HSF Analysis Facility Forum

Alessandra Forti, Manchester
Nicole Skidmore, Manchester
Diego Ciangottini, INFN
Lukas Heinrich, TUM

Challenges

We are all very aware of the upcoming data & computing challenges

Not only at the LHC, but more widely in HEP and fundamental physics

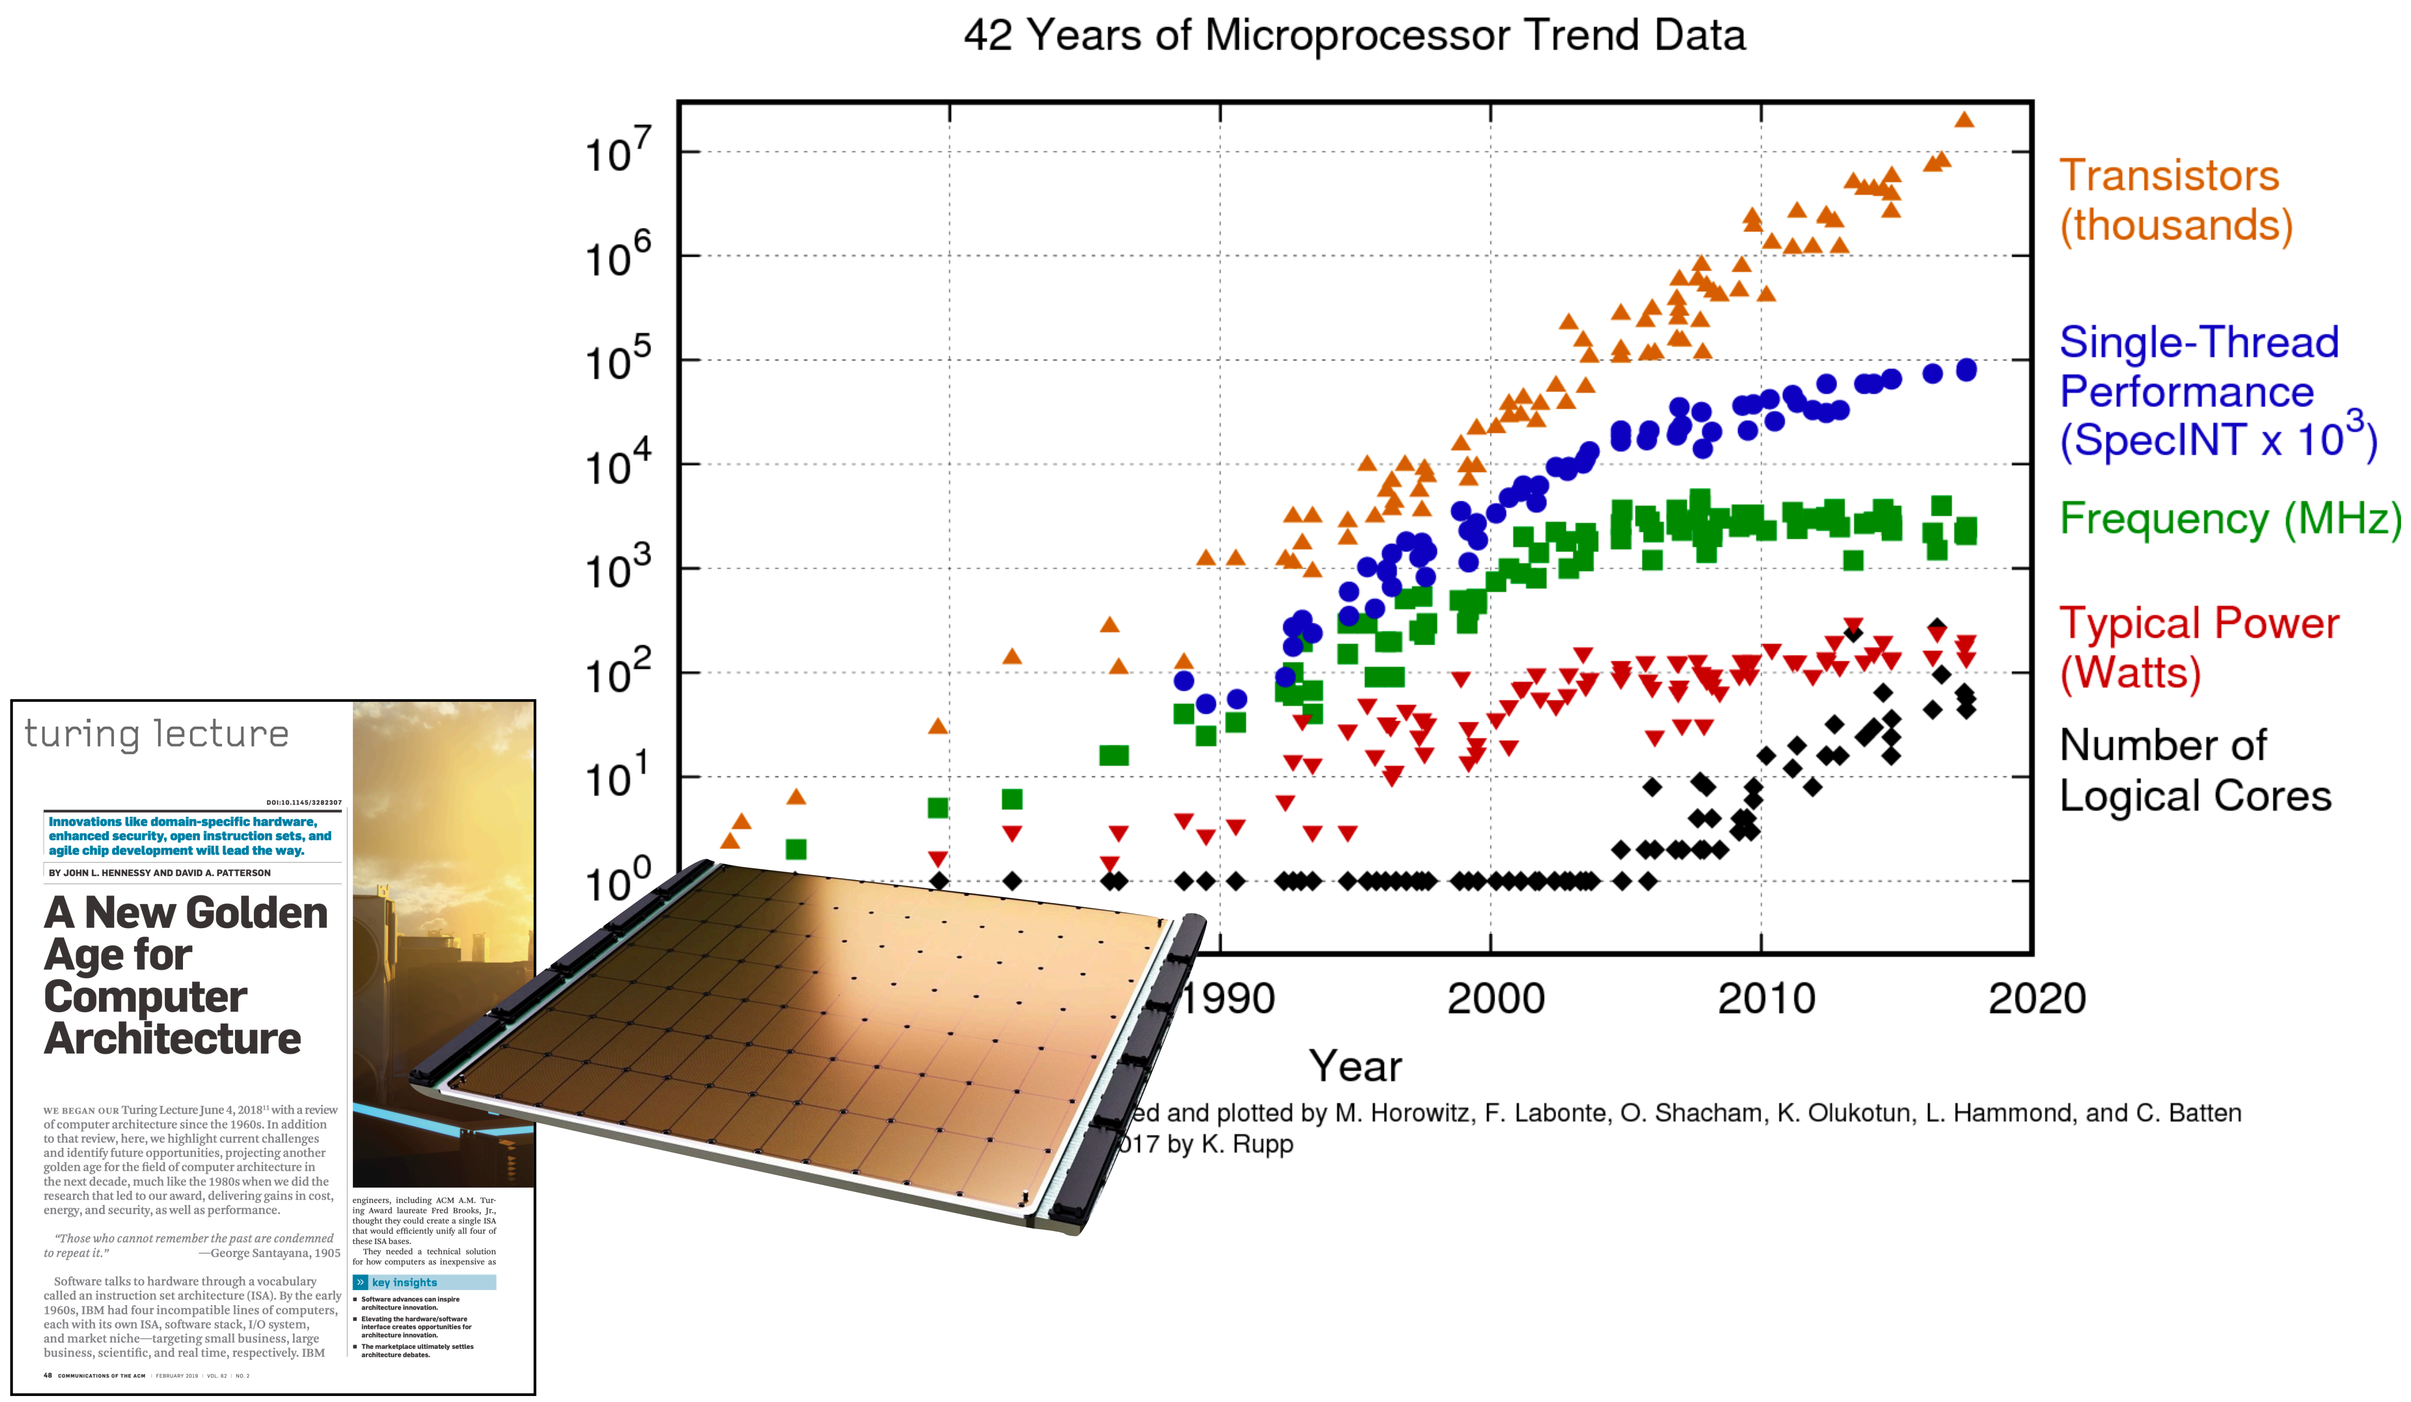

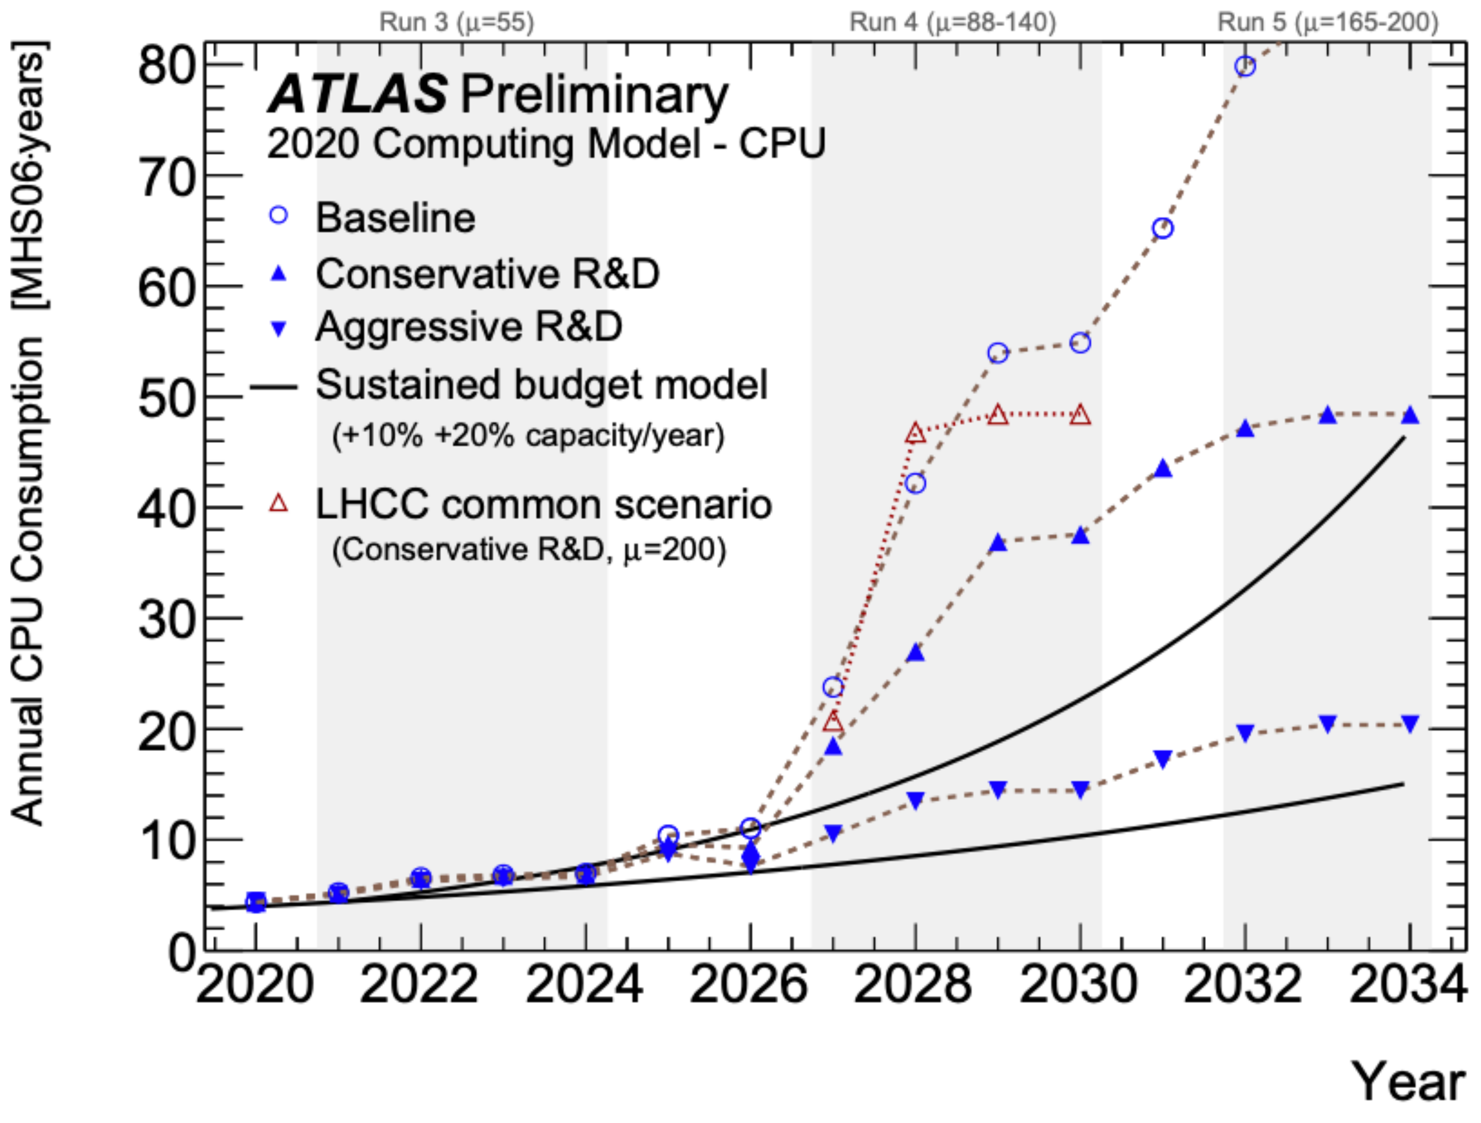

Not only a resource question, but a UX question as well

Analysis Facilities

The facilities folks use to do their physics determine the UX. AFs must play a major role in making physics analysis enjoyable in the future.

