CERNCOURIER

60 years 1959–2019

IN FOCUS COMPUTING

A retrospective of 60 years' coverage of computing technology

cerncourier.com 2019

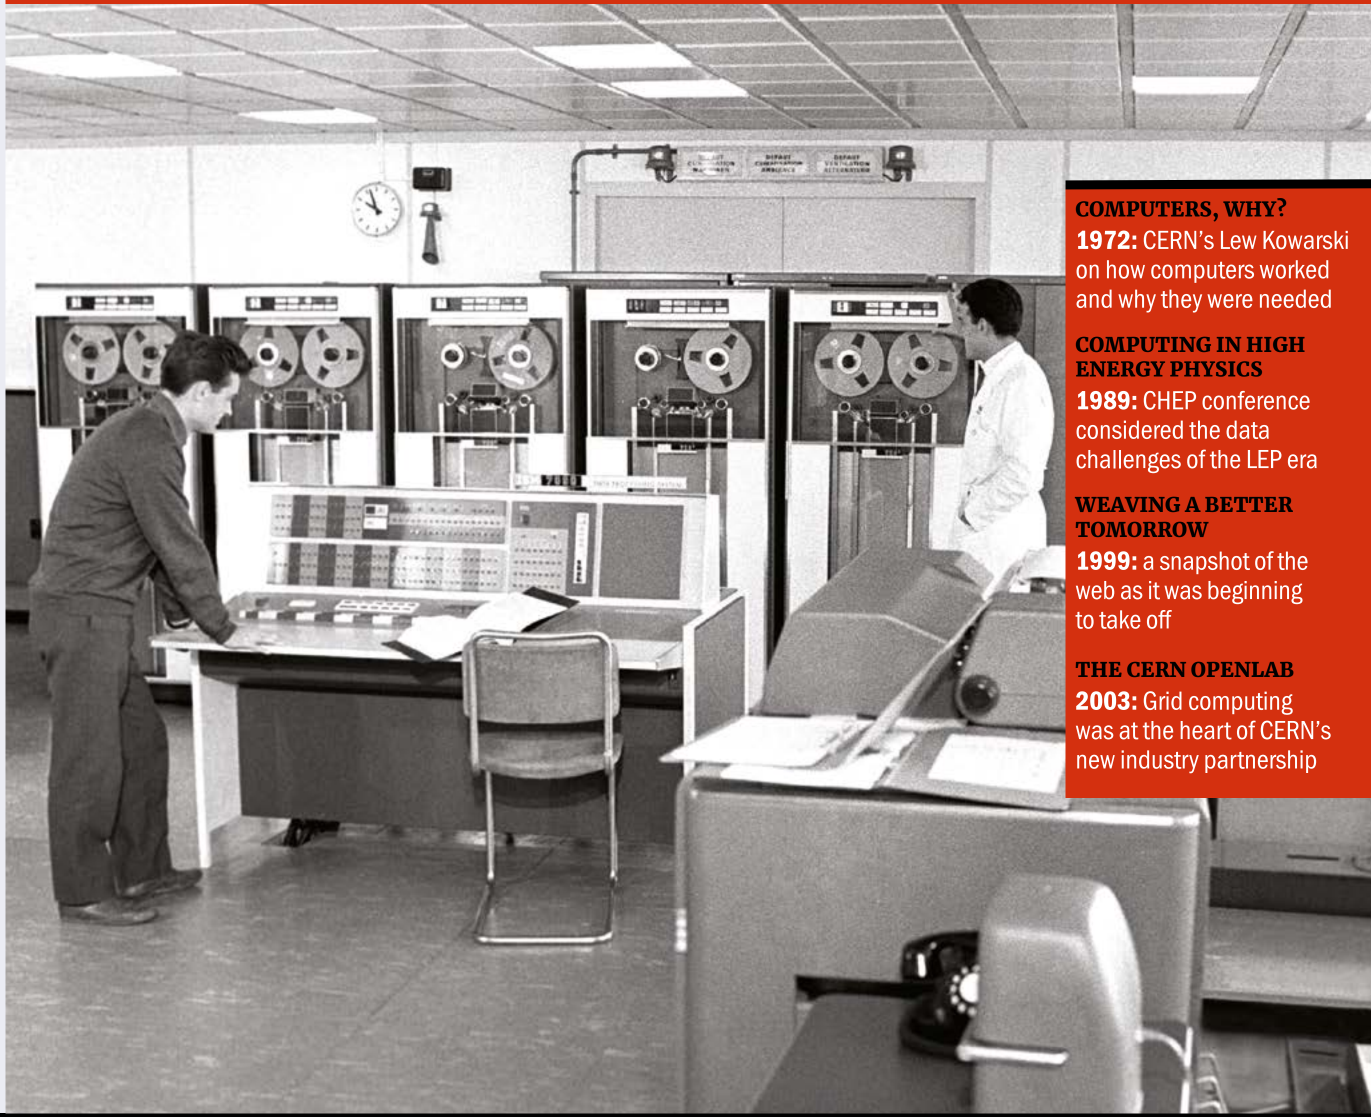

COMPUTERS, WHY?
1972: CERN's Lew Kowarski on how computers worked and why they were needed

COMPUTING IN HIGH ENERGY PHYSICS
1989: CHEP conference considered the data challenges of the LEP era

WEAVING A BETTER TOMORROW
1999: a snapshot of the web as it was beginning to take off

THE CERN OPENLAB
2003: Grid computing was at the heart of CERN's new industry partnership

IN THIS ISSUE

CERN COURIER IN FOCUS COMPUTING

From the Editor

Welcome to this retrospective supplement devoted to computing – the fourth in a series celebrating *CERN Courier*'s 60th anniversary. The *Courier* grew up during the computing revolution. CERN's first computer, the Ferranti Mercury, which had less computational ability than an average pocket calculator has today, was installed in 1958 and early issues of the magazine carried helpful articles explaining how computers work and why they are important. As the reproduced articles in this issue show, it did not take long for computing to become a pillar of high-energy physics. The field has inspired many advances, most famously the web and the Grid, and must now pursue the most advanced technology developed elsewhere to meet the demands of future experiments. In our special foreword, software-specialist Graeme Stewart sets out the need for physicists to prepare for a deluge of data in a fast-evolving technological and commercial landscape.

• All of this year's special supplements, including those devoted to detector, vacuum and magnet technology, can be read at cerncourier.com under the section "In Focus".

Matthew Chalmers

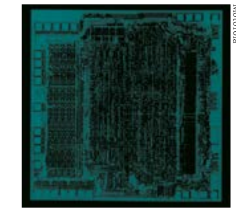

Motorola

Big impact *An early microprocessor chip.* **19**

CERN

Data shift *CERN's Computer Centre in 1998, with banks of computing arrays for the major experiments.* **25**

Editor Matthew Chalmers
Associate editor Mark Rayner
E-mail cern.courier@cern.ch

Produced for CERN by IOP Publishing Ltd Temple Circus, Temple Way, Bristol BS1 6HG, UK Tel +44 (0)117 929 7481

Head of Media Jo Allen
Head of Media Business Development Ed Jost
Art direction Frédérique Swist
Production editor Ruth Leopold

Advertising sales Tom Houlden, Chris Thomas
Advertisement production Katie Graham
Marketing and circulation Laura Gillham

Advertising Tel +44 (0)117 930 1026 or +44 (0)117 930 1164; e-mail sales@cerncourier.com

General distribution courrier-adressage@cern.ch

Published by CERN, 1211 Geneva 23, Switzerland Tel +41 (0) 22 767 61 11

Printed by Warners (Midlands) plc, Bourne, Lincolnshire, UK

IOP Publishing

SPECIAL FOREWORD

Adapting to exascale computing

CERN's Graeme Stewart surveys six decades of computing milestones in high-energy physics and describes the immense challenges in taming the data volumes from future experiments in a rapidly changing technological landscape.

Graeme Stewart *is a software specialist in the CERN EP-SFT group. He is a member of the ATLAS experiment and coordinator of the High-Energy Physics Software Foundation.*

S Bennett/CERN

Ever-evolving *The CERN data centre, photographed in 2016.*

It is impossible to envisage high-energy physics without its foundation of microprocessor technology, software and distributed computing. Almost as soon as CERN was founded the first contract to provide a computer was signed, but it took manufacturer Ferranti more than two years to deliver "Mercury", our first valve-based behemoth, in 1958. So early did this machine arrive that the venerable FORTRAN language had yet to be invented! A team of about 10 people was required for operations and the I/O system was already a bottleneck. It was not long before faster and more capable machines were available at the lab. By 1963, an IBM 7090 based on transistor technology was available with a FORTRAN compiler and tape storage. This machine could analyse 300,000 frames of spark-chamber data – a big early success. By the 1970s, computers were important enough that CERN hosted its first Computing and Data Handling School. It was clear that computers were here to stay.

By the time of the LEP era in the late 1980s, CERN hosted multiple large mainframes. Workstations, to be used by individuals or small teams, had become feasible. DEC VAX systems were a big step forward in power, reliability and usability and their operating system, VMS, is still talked of warmly by older colleagues in the field. Even more economical machines, personal computers (PCs), were also reaching a threshold of having enough computing power to be useful to physicists. Moore's law, which predicted the doubling of transistor densities every two years, was well established and PCs were riding this technological wave. More transistors meant more capable computers, and every time transistors got smaller, clock speeds could be ramped up. It was a golden age where more advanced machines, running ever faster, gave us an exponential increase in computing power.

Key also to the computing revolution, alongside the hardware, was the growth of open-source software. The GNU project had produced many utilities that could be used by hackers and coders on which to base their own software. With the start of the Linux project to provide a kernel, humble PCs became increasingly capable machines for scientific computing. Around the same time, Tim Berners-Lee's proposal for the World Wide Web, which began as a tool for connecting information for CERN scientists, started to take off. CERN realised the value in releasing the web as an open standard and in doing so enabled a success that today connects almost the entire planet.

LHC computing

This interconnected world was one of the cornerstones of the computing that was envisaged for the Large Hadron Collider (LHC). Mainframes were not enough, nor were local clusters. What the LHC needed was a worldwide system of interconnected computing systems: the Worldwide LHC Computing Grid (WLCG). Not only would information need to be transferred, but huge amounts of data and millions of computer jobs would need to be moved and executed, all with a reliability that would support the LHC's physics programme. A large investment in brand new grid technologies was undertaken, and software engineers and physicists in the experiments had to develop, deploy and operate a new grid system utterly unlike anything that had gone before. Despite rapid progress in computing power, storage space and networking, it was extremely hard to make a reliable, working distributed system for particle physics out of these pieces. Yet we achieved this incredible task. During the past decade, thousands of physics results from the four LHC experiments, including the Higgs-boson discovery, were enabled by the billions of jobs executed and the petabytes of data shipped around the world.

The software that was developed to support the LHC is equally impressive. The community had made a wholesale migration from the LEP FORTRAN era to C++ and millions of lines of code were developed. Huge software efforts in every experiment produced frameworks that managed data taking and reconstruction of raw events to analysis data.

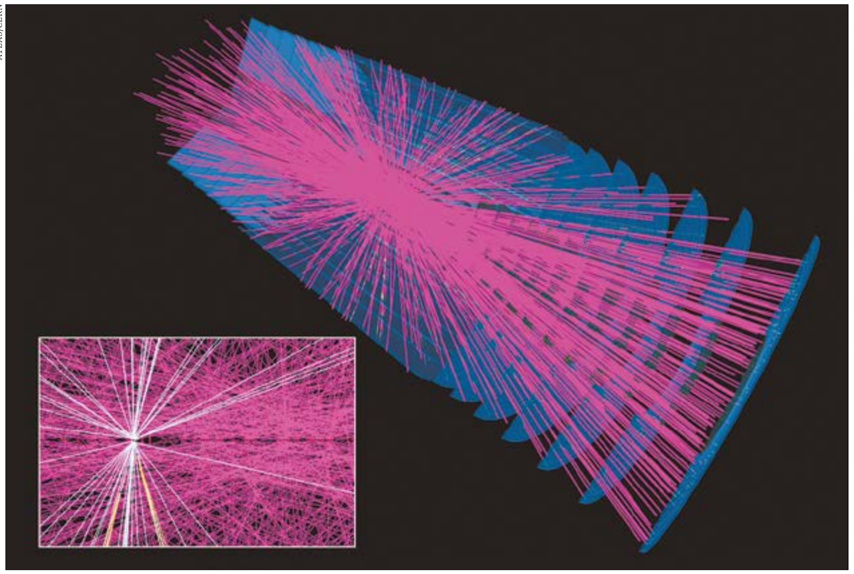

ATLAS/CERN

Close encounters *A simulated HL-LHC collision event in an upgraded ATLAS detector, which has an average of around 200 collisions per particle bunch crossing.*

In simulation, the Geant4 toolkit enabled the experiments to begin data-taking at the LHC with a fantastic level of understanding of the extraordinarily complex detectors, enabling commissioning to take place at a remarkable rate. The common ROOT foundational libraries and analysis environment allowed physicists to process the billions of events that the LHC supplied and extract the physics from them successfully at previously unheard of scales.

