

SLAC Mass Computing

Status Report - October 1st, 2020

Conveners

Cartaro, Kagan, Terao



U.S. DEPARTMENT OF
ENERGY

Stanford
University

SLAC NATIONAL
ACCELERATOR
LABORATORY

- SNOWMASS [Computational Frontier \(CF\)](#)
 - **Goal:** The CF will assess the software and computing needs of the High Energy Physics community emphasizing common needs and common solutions across the frontiers
 - Working groups in the CF:
 - [Experimental Algorithm Parallelization](#)
 - [Theoretical Calculations and Simulation](#)
 - [Machine Learning](#)
 - [Storage and processing resource access \(Facility and Infrastructure R&D\)](#)
 - [End user analysis](#)
 - [Quantum computing](#)
 - [Reinterpretation and long-term preservation of data and code](#)
- SLACmass Computing looked at all proposed aspects of computing, how they are viewed by the SLAC community, what are the needs, the challenges, and the possible solutions

- Computational Physics and Algorithms
- Machine and Deep Learning
- Statistical Techniques
- Distributed / Grid computing
- Large Scale High Performance Computing
- Networking
- Computing/software infrastructure and Storage systems
- Quantum Computing

- Survey sent to SLAC community focused on questions about future computing, HEP, and SLAC's role
- Followed survey with conversations with individuals who responded to survey to get a better picture of their thoughts
- Held additional conversations with computing / ML experts at SLAC
- SLACmass computer workshop ([Indico](#))
 - Focus on future computing needs at SLAC / in HEP

Feedback

- **The vast majority of feedback from survey and conversations focused on access to baseline computing resources, and future computing needs**
- **SLAC HEP groups share a strong culture of collaboration**
 - Shared computing (cores, storage, ...) resources
 - Easier to manage for computing division, possibility to better leverage existing resources
 - Huge advantage for smaller groups and individuals with limited resources
- **Continued need for local computing resources which offer more flexibility**
 - NERSC or other large facilities are not always the answer
 - Changes usually too aggressive for experiments lifetimes
 - Not ideal for certain tasks (online monitoring/processing)
 - Cloud is still expensive and science is not a priority in business models
 - Outsourcing from Centers of excellence only for small scale, limited scope projects (with uncertain results)
 - Access: should there be individual-level access for specialized co-processors? Should experiments have default allocation of HPC hours?
- **How to support both large & small experiments and interact w/ computing division**

- What is the expectation of technologies at SLAC in the near future, and how should we interface with the community to address our needs?
 - SDF (SLAC Data Facility) and new baseline for computing resources
 - Distributed file systems
 - Modern CPUs and co-processors (GPUs)
 - Is this (local cluster supported by the lab/PI's funding) the right model for SLAC?
 - How can our experience building this system inform the broader HEP community as they build similar systems?

- Several members of the community identified the need for training in software engineering for physicists / staff
- Face current challenges with new trends and new technologies
 - How to develop optimized code, and efficiently use computing resources
 - How to make use of modern workflow tools, and to develop ready-to-use workflows (event processing or infrastructure for experimental setups)
 - How to take full advantage of ML → training our community in modern ML methods
 - How to make use GPU resources
 - Even if not for ML, can current analysis strategies benefit?
 - How support/promote R&D activities on new computing technologies

Topics of Interest

ML in HEP has diverse set of physics drivers, but commonality exists for ML-HEP at SLAC

- Rare event searches, parameter estimation in high statistics samples, reconstruction in heterogeneous detectors, fast simulation for data generation or optimization, ...

Physics informed ML

- Adding physical constraints / knowledge in ML systems and pipelines
- Differentiable simulators, physics dynamics, symmetries, data informed kernels, ...

Uncertainty quantification

- Beyond prediction, how to get reliable and calibrated confidence estimates
- How to incorporate systematic uncertainties, and Simulator mismodeling in ML

Surrogate modeling / Fast simulation

- ML model to approximate simulation of physical system
- Generating data for analysis, design optimization, tuning..

System design and optimization

- ML optimization techniques provide new tools, combined with surrogate modeling
- Important for designing experiments, control and automation, simulator tuning...

