

SnowMass2021



Community Summer Study

SN  WMASS

July 17-26 2022, Seattle

The Future of HEP Software and Computing – the Snowmass report

V. Daniel Elvira (Fermilab) – CHEP 2023

Representing my co-conveners, Ben Nachman (LBNL) and Steven Gottlieb (Indiana University), and all participants

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Computational Frontier (CompF) conveners

Daniel Elvira


(Replaced Oli Gutsche in September 2021)
Fermilab

Steve Gottlieb

Indiana University

Ben Nachman

*Lawrence Berkeley
National Laboratory*



Most of the material referenced in this presentation is available at the Snowmass CompF web page: <https://snowmass21.org/computational/start>

(Links down the tree also provided in successive slides for your convenience)

Topical groups (TG)



CompF01 Experimental Algorithm Parallelization

Giuseppe Cerati (FNAL), Katrin
Heitmann (ANL), Walter Hopkins (ANL)



CompF02 Theory Calculations & Simulation

Peter Boyle (BNL), Kevin Pedro
(FNAL), Ji Qiang (LBNL)



CompF03 Machine Learning

Phiala Shanahan (MIT), Kazu Terao
(SLAC), Daniel Whiteson (Irvine)

**Mixed categorization, topics
and technologies**



CompF04 Storage and Processing Resource Access (Facility and Infrastructure R&D)

Wahid Bhimji (NERSC), Meifeng Lin
(BNL), Frank Würthwein (UCSD)



CompF05 End User Analysis

Gavin Davis (U. Mississippi),
Peter Onyisi (U. Texas at Austin),
Amy Roberts (UC Denver)

Topics: seeking solutions for
specific S&C challenges



CompF06 Quantum Computing

Travis Humble (ORNL), Gabriel Perdue
(FNAL), Martin Savage (U. Washington)