Changes in the wider world

While physicists were busy preparing for the LHC, the web became a pervasive part of people's lives. Internet superpowers like Google, Amazon and Facebook grew up as the LHC was being readied and this changed the position of particle physics in the computing landscape. Where particle physics had once been a leading player in software and hardware, enjoying good terms and some generous discounts, we found ourselves increasingly dwarfed by these other players. Our data volumes, while the biggest in science, didn't look so large next to Google; the processing power we needed, more than we had ever used before, was small beside Amazon; and our data centres, though growing, were easily outstripped by Facebook.

Technology, too, started to shift. Since around 2005, Moore's law, while still largely holding, has no longer been accompanied by increases in CPU clock speeds. Programs that ran in a serial mode on a single CPU core therefore started to become constrained in their performance. Instead, performance gains would come from concurrent execution on multiple threads or from using vectorised maths, rather than from faster cores. Experiments adapted by executing more tasks in parallel – from simply running more jobs at the same time to adopting multi-process and multi-threaded processing models. This *post hoc* parallelism was often extremely difficult because the code and frameworks written for the LHC had assumed a serial execution model.

It is both a challenge and also an opportunity to work with new scientific partners in the era of exascale science

The barriers being discovered for CPUs also caused hardware engineers to rethink how to exploit CMOS technology for processors. The past decade has witnessed the rise of the graphics processing unit (GPU) as an alternative way to exploit transistors on silicon. GPUs run with a different execution model: much more of the silicon is devoted to floating-point calculations, and there are many more processing cores, but each core is smaller and less powerful than a CPU. To utilise such devices effectively, algorithms often have to be entirely rethought and data layouts have to be redesigned. Much of the convenient, but slow, abstraction power of C++ has to be given up in favour of more explicit code and simpler layouts. However, this rapid evolution poses other problems for the code long term. There is no single way to programme a GPU and vendors' toolkits are usually quite specific to their hardware.

All of this would be less important were it the case that the LHC experiments were standing still, but nothing could be further from the truth. For Run 3 of the LHC, scheduled to start in 2021, the ALICE and LHCb collaborations are installing new detectors and preparing to take massively more data than they did up to now. Hardware triggers are being dropped in favour of full software processing systems and continuous data processing. The high-luminosity upgrade of the LHC for Run 4, from 2026, will be accompanied by new detector systems for ATLAS and CMS, much higher trigger rates and greatly increased event complexity. All of this physics needs to be supported by a radical evolution of software and computing systems, and in a more challenging sociological and technological environment. The LHC will also not be the only scientific big player in the future. Facilities such as DUNE, FAIR, SKA and LSST will come online and have to handle as much, if not more, data than at CERN and in the WLCG. That is both a challenge but also an opportunity to work with new scientific partners in the era of exascale science.

There is one solution that we know will not work: simply scaling up the money spent on software and computing. We will need to live with flat budgets, so if the event rate of an experiment increases by a factor of 10 then we have a budget per event that just shrank by the same amount! Recognising this, the HEP Software Foundation (HSF) was invited by the WLCG in 2016 to produce a roadmap for how to evolve software and computing in the 2020s – resulting in a community white paper supported by hundreds of experts in many institutions worldwide (*CERN Courier* April 2018 p38). In parallel, CERN open lab – a public–private partnership through which CERN collaborates with leading ICT companies and other research organisations – published a white paper setting out specific challenges that are ripe for tackling through collaborative R&D projects with leading commercial partners.

Facing the data onslaught

Since the white paper was published, the HSF and the LHC-experiment collaborations have worked hard to tackle the challenges it lays out. Understanding how event generators can be best configured to get good physics at minimum cost is a major focus, while efforts to get simulation speed-ups from classical fast techniques, as well as new machine-learning approaches, have intensified. Reconstruction algorithms have been reworked to take advantage of GPUs and accelerators, and are being seriously considered for Run 3 by CMS and LHCb (as ALICE makes even more use of GPUs since their successful deployment in Run 2). In the analysis domain, the core of ROOT is being reworked to be faster and also easier for analysts to work with. Much inspiration is taken from the Python ecosystem, using Jupyter notebooks and services like SWAN.

These developments are firmly rooted in the new distributed models of software development based on GitHub or GitLab and with worldwide development communities, hackathons and social coding. Open source is also vital, and all of the LHC experiments have now opened up their software. In the computing domain there is intense R&D into improving data management and access, and the ATLAS-developed Rucio data management system is being adopted by a wide range of other HEP experiments and many astronomy communities. Many of these developments got a shot in the arm from the IRIS–HEP project in the US; other European initiatives, such as IRIS in the UK and the IDT-UM German project are helping, though much more remains to be done.

All this sets us on a good path for the future, but still, the problems remain significant, the implementation of solutions is difficult and the level of uncertainty is high. Looking back to the first computers at CERN and then imagining the same stretch of time into the future, predictions are next to impossible. Disruptive technology, like quantum computing, might even entirely revolutionise the field. However, if there is one thing that we can be sure of, it's that the next decades of software and computing at CERN will very likely be as interesting and surprising as the ones already passed. ●

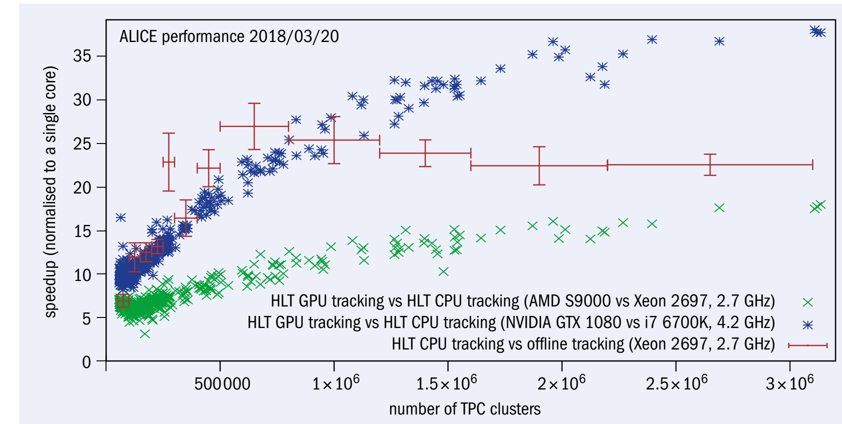

arXiv:1905.05515

Delivering data *Speedup of the ALICE TPC tracking on GPUs, normalised to a single CPU core, based on lead–lead collision data collected in 2015.*

When SCADA is the right solution

Throughout Cosylab's work in developing supervisory control and data acquisition (SCADA) solutions, we have often noticed that our customers have a lack of understanding of the role of SCADA in a system when we are asked to provide a solution. This poses a risk when it comes to whether or not our customers will find our solution useful for their work. As it is in our best interests to provide a useful resolution to our customers, and since this is also what we strive for, we would like to share what we see as the role of SCADA. In cases when SCADA really is the right solution for the given problem, then a number of questions must be asked in order to design a successful SCADA solution.

Our ultimate aim is to bring clarity to those thinking about whether they need SCADA or not.

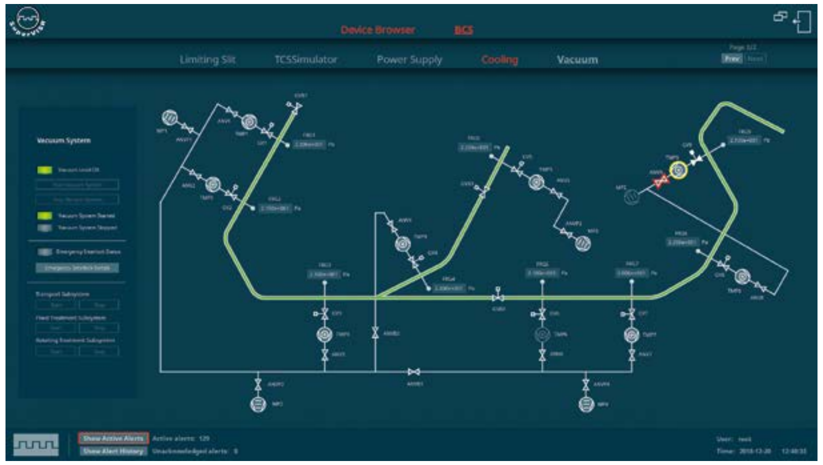

Vacuum

What is a SCADA system?

SCADA systems are used all over the world for supervisory control and data acquisition. Understanding a useful SCADA and how to make it begins with an understanding about the role of SCADA in a system.

The need for SCADA evolved over time after agrarian and handicraft economies shifted rapidly to industrial and machine-manufacturing economies during the Industrial Revolution in the 18th century. Initially, machines were developed that could perform repeatable processes faster, with more consistency and with greater precision than people. Much of this was also about eliminating human error. While this first step in the Industrial Revolution replaced many people with machines for work that was previously done by hand, it opened up questions about whether it was possible to completely eliminate the human factor from the whole process, including the management of connecting all the different aspects of the manufacturing process, as well as the hands-on aspects.

For example, there was a time when mechanised processes were made of multiple machines operated by many people, including at least one who would be the supervisor of the whole process. The process supervisor was there to make sure that the system worked accurately, and to ensure that this happened the machine operators had to report any problems that could hinder the operation to the supervisor. As the supervisor had an overview of the whole process, they could make the most appropriate decisions to keep it running.

This is the purpose of SCADA. It centralises control, congregates and exposes data to other system components on higher levels, and attempts to minimise human interaction with process control.

Determining the right SCADA solution

Only after we have decided whether SCADA is the correct solution to a problem or not can we start asking ourselves what we want SCADA to do. This can be achieved through the following set of questions:

1. What does the system need to be supervised and acted upon (by SCADA) to ensure that the process will do what we want it to do?
- o What are the building blocks of the process that need supervision and what exactly does the supervision mean for every single building block?
- o What are the supervision tasks that are common for all the building blocks of the process?

2. What are the possible events in the system that would negatively impact on what we want the system to do (e.g. machine failure, software bugs, human error) and which of these need be handled by SCADA?
- o How can these events be handled to avoid negative impacts on the system?
- o Is it possible to detect these events?
- o How will SCADA receive notifications about these events?
- o Are there any preventative actions that we can take to prevent these events from happening?

3. What supervisory information is required to enable improvements to the system?
- o Here is where different types of SCADA statistics come in, including data mining and key performance indicators.

4. Which parts of the work defined in the first three questions can be automated; which parts do we want to automate; and which parts require human intervention/supervision?
- o In the case of a fully automated process, you would need a graphical user interface with a single start button.

In conclusion

The key to getting a useful SCADA starts with a conversation with the customer about what they need and a clear conceptual design, guided by the above questions. Only after this process has been completed does it make sense to choose a specific SCADA technology (for example, WinCC OA, EPICS or TANGO) that best fits the job.

Author
Matevž Žugelj,
Project Manager,
WinCC OA Expert

COSYLAB

Cosylab., Control System Laboratory
Tel +386 1 477 66 76
Email info@cosylab.com
Web www.cosylab.com

COMPUTERS AT CERN

With CERN about to be equipped with the most powerful computer in Europe, F Beck of the data handling division considered the problems that it faced.

CERN PI 91.10.63

A corner of the IBM 7090 computer room at CERN. At the back are five of the magnetic-tape units, and in front of them the main control desk. In the right foreground is the on-line printer, which provides instructions to the operators and information on the progress of the calculations. Also in the picture (left to right): Richard Milan, Lucio Gourdiole and Eric Swoboda.

The popular picture of a biologist shows him wearing a white coat, peering through a microscope. If it is at all possible to draw a popular picture of an experimental high-energy physicist, he will probably be sitting at his desk and looking at his output from the computer. For the computer is increasingly becoming the tool by which raw experimental results are made intelligible to the physicist.

By one of those apparent strokes of luck that occur so often in the history of science, electronic digital computers became available to scientists at just that stage in the development of fundamental physics when further progress would otherwise have been barred by the lack of means to perform large-scale calculations. Analogue computers, which had been in existence for a number of decades previously, have the disadvantage for this work of limited accuracy and the more serious drawback that a more complicated calculation needs a more complicated machine. What was needed was the equivalent of an organized team of people, all operating desk calculating machines, so that a large problem in computation could be completed and checked in a reasonable time. So much was this need felt, that in England, for example, before digital computers became generally available, there was a commercial organization supplying just such a service of hand computation.

