Future computing strategy
Some considerations

Ian Bird

WLCG Overview Board

CERN, 28th September 2012

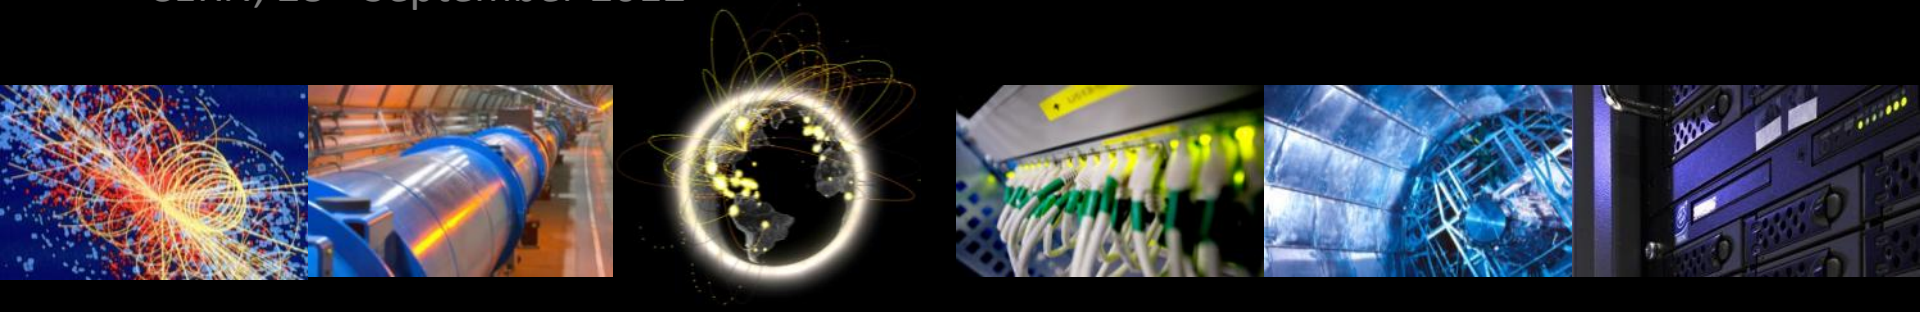

WLCG
Worldwide LHC Computing Grid

Introduction

- HEP needs for computing on ~2020 (post LS2) timescale are potentially very much larger than today
- Continuing the current computing models will not scale

☛Need to consider the next major change in the way we do computing (what and how)

- We also have an series of opportunities now:
  - Commonalities between experiments (not just LHC)
  - Common interests with other sciences
- We must also consider how major developments could be funded

Concerns

- WLCG has been successful in supporting LHC data analysis, and is already evolving based on experience with data and technology improvements, however, we must plan for the longer term
- Vital that the funding for the WLCG Tier 1 and Tier 2 centres be maintained at a level sufficient to ensure the full exploitation of the data produced by the LHC in the coming years
  - We are only at the beginning of the life of LHC and the experiments
- Anticipate significant higher trigger rates and data volumes (after LS1 and LS2)
  - Must continue to strive to be more efficient within the resources available

Data Management/Data Preservation

- Data Management: LHC has demonstrated very large scale data management. Must build on this success:
  - Improve our own efficiency – improved models of data placement, caching, and access
  - Work with wider scientific community to build community solutions or move towards open standards if possible
- Should drive/explore collaborations with other data intensive sciences and industry to consolidate data management tools in the long term
  - Build a broader support community, helps our sustainability
- Data preservation and open access to data is now very important and highly publicly visible
  - policy and technical investments are needed now
- CERN could provide the leadership to build a set of policies and technical strategies to address this for HEP, building on the work of DPHEP and collaborating with other big data sciences
- It is already urgent for LHC experiments to ensure the ability for long term use and re-use of the data

e-Infrastructures

- Grid computing has been very successful in enabling global collaboration and data analysis for LHC
- We should generalise this (with technical evolutions already under way) for a more general HEP-wide infrastructure
  - Several requests from other experiments to use the WLCG
  - Must understand how this would fit with regional infrastructures, national and international efforts etc., and how to integrate with other science communities
- We must plan for the next generation of e-infrastructure (e.g. on 5-10 yr timescale), making use of new technologies, e.g.:
  - Terabit networks – could enable true remote access
  - How does "cloud" computing fit? – need to understand costs of commercial clouds vs HEP (or science) clouds: cloud federations would be needed for our collaborations