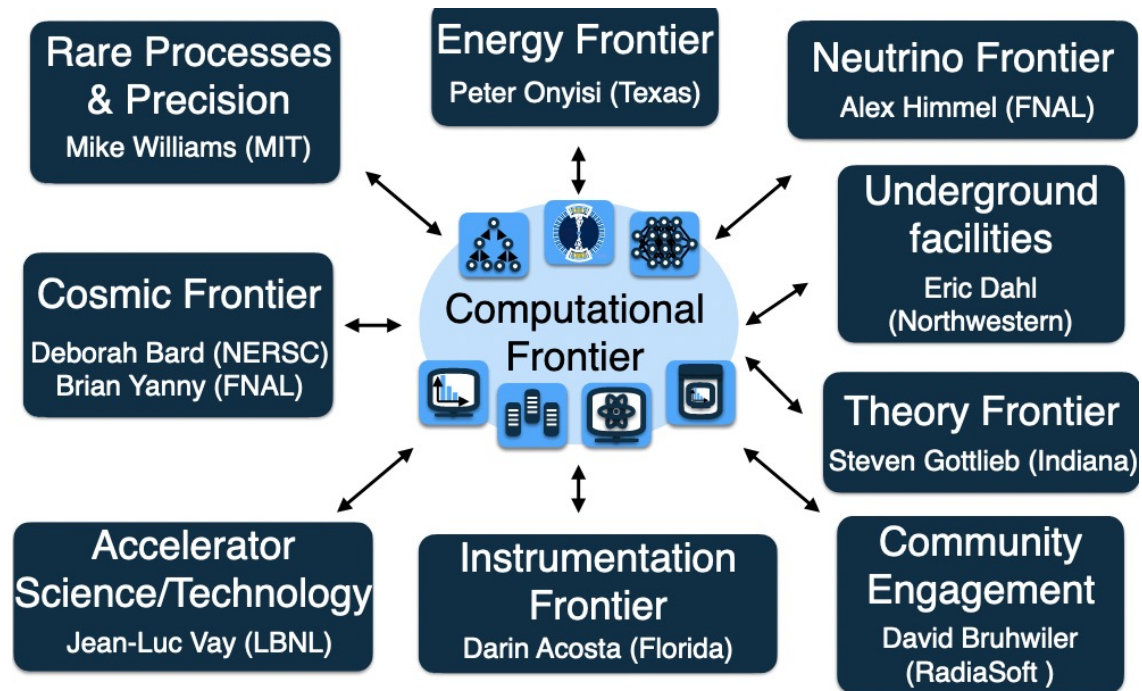


CompF07 Reinterpretation & Long-term Preservation of Data and Code

Stephen Bailey (LBNL), Kyle Cranmer (NYU),
Matias Carrasco Kind (Illinois/NCSA)

Technologies: recognize the
transformative impact of newly
established (Machine Learning)
and emerging (Quantum
Computing) technologies

Liaisons



Cross-frontier
communication a priority

Computing is an enabler of the
HEP physics programs

(Significant contributions to computer
science, shaping technologies and
methodologies while developing
solutions to HEP specific problems.)

Preparatory events and final reports

A series of events to organize the process and discuss progress

- **Biweekly conveners meetings** started in April 2020, a **Computational Frontier Workshop** was held in August 2020, many **TG specific meetings**: <https://indico.fnal.gov/category/1107>
- **Covid pause** from January to August, 2021
- **S&C for small HEP Experiments**: <https://indico.physics.lbl.gov/event/1756>
- **Letters of Intent (LOIs)**: 136 submitted to CompF as primary frontier, 248 in total
- **White Papers (WP)**: 71 submitted to CompF

The Community Summer Study (CSS), Seattle, July 2022

- **CSS was the final community-wide event**: <http://seattlesnowmass2021.net>
 - Sessions: all-CompF, individual TGs, small and large experiments parallels
 - AI, future of computing in HEP, quantum science and technology plenaries, CompF/industry session
- **First draft of CompF report** available for discussions and comments

“The Future of HEP S&C” (final CompF report): <https://arxiv.org/abs/2210.05822>

Outline


- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- **Newly established and emerging technologies**
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Hardware evolution: a brave new world in computing

A paradigm change in computing architecture

- **Dennard Scaling (DS):** power used by silicon device/volume independent on the number of transistors
- **Moore's Law:** transistor density doubles every two years
- **Clock speed (CS):** increased 1,000 times in 1970-2000

Computer speed
doubled every 2 years
while cost halved



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Break down of DS (leakage current), **Moore's Law** (atom sized devices), **CS** (too much power)



Evolution towards heterogeneous systems with multi-core machines using co-processors (e.g., GPUs) and complex memory configurations

AI: from emerging technology to mainstream

Machine Learning (ML) in particle physics

- **Multivariate analysis** commonplace since the 1990's
- **Modern ML** absent from the 2013 report, now making paradigm-shifting contributions to HEP
 - TG charge: understand development and deployment of methods, at large/small scales, online/offline setting

Dedicated HEP-ML Research

- Industry drives but **HEP requires dedicated solutions**
 - Level of precision and use of data-simulation hybrid methods calls for HEP-specific research in uncertainty quantification, validation, interpretability
 - Benefit from physics-aware learning with custom ML architectures
- **E.g., generative models** accelerate various steps in HEP simulation chain
 - Physics generators, detector simulation, lattice gauge theory
- **Unsupervised classification** is used in anomaly detection in searches for new physics
- **Gradient descent** is used in instrumentation to optimize detector designs
- **Ultra low-latency inference** is useful for control in particle accelerators

AI: from emerging technology to mainstream

Software and Hardware Needs

- **SciPy, TensorFlow, PyTorch, etc.**, more popular as (open source) ML platforms within HEP
 - Direct contributions would ensure HEP needs are met
- **GPU resources** are essential: dedicated hardware for prototyping, HPCs for deployment
 - Allocation process amenable to R&D and HEP schedules
- Industry is developing **advanced hardware accelerators** for ML training and inference
 - HEP should partner for R&D on custom solutions (stringent latency requirements, radiation hardness)
- **Cloud and on demand services** may reduce the needs for custom hardware on premises