At CERN the requirement for computing facilities in the Theory Division was at first largely satisfied by employing a calculating prodigy, Willem Klein. Mr. Klein is one of those rare people who combine a prodigious memory with a love of numbers, and it was some time before computers in their ever-increasing development were able to catch up with him. He is still with us in the Theory Division, giving valuable help to those who need a quick check calculation. It is of interest to note, however, that he has now added a knowledge of computer programming to his armoury of weapons for problem solution!

The first computer calculations made at CERN were also done in the very earliest days of the Organization. Even before October 1958, when the Ferranti Mercury computer was installed, computer work had been sent out to an English Electric Deuce in Teddington, an IBM 704 in Paris and Mercury computers at Harwell and Manchester. Much of this work was concerned with orbit calculations for the proton synchrotron, then being built.

We have now reached a stage at which there is hardly a division at CERN not using its share of the available computer time. On each occasion that a new beam is set up from the accelerator, computer programmes perform the necessary calculations in particle optics as a matter of routine; beam parameters are kept in check by statistical methods; hundreds of thousands of photographs from track-chamber experiments are 'digitized' and have kinematic and statistical calculations performed on them; the new technique of sonic spark chambers, for the filmless detection of particle tracks, uses the computer more directly. In addition, more than 80 physicists and engineers use the computer on their own account, writing programmes to solve various computational problems that arise in their day-to-day work.

There are now two computers at CERN, the original Ferranti Mercury and the IBM 7090, a transistorized and more powerful replacement for its predecessor, the IBM 709. The 7090, in spite of its great speed (about 100 000 multiplications per second!), is rapidly becoming overloaded and is to be replaced towards the end of this year by a CDC 6600, which at present is the most powerful computing system available in the world.

It is hoped that this new machine will satisfy the computing needs of CERN for upwards of five years. The Mercury computer is now being used more and more as an experimental machine and there is, for instance, a direct connexion to it at present from a spark-chamber experiment at the proton synchrotron. Calculations are performed and results returned to the experimental area immediately, giving great flexibility.

This article was adapted from text in *CERN Courier* vol. 4, June 1964, pp73–77

Track-chamber photographs

Among the biggest users of computer time are the various devices for converting the information on bubble-chamber and spark-chamber photographs, usually on

CERN/PI 2.10.63

Part of the Mercury computer at CERN, now being used for tests with 'on-line' experiments, in which the measuring equipment is connected more or less directly to the computer, so that the results can be worked out as the accelerator run progresses. Gildaz Auffret is at the control console, with Mrs. Peggy Minor (left) and Mrs. Ursula Franceschi in the background.

35-mm film, into a form in which the tracks of the particles can be fitted with curves and the entire kinematics of an event subsequently worked out. To this end, from the earliest days of CERN, IEPs (instruments for the evaluation of photographs) [rumour once had it that IEP stood for 'instrument for the elimination of physicists'!] have been built and put into use. These instruments enable accurately measured co-ordinates of points on a track, together with certain identifying information, to be recorded on punched paper tape. Their disadvantage is that measuring is done manually, requires skill and, even with the best operator, is slow and prone to errors. The paper tapes produced have to be copied on to a magnetic tape, checking for various possible errors on the way, and the magnetic tape is then further processed to provide in turn geometric, kinematic and statistical results.

It was recognized at an early stage, both in Europe and in the United States, that for experiments demanding the digitization of very large numbers of pictures, for example those requiring high statistical accuracy, some more-automatic picture-reading equipment would be needed. Two such devices are now coming into use at CERN. One, the Hough-Powell device (known as HPD), developed jointly by CERN, Brookhaven, Berkeley and the Rutherford Laboratory, is an electro-mechanical machine of high precision, which still requires a few pilot measurements to be made manually, on a measuring table named 'Milady', when used for bubble-chamber pictures. It has already been used in one experiment, for the direct processing of 200 000 spark-chamber photographs (for which the pilot measurements are unnecessary). The other device is called 'Luciole', a faster, purely electronic machine, although of lower precision, specially developed at CERN for digitizing spark-chamber photographs.

With the sonic spark chamber, the position of the spark between each pair of plates is deduced from the time intervals between its occurrence and the detection of the sound by each of four microphones. Arrays of such devices can be connected directly to the computer, thus dispensing with the taking, developing and examination of photographs.

The study of dynamics of particles in magnetic and electric fields gives rise to another important family of computer programmes. The electron storage ring, or beam-stacking model, required the writing of a programme that followed the motion of individual batches of particles during their acceleration and stacking in the ring. A previous study by the same group resulted in a series of programmes to examine the behaviour and stability of a proposed fixed-field, alternating-gradient stacking device. Various aspects of the performance of the linac (the linear accelerator that feeds the PS) have been studied and improved using the computer, and the later stages of the design of the PS itself involved a detailed computer simulation of the beam in the ring, including the various transverse or 'betatron' oscillations to which it is subjected.

There also exists a series of particle-optics programmes used for the design and setting up of particle beams, particularly the 'separated' beams producing particles of only one kind. These are 'production' calculations, in the sense that the programme is run with new parameters every time there is a major change in beam layout in any of the experimental halls.

Computer language

At first, a major bar to the use of computers for small, but important, calculations was the difficulty of programming them in their own special 'language' to solve the specific problem in hand. What is the use of a machine that can perform a particular calculation in a minute, the would-be-user asks, if it will take a month to provide the programme for the calculation? Given a hand calculating machine, a pencil and paper, and a quiet room, I can do it myself in three weeks! This valid argument limited the use of computers to two kinds of calculation: those too long to be performed by hand, and those that had to be carried out so often that the original effort of producing the programme was justified.

This situation was rectified by the use of 'programming languages', which make it possible to express one's problem in a form closely resembling that of mathematics. Such languages, if defined rigorously enough, express the problem unambiguously, and they can be translated automatically (by the computer) into the instructions for a particular computer. Two such languages have been used at CERN: Mercury Autocode, and Fortran. The use of Mercury Autocode has recently been discontinued, but until a short time ago many physicists used both languages with great success to express their computational problems in a form directly comprehensible by a computer. Courses in the Fortran language, given both in English and French, are held regularly, and usually last about three weeks. Such 'compiler languages', as they are called, used to be considered a rather inferior method for programming computers, as the translations obtained from them often used the machine at a low efficiency, but it is now recognized that the advantage of writing programmes in a language that can be translated mechanically for a number of different computers far outweighs a small loss in programme efficiency.

In the Theory Division, computer programmes are often written by individual theorists to check various mathematical models. Having devised a formula based on a novel theory, the physicist computes theoretical curves for some function that can be compared with experimental results.

Operation

The Data Handling Division, which actually has CERN's central computers under its charge, contains a number of professional programmers, mostly mathematicians by training. Some of these are responsible for the 'systems programmes', that is, the Fortran compiler and its associated supervisor programme. Some have the job of disseminating programming knowledge, helping individual users of the computer and writing programmes for people in special cases. Others are semi-permanently attached to various divisions, working on particular experiments as members of the team.

A small number of mathematicians are also engaged in what might be called 'specialist computer research', covering such things as list-programming languages and methods of translating from one programme language into another. Such work might be expected to yield long-term profits by giving increases in computing power and efficiency.

As at all large computing installations, computer programmers at CERN do not operate the machine themselves. Data and programmes are submitted through a 'reception office' and the results are eventually available in a 'computer output office', leaving the handling and organizing of the computer work-load, and the operating of the machine, to specialist reception staff and computer operators in the Data Handling Division.

What of the future of computers at CERN? In a field as new as this, predictions are even more dangerous than in others, but it is clear that the arrival of the new computer at the end of the year will cause a great change in the way computers are used. Having ten peripheral processors, each of which is effectively an independent computer, the machine may have many pieces of equipment for data input and output attached to it 'on-line'. The old concept that a computer waiting for the arrival of data is standing idle, and that this is therefore wasteful and expensive, need no longer be true. With the new system, a computer that is waiting for new data for one problem is never idle, but continues with calculations on others. Every moment of its working day is gainfully employed on one or other of the many problems it is solving in parallel. Even so, there will still be a need for a number of smaller computer installations forming part of particular experiments.

As M.G.N. Hine, CERN's Directorate Member for Applied Physics, pointed out at a recent conference, even with the growth of such facilities, the amount of computing time available may one day dictate the amount of experimental physics research done at CERN, in much the same way as the amount of accelerator time available dictates it now. ●

As at all large computing installations, computer programmers at CERN do not operate the machine themselves

COMPUTERS, WHY?

Kicking off a special issue on computing in March 1972 – the year Intel's 8008 processor was launched and the compact disc invented – CERN's Lew Kowarski explained why computers were here to stay.

CERN is the favourite showpiece of international co-operation in advanced scientific research. The public at large is, by now, quite used to the paradox of CERN's outstandingly large-scale electromagnetic machines (accelerators) being needed to investigate outstandingly small-scale physical phenomena. A visitor finds it natural that this, largest-in-Europe, centre of particle research should possess the largest, most complex and costly accelerating apparatus.

But when told that CERN is also the home of the biggest European collection of computers, the layman may wonder: why is it precisely in this branch of knowledge that there is so much to compute? Some sciences such as meteorology and demography appear to rely quite naturally on enormously vast sets of numerical data, on their collection and manipulation. But high energy physics, not so long ago, was chiefly concerned with its zoo of 'strange particles' which were hunted and photographed like so many rare animals. This kind of preoccupation seems hardly consistent with the need for the most powerful 'number crunchers'.

Perplexities of this sort may arise if we pay too much attention to the (still quite recent) beginnings of the modern computer and to its very name. Electronic digital computers did originate in direct descent from mechanical arithmetic calculators; yet their main function today is far more significant and universal than that suggested by the word 'computer'. The French term 'ordinateur' or the Italian 'elaboratore' are better suited to the present situation and this requires some explanation.

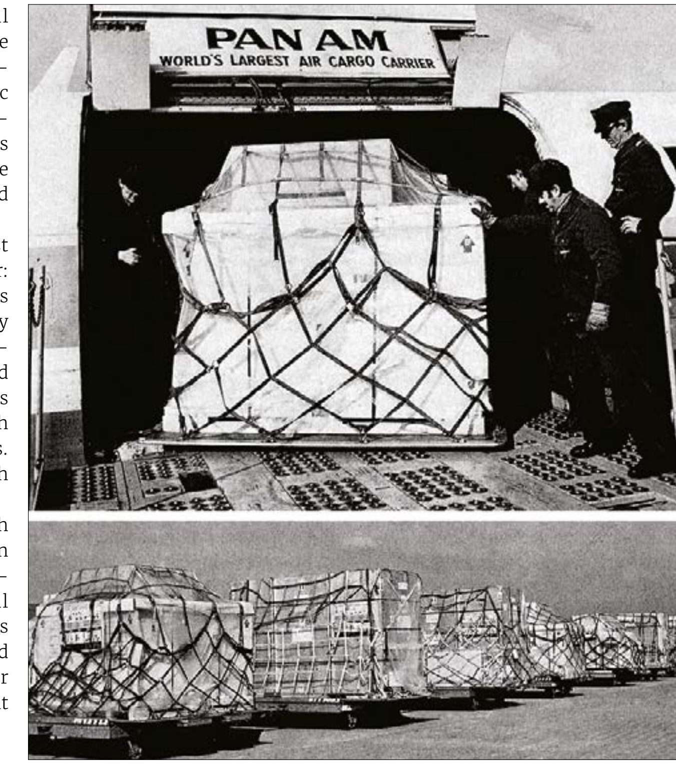

CERN's new central computer, a CDC 7600, was flown into Geneva airport mid-February and can be seen being unloaded from the plane and being wheeled across the tarmac.

What is a computer?

When, some forty years ago, the first attempts were made to replace number-bearing cogwheels and electro-mechanical relays by electronic circuits, it was quickly noticed that, not only were the numbers easier to handle if expressed in binary notation (as strings of zeros and ones) but also that the familiar arithmetical operations could be presented as combinations of two-way (yes or no) logical alternatives. It took some time to realize that a machine capable of accepting an array of binary-coded numbers, together with binary-coded instructions of what to do with them (stored program) and of producing a similarly coded result, would also be ready to take in any kind of coded information, to process it through a prescribed chain of logical operations and to produce a structured set of yes-or-no conclusions. Today a digital computer is no longer a machine primarily intended for performing numerical calculations; it is more often used for non-numerical operations such as sorting, matching, retrieval, construction of patterns and making decisions which it can implement even without any human intervention if it is directly connected to a correspondingly structured set of open-or-closed switches.

This article was adapted from text in a special issue of *CERN Courier* devoted to computing, vol. 12, March 1972, pp59–61

Automatic 'black boxes' capable of producing a limited choice of responses to a limited variety of input (for example, vending machines or dial telephones) were known before; their discriminating and logical capabilities had to be embodied in their rigid internal hardware. In comparison, the computer can be seen as a universally versatile black box, whose hardware responds to any sequence of coded instructions. The complication and the ingenuity are then largely transferred into the writing of the program.