- can I implement new ideas quickly or even interactively?
- can I reliably get new results without lots of babysitting?
- can I get my data and software to where I need it?
- can I share results / grant access with my team worldwide?
- can I run on the hardware I need (e.g. train ML models)
- can I preserve my analysis / port it to a new facility ?

Analysis Facilities

We have successful facilities: Grid, national facilities, Tier-3s, LXPLUS, … the primary mode over the last decades was

- Interactive Login Nodes + Local Batch Processing (data → code)
- Grid Computing (code → data)
- most on x86 CPU with SLC X Operating System

Under the current interest in “Analysis Facilities” there is a implicit assumption that things will evolve beyond these basic usage patterns

- various ideas on what exactly this looks like, a lot is in flux

A Sketch?

A future facility may need to provide data & infrastructure services targeted at multiple use-cases at once

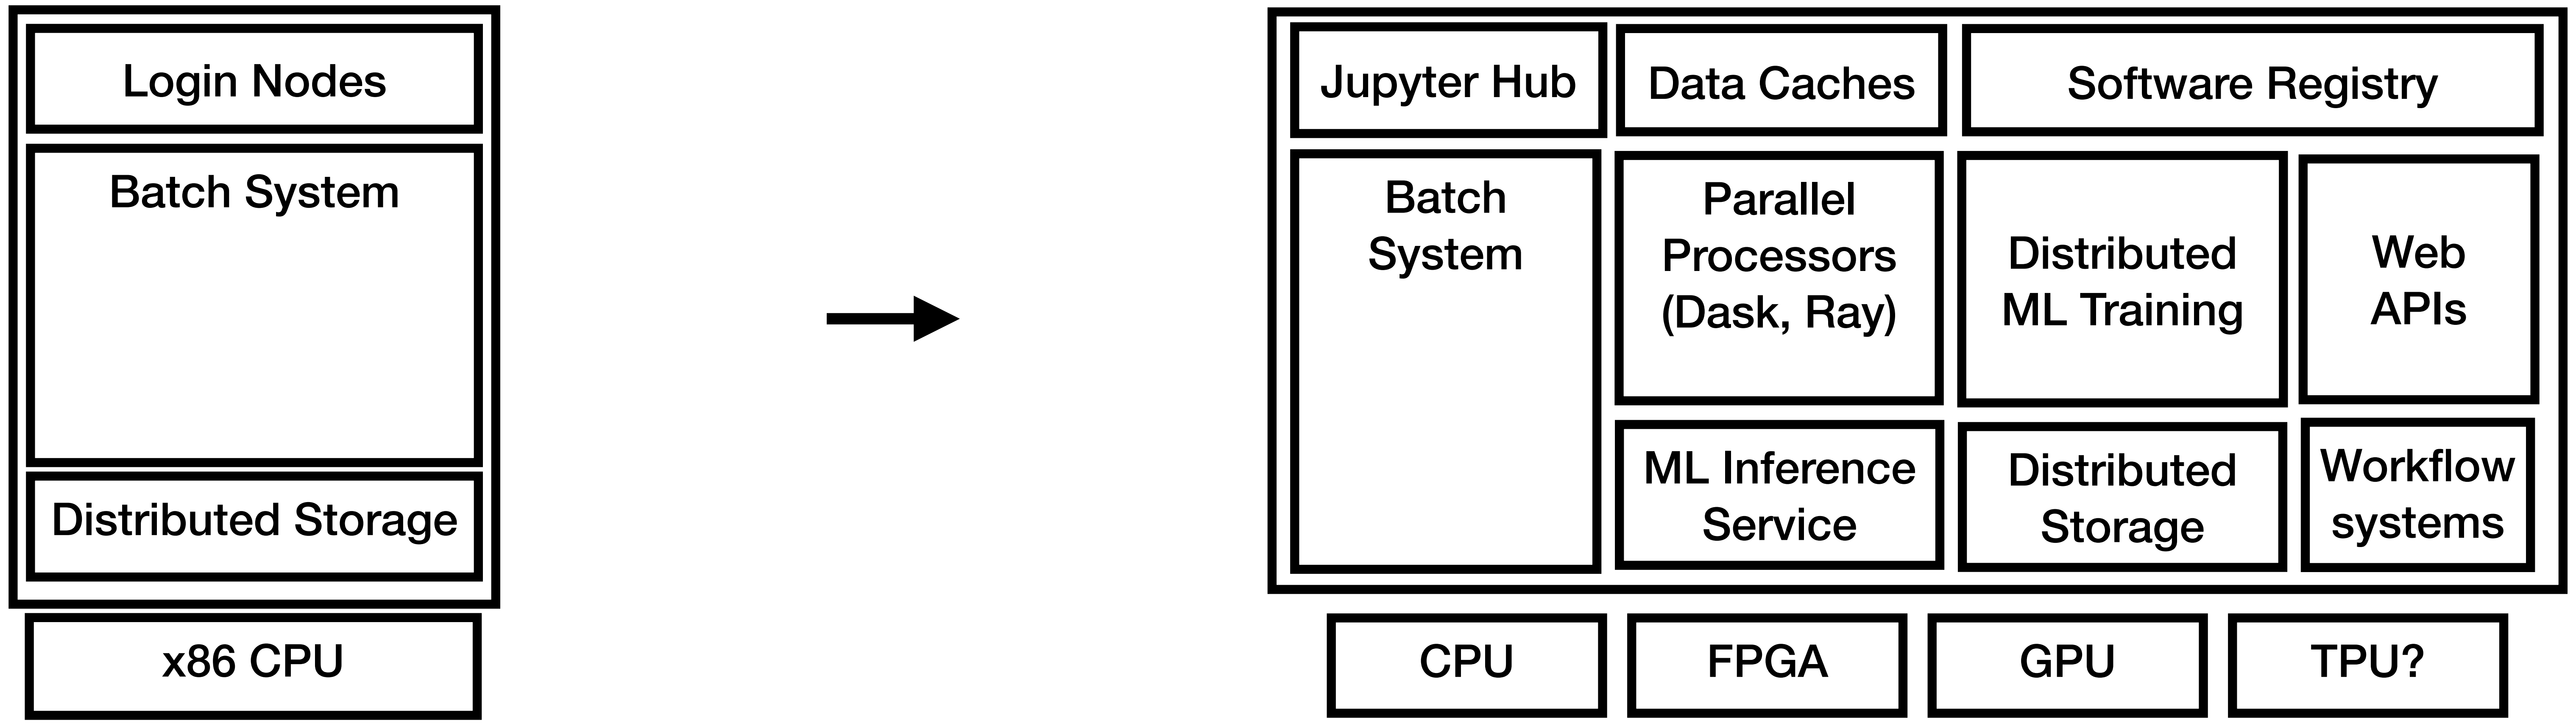

Key Questions:

- what components are really important / needed
- can we develop **shared infrastructure components** for the wider community like we did for the Grid?

Lots of interesting Work already ongoing

This Forum