Training and Personnel

- **ML tools part of standard training program:** faculty, graduate courses, summer schools, material
 - Level of precision and use of data-simulation hybrid methods calls for HEP-specific research in uncertainty
- Training and career paths for **development at the intersection of HEP and ML**
 - PhD physics programs with emphasis on statistical methods and data science
- **Cross-disciplinary collaborations, including industry**, present challenges (recruitment and retention, ethics, external collaborators)

Quantum Computing: a paradigm shift

Quantum computers (QCs) use interference and entanglement in calculations

- States represented with **qubits (bits in a two-dimensional Hilbert space)**
 - Even small QCs (# of qubits) outperform classical computers for certain questions
- **Quantum decoherence** (environmental noise affecting the quantum state) is a challenge

Near-term noisy intermediate-scale quantum (NISQ) era

- Rapid development of **software for QCs**
 - open-source tools to write algorithms and batch system to deploy them
- **Benchmark examples**: lattice gauge theory, event generation, data analysis
- **HEP participation** is essential
 - Develop motivational examples, tailor QC for HEP applications, build critical infrastructure for the future
 - Establish QC-HEP partnerships, including national QIS centers and industry

Develop training programs and a career path – QC expertise in high demand

The field should prepare for the **possibility that fault tolerance is achieved by the next Snowmass process, in about a decade**

Computing in HEP is not business as usual any more

Globalization

- Software and Computing for HEP is **a global endeavor**, advancing science across fields of research and across nations; efforts must be coordinated with partners

Complexity of experimental/observational instruments and size of data volumes

- Demand more **computing power, data storage, computationally expensive algorithms** in a funding context of **flat budgets**

S&C in HEP evolved to become an integral part of the measurement instrument

- **Not a "service" but an "element" of the "scientific apparatus" in HEP experiments/surveys**
- Complexity and physics content/impact of computing commensurable with that of detectors
- Exploit synergy between detector and software design
Co-development, with simulation to design/optimize detectors and detector parameters optimized for best physics and computing performance

Computing in HEP is not business as usual anymore

Hardware evolution calls for a redesign of the computing model for HEP

- **Experiment software frameworks** will need to use heterogeneous resources locally and remotely, including supercomputing centers and commercial facilities
- **Data Management Model** to handle data access, transfer, processing across a diverse set of computing systems
- **Adapt or re-engineer almost every piece of software**, including common software tools for event generation, detector simulation, end-user analysis, reconstruction algorithms
- **Portability tools** to avoid re-writing software for different computing hardware
- **Computing and software infrastructure for AI/ML** training and inference

Departure from stability of the past when the “same old software” would run faster and cheaper in future machines without adaptation or re-engineering

Research - to bridge the needed versus available resource gap and, development - to deliver production level software (heavy in labor - person power with rare and expensive talents)

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

A few notes about the Computational Frontier

Snowmass timeframe: 20-year global vision, 10-year execution plan

- The Snowmass book is mostly about post 2035 facilities
- The near-term programs, before 2035, are not a given, and the Snowmass process made a strong case for many of them
- The Computational Frontier received advice to focus on the upcoming 10-15 years
 - Secure funding to bridge the needs/resources gap of near-future programs
(Rapidly evolving technologies make predictions beyond 2035 very inaccurate)
- The near-term computational needs of feasibility studies and R&D for future facilities and instruments should also be addressed

Distinctive features of computing

- Computing technology and software paradigms change on a much faster scale than the Snowmass frequency, and the life of most experiments and surveys
- S&C efforts are neither funded nor managed as projects, unlike facilities and experimental devices

Main recommendation

We recommend the creation of a standing **Coordinating Panel for Software and Computing (CPSC)** under the auspices of DPF, mirroring the panel for advanced detectors ([CPAD](#)) established in 2012.