The new black box became virtually as versatile as the human brain; at the same time it offered the advantages of enormously greater speed, freedom from error and the ability to handle, in a single operation, any desired

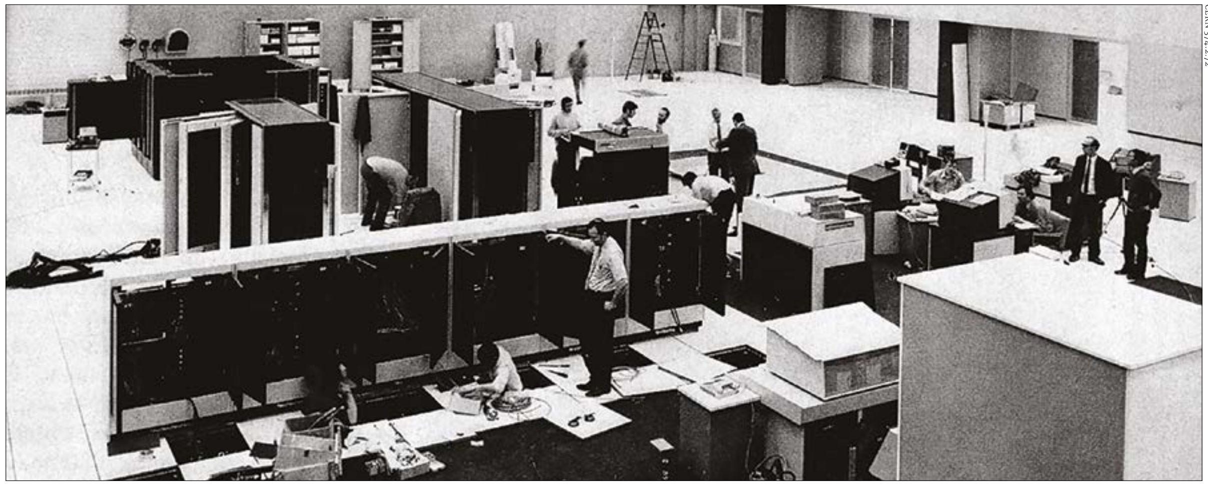

Installation of the 7600 well advanced in the large hall of the new computer centre. Top left, the shape of the main-frame can be picked out. Then comes the 6400 computer, the large core memory which will be shared by the 6400 and 6500 (Extended Core Store) and, nearest to the camera, a number of controllers which connect peripheral equipment to the 6400. On the right and in the foreground are the operator displays.

volume of incoming data and of ramified logical chains of instructions. In this latter respect the limit appears to be set only by the size (and therefore the cost) of the computer, i.e. by the total number of circuit elements which are brought together and interconnected in the same computing assembly.

High energy physics as a privileged user

We are only beginning to discover and explore the new ways of acquiring scientific knowledge which have been opened by the advent of computers. During the first two decades of this exploration, that is since 1950, particle physics happened to be the most richly endowed domain of basic research. Secure in their ability to pay, high energy physicists were not slow to recognize those features of their science which put them in the forefront among the potential users of the computing hardware and software in all of their numerical and non-numerical capabilities. The three most relevant characteristics are as follows:

1. Remoteness from the human scale of natural phenomena
Each individual 'event' involving sub-nuclear particles takes place on a scale of space, time and momentum so infinitesimal that it can be made perceptible to human senses only through a lengthy and distorting chain of larger-scale physical processes serving as amplifiers. The raw data supplied by a high energy experiment have hardly any direct physical meaning; they have to be sorted out, interpreted and re-calculated before the experimenter can see whether they make any sense at all – and this means that the processing has to be performed, if possible, in the 'real time' of the experiment in progress, or at any rate at a speed only a computer can supply.

2. The rate and mode of production of physical data
Accelerating and detecting equipment is very costly and often unique; there is a considerable pressure from the user community and from the governments who invest in this equipment that it should not be allowed to stand idle. As a result, events are produced at a rate far surpassing the ability of any human team to observe them on the spot. They have to be recorded (often with help from a computer) and the records have to be scanned and sifted – a task which, nowadays, is usually left to computers because of its sheer volume. In this way, experiments in which a prolonged 'run' produces a sequence of mostly trivial events, with relatively few significant ones mixed in at random, become possible without wasting the valuable time of competent human examiners.

3. High statistics experiments
As the high energy physics community became used to computer-aided processing of events, it became possible to perform experiments whose physical meaning resided in a whole population of events, rather than in each taken singly. In this case the need grew from an awareness of having the means to satisfy the need; a similar evolution may yet occur in other sciences (e.g. those dealing with the environment), following the currents of public attention and possibly de-throning our physics from pre-eminence in scientific computation.

Modes of application

In order to stress here our main point, which is the versatility of the modern computer and the diversity of its applications in a single branch of physical research, we shall classify all the ways in which the 'universal black box' can be put to use in CERN's current work into eight 'modes of application' (roughly corresponding to the list of 'methodologies' adopted in 1968 by the U.S. Association for Computing Machinery):

1. Numerical mathematics
This mode is the classical domain of the 'computer used as a computer' either for purely arithmetic purposes or for more sophisticated tasks such as the calculation of less common functions or the numerical solution of differential and integral equations. Such uses are frequent in practically every phase of high energy physics work, from accelerator design to theoretical physics, including such contributions to experimentation as the kinematic analysis of particle tracks and statistical deductions from a multitude of observed events.

2. Data processing
Counting and measuring devices used for the detection of particles produce a flow of data which have to be recorded, sorted and otherwise handled according to appropriate procedures. Between the stage of the impact of a fast-moving particle on a sensing device and that of a numerical result available for a mathematical computation, data processing may be a complex operation requiring its own hardware, software and sometimes a separate computer.

3. Symbolic calculations
Elementary logical operations which underline the computers' basic capabilities are applicable to all sorts of operands such as those occurring in algebra, calculus, graph theory, etc. High-level computer languages such as LISP are becoming available to tackle this category of problems which, at CERN, is encountered mostly in theoretical physics but, in the future, may become relevant in many other domains such as apparatus design, analysis of track configurations, etc.

4. Computer graphics
Computers may be made to present their output in a pictorial form, usually on a cathode-ray screen. Graphic output is particularly suitable for quick communication with a human observer and intervener. Main applications at present are the study of mathematical functions for the purposes of theoretical physics, the design of beam handling systems and Cherenkov counter optics and statistical analysis of experimental results.

5. Simulation
Mathematical models expressing 'real world' situations may be presented in a computer-accessible form, comprising the initial data and a set of equations and rules which the modelled system is supposed to follow in its evolution. Such 'computer experiments' are valuable for advance testing of experimental set-ups and in many theoretical problems. Situations involving statistical distributions may require, for their computer simulation, the use of computer-generated random numbers during the calculation. This kind of simulation, known as the Monte-Carlo method, is widely used at CERN.

6. File management and retrieval
As a big organization, CERN has its share of necessary paper-work including administration (personnel, payroll, budgets, etc.), documentation (library and publications) and the storage of experimental records and results. Filing and retrieval of information tend nowadays to be computerized in practically every field of organized human activity; at CERN, these pedestrian applications add up to a non-negligible fraction of the total amount of computer use.

7. Pattern recognition
Mainly of importance in spark-chamber and bubble-chamber experiments – the reconstruction of physically coherent and meaningful tracks out of computed coordinates and track elements is performed by the computer according to programmed rules.

8. Process control
Computers can be made to follow any flow of material objects through a processing system by means of sensing devices which, at any moment, supply information on what is happening within the system and what is emerging from it. Instant analysis of this information by the computer may produce a 'recommendation of an adjustment' (such as closing a valve, modifying an applied voltage, etc.) which the computer itself may be able to implement. Automation of this kind is valuable when the response must be very quick and the logical chain between the input and the output is too complicated to be entrusted to any rigidly constructed automatic device. At CERN the material flow to be controlled is usually that of charged particles (in accelerators and beam transport systems) but the same approach is applicable in many domains of engineering, such as vacuum and cryogenics.

We are only beginning to discover and explore the new ways of acquiring scientific knowledge which have been opened by the advent of computers

Centralization versus autonomy

The numerous computers available at CERN are of a great variety of sizes and degrees of autonomy, which reflects the diversity of their uses. No user likes to share his computer with any other user; yet some of his problems may require a computing system so large and costly, that he cannot expect it to be reserved for his exclusive benefit nor to be kept idle when he does not need it. The biggest computers available at CERN must perforce belong to a central service, accessible to the Laboratory as a whole. In recent years, the main equipment of this service has consisted of a CDC 6600 and CDC 6500. The recent arrival of a 7600 (coupled with a 6400) will multiply the centrally available computing power by a factor of about five.

For many applications, much smaller units of computing power are quite adequate. CERN possesses some 80 other computers of various sizes, some of them for use in situations where autonomy is essential (for example, as an integral part of an experimental set-up using electronic techniques or for process control in accelerating systems). In some applications there is need for practically continuous access to a smaller computer together with intermittent access to a larger one. A data-conveying link between the two may then become necessary.

Conclusion

The foregoing remarks are meant to give some idea of how the essential nature of the digital computer and that of high energy physics have blended to produce the present prominence of CERN as a centre of computational physics. The detailed questions of 'how' and 'what for' are treated in the other articles of this [special March 1972 issue of *CERN Courier* devoted to computing] concretely enough to show the way for similar developments in other branches of science. In this respect, as in many others, CERN's pioneering influence may transcend the Organization's basic function as a centre of research in high energy physics. ●

WEAVING A BETTER TOMORROW: THE FUTURE OF THE WEB

The Web was 10 years old and only just beginning to fulfil its potential, reported CERN's James Gillies from the eighth World Wide Web conference in Toronto in 1999.

Technical support at the World Wide Web conferences is traditionally supplied by volunteers. The WWW8 crew is seen here in front of the Web history wall, assembled during the conference to mark the Web's 10th anniversary. Volunteers pay their own way to get to the conference for a week's hard work, just for the pleasure of being there.

"It was a weird conference," said Ethernet inventor and self-styled technology pundit, Bob Metcalfe, summing up the eighth World Wide Web conference (WWW8). "Imagine," he continued, "sitting there listening to a senior executive of IBM wearing a tee-shirt and a beard." Appearances were not deceptive as Big Blue's vice president for Internet Technology, John Patrick, captured the spirit of the conference. "Power to the people," he said, would be the driving force behind the computing industry as we enter the new millennium. For if one thing is abundantly clear, it's that the political geography of information technology has been turned on its head by personal computing and the World Wide Web. "Stand aside, besuited corporate executives", came the message. Make way for the altruistic geeks: the future belongs to them.

It's rare to find such an optimistic bunch of people. The pony-tail count may have been way above average and the word "cool" still cool, but WWW8 delegates have their hearts in the right place. They are the ones who have made the Web, motivated only by the fun of playing with computers and the belief that the Web can make the world a better place. Some were concerned at Microsoft's conference sponsorship. There was grumbling that the delegates' pack included complementary Microsoft CD-ROMs (for Windows only). "Next year," one delegate was overheard to say, while tucking into a spring roll and sipping Chardonnay at the evening reception (courtesy of Bill Gates), "Microsoft will have bought the World Wide Web." However, his fears were not universally held. There is just too much grass-roots stuff going on out there for one company, however powerful, to take over completely.

The Web we're going to see emerging over the coming decade is none other than the one Tim Berners-Lee had running almost 10 years ago on an obscure computer called a Next cube at CERN

Information revolution

It may seem from the outside that the information revolution has arrived, but in John Patrick's view, "we're right at the beginning". The Web's inventor, Tim Berners-Lee, doesn't even go that far. The Web we're going to see emerging over the coming decade, he believes, is none other than the one he had running almost 10 years ago on an obscure computer called a Next cube at CERN. "Ask him about control-shift-N," said one delegate, referring to the combination of key strokes that instructed that early browser/editor to create a new document linked to the one you were already in. That simple manoeuvre encapsulates Berners-Lee's vision of what the Web should be, "a common space in which we could all interact", a medium in which we'd all be creators, not just consumers. Expediency prevented that reality from coming sooner as Berners-Lee and his team at CERN concentrated on providing Web services to the particle physics community, leaving the stage free for the entrance of Mosaic, a browser with no editing capacity, in 1994.

Even when the passive Web took off, Berners-Lee did not abandon his dream. To most users of the Web the choice of browsers comes down to two: Netscape Navigator and Internet Explorer. However, there's actually a lot more choice available. Many of the early browsers can still be found, and there are new companies turning out more. The Web consortium (W3C) has produced a browser/editor, called Amaya, that allows the kind of interactive Web use that Berners-Lee envisaged from the start. If you want to see what the Web was meant to be, open Navigator or Explorer for the last time, go to "http://www.w3.org" and click on "Amaya browser/editor".

