summer 1988, a CERN team went to Beijing to organize the very first academic connection between China and Western Europe. But more than the number of connections, it was the tremendous increase of the bandwidth that changed dramatically the resulting services. In 1988, most analogue links had been replaced by digital circuits, at multiples of 64 kbit/s. Furthermore, by funding the first European 2 Mbit/s academic link in Europe, between CNAF Bologna and CERN, INFN propelled external networking into the Megabit per second era. No ground links of that bit-rate were available, but experimental satellite channels were. From 1981 to 1983, the STELLA project evaluated protocol aspects of 2 Mbit/s satellite channels between CERN, Pisa and RAL.

CERN becomes the largest Network Hub in Europe

In 1990, the very first Internet transatlantic megabit link (1.5 Mbit/s) was set up between CERN and Cornell University. A full TCP/IP connection to the US National Science Foundation Network (NSFNet) was established, the link being the result of EASYNet, a new IBM initiative in Europe, three years after the end of their EARN funding.

In 1991, the star network around CERN became the largest internet hub in Europe: 80% of the Internet capacity in Europe for international traffic was installed at CERN (Fig. 9.5). It is no surprise that the performance of Tim Berners-Lee's first Web server impressed the world: it was at the heart of the European Internet, just a few hops from most destinations. Why did this occur? There was an absolute need for high bit-rate access to CERN from HEP centres so that physicists could work remotely almost as if they were on-site at CERN. In addition, networks grew like crystals: as one hub started "crystallizing", everyone wanted to connect to it being inspired by the novel opportunities offered, Therefore, non-HEP scientific centres started to set up links to CERN, often serving as relays to other domestic HEP sites. Finally, a networking hub, being a central point of failure, requires extremely reliable operation. Inherited from the accelerator culture, CERN had a tradition of excellence in 24 hours/7 days operation that was recognized worldwide.

CERN initiates IP address allocation in Europe

Another key contribution of CERN to the Internet infrastructure was the attribution of IP addresses. Any device connected to the Internet needs a unique identifier: the famous IP address. Before 1989, all addresses were allocated by the Internet Assigned Numbers Authority (IANA) in the USA. It was CERN, which convened at its site in December 1988 the first meeting to discuss the creation of a structure in Europe to allocate the European addresses. Six months later, the first meeting of RIPE (Réseaux IP Européens) took place in Amsterdam. RIPE is still, 25 years later, the authority that allocates Internet resources and services in Europe, in the Middle East and parts of Central Asia [36].

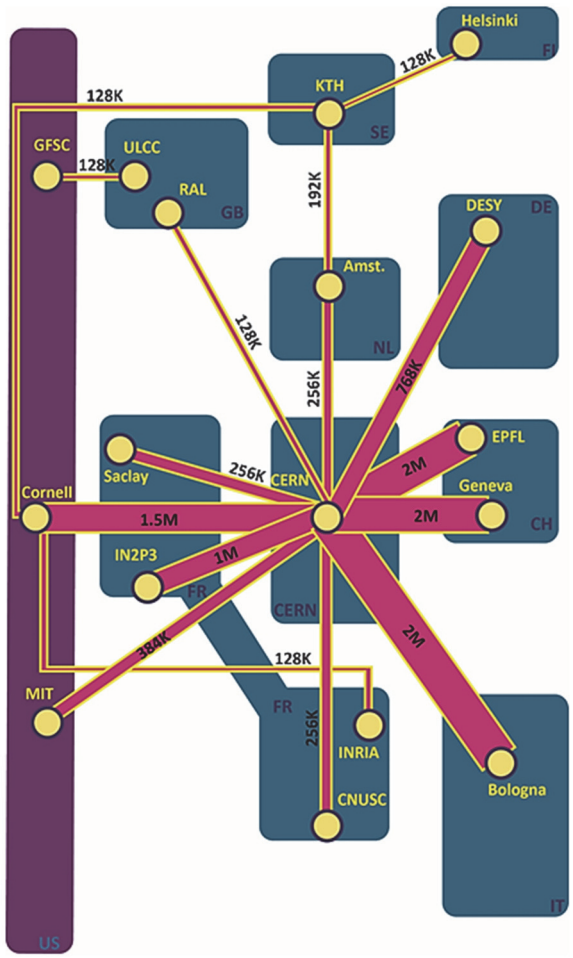

Fig. 9.5. The star network around CERN in 1991. Only high throughput data links are shown: for all links see http://cds.cern.ch/record/2054391.

