Software: challenges, opportunities, constraints

Michel Jouvin, IJCLab, michel.jouvin@ijclab.in2p3.fr

IN2P3 Scientific Council, June 23, 2022

Disclaimer

- I’m the head of IJCLab Computing Department
- I don’t pretend to know the situation for each IN2P3 community
  - My personal history is mainly with HEP
- IJCLab has a scientific scope similar to IN2P3
  - Significant involvement in HEP computing (grid) and software (AIDAinnova, ATLAS, ALICE, GEANT4…)
  - Significant involvement in software (and associated computing) for astrophysics/cosmology (LSST, SVOM, …) and multimessenger astronomy (GRANDMA, Fink)
  - Participation to starting effort, through Eurolabs project, to build a low-energy nuclear physics community on SW and computing

The Grand Challenge

- Data volume deluge in all communities
  - HEP: 10+ years of experience already with LHC experiments, soon Belle II and upcoming DUNE and HL-LHC
  - Astrophysics & cosmology: new experiments at LHC scale (LSST, CTA, SKA…)
  - Nuclear physics: EIC, GSI FAIR
- Data access: LHC demonstrated that the network is very reliable and not a technical challenge (funding challenge?) but data access is much more
  - Data locality vs. Remote access
  - Data federation: how many data stores seen/managed by the experiments
  - What is the optimal organisation/format for which type of processing
  - DOMA concept: Data Organisation, Management and Access → ESCAPE European project
- Data processing: how to efficiently (time, energy, price) process those data
  - Different answers for different type of data (ex: billion of independant events vs sky images)
  - How to efficiently use the new processor architectures that are very different from the generic ones
  - New algorithm approaches, e.g. machine learning