Improving the Web

Content that the Web is finally catching up with his original vision, Berners-Lee is now devoting his energies to improving it. The Web's biggest problem is caused by its success. There's so much information out there that it's often hard to find what you want. The answer, according to Berners-Lee, is what he calls the semantic Web. The kind of information on the Web today is understandable to humans but not to computers. If, for example, Berners-Lee wanted to buy a yellow car in Massachusetts and his neighbour wanted to sell a primrose automobile in Boston, how would his search engine know that what he wanted was right on his doorstep? If a current W3C project is successful, some kind of logical schema will tell the search engine that primrose is just a kind of yellow and that automobiles and cars are in fact the same thing.

Reminding delegates that there's nothing new under the Sun was IBM's John Patrick who spelled out his vision of how the Internet is poised to change our lives. Top of his list of next big things was instant messaging, which is just around the corner. Curiously familiar to anyone who used BITNET or DECNET in the 1980s, instant messaging is a sort of halfway house between e-mail and the telephone. Patrick demonstrated IBM's version by typing in "How is the weather in Heidelberg" to a colleague in Germany. Out boomed the mechanical words "*Wie ist das Wetter in Heidelberg*", followed, presumably after the Heidelberger had typed "*Es ist kalt und regnerisch*", by "It is cold and rainy". That's fine if all you want to do is discuss the weather, but IBM's translation software might have problems with more complex topics. Nevertheless, it served to show what's coming.

E-business

Symbiotic video came next on Patrick's list. That's clickable television to you and me. A coffee advertisement took us to a Web site where, you've guessed it, you could order coffee to be delivered to your door. This is an example of where e-business might be taking us and, as anyone who's looked at an IBM advert recently knows, e-business is IBM's next really big thing. Defined on their Web pages as "the transformation of key business processes through the use of Internet technologies", e-business, says Patrick, will force a new character onto the keyboard. (I don't have one, so to see what he means you'll have to look at IBM's Web site yourself, "http://www.ibm.com/") What it boils down to is businesses maximizing their potential through computers and the Web with the help, of course, of IBM.

Education is already benefiting from the Web. LEGO's high-tech programmable Mindstorms invites young engineers to submit their best designs and programs to a Web site ("http://www.legomindstorms.com/"). Mindstorms impressed Patrick so much that he felt inspired to submit his own, but was distressed to find that the "date of birth" choice only went back to 1970, so that's what he clicked. The Mindstorms design that most impressed him was posted by someone who had clicked on 1992.

Can the Internet handle all of these new big things? Yes, believes Patrick. Bandwidth is booming, and the much-touted address-space problem – simply running out of new addresses – will soon be a thing of the past. The next version of the Internet protocols will bring enough addresses for every proton, according to Patrick. "That ought to do it."

There were no surprises from Greg Papadopoulos, Chief Technology Officer at Sun Microsystems. He looks forward to the day when computers will not need to rely on complicated protocols to talk to each other and to peripheral devices. Instead, there will be just one simple protocol and it will be used for sending "objects" – executable programmes – around the Web. Coming from the company that filled our Web pages with Java applets, what else could he be expected to say? The example he gave was printing

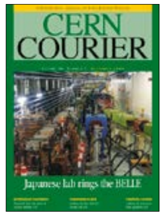

This article was adapted from text in *CERN Courier* vol. 39, September 1999, pp26–28

The Web's biggest problem is caused by its success. There's so much information out there that it's often hard to find what you want

from a Web phone, a device still far from most people's everyday reality but familiar to WWW8 conference-goers.

Patents

Conference co-chair Murray Maloney began the closing plenary session by presenting the Yuri Rabinsky award to Richard Stallman. Rabinsky was a pioneer of the Open Software movement making Stallman, founder of the Free Software Foundation and author of the operating system GNU, an appropriate recipient and Microsoft an unlikely sponsor of the award. This was an irony not lost on Stallman who graciously accepted the award while urging vigilance against those who would patent everything.

In the US, software can be patented and in Tim Berners-Lee's opinion, "the bar for what's patentable is far too low". As a consequence, a substantial part of W3C's energies are tied up in fighting patent applications covering things that the consortium believes should be standards. Stallman urges Europe not to succumb to American pressure to adopt software patents.

To sum up, Bob Metcalfe singled out the semantic Web of Tim Berners-Lee, whom he referred to as "The Duke of URL" (pronounced "Earl"), as the principal subject of the conference. A pundit's role is to stick his neck out, and Metcalfe is famous for doing that. Four years ago at WWW4, he predicted that the coming 12 months would see an Internet "Gigalapse" – a single network outage that would cost a billion man-hours. So confident was he that he promised to eat his words if there had not been one. Two years later, at WWW6, he took a copy of his column and ate it for all to see. The year's biggest outage had been estimated at a tenth of a Gigalapse.

Among his predictions at WWW8 was, again, the Gigalapse, but this time with no promises attached. Metcalfe also believes that microcharging is just around the corner. He surveyed his readers to find out how much they'd be prepared to pay to read his column. 0.2 cents came the reply, but, with half a million readers, he's quite happy about that. He had bad news for both Microsoft and Stallman, predicting that the former has peaked but the latter's Open Source will still never catch Mr Gates. The Internet stock bubble, he predicted, will burst on 8 November 1999.

How could he be so bold? He was recently invited to a meeting of venture capitalists and asked his opinion on this. After carefully explaining that he knew nothing of stock markets, he told them 8 November and was amused to see them all writing it down. Y2K will be a non-event. "Why?" he asked himself. "Are computers reliable?" he replied. "And anyway," he went on, "31 December is a Friday so we'll have the whole weekend to sort things out." So on that note, delegates were able to leave the conference looking forward to a restful last night of the millennium, whatever year they happen to believe that might occur. ●

THE MICROPROCESSOR BOOM

The early 1980s would see the microprocessor become a routine part of the high-energy physics toolkit, predicted the Courier in the summer of 1979.

During the past few years, electronic circuitry techniques have been developed which enable complex logic units to be produced as tiny elements or 'chips'. These units are now mass-produced, and are available relatively cheaply to anyone building data processing equipment.

Just a few years ago, the first complete processing unit on a single chip was produced. Now micro logic elements can be combined together to provide micro data processing systems whose capabilities in certain respects can rival those of more conventional computers. Commercially-available microcomputers are used widely in many fields.

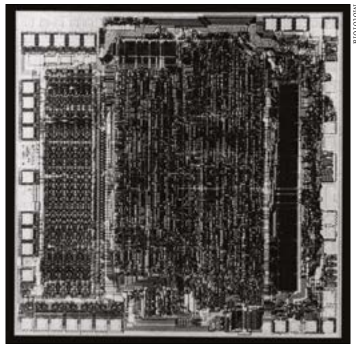
Enlargement of a microprocessor (in real life just a few millimetres across). Units like these could soon make a big impact in data processing for high energy physics.

Where an application requires special capabilities, it is preferable to take the individual micro logic units and wire them together on a printed circuit board to provide a tailor-made processing unit: If there is sufficient demand for the perfected design, the printed circuit board stage subsequently can be dispensed with and the processor can be mass-produced by large-scale integration (LSI) techniques as a single microprocessor.

With these processing units, there generally is a trade-off between speed and flexibility, the ultimate in speed being a hard-wired unit which is only capable of doing one thing. Flexibility can be achieved through programmable logic, but this affects the overall speed.

Programming micros is difficult, but one way of side-stepping these problems would be to design a unit which emulates a subset of an accessible mainframe computer. With such an emulator, programs could be developed on the main computer, and transferred to the micro after they have reached the required level of reliability. This could result in substantial savings in program development time. In addition, restricting the design to a subset of the mainframe architecture results in a dramatic reduction in cost.

High energy physics, which has already amply demonstrated its voracious appetite for computer power, could also soon cash in on this microcomputer boom and produce its own 'brand' of custom-built microprocessors.

According to Paolo Zanella, Head of CERN's Data Handling Division, now is the time to explore in depth the uses of microprocessors in high energy physics experiments. If initial projects now under way prove to be successful, the early 1980s could see microprocessors come into their own.

One of the biggest data processing tasks in any physics experiment is to sift through the collected signals from the various detecting units to reject spurious information and separate out events of interest. Therefore to increase the richness of the collected data, triggering techniques are used to activate the data collection system of an experiment only when certain criteria are met.

Even with the help of this 'hardwired' selection, a large proportion of the accumulated data has to be thrown away, often after laborious calculations. With experiments reaching for higher energies where many more particles are produced, and at the same time searching for rarer types of interaction, physicists continually require more and more computing power.

Up till now, this demand has had to be met by bringing in more and bigger computers, both on-line at the experiments and off-line at Laboratory computer centres. With the advent of microprocessors, a solution to this problem could be in sight. Micros could be incorporated into experimental set-ups to carry out a second level of data selection after the initial hard-wired triggering – an example of the so called 'distributed processing' approach where computing power is placed as far upstream as possible in the data handling process. In this way the demand on the downstream central computer would be reduced, and the richness of the data sample increased.

The micros would filter the readout in the few microseconds before the data is transferred to the experimental data collection system. Zanella is convinced that this could significantly improve the quality of the data and reduce the subsequent off-line processing effort to eliminate bad triggers.

As well as being used in the data collection system, micros would also be useful for control and monitoring functions. The use of off the-shelf microcomputers in accelerator control systems, for example, is already relatively widespread. Some limited applications outside the control area are already being made in experiments, a notable example being the CERN/Copenhagen/Lund/Rutherford experiment now being assembled at the CERN Intersecting Storage Rings.

Microcomputer projects are now being tackled at several Laboratories. At CERN three projects are under way in the Data Handling Division. Two of these are programmable emulators (one being based on the IBM 370/168 and the other on the Digital Equipment PDP-11), while the third is a very fast microprogrammable unit called ESOP.

High energy physics still has a lot to learn about microprocessor applications, and there is some way to go before their feasibility is demonstrated and practical problems, such as programming, are overcome.

However this year could see some of these initial projects come to fruition, and the early 1980s could live up to Zanella's expectations as the time when the microprocessor becomes a routine part of the high energy physicists' toolkit. ●

This article was adapted from text in *CERN Courier* vol. 19, July/August 1979, pp192–193