Want to provide a space for users and facility developers to figure out how to chart a path forward and share concrete implementations and experiences.

A place to share what's going on inside of the experiments / should be tightly coordinated with internal developments

Eventual Goal: Shared language / framework through which to built and interface future facilities ?

Possible Topics (just suggestions)

- **Deploying Kubernetes in academic settings:** k8s @ HPC, rootless kubernetes, k8s + batch,
- **Data Access:** Object Store APIs, Global Homedir/Drives, …
- **Software Distribution:** lazy-loaded container images, community registries, user containers, …
- **Authz/Authn, Federation:** token integration, multi-cluster APIs
- **Cloud Computing Tooling:** monitoring, cost control, …
- **Interactive Scale-out Systems:** multi-tenant dataframe processing
- …

Meetings & Coordinators

Bi-Weekly Meetings: Thursdays @ 18:00 CET

Mailing List: hsf-af-forum@googlegroups.com

Mattermost: https://mattermost.web.cern.ch/signup_user_complete/?id=ffib9ny91t8qbcdpeernf3sidy

Coordinators: hsf-af-forum-convenors@googlegroups.com

Website: https://hepsoftwarefoundation.org/activities/analysisfacilitiesforum.html

History

HEP has a long history in close collaboration, our AF work for HL-LHC should be another piece of that legacy