Having reached its zenith in 1991, the CERN share in the Internet infrastructure slowly decreased, for two reasons. First, though prepared to remain a major Internet Hub, CERN considered that it would be healthier for the Internet if other hubs emerged. Second, in many countries, National Research and Education Networks (NRENs) had emerged, to which most HEP centres were connected [37]. These NRENs were, and still are, interconnected via the pan-European academic backbones operated by the DANTE organization, initiated by the European Commission [38]. The connection of CERN to these successive backbones moved from 34 Mbit/s in 1997 to 10 Gbit/s in 2005. In that period, CERN played a similar role to that of the European NRENs. However, there was still a need to push the technological limits and in 2003, in the context of the EU-sponsored CERN DataTAG project, the world record for data transfer was broken by CERN, together with the California Institute of Technology (Caltech), the Los Alamos National Laboratory (LANL) and SLAC, with the successful transfer of 1 terabyte of data in less than an hour over the 10,037 km between CERN and California.

The Gigabit era

To meet the requirements of the Worldwide LHC Computing Grid (WLCG) [Highlight 9.2], CERN and its HEP partners returned to the approach of HEP-dedicated links between the CERN and the major HEP centres (Fig. 9.6).

These dedicated ultra-high bit-rate links are provided by GÉANT [37], a pan-European research and education network supported by the EU, and by other NRENs: they form the LHCOPN (LHC Optical Private Network). The minimum bit-rate is 10 Gbit/s and several circuits operate at higher speed including a dual 2×100 Gbit/s transatlantic link to the Energy Science Network (ESNet) in the US.

Through the deployment of the Internet infrastructure technology, CERN, together with the rest of the academic community, has also had an indirect impact. Academic organizations — be they universities or research centres — are sometimes compared to industrial companies in unfavourable terms regarding aspects such as efficiency, organization and rigour. This is perhaps viewed as a price to pay for creativity — a perception inherited from the old idea that inventiveness and organization are antagonistic qualities. The academic community was sometimes viewed as being only able to invent new technologies, but not to contribute to their development. The development of the academic and research networks has demonstrated that the academic community is capable of deploying globally extremely complex systems. They are capable of delivering professionally operated systems, which can revolutionize the economy in a few years and permeate all levels of society.

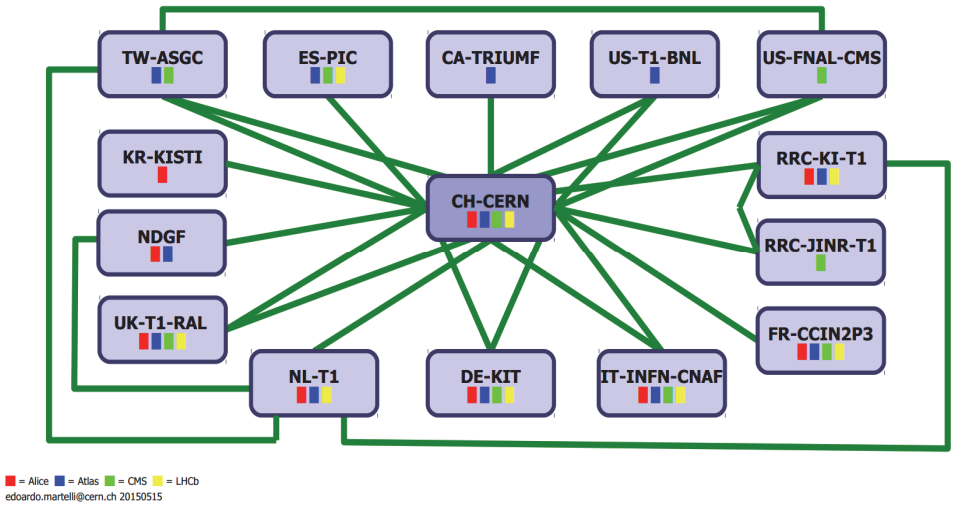

Fig. 9.6. LHCOPN topology. Details at http://cds.cern.ch/record/1997396?ln=en.