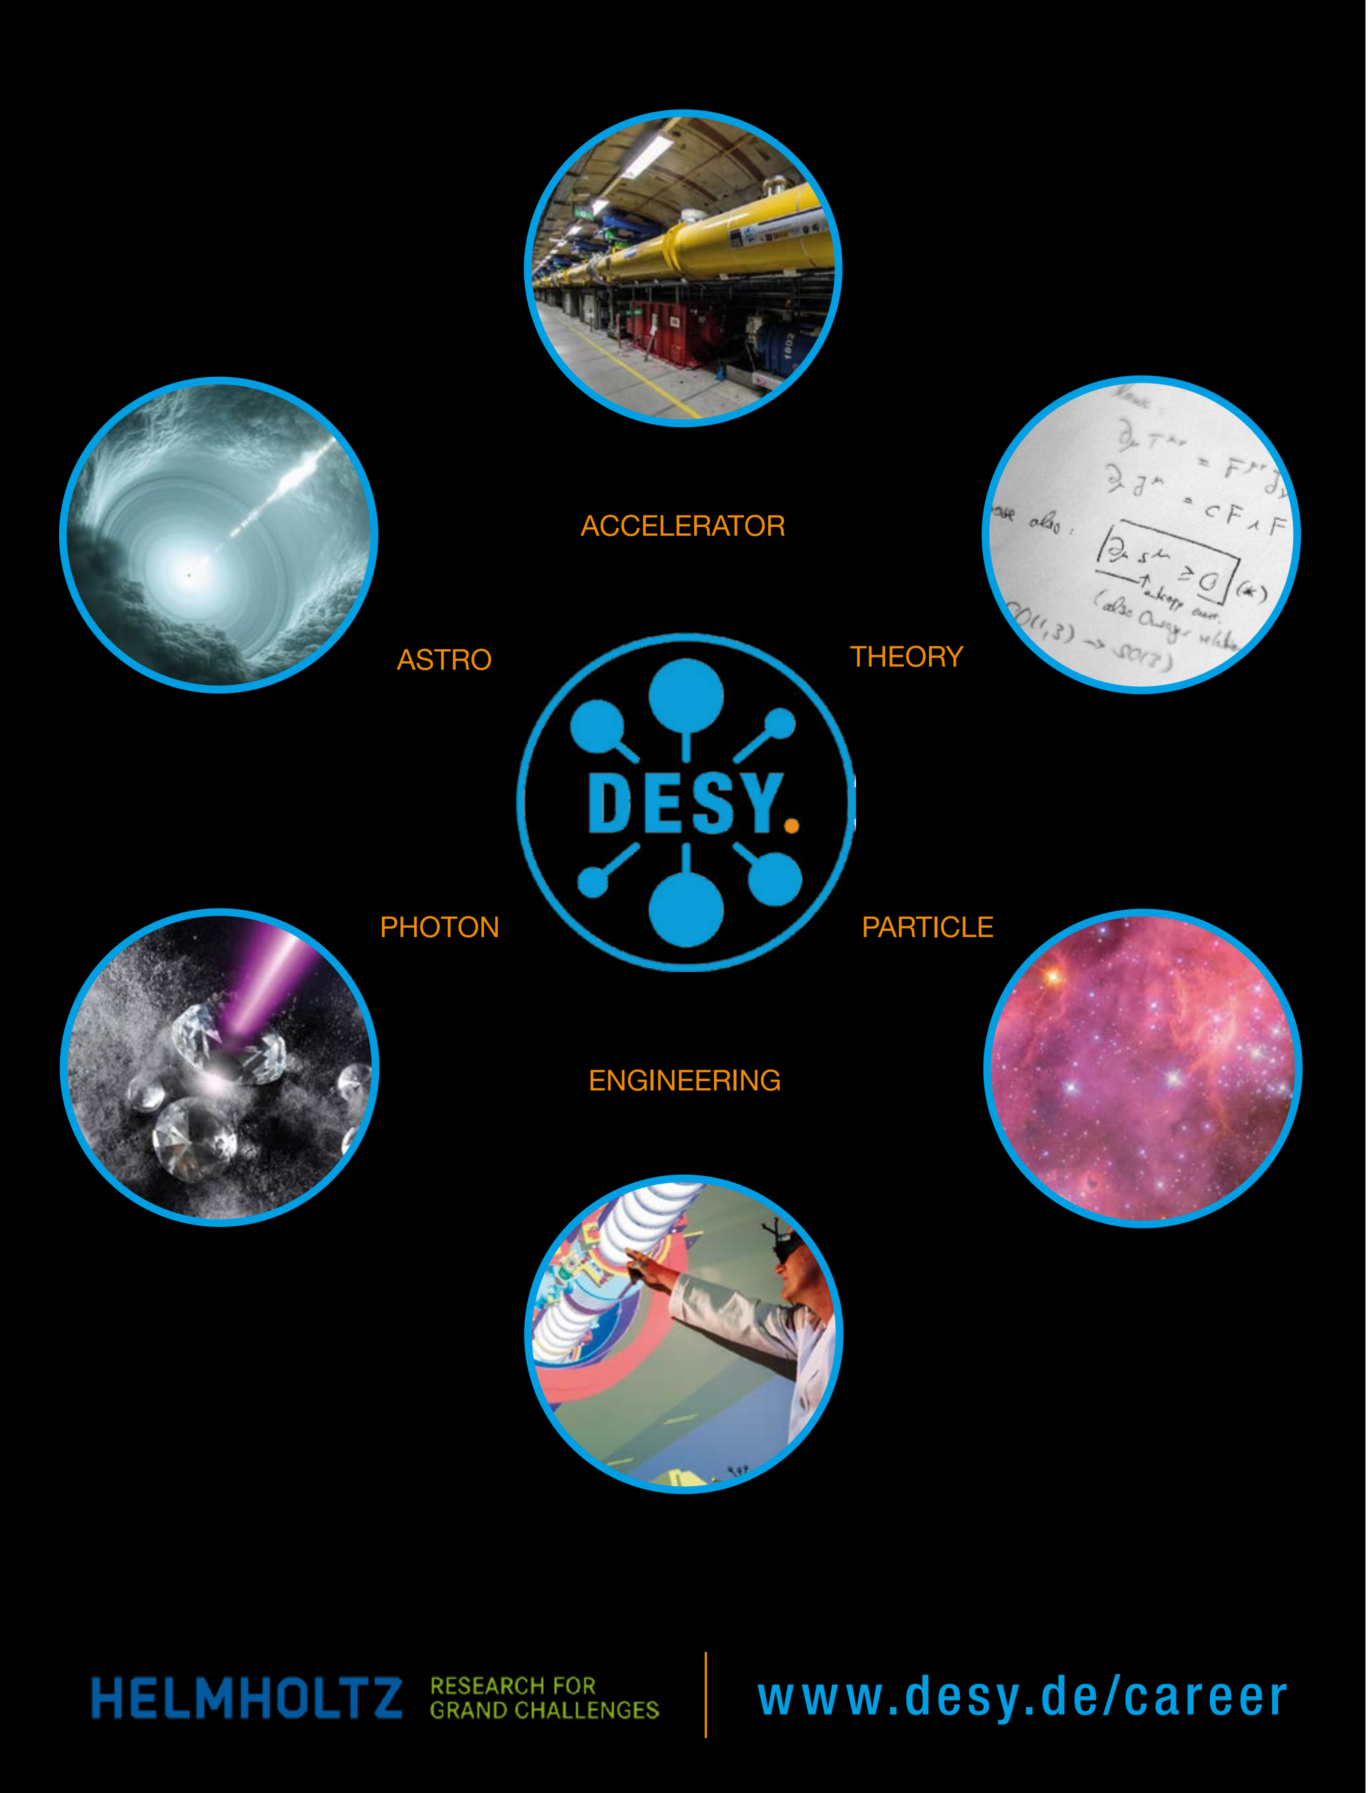

COMPUTING IN HIGH ENERGY PHYSICS

The 1989 Computing In High Energy Physics conference weighed up the challenges of analysing LEP and other data, reported Sarah Smith and Robin Devenish.

RAL

David Williams (CERN) sets the scene at the Oxford Conference on Computing in High Energy Physics.

Computing in high energy physics has changed over the years from being something one did on a slide-rule, through early computers, then a necessary evil to the position today where computers permeate all aspects of the subject from control of the apparatus to theoretical lattice gauge calculations.

The state of the art, as well as new trends and hopes, were reflected in this year's 'Computing In High Energy Physics' conference held in the dreamy setting of Oxford's spires. The 260 delegates came mainly from Europe, the US, Japan and the USSR. Accommodation and meals in the unique surroundings of New College provided a special atmosphere, with many delegates being amused at the idea of a 500-year-old college still meriting the adjective 'new'.

The conference aimed to give a comprehensive overview, entailing a heavy schedule of 35 plenary talks plus 48 contributed papers in two afternoons of parallel sessions. In addition to high energy physics computing, a number of papers were given by experts in computing science, in line with the conference's aim – 'to bring together high energy physicists and computer scientists'.

The complexity and volume of data generated in particle physics experiments is the reason why the associated computing problems are of interest to computer science. These ideas were covered by David Williams (CERN) and Louis Hertzberger (Amsterdam) in their keynote addresses.

The task facing the experiments preparing to embark on CERN's new LEP electron–positron collider is enormous by any standards but a lot of thought has gone into their preparation. Getting enough computer power is no longer thought to be a problem but the issue of storage of the seven Terabytes of data per experiment per year makes computer managers nervous even with the recent installation of IBM 3480 cartridge systems.

With the high interaction rates already achieved at the CERN and Fermilab proton–antiproton colliders and orders of magnitude more to come at proposed proton colliders, there are exciting areas where particle physics and computer science could profitably collaborate. A key area is pattern recognition and parallel processing for triggering. With 'smart' detector electronics this ultimately will produce summary information already reduced and fully reconstructed for events of interest.

Is the age of the large central processing facility based on mainframes past? In a provocative talk Richard Mount (Caltech) claimed that the best performance/cost solution was to use powerful workstations based on reduced instruction set computer (RISC) technology and special purpose computer-servers connected to a modest mainframe, giving maybe a saving of a factor of five over a conventional computer centre.

Networks bind physics collaborations together, but they are often complicated to use and there is seldom enough capacity. This was well demonstrated by the continuous use of the conference electronic mail service provided by Olivetti-Lee Data. The global village nature of the community was evident when the news broke of the first Z particle at Stanford's new SLC linear collider.

The problems and potential of networks were explored by François Fluckiger (CERN), Harvey Newman (Caltech) and James Hutton (RARE). Both Fluckiger and Hutton explained that the problems are both technical and political, but progress is being made and users should be patient. Network managers have to provide the best service at the lowest cost. HEPNET is one of the richest structures in networking, and while OSI seems constantly to remain just around the corner, a more flexible and user-friendly system will emerge eventually.

This article was adapted from text in *CERN Courier* vol. 29, July/August 1989, pp4–8

NPL Oxford

Hands-on experience at the DEC demonstration.

Harvey Newman (Caltech) is not patient. He protested that physics could be jeopardized by not moving to high speed networks fast enough. Megabit per second capacity is needed now and Gigabit rates in ten years. Imagination should not be constrained by present technology.

This was underlined in presentations by W. Runge (Freiburg) and R. Ruehle (Stuttgart) of the work going on in Germany on high bandwidth networks. Ruehle concentrated on the Stuttgart system providing graphics workstation access through 140 Mbps links to local supercomputers. His talk was illustrated by slides and a video showing what can be done with a local workstation online to 2 Crays and a Convex simultaneously! He also showed the importance of graphics in conceptualization as well as testing theory against experiment, comparing a computer simulation of air flow over the sunroof of a Porsche with a film of actual performance.

Computer graphics for particle physics attracted a lot of interest. David Myers (CERN) discussed the conflicting requirements of software portability versus performance and summarized with 'Myers' Law of Graphics' – 'you can't have your performance and port it'. Rene Brun (CERN) gave a concise presentation of PAW (Physics Analysis Workstation). Demonstrations of PAW and other graphics packages such as event displays were available at both the Apollo and DEC exhibits during the conference week. Other exhibitors included Sun, Meiko, Caplin and IBM. IBM demonstrated the power of relational database technology using the online *Oxford English Dictionary.*

Interactive data analysis on workstations is well established and can be applied to all stages of program development and design. Richard Mount likened interactive graphics to the 'oscilloscope' of software development and analysis.

Establishing a good data structure is essential if flexible and easily maintainable code is to be written. Paulo Palazzi (CERN) showed how interactive graphics would enhance the already considerable power of the entity-relation model as realized in ADAMO. His presentation tied in very well with a fascinating account by David Nagel (Apple) of work going on at Cupertino to extend the well-known Macintosh human/computer interface, with tables accessed by the mouse, data highlighted and the corresponding graphical output shown in an adjacent window.

The importance of the interface between graphics and relational databases was also emphasized in the talk by Brian Read (RAL) on the special problems faced by scientists using database packages, illustrated by comparing the information content of atmospheric ozone concentration from satellite measurements in tabular and graphical form – 'one picture is worth a thousand numbers'.

The insatiable number-crunching appetite of both experimental and theoretical particle physicists has led to many new computer architectures being explored. Many of them exploit the powerful new (and relatively cheap) RISC chips on the market. Vector supercomputers are very appropriate for calculations like lattice gauge theories but it has yet to be demonstrated that they will have a big impact on 'standard' codes.

An indication of the improvements to be expected – perhaps a factor of five – came in the talk by Bill Martin (Michigan) on the vectorization of reactor physics codes. Perhaps more relevant to experimental particle physics are the processor 'farms' now well into the second generation.

Paul Mackenzie (Fermilab) and Bill McCall (Oxford) showed how important it is to match the architecture to the natural parallelism of a problem and how devices like the transputer enable this to be done. On a more speculative note Bruce Denby (Fermilab) showed the potential of neural networks for pattern recognition. They may also provide the ability to enable computers to learn. Such futuristic possibilities were surveyed in an evening session by Phil Treleavan (London). Research into this form of computing could reveal more about the brain as well as help with the new computing needs of future physics.

The importance of UNIX

With powerful new workstations appearing almost daily and with novel architecture in vogue, a machine-independent operating system is obviously attractive. The main contender is UNIX. Although nominally machine independent, UNIX comes in many implementations and its style is very much that of the early 70s with cryptic command acronyms – few concessions to user-friendliness! However it does have many powerful features and physicists will have to come to terms with it to exploit modern hardware.

Dietrich Wiegandt (CERN) summarized the development of UNIX and its important features – notably the directory tree structure and the 'shells' of command levels. An ensuing panel discussion chaired by Walter Hoogland (NIKHEF) included physicists using UNIX and representatives from DEC, IBM and Sun. Both DEC and IBM support the development of UNIX systems.

David McKenzie of IBM believed that the operating system should be transparent and looked forward to the day when operating systems are 'as boring as a mains wall plug'. (This was pointed out to be rather an unfortunate example since plugs are not standard!)

Two flavours of UNIX are available – Open Software Foundation version one (OSF1), and UNIX International release four (SVR4) developed at Berkeley. The two implementations overlap considerably and a standard version will emerge through user pressure. Panel member W. Van Leeuwen (NIKHEF) announced that a particle physics UNIX group had been set up (contact HEPNIX at CERNVM).

