Fundamentals of an HL-LHC Computing Model

Frank Wuerthwein
UCSD

The purpose of this white paper is to write down some fundamentals that any detailed HL-LHC computing model ought to have because they are to good approximation technology independent.

Science Driven by Colliding beams:

Fundamentally, the LHC is colliding bunches of protons on bunches of protons. The most fundamental unit of computation is that of an event, i.e. the crossing of two bunches of protons that happens every 25ns in each of the ATLAS and CMS detectors. As the brightness of the beams increases, more proton-proton collisions occur on average per crossing of the beams, i.e. per event. More collisions per event increases the occupancy of the detector, and thus the CPU time required in the reconstruction.

We make the fundamental approximation that any two events are independent. We may decide to process multiple events simultaneously, or fragments of multiple events simultaneously, but we shall do so only for convenience. Nothing in the science requires it.

The LHC computing problem is thus fundamentally one that lends itself to High Throughput Computing (HTC) perfectly. It includes hundreds of billions of events, all of which could be analyzed independently. It requires a high degree of automation due to the very nature of the large scale processing requirements. HTC defined as "automated processing of a large number of independent entities" fits the HL-LHC perfectly.

Science as a team sport:

The intellectual property of each of the LHC detector collaborations is to a good approximation preserved in their software and data, broadly defined. In one way or the other, the vast majority of the collaboration contributes to its intellectual property. The scientific results are unthinkable without the collaboration, leading to author lists that include a couple thousand authors from close to 200 institutions in more than 40 countries per collaboration. As people have lives, and it is both impractical and ineffective to assume that all of people's lives can converge in one point, it is a reasonable assumption to consider a "globally distributed" computing model to be fundamental to the success of the HL-LHC. Even if all countries were willing to send checks to one place to host all of the HL-LHC computing, we assume that this would be less effective in enabling the globally distributed workforce of the HL-LHC than a well designed globally distributed computing model.

Central vs distributed control and decision making:
The path from instrument to discovery to publication has fundamentally three aspects to it. First, the collaboration produces data and software for the collective. We shall call this “public data products” in the spirit that 2000 collaborators having equal rights to all of the data and software to analyze it makes the data a public good within the collaboration. Second, the individuals within the collaboration assemble in groups that both collaborate and compete against each other for the “best science” to be derived from the public data. In the process, secondary data gets derived using a mix of private and public software that potentially infers their progenitors a competitive advantage. Third, a centrally organized review process determines which results are ready to be released to the general public. We acknowledge that there is a priori some potential for tension between steps two and three. Any discovery, or any other science from each of the HL-LHC experiments will have to go through all three of these high level steps. This implies an interplay of central and distributed control and decision making, as well as traceability and reproducibility of the results that is fundamental to how we have done, and will continue to do our science. While we may decide to move the line with regard to what level of data shall be the common starting point for all (more “refined” vs more “flexible”), there will always be a competitive piece to it. A successful computing model for the HL-LHC must allow for such creative competition between individual groups for the benefit of the science that comes out.

LHC Science is cyclical:
The lion share of the production of the “public data products” is release/campaign driven. Each experiment produces a software release, then validates it, possibly going back and forth a few times, before the collaboration is ready for large scale production of simulation and detector data with that release. Here “software release” refers to the combination of software, calibrations, etc., i.e. anything required to produce a consistent and complete set of public data products. If the HL-LHC computing model were to allow the collaborations the means to shrink the time between sign-off on a release and having a consistent and complete set of data products with that release then this would accelerate the science, and thus be a desirable feature. The concept of “elastic scale out” is thus fundamentally desirable. Given the previous discussion, we assume that this is desirable at all scales. An individual graduate student potentially benefits if her or his processing can elastically scale out into resources she or he controls as much as a research group, a University, a national laboratory, or a country, or the entirety of a collaborations global operations. Targets for elastic scale out may be as varied as business models for computing, and include commercial, institutional, and nationally allocatable resources. To be perfectly clear. While it would be potentially disastrous for the HL-LHC science to depend on resources the collaborations have no control over, it would be very short-sighted not to plan for a computing model that is agile enough to dynamically add additional resources in order to shrink turn-around time via elastic scale out.