9.5 Detector Simulation: Events Before the Event

René Brun

Simulation[d] of a particle physics experiment relies on a chain of models which follow the evolution of the particles produced in the interaction up to the signals generated in the detectors. These can be:

- Models of the physical process representing the primary goal of the experiment and of all the known processes that could be considered as background;
- A model of the detector (geometry, materials, magnetic and electric fields) and of the path and interactions of each particle through the detector elements;
- Models of the signals generated by the particles in each detector type, followed by the electronic processing to produce digital information;
- A model of the selection procedures (trigger) which retain the wanted topology and reject background.

All models in the various steps are complex and contain many variables: simulation of an experiment is possible only by using Monte Carlo methods.

In the end, the simulation chain produces "data" in the same format as the real experiment. Models of the physical processes are needed to test the viability of the initial idea. Geometry and acceptance go in step with the overall design of the

[d]A simulation uses a mathematical description, or model, of a real system in the form of a computer program (Encyclopaedia Britannica).

apparatus. Signal generation is useful to understand in detail the behaviour of individual detectors. Simulated signals are used to train and tune selection and pattern recognition procedures. Simulated raw "data" help testing the reconstruction software.

Once real data become available and are processed, results are compared with the output of the simulation chain. In order to extract efficiently a potential signal from the background, one must identify the most efficient (uncorrelated) variables. Multivariate analysis (MVA) treats many statistical variables simultaneously, taking into account the effects of all variables on the response of interest. The tools available, such as neural nets, "boosted" decision trees, etc., go under the name of Machine Learning (ML). These tools make extensive use of simulated "data" to train their decision algorithms. TMVA [39] (Toolkit for Multivariate Data Analysis with ROOT) is used by the LHC experiments for this task.

As the understanding of the experiment improves, more refined simulations may be desirable, requiring running the whole cycle several times.

The full simulation cycle is very demanding of computer resources. Different methods of "fast" simulation have been devised to replace the full simulation of sub-systems by introducing simplified assumptions on detector geometry or by skipping some Monte Carlo calculations and using the results from formulae or from a stored catalogue. Computing time can thus be reduced by a large factor.

Fast simulation is not a replacement of full detector simulation. It must be tuned/validated with full simulation results (while full simulation is itself tuned/validated with data containing well-known reference signals). Fast simulation is used primarily for quick and approximate estimates of signal and background rates. It is also useful for an initial survey of complex, multi-dimensional parameter spaces which require a very large (simulated) data set.

Simulation is essential in designing experiments and interpreting the results, not only in high energy physics (HEP), but also in many other branches of science and technology. Simulating the known phenomena with good precision is a necessary prerequisite to the discovery of new effects, by comparing simulation with data.

Event generators have been developed for simulating interactions of different primary particles and a wide range of energies. An interaction process is described mathematically by a "matrix element" which must be integrated over the multidimensional internal parameters phase space to provide the total and differential cross-sections. Each point of this phase space is associated with a probability. According to this probability events are randomly generated.

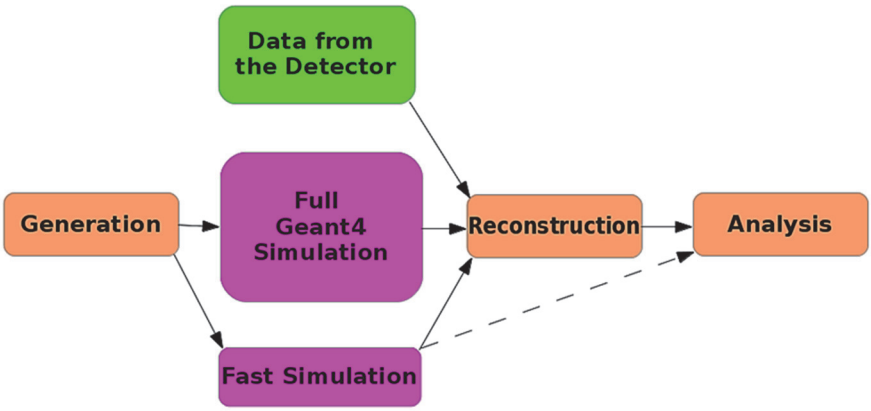

Fig. 9.7. Simulation in the LHC experiments.

