UiO : University of Oslo

HEP Software Foundation

James Catmore (UiO)

Please note: this is a summary of my understanding of the current status - I am not speaking on behalf of the HSF!

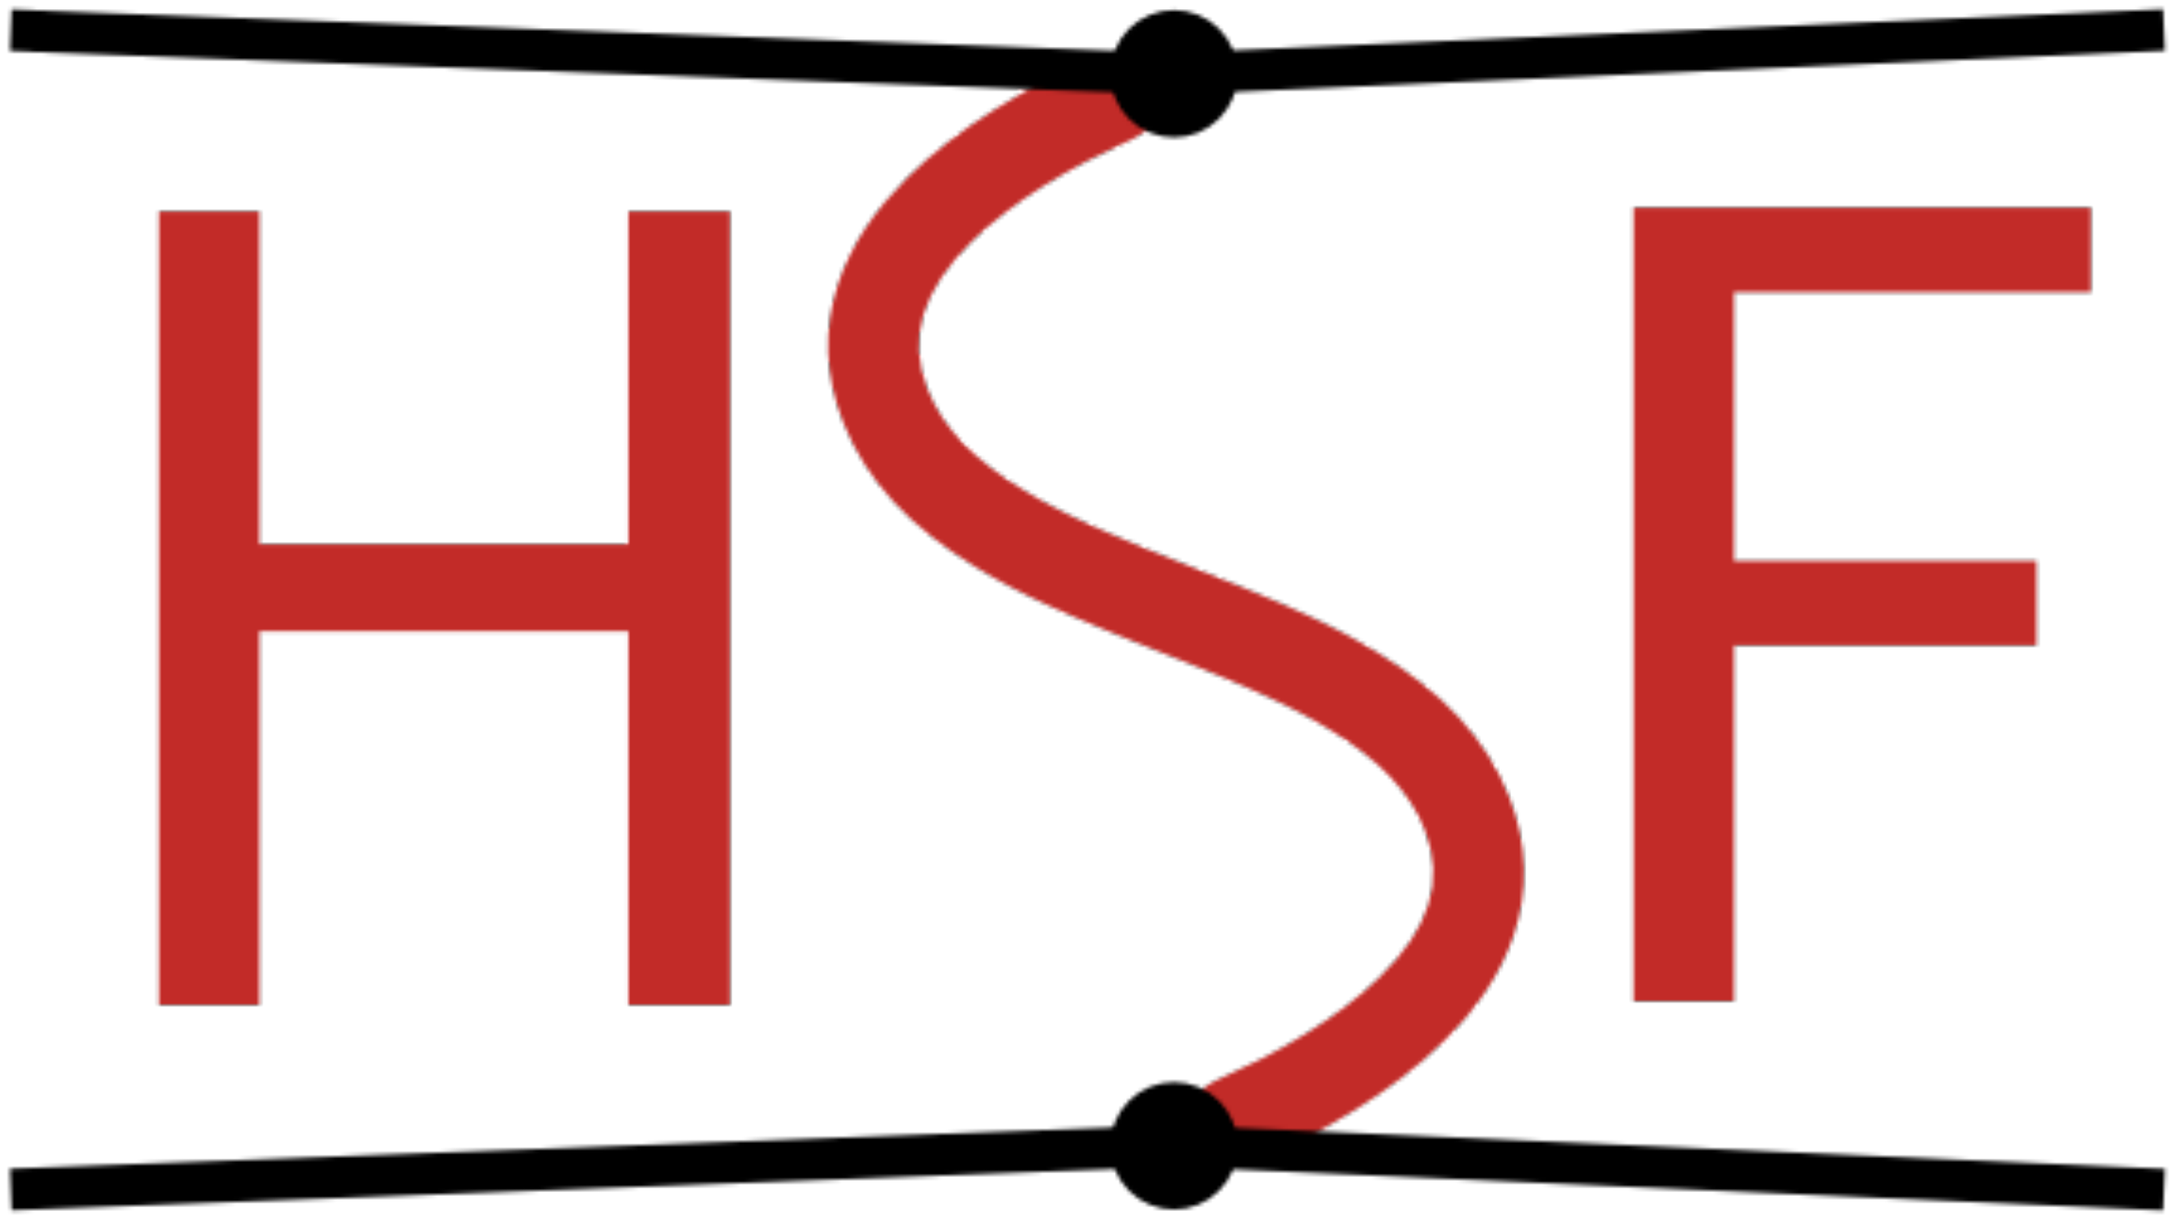
HSF
HEP Software Foundation

The HEP Software foundation

- http://hepsoftwarefoundation.org
- Facilitates coordination and common efforts in high energy physics software and computing, internationally
- Loose collaboration, "do-ocracy"
  - Open to new projects, as long as they abide by certain standards: http://hepsoftwarefoundation.org/project_guidelines.html
- Several working groups and discussion forums
- Frequent open meetings, managed via Vidyo and Indico
- Larger workshops organised on a regular basis
  - e.g. one taking place this week in LAP (Annecy)