Promote, coordinate, and assist the HEP community on Software and Computing, working with scientific collaborations, grassroots organizations, institutes and centers, community leaders, and funding agencies on the evolving HEP Software and Computing needs of experimental, observational, and theoretical aspects of the HEP programs. The scope should include research, development, maintenance, and user support.

(There is also a recommendation for the CPSC to setup a study group on DEI in HEP computing – See backup slide for details.)

In addition we have identified four key areas of need where increased investment would significantly enhance the physics output of the US HEP community:

Findings and recommendations

1. Long-term development, maintenance, and user support of essential software packages cutting across project or discipline boundaries are largely unsupported.

- A new structure is needed to fund modernization, maintenance, and user support of existing tools
 - Grants typically only fund ground-breaking R&D of development of new software
- Examples include:
 - Event generators, and simulation tools such as Geant4, which do not belong to a particular facility, experiment, or survey
 - S&C tools associated to one or more experiments
 - Data and software preservation for utilization after an experiment has ended

Recommendation: the US HEP community should take a leading role in long-term development, maintenance, and user support of essential software packages with targeted investment.

Findings and recommendations

2. Research and development (R&D) for software and computing cutting across project or discipline boundaries receive insufficient support.

- Computational HEP is a vehicle for cross-cutting R&D.
Supporting research in this area at a variety of scales would be broadly impactful
- Examples include S&C for theoretical calculations/generators; cosmological, accelerator, and detector modeling; machine learning methodology and hardware ecosystems; and algorithms and packages across experiment boundaries
 - DOE-CCE, NSF IRIS-HEP, and NSF AI institutes are great examples of programs, institutions, organizations that fund interdisciplinary R&D in these areas
 - The HEP Software Foundation (HSF) is a successful example of a community organization to facilitate cooperation in common efforts internationally

Recommendation: through existing, reshaped and expanded programs, **R&D efforts cutting across project or discipline boundaries should be supported from proof of concept to prototype to production.**

Findings and recommendations

3. Scarcity of personnel and expertise jeopardizes full and optimal use of **heterogeneous and high-performance computing (HPC) resources**

- Most HEP software runs on a single computing platform, making it difficult to use the diversity of hardware accelerators and computing resources, like cloud, HPC, etc.
 - Many efforts within experiments and surveys, as well as dedicated programs: DOE-ECP, HEP-CCE, SciDAC, CompHEP, NSF IRIS-HEP
- To satisfy the needs of inherently serial algorithms that are still transitioning towards computing accelerators or are not cost-effective to port, an appropriate level of traditional CPU-based hardware should coexist with heterogeneous resources

Recommendation: **support computing professionals/researchers, and physicists to conduct code re-engineering and adaptation** to enable use of heterogeneous resources effectively.

Findings and recommendations

4. Investment in **training and career paths** within HEP for S&C researchers is insufficient

- Sustainable efforts in HEP computation require continual **recruitment and training in the context of an environment that is diverse, inclusive, supportive, and welcoming**
- Successful training events have been carried out through HEP experiments, institutes/organizations, and a growing number of university courses
 - Need to grow efforts for documentation and training at multiple levels
- Faculty/staff positions for physicists with expertise in S&C for HEP are scarce and person-power shortfall endemic.
 - Faculty-level appointments in S&C with joint appointments at national laboratories

Recommendation: strong investment in **career development** for HEP S&C researchers to ensure future success.