As an example, consider the steps involved in the generation of events in proton-proton interactions as they happen at the LHC. The initial state consists of colliding protons, but the most interesting interactions are the hard collisions between constituent partons [Boxes 6.4 and 8.3]: the simulation must first take into account the distribution of partons inside a proton, as described by its structure functions. Then the specific "hard interaction" between partons has to be selected. As the most interesting processes (e.g. Higgs physics) are very rare compared to the dominant QCD background, the latter must be correctly simulated, which implies considering many relevant loop diagrams [Boxes 4.2 and 5.1]. Other partons belonging to the initial protons may interact "softly", and must be also simulated as they contribute to the complexity of the event. Following the initial hard interaction, the bare partons do not exist in nature (they are confined inside the hadrons) and must be "dressed", to form the known hadrons or mesons. This goes in steps: parton shower (describing their radiation), hadronization (grouping partons into particles) and decay of unstable particles. The task of comparing simulated events with data is further complicated by the fact that at the LHC several proton-proton interactions take place at each beam crossing.

Event generators must consider the physics complexity at the appropriate level together with the most efficient simulation methods. The goal is to generate enough simulated events to match the number of real events with affordable computing resources. An event generation package widely used for high energy collisions of electrons, positrons, protons and antiprotons is PYTHIA [40]; other packages have been developed for specific processes or initial states (e.g. heavy ion collisions).

Up to 1975 the major software development for HEP was associated with bubble chamber experiments. Event kinematics was reconstructed from the

information obtained by scanning and measuring pictures. With early electronics-based experiments, the reconstruction software was tested and run only after data taking. Events consisted of straight or curved tracks hitting detector "planes" (digitized spark chambers or multiwire proportional chambers) and showers in calorimeter "blocks" (sodium iodide or lead-glass). EGS [41] from SLAC was the state-of-the-art program for simulating an electromagnetic shower in the simple geometry of these calorimeters. EGS users had to implement essentially two functions: to tell the system "where I am" or "where do I go". With the growing complexity of the detectors one soon realized that full simulation (tracking + calorimetry) was essential for testing the reconstruction programs and computing the geometrical acceptance. In 1975 a combined effort between the Omega, NA3 and NA4 experiments at the SPS led to the Geant1 and later Geant2 simulation tools (Generation of Events ANd Tracks). These provided a simple environment to propagate particles produced by an event generator, specific to each experiment, through the various detector modules, taking into account magnetic fields and multiple scattering. In 1981, a major upgrade to Geant3 [42] provided a tool to describe geometry with an automatic transport mechanism. Geant3 was interfaced to EGS3, then EGS4 and several hadronic shower packages. Geant3 included major developments and became the main simulation tool for the LEP detectors, for experiments at the Tevatron and RHIC and for the design of LHC detectors.

In 1994, the Geant4 [43] system was implemented in C++, based on the same principles as Geant3 as concerns detector description and transport. Today it is the cornerstone for all detector simulations in the HEP world. During phase1 of the LHC, about 55% of the total CPU time on the GRID machines was dedicated to detector simulation. It has become essential to improve the performance of the simulation tools, and a new project GeantV was launched in 2012 with as objective to gain possibly a factor 10 in speed. Both Geant3 and Geant4 had and still have a strong impact outside the HEP world, contributing to the design and analysis of many devices in astrophysics, medicine, hadron therapy, materials science.

Another development should be mentioned in the field of simulation: FLUKA [44]. Started as a program to calculate radiation shielding at the PS in the early 1960s, it was refined for the same purpose to cope with the higher energies of the SPS. It went through a major upgrade which made it a general purpose tool for calculations of particle transport and interactions with matter, covering an extended range of applications from proton and electron accelerator shielding to target design, calorimetry, activation, dosimetry, detector design, cosmic rays, neutrino physics and radiotherapy.