

# 7<sup>th</sup> Scientific Computing Forum

CERN, 02 October 2019, 10:00 h

## Minutes

Indico event page: <https://indico.cern.ch/e/Scientific-Computing-Forum-2019-10-02>

### Introduction and approval of the minutes (E Elsen)

E Elsen welcomed participants and reminded them that the meeting serves as an informal forum for exchange of information and views on scientific computing between CERN, the Member States and other countries.

On the minutes of the 6<sup>th</sup> meeting on 25 February 2019, a comment had been received by Bob Jones concerning the cost of commercial cloud offerings. The meeting agreed to change the minutes as requested.

### WLCG preparations for Run 3 and beyond (I Bird)

I Bird presented the slides attached to the agenda, highlighting that the previous assumptions on the cost evolution of computing equipment are probably too optimistic now; nonetheless for Run 3, the detailed timing of which will become clear only at the end of November, it is expected that the requirements can be fulfilled. However, constant investment by all funding agencies will be needed in order to ramp up to the required level for Run 4.

P Clarke commented that significant gains could be achieved from porting the GEANT framework to GPUs, as significant GPU capacity could be expected from HPC installations; the port requires a complete re-write, which is a major effort for which a call is open in the UK; simply adapting the current code will probably not suffice. E Lancon replied that the community needs to more generally consider the heterogeneity of the platforms that can be made available. E Elsen pointed to the information that would be provided in the subsequent talk in this SCF meeting.

E Elsen commented that no major changes to LS2 are expected from the LHC scheduling discussion in November, and advised to base the resource planning on the “typical” luminosity expectations rather than the “optimistic” ones.

E Elsen expressed satisfaction over the progress on data storage, the concentration of which in fewer and larger sites will help reduce investments as well as operational costs.

Replying to a question by J Templon, I Bird explained that it should be decided at the national level how to implement the sustained investments in preparation of Run 4, when possibly much higher levels of resources will be required. If possible, funds could be moved forward; otherwise experiments will need to request more capacity than they really require during LS3, or sites will have to provide more than has been requested.

E Elsen concluded by underlining CERN’s priority of ensuring that all data are being recorded, and the importance of the new data centre (PCC) for Run 4, which is as high a priority as the high-luminosity upgrade of the LHC machine and the Phase II upgrades of the experiments.

### GPUs in CMS: status and plans (D Piparo)

D Piparo presented the slides attached to the agenda.

E Elsen commented that there are similar efforts to CMS in the other experiments; ALICE even critically depends on running code efficiently on GPUs for their O2 project.

Replying to a question by E Elsen, D Piparo explained that currently, the GPU resources are dominated by Nvidia cards, which can be programmed in C++ and CUDA; code using CUDA can probably be ported to different GPU architectures. It is hence adequate to go that way for the time being; the issue is how to attract young, productive programmers to work on this. A Di Meglio added that many companies now release programming interfaces as open source without any particular focus on HEP; it is probably too early to choose one interface or the other now.

E Elsen commented that there had been attempts, such as GeantV, to “vectorise” the Geant code. The successful parts of that endeavour have largely been integrated into the production framework (Geant4). HPC centres have software support groups, which provide assistance in adapting the code. D Piparo replied that for Run 3, CMS estimates that at most 20% of the total CPU capacity will be used for simulation, to which event generation contributes in a minor way. J Catmore agreed that simulation should be considered very seriously, as in Run 4, the HLT farms possibly provide significant GPU capacity, which we must be prepared to take advantage of. E Elsen added that ALICE is in this situation in Run 3 already, as the O2 facility will be used for heavy-ion data taking only during a minor fraction of the run. C Bozzi commented that LHCb will use a much larger fraction of the CPU resources for simulation than indicated for CMS.

Replying to a question by C Bozzi, D Piparo said that CMS intends to equip all CPU nodes of the HLT farm with GPUs.

E Lancon commented that capacity for machine learning is available now on supercomputers; it is vital for HEP to create strong links with this community now, as it may be much more difficult later.

A Streit commented that the performance comparison with Intel CPUs is interesting, but AMD CPUs with high core counts should be added to it, as these offer some advantages over Intel that are also found on GPUs. Replying to a question by A Streit, D Piparo explained that CMS have looked at the POWER architecture, which was found to be less cost-effective than the x86 one; NEC vector cards have not been considered.

## CERN School of Computing (S Lopienski)

S Lopienski presented the slides attached to the agenda.

Replying to a question by E Lancon, S Lopienski explained that the CERN School of Computing (CSC) tries to link with other science domains by not focusing too much on HEP in general or LHC in particular; the organisers actively look out for potential lecturers from other communities. There is no formal link with photon science. Still probably about 90% of the students come from HEP, of which not all are from LHC experiments.