Software engineering became a heated talking point as the conference progressed. Two approaches were suggested: Carlo Mazza, head of the Data Processing Division of the European Space Operations Centre, argued for vigorous management in software design, requiring discipline and professionalism, while Richard Bornat ('SASD- All Bubbles and No Code') advocated an artistic approach, likening programming to architectural design rather than production engineering. A. Putzer (Heidelberg) replied that experimenters who have used software engineering tools such as SASD would use them again.

Software crisis

Paul Kunz (SLAC) gave a thoughtful critique of the so-called 'software crisis', arguing that code does not scale with the size of the detector or collaboration. Most detectors are modular and so is the code associated with them. With proper management and quality control good code can and will be written. The conclusion is that both inspiration and discipline go hand in hand.

A closely related issue is that of verifiable code – being able to prove that the code will do what is intended before any executable version is written. The subject has not yet had much impact on the physics community and was tackled by Tony Hoare and Bernard Sufrin (Oxford) at a pedagogical level. Sufrin, adopting a missionary approach, showed how useful a mathematical model of a simple text editor could be. Hoare demonstrated predicate calculus applied to the design of a tricky circuit.

Less high technology was apparent at the conference dinner in the candle-lit New College dining hall. Guest speaker Geoff Manning, former high energy physicist, one-time Director of the UK Rutherford Appleton Laboratory and now Chairman of Active Memory Technology, believed that physicists have a lot to learn from advances in computer science but that fruitful collaboration with the industry is possible. Replying for the physicists, Tom Nash (Fermilab) thanked the industry for their continuing interest and support through joint projects and looked forward to a collaboration close enough for industry members to work physics shifts and for physicists to share in profits!

Summarizing the meeting, Rudy Bock (CERN) highlighted novel architectures as the major area where significant advances have been made and will continue to be made for the next few years. Standards are also important provided they are used intelligently and not just as a straitjacket to impede progress. His dreams for the future included neural networks and the wider use of formal methods in software design. Some familiar topics of the past, including code and memory management and the choice of programming language, could be 'put to rest'. ●

The insatiable number-crunching appetite of both experimental and theoretical particle physicists has led to many new computer architectures being explored

A MAJOR SHIFT IN OUTLOOK

In the summer of 2001, computer specialist Ben Segal described how distributed UNIX boxes took over from CERN's all-powerful IBM and Cray mainframe workhorses.

CERN

The changing landscape at CERN's Computer Centre: 1988, with the Cray supercomputer (in blue and yellow) in the background.

CERN

Ten years on: CERN's Computer Centre in 1998, with banks of SHIFT-distributed computing arrays for major experiments.

I don't remember exactly who first proposed running physics batch jobs on a UNIX workstation, rather than on the big IBM or Cray mainframes that were doing that kind of production work in 1989 at CERN. The workstation in question was to be an Apollo DN10000, the hottest thing in town with reduced instruction set (RISC) CPUs of a formidable five CERN Units (a CERN Unit was defined as one IBM 370/168, equivalent to four VAX 11-780s) each and costing around SwFr 150 000 for a 4-CPU box.

It must have been the combined idea of Les Robertson, Eric McIntosh, Frederic Hemmer, Jean-Philippe Baud, myself and perhaps some others who were working at that time around the biggest UNIX machine that had ever crossed the threshold of the Computer Centre – a Cray XMP-48, running UNICOS.

At any rate, when we spoke to the Apollo salespeople about our idea, they liked it so much that they lent us the biggest box they had, a DN10040 with four CPUs plus a staggering 64 Mb of memory and 4 Gb of disk space. Then, to round it off, they offered to hire a person of our choice for three years to work on the project at CERN.

In January 1990 the machine was installed and our new "hireling", Erik Jagel, an Apollo expert after his time managing the Apollo farm for the L3 experiment, coined the name "HOPE" for the new project. (Hewlett-Packard had bought Apollo and OPAL had expressed interest, so it was to be the "HP OPAL Physics Environment").

We asked where we could find the space to install HOPE in the Computer Centre. We just needed a table with the DN10040 underneath and an Ethernet connection to the Cray, to give us access to the tape data. The reply was: "Oh, there's room in the middle" – where the recently obsolete round tape units had been – so that was where HOPE went, looking quite lost in the huge computer room, with the IBM complex on one side and the Cray supercomputer on the other.

Soon the HOPE cycles were starting to flow. The machine was surprisingly reliable, and porting the big physics FORTRAN programs was easier than we had expected. After around six months, the system was generating 25 per cent of all CPU cycles in the centre. Management began to notice the results when we included HOPE's accounting files in the weekly report we made that plotted such things in easy-to-read histograms.

We were encouraged by this success and went to work on a proposal to extend HOPE. The idea was to build a scalable version from interchangeable components: CPU servers, disk servers and tape servers, all connected by a fast network and software to create a distributed mainframe. "Commodity" became the keyword – we would use the cheapest building-blocks available from the manufacturers that gave the best price performance for each function.

On how large a scale could we build such a system and what would it cost? We asked around, and received help from some colleagues who treated it as a design study. A simulation was done of the workflow through such a system, bandwidth requirements were estimated for the fast network "backplane" that was needed to connect everything, prices were calculated, essential software was sketched out and the manpower required for development

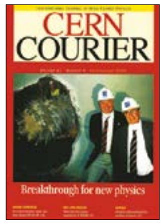

This article was adapted from text in *CERN Courier* vol. 41, July/August 2001, pp19–20

PC power: CERN's Computer Centre in 2001 has wall-to-wall PCs.

and operation was predicted.

Software development would be a challenge. Fortunately, some of us had been working with Cray at CERN, adding some facilities to UNIX that were vital for mainframe computing: a proper batch scheduler and a tape-drive reservation system, for example. These could be reused quite easily.

Other new functions would include a distributed "stager" and a "disk-pool manager". These would allow the pre-assembly of each job's tape data (read from drives on tape servers) into efficiently-managed disk pools that would be located on disk servers, ready to be accessed by the jobs in the CPU servers. Also new would be the "RFIO", a remote file input-output package that would offer a unified and optimized data-transfer service between all of the servers via the backplane. It looked like Sun's networking filing system, but was much more efficient.

SHIFT in focus

Finally, a suitable name was coined, again by Erik Jagel: "SHIFT", for "Scalable Heterogeneous Integrated FaciliTy", suggesting the paradigm shift that was taking place in large-scale computing: away from mainframes and towards a distributed low-cost approach.

The "SHIFT" proposal report was finished in July 1990. It had 10 names on it, including the colleagues from several groups that had offered their ideas and worked on the document.

"Were 10 people working on this?" and "How many Cray resources were being used and/or counted?" came the stern reply. In response, we pointed out that most of the 10 people had contributed small fractions of their time, and that the Cray had been used simply as a convenient tape server. It was the only UNIX machine in the Computer Centre with access to the standard tape drives, all of which were physically connected to the IBM mainframe at that time.

Closer to home, the idea fell on more fertile ground, and we were told that if we could persuade at least one of the four LEP experiments to invest in our idea, we could have matching support from the Division. The search began. We spoke to ALEPH, but they replied, "No, thank you, we're quite happy with our all-VAX VMS approach." L3 replied, "No thanks, we have all the computing power we need." DELPHI replied, "Sorry, we've no time to look at this as we're trying to get our basic system running."

Only OPAL took a serious look. They had already been our partner in HOPE and also had a new collaborator from Indiana with some cash to invest and some small computer system interface (SCSI) disks for a planned storage enhancement to their existing VMS-based system. They would give us these contributions until March 1991, the next LEP start-up – on the condition that everything was working by then, or we'd have to return their money and disks. It was September 1990, and there was a lot of work to do.

Our modular approach and use of the UNIX, C language, TCP/IP and SCSI standards were the keys to the very short timescale we achieved. The design studies had included technical evaluations of various workstation and networking products.

By September, code development could begin and orders for hardware went out. The first tests on site with SGI Power Series servers connected via UltraNet took place at the end of December 1990. A full production environment was in place by March 1991, the date set by OPAL.

And then we hit a problem. The disk server system began crashing repeatedly with unexplained errors. Our design evaluations had led us to choose a "high-tech" approach: the use of symmetric multiprocessor machines from Silicon Graphics for both CPU and disk servers, connected by the sophisticated "UltraNet" Gigabit network backplane. One supporting argument had been that if the UltraNet failed or could not be made to work in time, then we could put all the CPUs and disks together in one cabinet and ride out the OPAL storm. We hadn't expected any problems in the more conventional area of the SCSI disk system.

Our disks were mounted in trays inside the disk server, connected via high-performance SCSI channels. It looked standard, but we had the latest models of everything. Like a performance car, it was a marvel of precision but impossible to keep in tune. We tried everything, but still it went on crashing and we finally had to ask SGI to send an engineer. He found the problem: inside our disk trays was an extra metre of flat cable which had not been taken into account in our system configuration. We had exceeded the strict limit of 6 m for single-ended SCSI, and in fact it was our own fault. Rather than charging us penalties and putting the blame where it belonged, SGI lent us two extra CPUs to help us to make up the lost computing time for OPAL and ensure the success of the test period!

At the end of November 1991, a satisfied OPAL doubled its investment in CPU and disk capacity for SHIFT. At the same time, 16 of the latest HP 9000/720 machines, each worth 10 CERN Units of CPU, arrived to form the first Central Simulation Facility or "Snake Farm".

The stage was set for the exit of the big tidy mainframes at CERN, and the beginning of the much less elegant but evolving scene we see today on the floor of the CERN Computer Centre. SHIFT became the basis of LEP-era computing and its successor systems are set to perform even more demanding tasks for the LHC, scaled this time to the size of a worldwide grid. ●

Like a performance car, it was a marvel of precision but impossible to keep in tune

THE CERN OPENLAB: A NOVEL TESTBED FOR THE GRID

Grid computing was the computer buzzword of the decade, wrote CERN's François Grey in this 2003 feature about a new model for partnership between CERN and industry.

Sverre Jarp, chief technology officer of the CERN openlab, with equipment from Enterasys Networks.

Grid computing is the computer buzzword of the decade. Not since the World Wide Web was developed at CERN more than 10 years ago has a new networking technology held so much promise for both science and society. The philosophy of the Grid is to provide vast amounts of computer power at the click of a mouse, by linking geographically distributed computers and developing "middleware" to run the computers as though they were an integrated resource. Whereas the Web gives access to distributed information, the Grid does the same for distributed processing power and storage capacity.

There are many varieties of Grid technology. In the commercial arena, Grids that harness the combined power of many workstations within a single organization are already common. But CERN's objective is altogether more ambitious: to store petabytes of data from the Large Hadron Collider (LHC) experiments in a distributed fashion and make the data easily accessible to thousands of scientists around the world. This requires much more than just spare PC capacity – a network of major computer centres around the world must provide their resources in a seamless way.

CERN and a range of academic partners have launched several major projects in order to achieve this objective. In the European arena, CERN is leading the European DataGrid (EDG) project, which addresses the needs of several scientific communities, including high-energy particle physics. The EDG has already developed the middleware necessary to run a Grid testbed involving more than 20 sites. CERN is also leading a follow-on project funded by the European Union, EGEE (Enabling Grids for E-Science in Europe), which aims to provide a reliable Grid service to European science. Last year, the LHC Computing Grid (LCG) project was launched by CERN and partners to deploy a global Grid dedicated to LHC needs, drawing on the experience of the EDG and other international efforts. This project has started running a global Grid, called LCG-1.

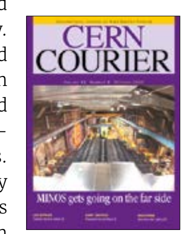

This article was adapted from text in *CERN Courier* vol. 43, October 2003, pp31–34

Enter the openlab

The CERN openlab for DataGrid applications fits into CERN's portfolio of Grid activities by addressing a key issue, namely the impact on the LCG of cutting-edge IT technologies that are currently emerging from industry. Peering into the technological crystal ball in this way can only be done in close collaboration with leading industrial partners. The benefits are mutual: through generous sponsorship of state-of-the-art equipment from the partners, CERN acquires early access to valuable technology that is still several years from the commodity computing market the LCG will be based on.

In return, CERN provides demanding data challenges, which push these new technologies to their limits – this is the "lab" part of the openlab. CERN also provides a neutral environment for integrating solutions from different partners, to test their interoperability. This is a vital role in an age where open standards (the "open" part of openlab) are increasingly guiding the development of the IT industry.

The CERN openlab for DataGrid applications was launched in 2001 by Manuel Delfino, then the IT Division leader at CERN. After a hiatus, during which the IT industry was rocked by the telecoms crash, the partnership took off in September 2002, when HP joined founding members Intel and Enterasys Networks, and integration of technologies from all three led to the CERN opencluster project.