Example : the HL-LHC Challenge

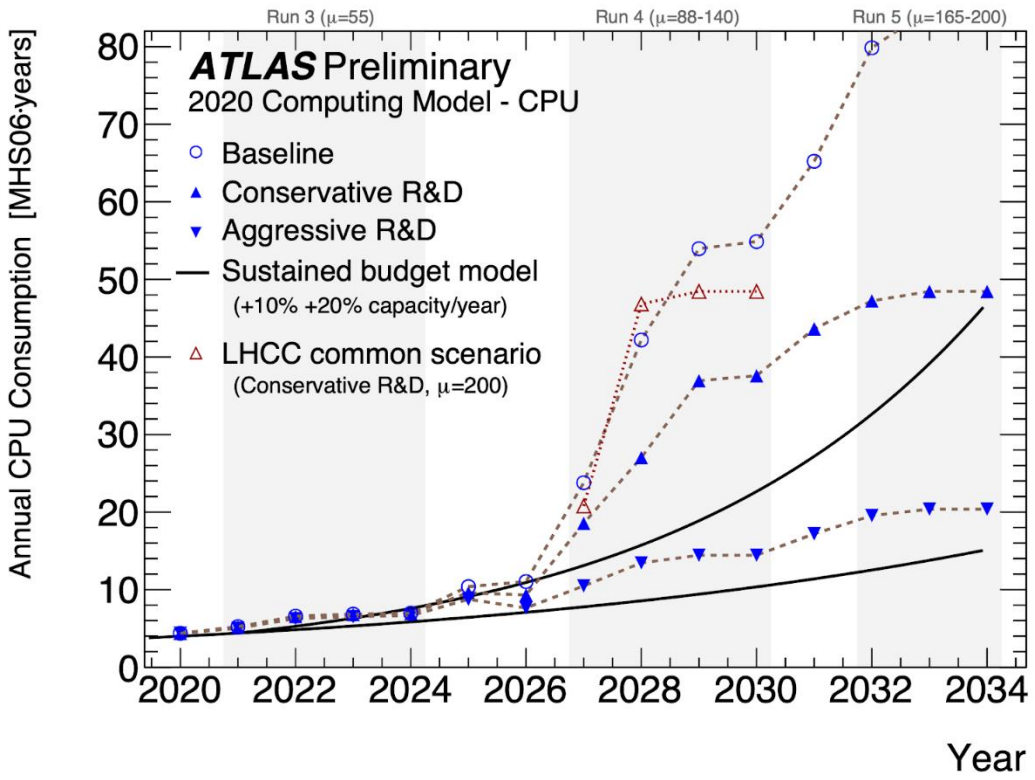

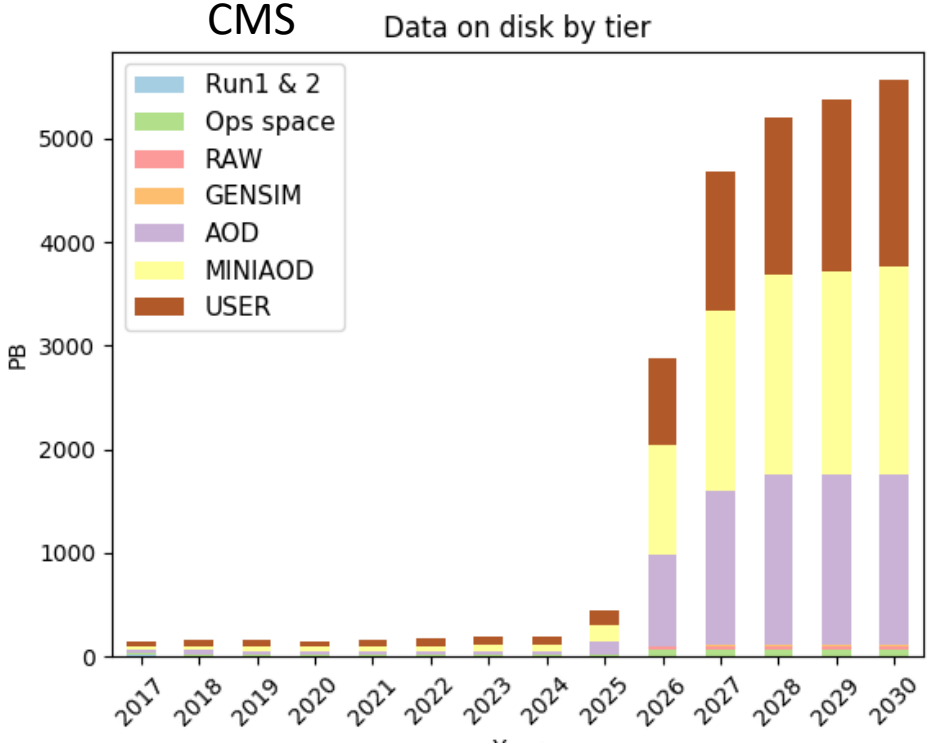

HL-LHC challenge: both rate and complexity rise, pushing resource needs far above « natural » technology growth

A game changer : HEP Software Fundation (HSF)

- First multi-experiment international collaboration around SW development in our communities
  - Bootstrapped in 2016 to address the challenges of HEP software but has always been open to others
  - Well recognized since 2018 after publishing the Community White Paper « A Roadmap for HEP Software and Computing R&D in the 2020s »
  - Sollicited by various HEP bodies to report on SW&C: LHCC, ESPP, Snowmass…
  - Joined by EIC community in 2021
- Main idea: many challenges common to several experiments or communities, not reasonable to reinvent the wheel each time
  - Several tools already widely adopted: GEANT4, ROOT, Rucio, Gaudi…
  - Promote the adoption of existing tools by new/future experiments (LC, FCC…) : Turnkey SW stack
- Name probably misleading: no money to distribute!
  - Derived from Apache Software Fundation (httpd server and many other open-source projects) used as a model which is not funding projects either
  - Organize the collaboration between projects that remain independent of HSF
  - Several topical WGs where are discussed challenges, presented existing solutions and R&D, and designed common projects for future

HSF in a Nutshell

- HSF is an empty shell without projects
  - Projects remain autonomous, have their own governance and funding
  - Allow to build an ecosystem integrating all the projects
  - Bootstrap new projects around common challenges
- Covering a wide range of software, organised in WG
  - Detector simulation: GEANT4, R&D projects (AdePT, Celeritas)…
  - Data analysis: ROOT, Coffea…
  - Frameworks: Gaudi, CMSSW, AliEN…
  - PyHEP, Reconstruction & tiggers, Physics generators, SW Tools and packaging…
- A recent example: Analysis Ecosystem workshop (IJCLab, May 2022)
  - https://indico.cern.ch/event/1125222/: 120 participants (80 in-person), 2nd edition (1st in 2017)
  - Covered several analysis challenges: ML, analysis facilities, metadata & systematics, real-time analysis, declarative languages…
  - Build a vision for the next 5 years: different focus from the last 5 years