Motivation

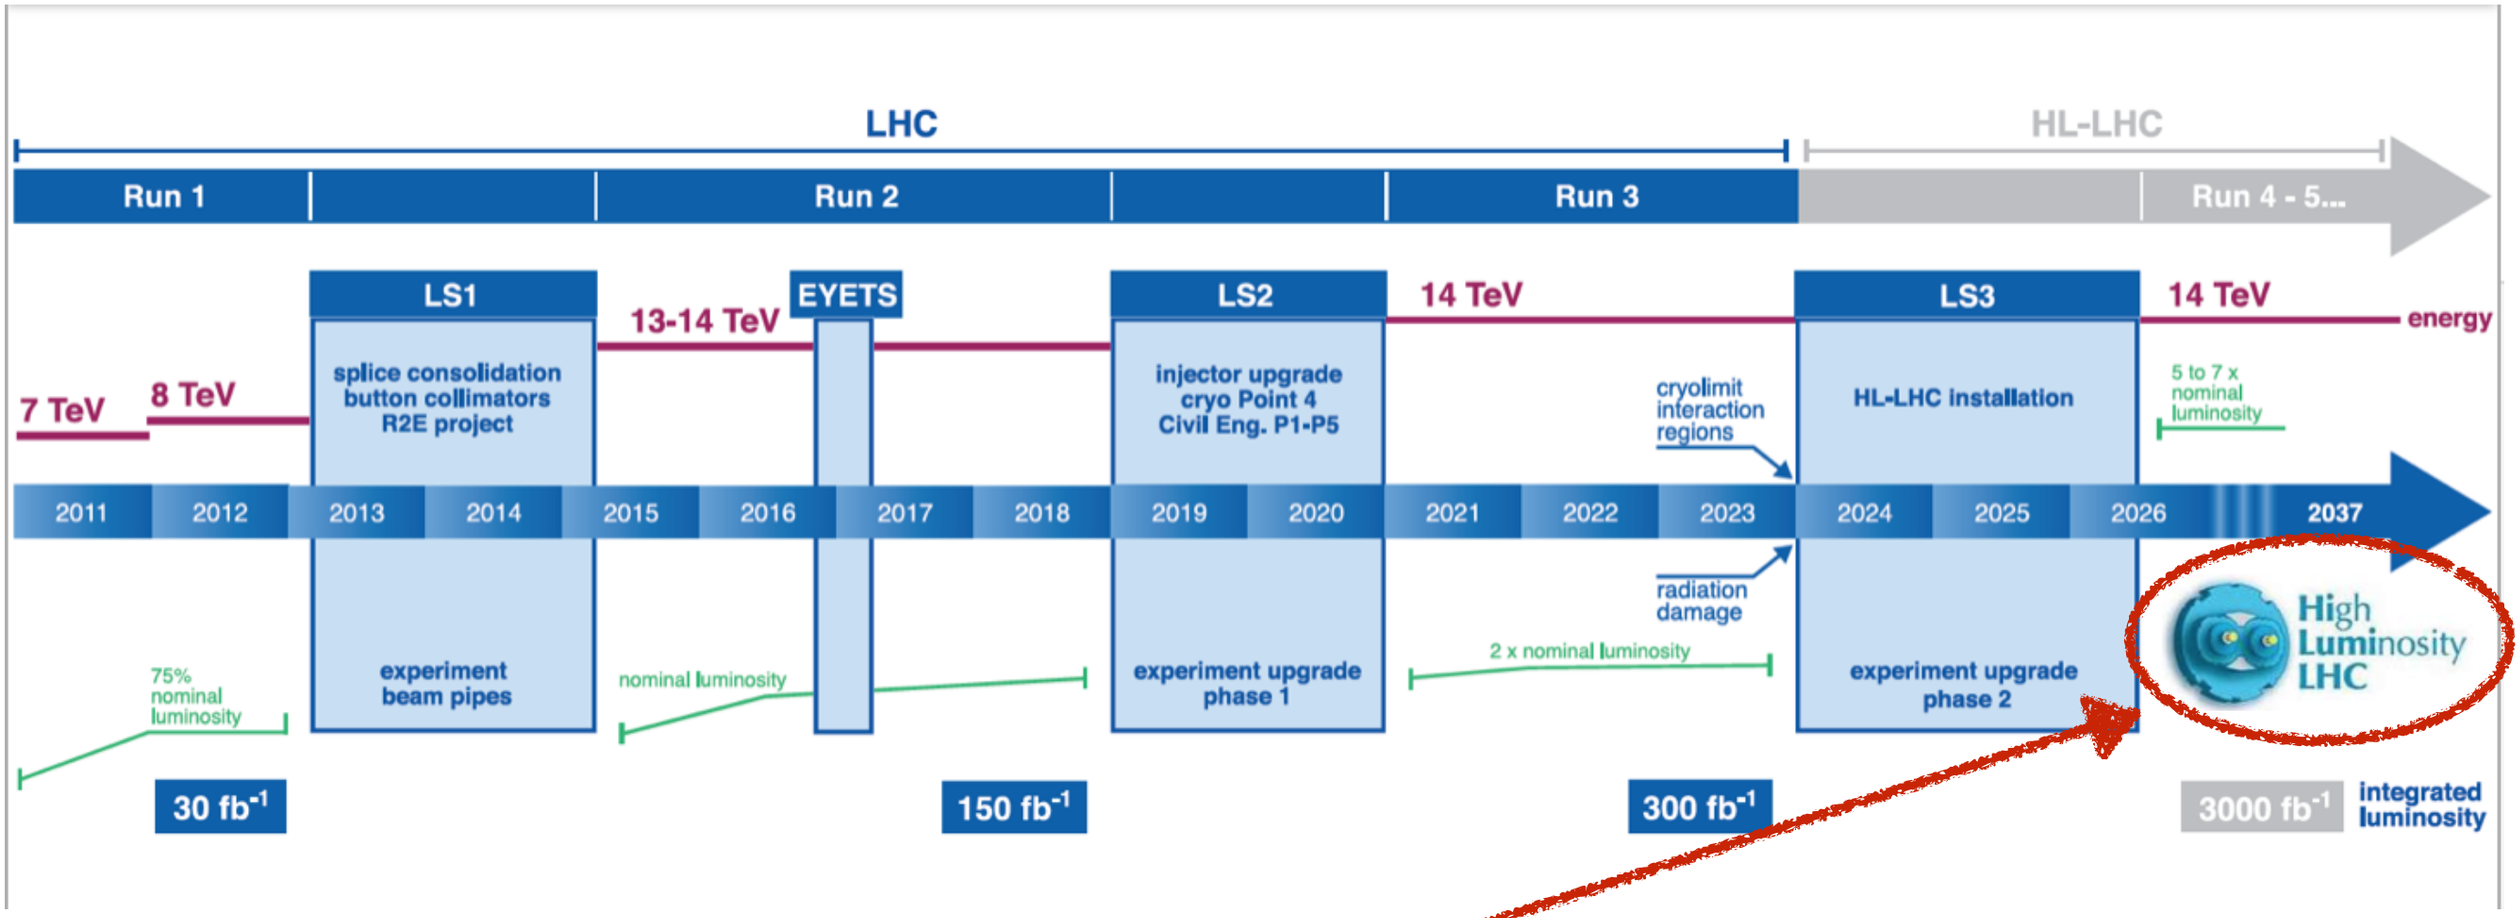

10x more events

4x more complex

Motivation

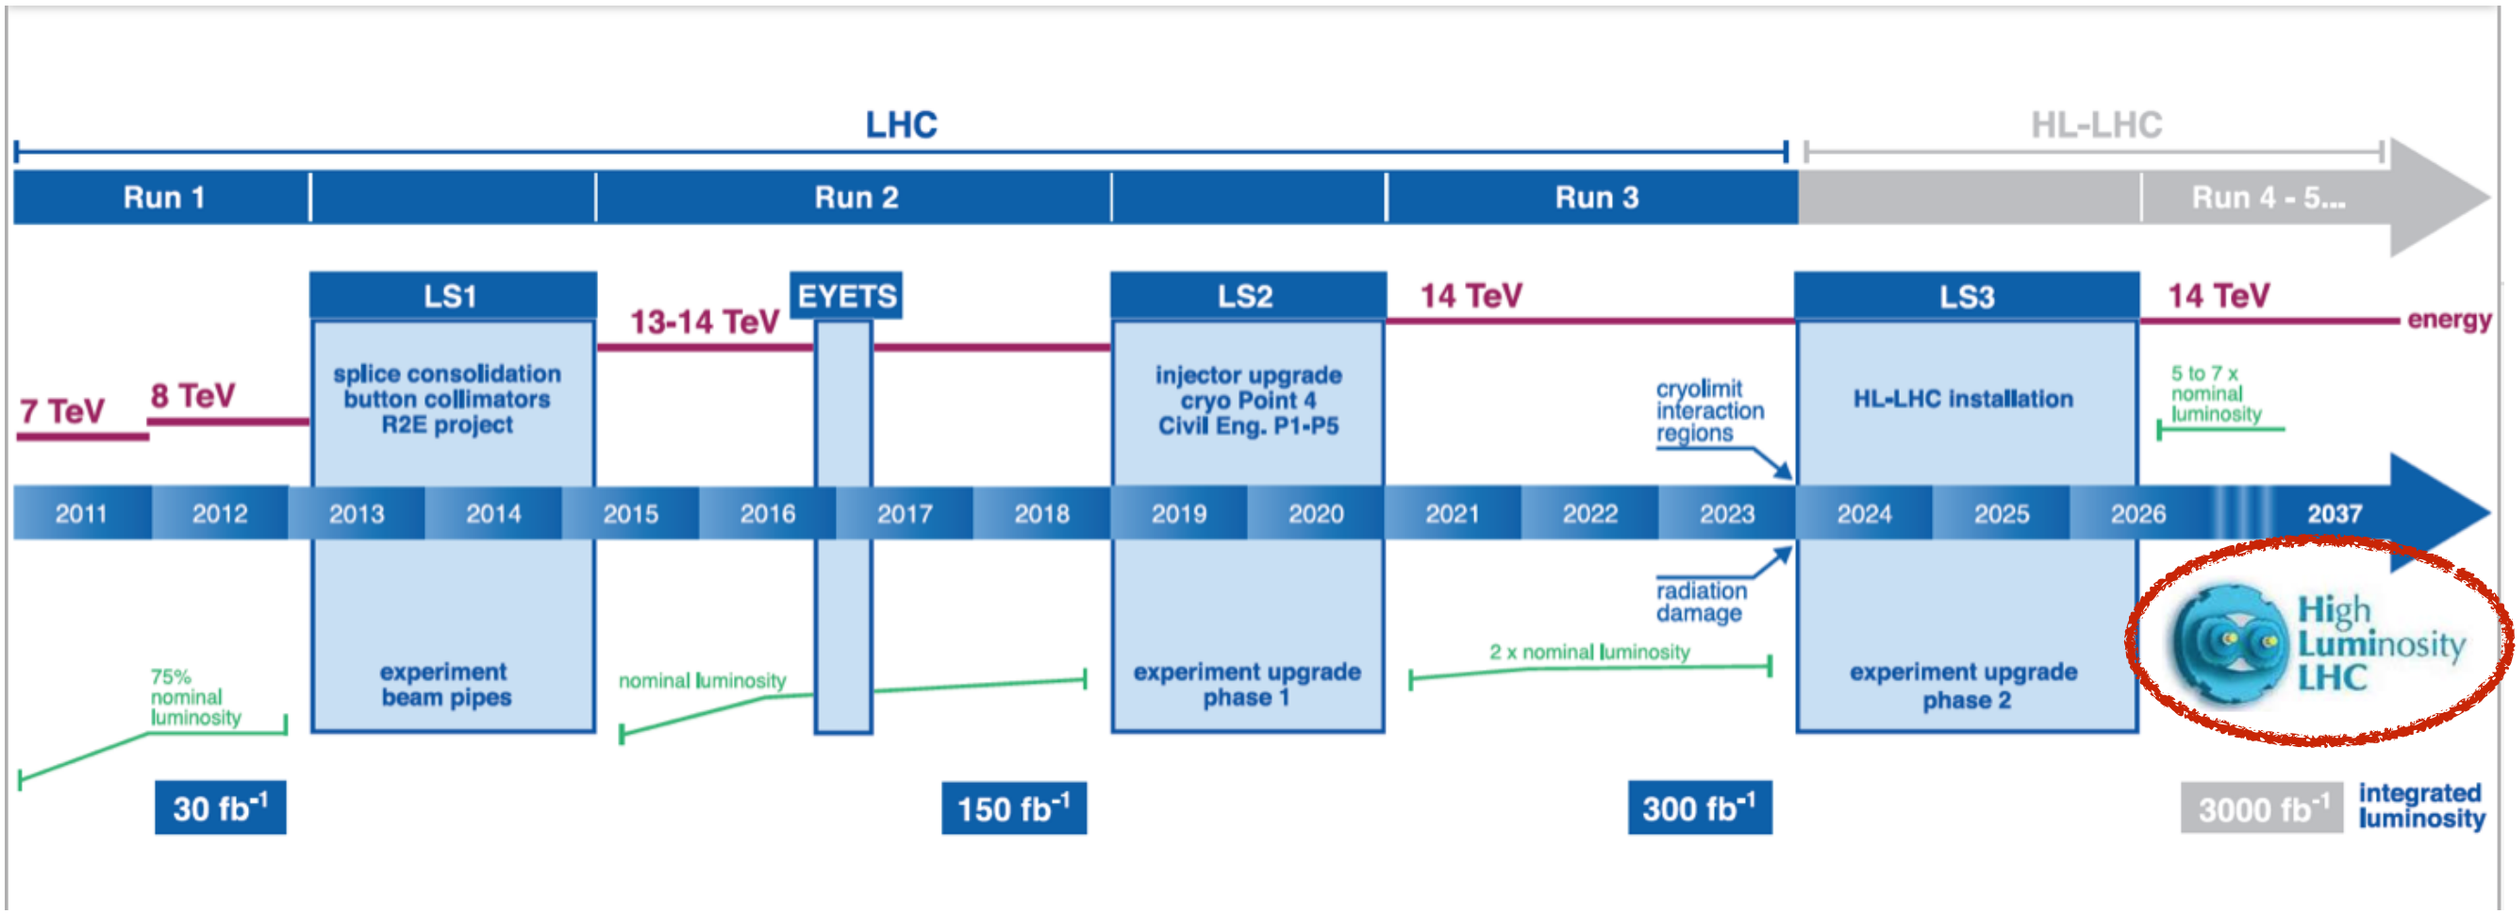

- The volume of data our computing/software will be expected to process and store is growing very quickly
- In order to make full use of new processor architectures software will need to be flexible and adaptable; also making use of non-HEP HPCs and commercial clouds
- Resources (both computing and human) are limited and not growing
- Some sharing of the development burden between experiments is going to be needed
- See talk tomorrow by Borut…

Objectives

- Catalyze new common projects