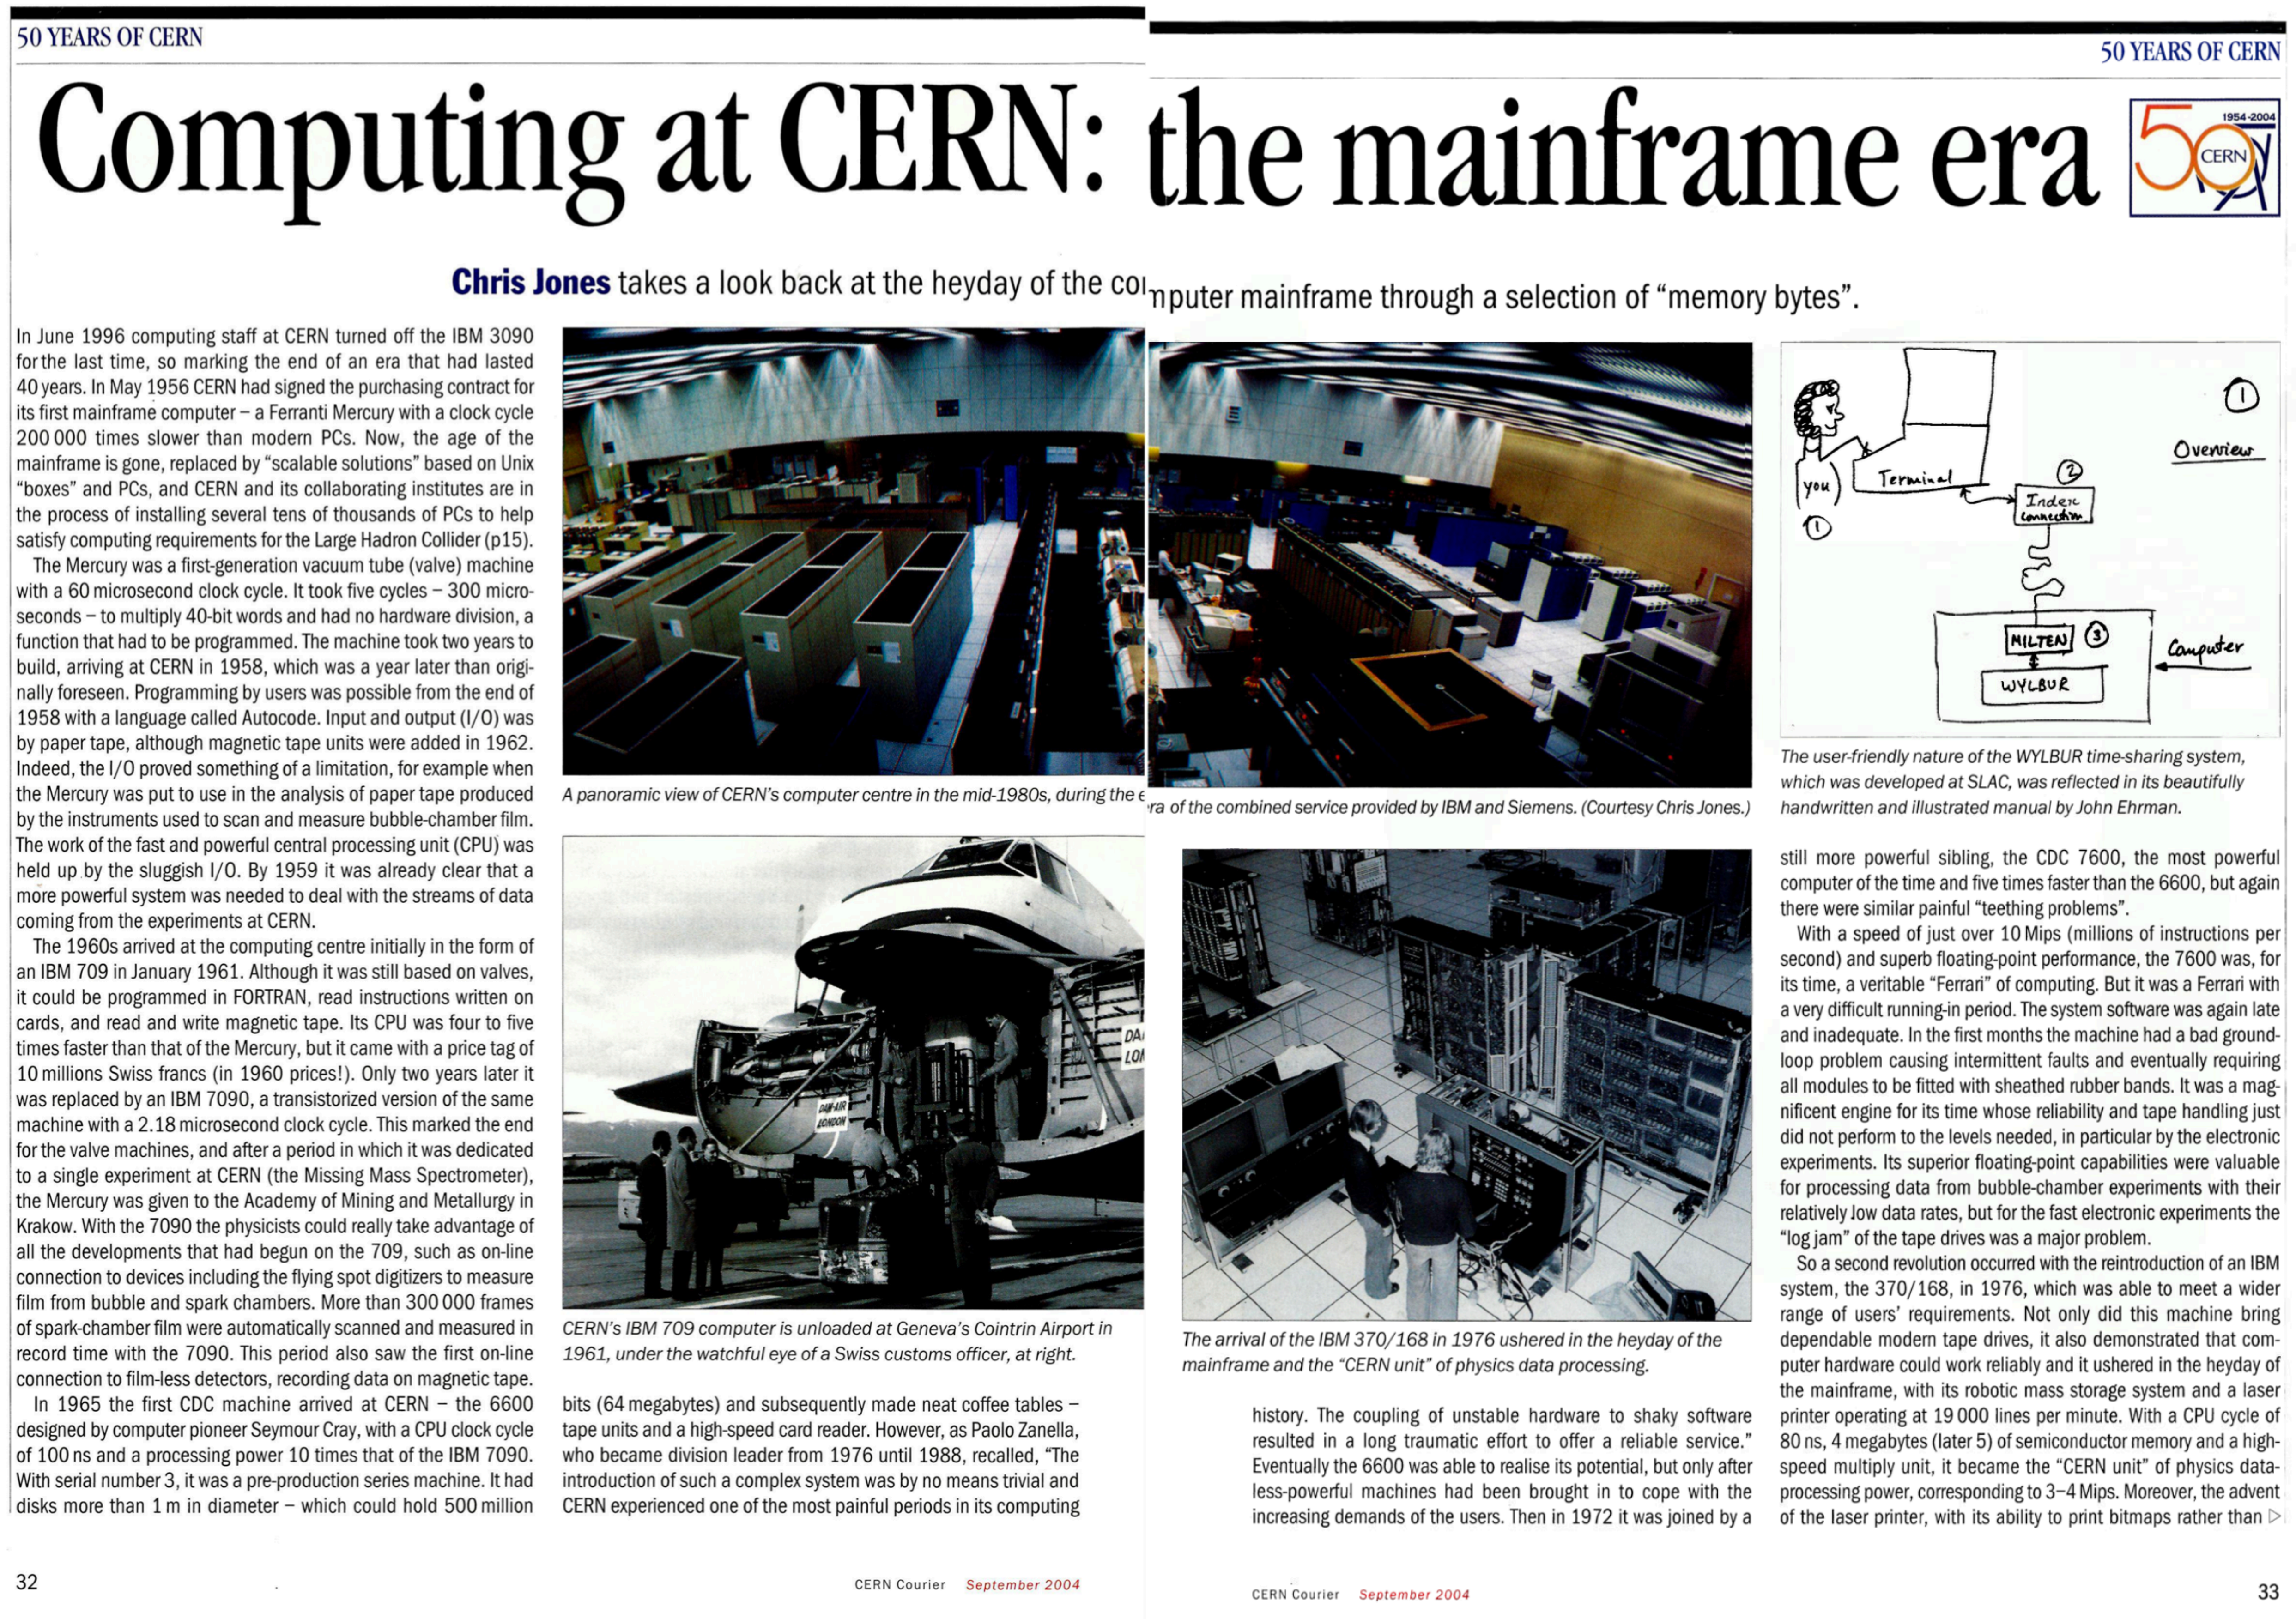

50 YEARS OF CERN

Computing at CERN: the mainframe era

Chris Jones takes a look back at the heyday of the computer mainframe through a selection of "memory bytes".

In June 1996 computing staff at CERN turned off the IBM 3090 for the last time, so marking the end of an era that had lasted 40 years. In May 1956 CERN had signed the purchasing contract for its first mainframe computer – a Ferranti Mercury with a clock cycle 200 000 times slower than modern PCs. Now, the age of the mainframe is gone, replaced by "scalable solutions" based on Unix "boxes" and PCs, and CERN and its collaborating institutes are in the process of installing several tens of thousands of PCs to help satisfy computing requirements for the Large Hadron Collider (p15).

The Mercury was a first-generation vacuum tube (valve) machine with a 60 microsecond clock cycle. It took five cycles – 300 microseconds – to multiply 40-bit words and had no hardware division, a function that had to be programmed. The machine took two years to build, arriving at CERN in 1958, which was a year later than originally foreseen. Programming by users was possible from the end of 1958 with a language called Autocode. Input and output (I/O) was by paper tape, although magnetic tape units were added in 1962. Indeed, the I/O proved something of a limitation, for example when the Mercury was put to use in the analysis of paper tape produced by the instruments used to scan and measure bubble-chamber film. The work of the fast and powerful central processing unit (CPU) was held up by the sluggish I/O. By 1959 it was already clear that a more powerful system was needed to deal with the streams of data coming from the experiments at CERN.