Cosmology: LSST data deluge

- Vera Rubin Observatory will scan ½ sky in 4 days, generating 10 TB of raw data every night
  - Real time processing of these data to identify what changed and generate alerts with the relevant information sent downstream to alert brokers
  - 10 millions of alerts (1 TB) per night that must be analysed, classified, according to science drivers (MMA, GW, microlensing…)
- Alert brokers are not a new service but the amount of data to process is order of magnitudes bigger than in previous experiments
  - Not just a matter of scaling up previous brokers: require new approaches both in the processing technology and in algorithms
  - Fink: a French-led broker chosen by LSST in July 2021 to be one of the 7 official brokers that will received and process the full stream of alerts, supporting most of the LSST science drivers. Based on big data technologies (Kafka, Spark) and machine learning.

Multi-Messenger Astronomy

- Extends the alert broker concept to be able to correlate multiple sources of events or to notify one instrument when an event is detected by another one
  - Exemple: look at EM counterpart of gravitational waves with instruments like SVOM → GRANDMA project
  - Need to both analyse alerts in realtime and build an observation plan for associated instruments like telescopes (need to take into account where the telescope is located, when it can observe…)
- An area for collaboration between experiments around scalable tools but no formal collaboration to do it
  - Link between projects currently done by people/laboratories involved in several ones. Example: IJCLab with Fink and GRANDMA
  - Trying to leverage expertise developed in the different projects and to build new features on existing platforms, like the next generation GRANDMA that will be built on Berkeley astronomical data platform SkyPortal

Cherenkov Astronomy

- Another area where consolidation and convergence on common solutions started because of the challenges posed by the next generation experiment, CTA
- Data management tools and computing infrastructure services shared with HEP
  - DIRAC, Rucio…
- Significant effort to promote common solutions in the community with the ASTERICS and H2020 ESCAPE projects
- A common framework, Gammapy, adopted by CTA in 2021 as the official science tool library
  - Already adopted/recommended by most of the existing Cherenkov telescopes (H.E.S.S., VERITAS, MAGIC, HAWC)
  - Modern, Python-based, library developed in the last 5 years
  - Based on widely used and very efficient scientific packages from the Python ecosystems: Numpy, Scikit, Astropy…
  - Community-developed framework aggregating contributions from many different parties using the now standard software development tools and workflows (SW forge, integration tests…)

Nuclear physics: the start

- Up to now, nuclear physics has been dominated by independent experiments that tended to develop their own specific tools
  - Still using a few well established frameworks like GEANT4 or ROOT but mainly consuming them…
  - Started to change: EIC in the USA, Eurolabs European project
- EIC: after discussions in the 2020, joined the HSF in 2021
  - SW co-developed by JLab and BNL: BNL is well inserted in the HSF, some members among the founders
  - Several EIC members are co-conveners of HSF WGs (e.g. event generators, reconstruction and SW triggers)
- EURO-LABS: EUROpean Laboratories for Accelerator-Based Sciences
  - More than nuclear physics but among the explicit goals: implementation of good practices for data management and activities relating to targeted service improvement at RIs, build the foundations to create synergies and collaborations between the RIs of the Nuclear and High Energy communities
  - ICJLab involved in this project as one of the possible bridges with HEP as far as computing and SW are concerned

Publishing SW & Computing R&D and work

- Sharing ideas and identifying possible collaborations require a well-known place where the existing work can be published and archived
  - Used to be missing for software and computing work which is neither physics work nor computing research
  - General approach was to embed the computing work description in physics-related papers, making difficult to identify the potential for sharing or generalization
- Changed in 2017 with the creation of the Computing and Software for Big Science journal
  - Partnership with Springer, open-access policy promoted
  - IN2P3 well involved in the journal: one of the editor-in-chief (V. Beckmann), several associate editors
  - Already well indexed by most of the main journal indexes and with a high citation score… but not so many publications involving French authors!