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Snowmass process preceded by European Strategy Report

Unsurprisingly, the recommendations of the Computational Frontier resonate with the European Strategy, spelled out in their recent report:

[Physics Briefing Book](#): European Strategy for Particle Physics Preparatory Group

Summary of key points on computing:

- It is critically important, more than ever, for the community to maintain a strong focus on instrumentation R&D and to foster an environment that stimulates innovation, with the primary goal of addressing the well-defined technological challenges of future experimental programmes.
- Computing challenges are immediate and need to be addressed now through a vigorous R&D programme. A cross-computing view tensioning new hardware costs with the costs of software development is required.
- It is becoming increasingly vital to take a more holistic approach to detector design which includes impact on computing resources; however, the detector and the computing/software communities have been drifting apart, and individuals that can bridge the growing gap are rare. This is a challenge to the community.
- Both detector and computing development efforts benefit from the existence of networks and consortia, within particle physics as well as extending to other disciplines and with industry. There is a clear need to strengthen existing R&D collaborative structures, and create new ones, to address future experimental challenges of the field post HL-LHC.
- While producing experts who then go on to shape the wider world is a significant benefit to society, a limited amount of success in attracting, developing and retaining instrumentation and computing experts poses a growing risk to the field of particle physics. It is of utmost importance that both instrumentation and computing development activities be recognized correctly as fundamental research areas bearing a large impact on the final physics results.

European initiatives to address challenges

EOSC - European Open Science Cloud

- A cloud for research data in Europe allowing seamless access to data and interoperable services
- Federate existing infrastructures, National data centers and research infrastructures, allowing researchers and citizens to publish, access, re-use data



**EUROPEAN OPEN
SCIENCE CLOUD**

<https://eosc-portal.eu>

ESCAPE - European Science Cluster of Astronomy & Particle Physics Research infrastructures (ESFRI projects: particle and nuclear physics, gravitational waves, etc.)

- Develop common infrastructure solutions for particle physics and astronomy facilities
- Push the state-of-the-art in data mgmt. and computing R&D
- Address common challenges for data preservation, sustainability, open access to data



<https://projectescape.eu>

(Special thanks to Graeme Stewart for pointers to European Strategy material.)

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- **Final thoughts**
- Backup Slides
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Final thoughts

- The Snowmass CompF process was a successful participative experience
 - The conveners hope the report represents the community faithfully
- Computing plays a fundamental role in the execution of current and future HEP theoretical, experimental, observational programs
 - Not a "service" but an "element" of the "scientific apparatus" in HEP experiments/surveys
- Modern computing architectures, newly established and emerging technologies are changing the way we do particle physics but also bringing transitional challenges
 - Balanced and timely support for software, from prototype all the way to deployment
 - Education, training, career paths in an inclusive, diverse and welcoming environment
- Strike a balance between R&D and improvement/maintenance of existing tools
 - Simulation common tools are underfunded
- Coordinating panel (CPSC) would play a critical role in achieving coordinated, timely, balanced, effective, sustainable investments in S&C for HEP
 - The DPF leadership plans to establish a task force with broad community representation to write a report on a proposed CPSC creation process

Backup Slides

Outline

- The Computational Frontier (CompF) within the Snowmass process
 - Organization, preparatory events, the Seattle Community Summer Study (CSS), the report
- Newly established and emerging technologies
 - Hardware, artificial intelligence (AI), quantum computing (QC)
- Findings and recommendations
 - Coordinating Panel on Software and Computing (CPSC)
- The CompF report in the global context
- Final thoughts
- **Backup Slides**
 - More detail on diversity and climate within computing in HEP
 - More detail on challenges and opportunities by S&C topic

Diversity and climate within computing in HEP

The percentage of Underrepresented Minorities (URMs) obtaining a Ph.D. in physics in the US is < 5% (Hispanics) and < 2% (Blacks), the percentage of women is ~20%

- **In computer science** Hispanics earn ~1.5% of Ph.D. degrees and Blacks ~1%
 - Stark contrast with numbers as a fraction of the US population, 19% (Hispanics) and 14% (Blacks)
 - Beyond race and gender, **individuals face adversity based on other identity elements**
- **The situation was not always so dire for women in the case of computer science**
 - Women majoring in computer science were 12% (1970), increasing to 37% (mid eighties), only to plummet and flatten to a 17% (2010)
 - 7% of the Ph.D.s in physics awarded to women (1970), increasing linearly to flatten to a 20% in the early 2000's
- **Several compounding factors may contribute to persistent diversity gaps in computer science**
 - E.g., stereotype threat, implicit and explicit bias in higher education and society, isolation, lack of community
- **Lack of recognition for S&C work** affects all demographics
 - Instrumentation work better rewarded career-wise, may come from a perception of having higher physics content