Computer Vision

- Reconstructing / analyzing image like data, adaptable to many experiments

Graph Neural Networks and analyzing heterogeneous detector data

Simulators are a key driver to perform the physics we want to do - we can not do our science without precision simulators (at least in the way we do science now)

Across HEP, we rely heavily on simulators for experiment design and control, predicting outcomes of experiments, statistical analysis, systematic uncertainty estimations...

Large interest from SLAC community in developing and improving simulators

- Implementing new simulations using existing tools, for new purposes in Neutrino physics, DM searches, accelerator modeling, colliders, ...
- Improve simulator speed and efficacy - take advantage of modern SW parallelism and modern hardware such as GPU / TPU / FPGA and facilities like HPCs
- Increasing simulator capabilities, with tools such as automatic differentiation

Interests/Needs around “data”

- Long term preservation for experimental data
 - Findable, Accessible, Interoperable, Reusable (FAIR) data
- Public dataset for cross-cut development and collaboration
 - Policy, technical designs, interface to cross-cut community
- Data storage and file format that suits high performance/throughput computing

Software landscape

- How to bridge experimental software in HEP (~a decade scale) and modern computing (specialized at hardware/firmware level, shorter life cycle) efficiently?
- A systematic, community-wide approach for education about modern software concepts/tools and seeding of cross-cut development
 - Not only for young folks! Target everyone (your boss struggling to learn Python and ML)

SNOWMASS Computing LOI's with SLAC Members

SNOWMASS Computing LOI's with SLAC Members (1)

Large interest and focus on developing simulators in many subfields

1. [End-to-End Virtual Accelerators \(EVA\)](#) - Zhirong Huang, Cho-Kuen Ng, Adi Hanuka, Auralee Edelen, et al.
2. [Center\(s\) for Accelerator and Beam Physics Modeling](#) - Cho-Kuen Ng, et. al.
3. [Interdisciplinary simulations: Integrating accelerator RF and particle-matter interaction codes](#) - Cho-Kuen Ng, Zenghai Li, et. al.
4. [Modeling of structured plasmas for next generation accelerators](#) - Petros Tzeferacos, et. al.
5. [Modeling Needs for Structure Wakefield Accelerators](#) - Cho-Kuen Ng, et. al.
6. [Neutrino Event Generators](#) - A. M. Ankowski, Y.-T. Tsai, et. al.
7. [Simulations of Low-Energy Crystal Physics for Dark Matter Detectors](#) - Michael Kelsey, et. al.
8. [Computing Challenges for Event Generators](#) - F. Kling, et. al.
9. [Surface Methods for Precision Accelerator Design and Virtual Prototyping of Accelerator Systems](#) - Cho-Kuen Ng, et. al.
10. [Fast Simulations for Noble Liquid Experiments](#) - Maria Elena Monzani, et. al.

SNOWMASS Computing LOI's with SLAC Members (1)



Improving software tools, ecosystems, access, data analysis frameworks, and push for open software and science

1. [*A modular community ecosystem for multiphysics particle accelerator modeling and design*](#) - Cho-Kuen Ng, Spencer Gessner, et. al.
2. [*Embracing modern software tools and user-friendly practices, when distributing scientific codes*](#) - Adi Hanuka, Auralee Edelen, et al.
3. [*Develop/integrate data standards & start-to-end workflows for Accelerator Physics*](#) - C. Mayes, S. J. Gessner, C. H. Yoon, A. L. Edelen, et. al.
4. [*Aspiration for Open Science in Accelerator & Beam Physics Modeling*](#) - C. Mayes, S. J. Gessner, A. L. Edelen, et. al.
5. [*Link the Future And Past User Data Analysis with Jupyter and Xcache*](#) - Wei Yang, et. al.
6. [*Facilitating cosmic probes and dark matter searches with improved data access and software tools for multi-wavelength and multi-messenger analyses.*](#) - E. Charles
7. [*Jas4pp - a Data-Analysis Framework for Physics and Detector Studies*](#) - N. A. Graf, et. al.