Career Recognition

- Computing and software challenges faced by our disciplines require a collaboration between domain scientists and computing experts
  - A key role is played by people that have one foot on each side: Research Software Engineer (RSE)
  - they have been instrumental in the recent evolution of many frameworks
  - Unfortunately career recognition for people with these profiles are often problematic
- CNRS distinction between researchers and engineers is not helping
  - Such a profile has almost no chance to get a researcher position
  - Often not seen as expert enough in computing to get a Research Engineer (IR) position
  - If he/she gets a Research Engineer (IR) position, will probably loose over time its domain expertise
- This requires some attention to ensure that we have enough RSE to play a long-term role in software collaborations
  - Currently this role is often filled by postdocs: an expertise which is lost if the persons don’t get permanent positions

Training

- Another key activity for allowing a large contribution to the various tools
- Not only for young people and beginners: technology is evolving rapidly
  - New languages: even a well-established language like C++ is evolving a lot at each new version
  - New paradygms and frameworks
  - Ecosystems: modern software is relying on rich ecosystems
  - Packaging and build tools: a key component to allow interoperability and reuse
- A model emerged, community-driven: Software Carpentry
  - Workshop with progressive exercises in terms of complexity done by students helped by mentors
  - Courses freely available as Git repositories: can be used outside courses
  - Model successfully implemented by the HSF on several topics: advanced C++, Python, SW development tools… more coming
  - Requires involvement by community experts to build the courses and by community members with some experience to act as mentors: also a potential question of recognition

National/topical projects

- Existing software collaborations are mostly unfunded
  - Volunteer efforts by people convinced that we need to promote common solutions
  - Some projects are recognized enough to allow some people dedicated to them (e.g. GEANT4)
  - Good starting point but not enough to address all the challenges we are facing
- Model that seems to work well is national/topical projects collaborating with the software collaborations and focused on some specific issues
  - Successful example: IRIS-HEP (USA/NSF, https://iris-hep.org) hired ~10 people working on various aspects of the HSF program, especially on analysis tools
  - UK has a similar program being finalized
  - Europe tends to be more difficult to support our SW collaborations: any funded work has to be part of a program with broad goals and not specific to one (of our) communities. NSF model seems the opposite.
  - France: ANR would be the main source of funding, not clear how to get support from it as it is too much an engineering (as opposed to research) work

SIDIS

- Software Institute for Data-Intensive Sciences
  - Permanent structure to organize/fund work on key areas and R&Ds: complement HSF
  - https://sidis.web.cern.ch
- Objectives
  - *Enable concrete R&D work* resulting from the collaboration of computer/data science and natural science.
  - *Establish a career path* for scientists and engineers working in software and computing in natural science.
  - *Gather, curate and disseminate the acquired knowledge* across science domains.
  - *Act as lobbying organisation* and raise awareness for the need of software and computing in natural science
- A project currently in “conceptualization phase” building a network of motivated people, refining the idea and elaborating a possible structure, including the funding to support it
  - Funding: applying to European COST program (failed in 2021, planned again in 2022)
  - Suffered a lot from COVID crisis as the initiative started just before…
  - 2021 achievement: partners from several domains (biology, chemistry…) joined the project no longer HEP

French (IN2P3) Situation

- Several IN2P3 members involved in major software projects
  - Software contributors in many projects
  - Leading position in several of them, e.g. GEANT4, Gammapy, Fink, DIRAC
- A dedicated masterproject on the challenges of new hardware architectures
  - A French project but members connected to other community efforts
  - See David Chamont's talk
- Less involvement in collaboration efforts like HSF
  - Probably not seen as useful as the projects
- Suffer from landscape fragmentation in France
  - HPC / HTC competition: difficult to establish productive collaborations on SW challenges with the HPC world for event-based physics (as opposed to images) compared to some other countries (not only USA but Switzerland or Germany)
- Collaboration with computer/data science improving
  - 3 comanaged PhD at IJCLab, 1 at LUPM
  - Performance portability and precision seems a promising area

(My)Recommandations

- Recognition of the role of software collaborations
  - HSF is probably the most advanced one but not the only model
  - Astrophysics has some strong initiatives/projects: is a more global collaboration feasible?
  - Several computing challenges are cross-cutting issues, not specific to one particular science field or community: exchange of ideas and practices are beneficial
- Promote collaboration with computer scientists and computer engineering
  - SIDIS initiative launched by CERN and a few others to foster collaboration: may want to liaise with them
- Encourage participation to collaborative software efforts, in addition to projects
  - Training activities deserve a significant effort: both seniors and juniors are needed
  - Coordination, e.g. HSF, requires that some persons dedicate a significant fraction of their time to it
- Identify funding opportunities for our challenges rather than address our challenges as by-product of funding opportunities for others
  - Situation seems much better in the USA with NSF (e.g. IRIS-HEP), UK following the same approach