Replying to a question by E Elsen, S Lopienski explained that the programme committee consists of representatives of the experiments, EP-SFT, IT and other units, which explain what their respective priorities are.

O Smirnova commented that in her experience, obtaining a student slot for a school was not easy; in addition, not all potential applicants are able to fund the travel. She asked whether on-line courses had been considered. S Lopienski replied that there are usually more than 100 applications for the 70 student slots (the latter number is limited for logistical reasons), most of which are very well justified, even though some are somewhat surprising. With a fee of 1’200 EUR for the all lectures, lodging and meals, the cost is very moderate; some stipends are available for students in need. The lectures of the Inverted CERN School of Computing (iCSC) are webcast and recorded.

E Elsen suggested to consider whether training on new programming languages should be offered at the forthcoming schools.

## CERN openlab (A Di Meglio)

A Di Meglio presented the slides attached to the agenda.

Replying to a question by O Smirnova, A Di Meglio explained that the data formats used in HEP have been highly optimised and are hence very HEP-specific; Hadoop, Spark and other candidate solutions have been investigated, but do not outperform the HEP-specific solutions.

Replying to a question by C Grab, A Di Meglio said that there are a few formal collaborations with computer science departments on machine learning, but there is room for more.

## CERN's MALT project (T Smith)

T Smith presented the slides attached to the agenda.

Replying to a question by O Smirnova, T Smith explained that very little in-house development work has been done by CERN, which is mostly a one-off effort. The longer-term support effort is expected to be largely comparable with the one for the previous tools. In some cases, more standard solutions offer savings over the previous tools; for example, Kopano runs on very standard hardware, while Exchange required specially configured servers.

P Clarke expressed high appreciation of the work already done and still ongoing. E Elsen replied that the most difficult part will probably be rolling out the new tools in mass, which is still ahead.

## Round table discussion

Under this topic, another discussion on GPUs and heterogeneous architectures was held with contributions by E Elsen, P Clarke, C Grab, D Britton, A Di Meglio, O Smirnova, J Templon and I Bird, during which the following points were made:

- The challenges of large-scale GPU installations on the one hand, and heterogeneous architectures on the other hand, may be similar, but are not identical.
- A lot can be achieved via collaborations with supercomputer centres and universities, without spending undue amounts of money.
- There is a general tendency towards providing pledges and unpledged resources at supercomputer centres. These centres are investing much more massively in GPU resources than in CPU ones. HEP must act. Enabling Geant4 to take advantage from GPUs and heterogeneous platforms should be the highest priority.
- The required work on software can be split in developing or adapting algorithms, which can be done in isolation, and making frameworks compatible, which cannot.
- In view of the possibly significant costs of the software developments, careful cost-benefit analyses are required in order to find the right balance. Work on other architectures should be included in these analyses.
- A possible way of attracting effort would be to allow funding agencies to pledge for development resources, which could be helpful for the career of the persons concerned. WLCG has done this in the first phase. However finding the right metrics is challenging, and the “creative chaos” we benefit from now could be unduly constrained; it may be more of an option for the implementation phase, once the direction to take is clear.
- Some of the developments towards GPUs and heterogeneous platforms are expected to impact the performance on CPUs positively as well.

E Elsen suggested that a report on the status and plans of Geant4 be a topic in the next meeting. M Girone pointed out that on 15 October 2019, an HSF meeting will be held on the performance of the GeantV prototype.

Replying to a question by F Meijers, I Bird explained that following experiments' requests, a limited service of GPU capacity is being prepared in IT, carefully considering the choice of GPU cards. F Hemmer added that such service could be complemented at other sites.

### **Any other business and date of next meeting**

It was tentatively agreed to hold the next meeting on Thursday 13 February 2020 at 10:00 h at CERN (09:00 h UTC). E Elsen and H Meinhard invited participants to suggest topics and/or speakers.

E Elsen thanked participants and closed the meeting stating that there are still many more questions than answers.

## Annex: Attendance

### At CERN:

Maite Barroso-Lopez  
Ian Bird  
Concezio Bozzi  
David Britton  
Simone Campana  
James Catmore  
Alberto Di Meglio  
Mihnea Dulea  
Eckhard Elsen (chair)  
Maria Girone  
Christophorus Grab  
Frederic Hemmer  
Andreas Hoecker  
Eric Christian Lancon  
Sebastian Lopienski  
Frans Meijers  
Helge Meinhard (scientific secretary)  
Emmanuel Ormancey  
Danilo Piparo  
Oxana Smirnova  
Tim Smith

### Via teleconferencing:

Sadaf Alam  
Volker Beckmann  
Peter Clarke  
Josep Flix  
Thorsten Kollegger  
Gonzalo Merino  
Jean-Pierre Meyer  
Pepe Salt  
Achim Streit  
Jeff Templon

### Apologies:

Karl Jakobs  
Kajari Mazumdar