SNOWMASS Computing LOI's with SLAC Members (2)

Drive for ML methods - with overlaps in our drive for better simulation and software

1. [*Machine learning and surrogate models for simulation-based optimization of accelerator design*](#) - Adi Hanuka, Auralee Edelen, Claudio Emma, Spencer Gessner, et. al.
2. [*Graph Data Structures and Graph Neural Networks for High Energy Physics*](#) - Kazuhiro Terao, et. al.
3. [*Machine Learning Techniques and Software for Neutrino Physics*](#) - Kazuhiro Terao, et. al.
4. [*Scalable, End-to-End Optimizable Data Reconstruction and Physics Inference Techniques for Large-scale Particle Imaging Neutrino Detectors*](#) - Kazuhiro Terao, et. al.
5. [*Differentiable Simulators for HEP*](#) - Michael Kagan, Kazuhiro Terao, et. al.
6. [*The Future of Machine Learning in Rare Event Searches*](#) - Maria Elena Monzani, et. al.
7. [*Differentiable Programming in High-Energy Physics*](#) - Jannicke Pearkes, et. al.

Quantum Computing

1. [*Front-form calculations on near-term and far-future quantum computers*](#) - Stanley J. Brodsky, et. al.

Some insights from LOIs

- Largest number of SLAC LOI contributions were targeting simulator development
 - We are a simulation driven science - developing and improving simulators is key, and needs to be addressed across all physics goals
- Benefit from acceleration / parallelization with modern co-processors (e.g. GPU, etc.)
 - ML on FPGA for very low latency applications
 - GPU for particle tracking (e.g. ray tracing) and more generally any parallelizable algorithms
- Develop and provide open and common software/data ecosystems with modern SW tools
 - Not only for increasing efficiency/productivity, but also for reproducibility and transparency of science
 - Especially, but not only, in fields where software tools are developed by many groups independently
- Machine learning goals are varied, but find commonality:
 - Utilize simulator prediction within ML pipelines, not just for generating training datasets
 - Finding mechanism to incorporate physics and detector knowledge in ML model design
 - Provide fast and differentiable simulation and prediction engines
 - Use ML to help design, control, and tune experiments and simulators

Snowmass Whitepaper Topics We Would Like to See

Snowmass Whitepaper Topics We Would Like to See

Accelerated Simulation Tools

- Benefitting from specialized co-processor (e.g. GPU, etc.) capabilities for future simulation tools including event generators, particle tracking, and detector modeling

Next generation analysis ecosystems

- Capable of benefiting from GPU / TPU capabilities for future generators
- Seamlessly integrated with ML and Automatic Differentiation tools

Public dataset and software distribution for open, cross-cut development ecosystem

- Dataset with defined scientific goals and metrics, well designed structure and necessary software stack for distribution and access to catalyze cross-cut R&D

Data v.s. simulation domain discrepancy mitigation techniques

- In-line with R&D of simulators that can be automatically optimized against real data
- ML methods to make data reconstruction and analysis robust against mismodeling

Snowmass Whitepaper Topics We Would Like to See (II)

Simulator Driven ML

- Bringing physics domain knowledge into ML directly with simulators
- Differentiable simulators in ML pipelines capable of using these tools

ML Driven optimization, control, design, and automation for experiments

- The learning and optimization capability of ML can drive the next generation of designing and operating experiments
- Key interplay between the optimization method and the simulator capability

End-to-end ML at HEP experiment scale

- Design of ML models for the full data reconstruction chain with inductive bias and causal structures to address dependencies and interplays across different stages
- Demonstrate an end-to-end optimization capability for the reconstruction and analysis chain and reduce months/years of “by-hand/work-hard” tuning need

SLACMass Takeaways

- SLAC groups need baseline resources
 - SDF is a step in right direction technologically!
 - More support is needed for the baseline resource (=available to everyone by default)
 - Improve policy to accommodate needs of small projects (i.e. “collaborative” culture)
- Wide interest in new technologies and methodologies
 - AI/ML, simulation, optimization of workflows, software management
- Recognized critical need for lab-wide training
 - Programming language, software design & distribution, infrastructure-specific workflow

Toward Snowmass

- Many LOIs from SLAC is a sign of thriving interest in the Computing Frontier
 - Meet among members of SLAC with similar interests! ... and even write a white paper together :)
 - We, at SLAC, need to build strong foundations (training and baseline resources) and develop sound strategies (planning) in order to face the **HEP community challenges** of the next years