Prospects for HEP software on new CPUs

- HEP code runs efficiently on today's CPUs and large efforts continue to be made to improve algorithms to achieve better performance
- However, analysis of software performance also show that further big gains can be made by exploiting features of the CPUs' microarchitecture
  - by making use of vector registers, instruction pipelining, multiple instructions per cycle
  - by improving *data and code locality* and making use of *hardware threading*
- Moreover, new architectures can be useful for off-loading large computations to accelerators (GPGPUs, Intel MIC)
- Today multi-core architectures employing O(10) cores are well exploited using a multi-process model (1 job/core). However this performance will not scale to future generations of many-core architectures employing O(100) cores due to memory issues
  - there are technical issues related to connecting many cores to shared memory that will reduce the amount of memory available to each core
  - whereas the memory footprint of HEP code is increasing due to increasing event complexity as the energy and luminosity of the LHC is increased
- There is a growing consensus that HEP code needs to be re-engineered to introduce parallelism at all levels
  - the event level, the algorithm level, and the sub-algorithm level
- We need to act now or run the risk that HEP software will not run at all in the future

The challenge

- Concrete algorithms can be run in parallel by making use of *threads* but integrating them to run in a single application is highly non-trivial
- It will require new levels of expertise that need to be acquired by the community
- Tasks will involve a major re-engineering of frameworks, data structures and algorithms e.g.
  - making code thread-safe to exploit multi-threading
  - developing new framework services for scheduling parallel execution of threads on all available cores
  - targeting libraries/toolkits where the biggest impact can be achieved e.g. GEANT
- Making changes in the code of running experiments must be done gradually whilst preserving the correctness of the physics output
- A fast response is needed since technology is evolving quickly
  - At the LHC, LS1 provides a window of opportunity for introducing first changes
- The scale of the changes requires a very large effort to be invested
- The adoption of a collective response will help to meet the challenges using available expertise and resources and within the required timescale

Initiatives taken so far

- A new forum was established at the start of this year, the Concurrency Forum, with the aim of :
  - sharing knowledge amongst the whole community
  - forming a consensus on the best concurrent programming models and on technology choices
  - developing and adopting common solutions
- The forum meets bi-weekly and there has been an active and growing participation involving many different laboratories and experiment collaborations
- A programme of work was started to build a number of demonstrators for exercising different capabilities, with clear deliverables and goals
  - 16 projects are in progress started by different groups in all corners of the community
- In the longer term this may need to evolve into other means for measuring progress and steering the future work programme

Long Term strategy

- HEP computing needs a forum where these strategic issues can be coordinated since they impact the entire community:
  - Build on leadership in large scale data management & distributed computing – make our experience relevant to other sciences – generate long term collaborations and retain expertise
  - Scope and implementation of long term e-infrastructures for HEP – relationship with other sciences and funding agencies
  - Data preservation & reuse, open and public access to HEP data
  - Significant investment in software to address rapidly evolving computer architectures is necessary
  - HEP must carefully choose where to invest our (small) development effort – high added value in-house components, while making use of open source or commercial components where possible
  - HEP collaboration on these and other key topics with other sciences and industry
- CERN plans a workshop to discuss these topics and to kick-start such an activity

Long term – LHC experiments

- Recognise the importance of re-thinking computing models for post-LS2 timeframe
  - Could easily imagine order(s) of magnitude more data
  - Current models simply will not continue to scale
- Recognition that experiments' solutions today are not so different
  - Must devote effort on effective areas:
    - CPU-heavy algorithms, simulation, data management, etc
  - Frameworks and data serving mechanisms could be common
  - Willingness to bring these ideas together and start to plan what needs to be done

Summary

- WLCG/LHC experiments recognise the need to plan for the future
- HEP has significant expertise/experience now on core technologies and has an opportunity to engage with other sciences and industry
  - Essential if we are to maintain funding and attract people
- Plan community-wide discussion/engagement on key themes

LHCC: Update of computing models

- For the period LS1-LS2
- Agreed to document:
  - Concepts of the computing models now, anticipated evolution in this period
  - Explain the use of new technologies and consequent improved use of resources, common tools and strategies as well as differences
    - E.g. Disk strategies, data popularity, data federation
    - Analysis strategies – prompt vs delayed analysis with high trigger rates
    - Use common templates, tables, etc – common understanding of concepts
    - Also explain strategies for better use of current and potential CPU architectures
- Timescale:
  - (Draft) report in time for C-RRB in October 2013 ➔ need to have a good draft for discussion by September 2013