Recommendation of a study group to explore the challenges in the area of diversity, inclusion and climate that are particular to the computing sub-field within HEP, produce recommendations — in the context of the Coordinating Panel for S&C (CPSC)

Frameworks and experimental algorithms

Experimental algorithms are code to reconstruct physics objects, process observational data, perform their calibration, etc.

- Challenge to **improve physics performance while minimizing computing cost**
 - Algorithm re-design, ML alternatives, parallelization, platform portability, effective use of accelerators and HPCs

Software frameworks need to be adapted to modern computing paradigms

- Support **parallelism** in multi-core CPUs and execution on heterogeneous platforms
- **A common framework and reconstruction tools, a platform for R&D** work including AI/ML, algorithm parallelization and HPC deployment would reduce duplication, benefit (small) experiments

Examples in various HEP sub-fields

- **Hadron colliders:** physics/computing performance in the presence of pileup, increased detector complexity, trajectory reconstruction, faster trigger algorithms, etc.
- **Cosmology (static/time domain science):** collect and analyze big data, identify/classify events quickly
- **Direct dark matter detection:** process large data volumes ~ 1 PB/year
- **Neutrino experiments (LArTPC detectors):** (a) signal processing and (b) high-level algorithms use traditional algorithms combined with deep learning. Challenge: memory/disk for large events

Theoretical calculations and simulation

Next-generation cosmological simulations focus on high-fidelity modeling of observational outputs from multiple surveys

- **N-body simulations** model dynamics of galaxies at different scales, describe dark matter fluctuations. **Hydrodynamical simulations** provide descriptions of distributions of baryons, their effect on probes of large-scale structure, results for distribution and properties of galaxies and clusters
 - Challenges are computational cost of including multiple probes with high precision, computationally expensive hydrodynamical modeling of the distribution of gas in the universe, correlations between multi-wavelength observables in galaxy formation models, modeling of massive neutrinos and supernovae
 - Understanding, mitigation, control of systematic uncertainties using “virtual universes”

Particle accelerator modeling

- **Next-generation accelerator modeling tools** must include all types of accelerator components and address grand challenges on intensity, quality, control, safety, and prediction
 - From source to interaction region, physics models for high precision, collective effects, beam transfers, etc.
- **Accelerator modeling community** pioneer in advanced algorithms, parallelization, utilization of heterogeneous computing and HPCs
 - Support needed for code portability, improvement of workload balance and parallel efficiency, optimal use of supercomputing facilities, integration of AI/ML to physics modeling, control systems, uncertainty quantification

Theoretical calculations and simulation

Physics generators

- **Many generator packages** requiring solutions for matrix element calculation, hadronization and parton showering, underlying event modeling, parton matching/merging, decays, cross sections
 - Negative weights, inefficiencies in phase space sampling
- **Increased use of NLO calculations** puts pressure on computing resources (up to 20% for HL-LHC). Neutrino experiments require similar percentual increases in computing for generators
- **Demand for coordinated activity (e.g., MCNet)** to face the challenges of adapting to new computing platforms and integrating AI solutions

Detector simulation

- Experiments demand **increased physics fidelity and significantly larger simulated data volumes**
 - Already consume a large fraction of HEP computing resources (>50% for LHC experiments in Run 1,2)
 - R&D for improved physics, geometry, navigation algorithms, AI/ML options
- **Geant4 primary tool** for detector modeling, from instrument design to data analysis
 - Resources for Geant4 physics models, software maintenance and support severely reduced in the USA
- **Geant4 adaptation/re-engineering for GPUs and HPCs** is underway
 - Extremely challenging, simulation is intrinsically a High Throughput Computing (HTC) problem
 - Person-power intensive, requiring scarce and expensive software skills