The 1960s arrived at the computing centre initially in the form of an IBM 709 in January 1961. Although it was still based on valves, it could be programmed in FORTRAN, read instructions written on cards, and read and write magnetic tape. Its CPU was four to five times faster than that of the Mercury, but it came with a price tag of 10 millions Swiss francs (in 1960 prices!). Only two years later it was replaced by an IBM 7090, a transistorized version of the same machine with a 2.18 microsecond clock cycle. This marked the end for the valve machines, and after a period in which it was dedicated to a single experiment at CERN (the Missing Mass Spectrometer), the Mercury was given to the Academy of Mining and Metallurgy in Krakow. With the 7090 the physicists could really take advantage of all the developments that had begun on the 709, such as on-line connection to devices including the flying spot digitizers to measure film from bubble and spark chambers. More than 300 000 frames of spark-chamber film were automatically scanned and measured in record time with the 7090. This period also saw the first on-line connection to film-less detectors, recording data on magnetic tape.

In 1965 the first CDC machine arrived at CERN – the 6600 designed by computer pioneer Seymour Cray, with a CPU clock cycle of 100 ns and a processing power 10 times that of the IBM 7090. With serial number 3, it was a pre-production series machine. It had disks more than 1 m in diameter – which could hold 500 million bits (64 megabytes) and subsequently made neat coffee tables – tape units and a high-speed card reader. However, as Paolo Zanella, who became division leader from 1976 until 1988, recalled, "The introduction of such a complex system was by no means trivial and CERN experienced one of the most painful periods in its computing history. The coupling of unstable hardware to shaky software resulted in a long traumatic effort to offer a reliable service." Eventually the 6600 was able to realise its potential, but only after less-powerful machines had been brought in to cope with the increasing demands of the users. Then in 1972 it was joined by a still more powerful sibling, the CDC 7600, the most powerful computer of the time and five times faster than the 6600, but again there were similar painful "teething problems".

With a speed of just over 10 Mips (millions of instructions per second) and superb floating-point performance, the 7600 was, for its time, a veritable "Ferrari" of computing. But it was a Ferrari with a very difficult running-in period. The system software was again late and inadequate. In the first months the machine had a bad ground-loop problem causing intermittent faults and eventually requiring all modules to be fitted with sheathed rubber bands. It was a magnificent engine for its time whose reliability and tape handling just did not perform to the levels needed, in particular by the electronic experiments. Its superior floating-point capabilities were valuable for processing data from bubble-chamber experiments with their relatively low data rates, but for the fast electronic experiments the "log jam" of the tape drives was a major problem.

So a second revolution occurred with the reintroduction of an IBM system, the 370/168, in 1976, which was able to meet a wider range of users' requirements. Not only did this machine bring dependable modern tape drives, it also demonstrated that computer hardware could work reliably and it ushered in the heyday of the mainframe, with its robotic mass storage system and a laser printer operating at 19 000 lines per minute. With a CPU cycle of 80 ns, 4 megabytes (later 5) of semiconductor memory and a high-speed multiply unit, it became the "CERN unit" of physics data-processing power, corresponding to 3–4 Mips. Moreover, the advent of the laser printer, with its ability to print bitmaps rather than ▷

A panoramic view of CERN's computer centre in the mid-1980s, during the era of the combined service provided by IBM and Siemens. (Courtesy Chris Jones.)

CERN's IBM 709 computer is unloaded at Geneva's Cointrin Airport in 1961, under the watchful eye of a Swiss customs officer, at right.

The arrival of the IBM 370/168 in 1976 ushered in the heyday of the mainframe and the "CERN unit" of physics data processing.

The user-friendly nature of the WYLBUR time-sharing system, which was developed at SLAC, was reflected in its beautifully handwritten and illustrated manual by John Ehrman.

32 CERN Courier September 2004

CERN Courier September 2004 33

do this better than the others. In order to defend the rights of the itinerant physicist, in 1983 Norman McCubbin from the Rutherford Appleton Laboratory made the radical but irresistible proposal: "don't do it better, do it the same!"
The HEPVM collaboration comprised most of the sites who ran

(after exploring options…)

cussion and explanation and working at the issues. Two important products resulted from this collaboration. A HEPVM tape was distributed to more than 30 sites, containing all the code necessary for producing a unified HEP environment, and the "concept of collaboration between sites" was established as a normal way to proceed. The subsequent off-spring, HEPiX and HEPNT, have continued the tradition of collaboration and it goes without saying that such collaboration will have to take a higher level again in order to make Grid computing successful.

Agenda

Analysis Facilities Forum

Thursday 21 Apr 2022, 18:00 → 19:00 Europe/Zurich

Videoconference zoom Analysis Facilities Forum — Join

18:00 → 18:10 **Introduction** — 10m

Speakers: Alessandra Forti (University of Manchester (GB)), Diego Ciangottini (INFN, Perugia (IT)), Lukas Alexander Heinrich (CERN), Nicole Skidmore (University of Manchester)

18:10 → 18:40 **kubernetes existing activities update** — 30m

Speaker: Ricardo Rocha (CERN)

Kubernetes and Ba...

18:40 → 19:00 **Discussion** — 20m