- Promote commonality and collaboration in new developments to make the most of limited resources
- Provide a framework for attracting effort and support to S&C common projects (new resources!)
- Provide a structure to set priorities and goals for the work
- Currently working on a Community White Paper (CWP)
  - CWP workshop this week: https://indico.cern.ch/event/613093/overview

CWP working groups

- Software Trigger and Event Reconstruction
- Machine Learning
- Data Access, Organisation and Management
- Software Development, Deployment and Validation/Verification
- Data Analysis and Interpretation
- Conditions Database
- Data and Software Preservation
- Event Processing Frameworks
- Physics Generators
- Workflow and Resource Management
- Visualisaton
- Computing Models, Facilities, and Distributed Computing
- Careers, Staffing and Training

CWP Highlights: analysis ecosystems

- Workshop in Amsterdam 22-24 May
  - https://indico.cern.ch/event/613842/
- Review of the experiments' analysis models
- The central role of ROOT and its future
- Data formats
  - Beyond files? Database-style access for plotting and statistical analysis?
- Analysis in a multi-threaded / massively parallel environment
- Programming languages
- Interface with non-HEP tools (machine learning…)

Database-style interaction

Rapid queries on big data are possible; in fact, it's a big field:

- Instead of users maintaining private skims, running local processes, they share a distributed query server.
- Instead of fetching data from disk, cache in RAM/SSD/X-Point.
- Instead of evaluating user's code as-is, translate it into a form that optimizes memory bandwidth.
- Instead of executing operations on nested objects, execute on flat arrays of numbers.

Jim Pivarski (https://indico.cern.ch/event/613842/contributions/2585803/attachments/1464053/2264325/pivarski-femtocode-updated.pdf)

CWP Highlights: reconstruction

- Challenges include:
  - How to cope with increased event complexity/instantaneous luminosity
  - How to adapt to commodity CPUs moving towards many-core,
  - Which architectures (CPU+coprocessor hybrids, GPUs and/or FPGAs) will provide the best throughput per cost in the future?
  - Can we exploit novel techniques to improve physics output or to improve throughput (e.g., real-time analysis, machine learning algorithms)
- This means
  - Going faster (faster algorithm implementations, making full use of hardware, e.g. multithreading)
  - Doing things right the first time (eg, solid validation/verification procedures, deriving "good enough" conditions derived quickly)

Software developments in trigger/reco can benefit physics output of future experiments

Doing more with less?

