Thoughts from ATLAS

On the Snowmass CompF Summary Document

Snowmass Summer Meeting: EF Cross-Frontier "Big Experiments" Session

21 July 2022

Alessandro Di Girolamo (CERN) and Zach Marshall (LBNL)

With thanks to James Catmore, Andrea Dell'Acqua, Johannes Elmsheuser, David South, and Nick Styles

Brief Disclaimer

- We received this 30-page document on Monday evening
- There's been no real time for any discussion within the experiment, or even amongst the coordination team
- We suggest that at least 2 weeks time should be given to the experiments (and the community) to review the document when finished and provide final feedback
- We have endeavored to identify some issues of relevance to our experiment

- Two missing / "incomplete" sections are quite important
  - Personnel and Training
  - Conclusions, Outlook, and Recommendations (particularly Recommendations)

General Remark

- It should be emphasized that this document is not in isolation
- For example, in the Introduction, L51-L59: it would be useful to mention the ATLAS S&C HL-LHC Roadmap (currently Ref [39], but referenced only in L413). There might be also other experiments public documents with HL-LHC plans (e.g. LHCb TDR)
- Very important to clarify that (some of) the experiments *do* have a clear strategy and roadmap defined for the next 10y – Snowmass should complement those plans (cf. L611)

Executive Summary Section: “Four Critical Challenges”

- Continuous development (support for Dev, Ops and Integration, not just R&D)
  - We agree this is a major issue. We are pleased to see some grant applications asking/commenting directly about long-term support for R&D outcomes.
  - The examples might be more explicit on the experiment specific SW: e.g. Athena for ATLAS.
  - Not clear whether data and software presentation belong here. How are CDF / D0 doing in that area? How is BaBar?
- Support for cross-cutting development
  - This seems to get some support in the US, particularly for “computing” projects (which, incidentally, deserve a mention here)
  - It is certainly true that more is needed
- Heterogeneous resources
  - While the other “challenges” are (basically) requests for funding, it is not clear what message is being sent in the phrasing of this challenge – perhaps it is a mixture of Challenges 2 and 4?
  - The text seems to conflate non-standard architectures (e.g. GPU / ARM) with non-standard sites (HPC, Cloud); the rest of the doc often conflates “challenges of using an accelerator” (e.g. GPU programming) with “challenges of heterogeneous resources” (i.e. portability)
- Training and career paths
  - We agree this is a challenge but want to emphasize that we need support for *Physicists* who are Software and/or Computing experts, as well as S&C professionals (who likely would not be professors in a physics department)

Overall: challenge 3 is some sort of convolution of 1, 2 and 4……

- Thinking aloud: do we need 4 challenges, or are 3 “enough”?

A few additional comments

- The discussion of QC and ML development seem to contradict the first critical challenge, “Grants typically fund ground-breaking R&D or development of new software, but not modernization, maintenance, and user support of existing tools.”
  - Machine learning is not an “emerging technology”, it is pervasive, and QC should not be put on the same level (or even in the same sentence)
  - Difficult to see QC developments having practical HEP impact for the next 5+ years, but obviously “sexy” from the point of view of funds/grants requests
  - We might need to/be able to participate in such proposals, but we have to be careful to not just follow the hype.
- Seemingly outsized discussion of GPUs (e.g. L650) compared to different CPU architectures (ARM/SoC are not mentioned)
- “Cloud computing is useful but at what cost?” (L659/60): it has a too negative spin.
- Fig 0-1: it is an interesting attempt to illustrate the boundaries between “online” and “offline” but the result is dubious, the “Edge” numbers might be misleading.
  - Maybe just have the “Offline” part?
  - Also, “Online” have often their own detailed documents (it is definitely true for both ATLAS and CMS)
- A few things are missing that might deserve some mention
  - Environmental impact of computing, infrastructure (build, code distribution, etc)
  - Distributed Computing infrastructure: it is only “gently” touched in chapter 5, but it might deserve more prominence taken into consideration that it is a relevant financial investment (both personpower and HW).

Extras