At present, the CERN opencluster consists of 32 Linux-based HP rack-mounted servers, each equipped with two 1 GHz Intel Itanium 2 processors. Itanium uses 64-bit processor technology, which is anticipated to displace today's

IBM joins CERN openlab to tackle the petabyte challenge

The LHC will generate more than 10 petabytes of data per year, the equivalent of a stack of CD ROMs 20 km high. There is no obvious way to extend conventional data-storage technology to this scale, so new solutions must be considered. IBM was therefore keen to join the CERN openlab in April 2003, in order to establish a research collaboration aimed at creating a massive data-management system built on Grid computing, which will use innovative storage virtualization and file-management technology.

Rainer Többicke of CERN with the IBM-sponsored 28 terabyte storage system.

IBM has been a strong supporter of Grid computing, from its sponsorship of the first Global Grid Forum in Amsterdam in 2001 to its participation in the European DataGrid project. The company sees Grid computing as an important technological realization of the vision of "computing on demand", and expects that as Grid computing moves from exclusive use in the scientific and technical world into commercial applications, it will indeed be the foundation for the first wave of e-business on demand.

The technology that IBM brings to the CERN openlab partnership is called Storage Tank. Conceived in IBM Research, the new technology is designed to provide scalable, high-performance and highly available management of huge amounts of data using a single file namespace, regardless of where or on what operating system the data reside. (Recently, IBM announced that the commercial version will be named IBM TotalStorage SAN File System.) IBM and CERN will work together to extend Storage Tank's capabilities so it can manage the LHC data and provide access to it from any location worldwide.

Brian E Carpenter, IBM Systems Group, and Jai Menon, IBM Research.

32-bit technology over the next few years. As part of the agreement with the CERN openlab partners, this cluster is planned to double in size during 2003, and double again in 2004, making it an extremely high-performance computing engine. In April this year, IBM joined the CERN openlab, contributing advanced storage technology that will be combined with the CERN opencluster (see box above).

For high-speed data transfer challenges, Intel has delivered 10 Gbps Ethernet Network Interface Cards (NICs), which have been installed on the HP computers, and Enterasys Networks has delivered three switches equipped to operate at 10 gigabits per second (Gbps), with additional port capacity for 1 Gbps.

Over the next few months, the CERN opencluster will be linked to the EDG testbed to see how these new technologies perform in a Grid environment. The results will be closely monitored by the LCG project to determine the potential impact of the technologies involved. Already at this stage, however, much has been learned that has implications for the LCG.

For example, thanks to the preinstalled management cards in each node of the cluster, automation has been developed to allow remote system restart and remote power control. This development confirmed the notion that – for a modest hardware investment – large clusters can be controlled with no operator present. This is highly relevant to the LCG, which will need to deploy such automation on a large scale.

Several major physics software packages have been successfully ported and tested on the 64-bit environment of the CERN opencluster, in collaboration with the groups responsible for maintaining the various packages. Benchmarking of the physics packages has begun and the first results are promising. For example, PROOF (Parallel ROOT Facility) is a version of the popular CERN-developed software ROOT for data analysis, which is being developed for interactive analysis of very large ROOT data files on a cluster of computers. The CERN opencluster has shown that the amount of data that can be handled by PROOF scales linearly with cluster size – on one cluster node it takes 325 s to analyse a certain amount of data, and only 12 s when all 32 nodes are used.

One of the major challenges of the CERN opencluster project is to take maximum advantage of the partners' 10 Gbps technology. In April, a first series of tests was conducted between two of the nodes in the cluster, which were directly connected (via a "back-to-back" connection) through 10 Gbps Ethernet NICs. The transfer reached a data rate of 755 megabytes per second (MB/s), a record, and double the maximum rate obtained with 32-bit processors. The transfer took place over a 10 km fibre and used very big frames (16 kB) in a single stream, as well as the regular suite of Linux Kernel protocols (TCP/IP).

The best results through the Enterasys switches were obtained when aggregating the 1 Gbps bi-directional traffic involving 10 nodes in each group. The peak traffic between the switches was then measured to be 8.2 Gbps. The next stages of this data challenge will include evaluating the next version of the Intel processors.

Data challenge

In May, CERN announced the successful completion of a major data challenge aimed at pushing the limits of data storage to tape. This involved, in a critical way, several components of the CERN opencluster. Using 45 newly installed StorageTek tape drives, capable of writing to tape at 30 MB/s, storage-to-tape rates of 1.1 GB/s were achieved for periods of several hours, with peaks of 1.2 GB/s – roughly equivalent to storing a whole movie on DVD every four seconds. The average sustained over a three-day period was of 920 MB/s. Previous best results by other research labs were typically less than 850 MB/s.

The significance of this result, and the purpose of the data challenge, was to show that CERN's IT Division is on track to cope with the enormous data rates expected from the LHC. One experiment alone, ALICE, is expected to produce data at rates of 1.25 GB/s.

In order to simulate the LHC data acquisition procedure, an equivalent stream of artificial data was generated using 40 computer servers. These data were stored temporarily to 60 disk servers, which included the CERN opencluster servers, before being transferred to the tape servers. A key contributing factor to the success of the data challenge was a high-performance switched network from Enterasys Networks with 10 Gbps Ethernet capability, which routed the data from PC to disk and from disk to tape.

An open dialogue

While many of the benefits of the CERN openlab for the industrial partners stem from such data challenges, there is also a strong emphasis in openlab's mission on the opportunities that this novel partnership provides for enhanced communication and cross-fertilization between CERN and the partners, and between the partners themselves. Top engineers from the partner companies collaborate closely with the CERN openlab team in CERN's IT Division, so that the inevitable technical challenges that arise when dealing with new technologies are dealt with rapidly and efficiently. Furthermore, as part of their sponsorship, HP is funding two CERN fellows to work on the CERN opencluster. The CERN openlab team also organizes thematic workshops on specific topics of interest, bringing together leading technical experts from the partner companies, as well as public "First Tuesday" events on general technology issues related to the openlab agenda, which attract hundreds of participants from the industrial and investor communities.

A CERN openlab student programme has also been created, bringing together teams of students from different European universities to work on applications of Grid technology. And the CERN openlab is actively supporting the establishment of a Grid café for the CERN Microcosm exhibition – a Web café for the general public with a focus on Grid technologies, including a dedicated website that will link to instructive Grid demos.

Efforts are ongoing in the CERN openlab to evaluate other possible areas of technological collaboration with current or future partners. The concept is certainly proving popular, with other major IT companies expressing an interest in joining. This could occur by using complementary technologies to provide added functionality and performance to the existing opencluster. Or it could involve launching new projects that deal with other aspects of Grid technology relevant to the LCG, such as Grid security and mobile access to the Grid.

In conclusion, the CERN openlab puts a new twist on an activity – collaboration with leading IT companies – that has been going on at CERN for decades. Whereas traditionally such collaboration was bilateral and focused on "here-and-now" solutions, the CERN openlab brings a multilateral long-term perspective into play. This may be a useful prototype for future industrial partnerships in other high-tech areas, where CERN and a range of partners can spread their risks and increase their potential for success by working on long-term development projects together. ●

The purpose of the data challenge was to show that CERN's IT Division is on track to cope with the enormous data rates expected from the LHC

CERN'S ULTIMATE ACT OF OPENNESS

The seed that led CERN to relinquish ownership of the web in 1993 was planted when the Organization formally came into being, wrote Maarten Wilbers and Jonathan Drakeford of CERN legal service on the web's 30th anniversary.

At a mere 30 years old, the World Wide Web already ranks as one of humankind's most disruptive inventions. Developed at CERN in the early 1990s, it has touched practically every facet of life, impacting industry, penetrating our personal lives and transforming the way we transact. At the same time, the web is shrinking continents and erasing borders, bringing with it an array of benefits and challenges as humanity adjusts to this new technology.

This reality is apparent to all. What is less well known, but deserves recognition, is the legal dimension of the web's history. On 30 April 1993, CERN released a memo (part of which is pictured right) that placed into the public domain all of the web's underlying software: the basic client, basic server and library of common code. The document was addressed "To whom it may concern" – which would suggest the authors were not entirely sure who the target audience was. Yet, with hindsight, this line can equally be interpreted as an unintended address to humanity at large.

The legal implication was that CERN relinquished all intellectual property rights in this software. It was a deliberate decision, the intention being that a no-strings-attached release of the software would "further compatibility, common practices, and standards in networking and computer supported collaboration" – arguably modest ambitions for what turned out to be such a seismic technological step. To understand what seeded this development you need to go back to the 1950s, at a time when "software" would have been better understood as referring to clothing rather than computing.

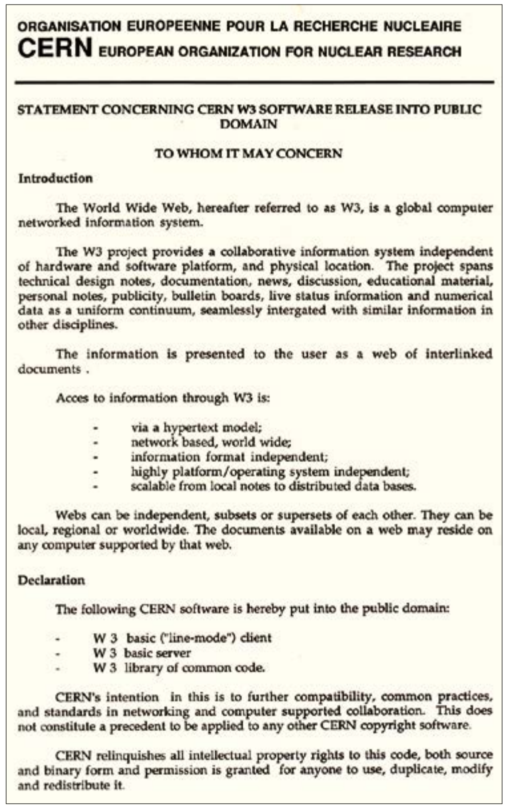

ORGANISATION EUROPEENNE POUR LA RECHERCHE NUCLEAIRE
CERN EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

STATEMENT CONCERNING CERN W3 SOFTWARE RELEASE INTO PUBLIC DOMAIN

TO WHOM IT MAY CONCERN

Introduction

The World Wide Web, hereafter referred to as W3, is a global computer networked information system.

The W3 project provides a collaborative information system independent of hardware and software platform, and physical location. The project spans technical design notes, documentation, news, discussion, educational material, personal notes, publicity, bulletin boards, live status information and numerical data as a uniform continuum, seamlessly intergated with similar information in other disciplines.

The information is presented to the user as a web of interlinked documents .

Acces to information through W3 is:

- via a hypertext model;
- network based, world wide;
- information format independent;
- highly platform/operating system independent;
- scalable from local notes to distributed data bases.

Webs can be independent, subsets or supersets of each other. They can be local, regional or worldwide. The documents available on a web may reside on any computer supported by that web.

Declaration

The following CERN software is hereby put into the public domain:

- W 3 basic ("line-mode") client
- W 3 basic server
- W 3 library of common code.

CERN's intention in this is to further compatibility, common practices, and standards in networking and computer supported collaboration. This does not constitute a precedent to be applied to any other CERN copyright software.

CERN relinquishes all intellectual property rights to this code, both source and binary form and permission is granted for anyone to use, duplicate, modify and redistribute it.

European project

CERN was born out of the wreckage of World War II, playing a role, on the one hand, as a mechanism for reconciliation between former belligerents, while, on the other, offering European nuclear physicists the opportunity to conduct their research locally. The hope was that this would stem the "brain drain" to the US, from a Europe still recovering from the devastating effects of war.

In 1953, CERN's future Member States agreed on the text of the organisation's founding Convention, defining its mission as providing "for collaboration among European States in nuclear research of a pure scientific and fundamental character". With the public acutely aware of the role that destructive nuclear technology had played during the war, the Convention additionally stipulated that CERN was to have "no concern with work for military requirements" and that the results of its work, were to be "published or otherwise made generally available".

In the early years of CERN's existence, the openness resulting from this requirement for transparency was essentially delivered through traditional channels, in particular through publication in scientific journals. Over time, this became the cultural norm at CERN, permeating all aspects of its work both internally and with its collaborating partners and society at large. CERN's release of the WWW software into the public domain, arguably in itself a consequence of the openness requirement of the Convention, could be seen as a precursor to today's web-based tools that represent further manifestations of CERN's openness: the SCOAP3 publishing model, open-source software and hardware, and open data.

Perhaps the best measure for how ingrained openness is in CERN's ethos as a laboratory is to ask the question: "if CERN would have known then what it knows now about the impact of the World Wide Web, would it still have made the web software available, just as it did in 1993?" We would like to suggest that, yes, our culture of openness would provoke the same response now as it did then, though no doubt a modern, open-source licensing regime would be applied.

This, in turn, can be viewed as testament and credit to the wisdom of CERN's founders, and to the CERN Convention, which remains the cornerstone of our work to this day. ●

This article was adapted from text in *CERN Courier* vol. 59, March/April 2019, pp39–40