Theoretical calculations and simulation

Theory calculations

- **High-precision perturbative calculations** have become possible with modern computing systems
 - Requirements for electroweak processes are moving from percent to per-mille accuracy
 - QCD calculations at N³LO for 2-to-1, NNLO for 2-to-3 and 2-to-2 (with loops) are now state-of-the-art
- **Conformal bootstrap to solve conformal field theories**, is a non-perturbative approach
- **Significant computing challenges** arising from large memory/core requirement and CPU reliance
 - Perturbative calculations have sizable barrier to use on GPUs, conformal bootstrap may be amenable to GPUs
- **Machine learning** still not playing a large role

Lattice QCD

- **Numerical evaluation of Feynman-path integrals** for first-principles predictions of QFTs
- **Worldwide lattice gauge theory efforts** to calculate properties of hadrons – vital for many experiments
 - Muon magnetic moment ($g-2$), nucleon/parton structure, neutron electric dipole, form factors in ν -nucleus, etc.
- **Finer lattice spacing and larger volumes** requires algorithmic research, computing beyond exascale, improved software engineering
 - Massive vector parallelism amenable to GPUs and the use of ML techniques

Storage and processing resource access

Diversity of resources across technologies/countries is a challenge – R&D

- **Storage** needs are vast, major concern in the context of current budget projections
 - R&D on data formats more efficient for modern processors, explore data compression
 - Tape continues as the archival medium, spinning disk (favored by cost) and solid state will co-exist
- **Data processing** would benefit from data management framework across different types of facilities
 - More flexible access policies at HPCs, R&D for efficient utilization of diverse resources
 - Evaluate trade off between HEP computing facilities, HPC centers, Grid, Cloud resources
- **Analysis facilities** with infrastructure and services to execute analysis workflows
 - Prototype analysis facilities, analysis benchmark to demonstrate use of facility at scale of different analyses,
 - Integrate AI/ML tools and security in distributed computing environments
- **AI/ML hardware** directly connected to (local) computer systems or as a service
 - Development of a specialized ML benchmark suite focused on HEP to select best hardware for each need
- **Network traffic for HEP** expected to grow by a factor of 10 by the end of the decade
 - Issues: network interaction optimization, resource orchestration/automation, traffic visibility, data movement optimization

End-user analysis

Run analysis code on full datasets for timely delivery of physics results

- **Big HEP data** demand novel data format and access tools, as well as sufficient computing infrastructure
 - Software ecosystems for fast, easy, user-friendly access
- **Scaling of paramount importance** to cover small tests (iterative) and full datasets (distributed execution) infrastructure
- **ROOT ecosystem:** set of libraries covering I/O, event loop execution, histogramming, statistical analysis, and visualization. It is hosted by CERN with a leading contribution from the US to I/O
 - Using alternative software for specific functionalities is difficult (tightly-bound system)
 - Planned evolution to columnar data format (RNTuple), mapping to vectorized/parallel hardware for speed
- **Python ecosystem:** set of tools, with Python as the primary language interface, to enable the use of software developed outside of HEP, including for machine learning
 - Coffea-casa is an analysis system prototype providing services for low latency columnar analysis
 - Interactive experience, scales to full datasets, exploitation of coprocessors and accelerators
- **Collaborative software** including messaging between users, discussion forums, software version control, bug tracking, document workflow management

Reinterpretation and long-term preservation

Preserving data, simulations, analyses and the codes can increase scientific output from the investments made in the original inquiry. Requires resources and effort!

- **All running and in-preparation experiments and surveys** must have a strategy and resources for long-term preservation of data and analysis capabilities beyond their lifetime
- **Shared cyberinfrastructure** to preserve data and perform comprehensive analysis
- Setup a **facility to preserve Cosmic Frontier datasets and simulations**
 - To facilitate joint analysis across different computing centers