Constraint	Implication	Opportunity
Fixed CPU budget	CHF = (Time/event) * (# of events)	Keep more events (higher trigger rate) or do more per event (lower pt thresholds) with faster algorithms
Fixed storage budget	CHF = (Size/event) * (# of events)	Keep more events (higher trigger rate) with less information

D. Lange, V. Gligorov, https://indico.cern.ch/event/613093/contributions/2615321/attachments/1482729/2300054/TriggerRecoAnnecy.pdf

Evolutionary vs Revolutionary approaches

- The steady improvement of the performance our hungriest algorithms (eg, tracking implementations) may be sufficient to keep up with demands of future facilities (eg, event rate, pileup)?
  - In which areas are new approaches needed to complement current algorithms?
- New detector concepts require R&D (eg, high-granularity calorimetry)
- Explore breaking the “event” paradigm in places to facilitate the use of GPUs or other “new” architectures
- Determine the “right” mix of “online” vs “offline” processing capabilities for physics. When does reusing online derived quantities help?
- The evolution of analysis facilities will impact data structure requirements
- Reducing the raw data saved per event requires big changes to calibration procedures (likely not possible in all cases)

CWP Highlights: distributed computing

Computing Models, Facilities, Distributed Computing

- Draft Chapter
- Scope
  - Consider changes in the computing models, and the distributed infrastructure that can help to mitigate the disparity between the anticipated requirements for storage & computing and what is probably affordable. How can the models easily adapt to changing technology, and resources, without need for continual major porting efforts? We must be able to easily use any and all resources that may be offered to us. NB. This is a cross-cutting and high level view of the overall computing models
- Charge
  - This working group focuses on how computing models might evolve and which elements might be key to developing different possibilities over the next years. Contributions originating from many other CWP WGs will of course contribute to the overall computing models.
    - How will (or should) computing models evolve in the next 5-10 years to meet the challenges for HEP in the 2020s and which elements are key to developing these possibilities?
    - Which elements of the current computing models will ***not*** need to change over the next 5-10 years?

I. Bird: https://indico.cern.ch/event/613093/contributions/2615330/attachments/1482552/2299688/Annecy_Workshop_WG_summaries.pdf

CWP Highlights: distributed computing

Key challenges and opportunities (1 slide)

- Fitting requirements (>>x10-50) within constrained funding
  - Requirements significantly exceed likely hardware evolution expectations – but that is a big unknown
  - Reduce operational costs
  - Understand cost models and trade-offs
- Potential for common solutions between experiments – at all levels of the computing models
- Facilities – consolidation of resources in regions
  - Efficiency of scale and operations costs
  - Benefit from technology advances e.g. as demonstrated by large internet companies
  - Agile use of heterogenous resources, workflow-specific facilities, opportunistic resources, etc.
  - Create opportunities for very different models of processing and analysis

I. Bird: https://indico.cern.ch/event/613093/contributions/2615330/attachments/1482552/2299688/Annecy_Workshop_WG_summaries.pdf

Other HSF activities

- Google Summer of Code @ CERN: http://hepsoftwarefoundation.org/activities/gsoc.html
- Licensing: http://hepsoftwarefoundation.org/activities/licensing.html
- Packaging: http://hepsoftwarefoundation.org/activities/packaging.html
- Common tracking software forum: http://hepsoftwarefoundation.org/activities/tracking.html
- Training: http://hepsoftwarefoundation.org/activities/training.html
- Visualisation: http://hepsoftwarefoundation.org/activities/visualization.html

Summary

- Demands on HEP computing & software are growing and growing and growing
  - We don't have a magic money tree
  - We must *do* much more, without *getting* a lot more
  - We must make use of non-HEP resources and full use of new technologies
- The HEP community will need to work together to succeed - much more than in the past
  - The HEP Software Foundation provides an excellent basis for this
  - Please read the White Paper when the draft is presented to the community
  - … and get involved: http://hepsoftwarefoundation.org/get_involved.html