BROOKHAVEN NATIONAL LABORATORY
PROPOSAL INFORMATION QUESTIONNAIRE
LABORATORY DIRECTED RESEARCH AND DEVELOPMENT PROGRAM

PRINCIPAL INVESTIGATOR	**Torre Wenaus, Meifeng Lin**	**PHONE**	**x4755, x4379**
DEPARTMENT/DIVISION	**PO/CC - NPP/CSI**	**DATE**	**06/16/2021**
OTHER INVESTIGATORS	**Brett Viren (PO)** *Collaborator:* **Elke C. Aschenauer (EIC)**		

TITLE OF PROPOSAL **TYPE A**	**Towards a Scalable and Distributed Machine Learning Service for Data-Intensive Applications**		
PROPOSAL TERM (month/year)	From **06/01/2022**	Through **05/31/2025**	

SUMMARY OF PROPOSAL

Description of Project: Machine Learning (ML) plays an increasing role in the scientific research areas critical to the missions of Department of Energy. While new ML algorithms are being developed at a fast rate, deploying them in large-scale data-intensive scientific applications for production will still require significant improvement to the existing software infrastructure. There is an urgent need for an ML software service infrastructure that can assist the users in utilizing the diverse computing resources efficiently for a better integration with the scientific workflows. We propose to combine Computational Science Initiative (CSI)'s expertise in high performance computing (HPC) and Machine Learning (ML), and Physics Department Nuclear and Particle Physics Software (NPPS) group's expertise in providing large-scale distributed computing services, to investigate and develop a scalable hardware-and-application-agnostic distributed ML service environment. To achieve this, key issues related to portability (to increase accessible re-sources), scalability (to ensure high performance) and agility (to reduce deployment latency) will be systematically studied. Issues related to cyber security such as user authentication will be addressed. To accomplish this, we will develop an integrated software and hardware test bed based on the Institutional Cluster resources available at Brookhaven National Laboratory. Throughout the project, we will use ATLAS and DUNE ML applications as driving examples, such as the ATLAS FastCaloGAN training mentioned previously. Performance benchmarks (in terms of time to solution, coding efforts and other metrics) will be documented periodically to track the progress of the project.

Expected Results: The outcome of the project will be the evaluation of different approaches for user interface, intermediate provisioning and backend execution along with a set of prototype tools that can be used to guide the design and development of a production-level ML service framework in the future. This will be an enabler for expanding scientific creativity in AI/ML applications towards models of a scale and complexity otherwise impractical, or with prohibitive development/training time requirements. The project aligns with and is an enabler for BNL's strategic growth plan to use AI/ML to create solutions to experiment-driven computing challenges across the Laboratory, leveraging BNL and wider computing resources to do so.

PROPOSAL

Abstract

We propose to combine Computational Science Initiative (CSI)'s expertise in high performance computing (HPC) and Machine Learning (ML), and Physics Department Nuclear and Particle Physics Software (NPPS) group's expertise in providing large-scale distributed computing services, to investigate and develop a scalable hardware-and-application-agnostic distributed ML service environment. To achieve this, key issues related to portability (to increase accessible resources), scalability (to ensure high performance) and agility (to reduce deployment latency) will need to be systematically studied, and are the key research questions of our proposal. Containerization and portable ML frameworks will be investigated as potential portability solutions, and performance analysis and modeling of the application behaviors will be studied to help intelligently select the best platform to execute the ML workflows in a scalable manner. Lastly workflow management systems such as PanDA will be investigated for fast deployment to the target large-scale distributed systems. Issues related to cyber security such as user authentication will also need to be addressed. To accomplish this, we will set up an integrated software and hardware test bed based on the Institutional Cluster resources available at Brookhaven National Laboratory (BNL). The success of the proposed research will pave the way for the development of a scalable distributed ML service to provide large-scale processing resources to ML applications across the BNL science programs and beyond.

1 Introduction

We propose to combine Computational Science Initiative (CSI)'s expertise in high performance computing (HPC) and Machine Learning (ML), and Physics Department Nuclear and Particle Physics Software (NPPS) group's expertise in providing large-scale distributed computing services, to investigate and develop a scalable hardware-and-application-agnostic distributed ML service environment. While new ML algorithms are being developed at a fast rate, deploying them in large-scale scientific applications for production will still require significant improvement to the existing software infrastructure. This is especially crucial when it comes to data-intensive applications where training and hyperparameter optimization often take a long time, and require large-scale computational resources. As ML plays an increasing role in the scientific research areas critical to the missions of Department of Energy, there is an urgent need for an ML software service infrastructure that can assist the users in utilizing the diverse computing resources efficiently for a better integration with the scientific workflows. This service should aim to take all the manual optimizations and orchestration out of the users' hands, ideally through one coherent suite of services and scientist-friendly interfaces. Such a need is outlined in the Artificial Intelligence (AI) for Science Townhall Meeting report [1], where challenges and research needed are described. Our research will specifically address the challenge of "a software stack that facilitates efficient use of a broader range of algorithmic and mathematical techniques." and "Simultaneously, given the current trends, data sizes, and the role of hardware in ML, it is reasonable to expect the evolved software stack to take even greater advantage of the new hardware accelerators, as well as to target distributed computing architectures". This project is also in line with BNL's strategy to employ ML and AI for accelerated scientific discovery.

To achieve this, key issues related to portability (to increase accessible resources), scalability (to ensure high performance) and agility (to reduce deployment latency) will need to be systematically studied, and are the key research questions of our proposal. They thus constitute the three main tasks of the proposed project:

- **Portability:** There are two components to portability in this context: software environment portability and hardware portability. The former may be addressed by containerization, while the latter may be embedded in the ML frameworks under consideration (such as GPU/TPU support in TensorFlow). How to efficiently marry the two will be a topic for investigation, as well other potential portability solutions.
- **Scalability:** Scalability of ML training can sometimes be inherent to the algorithms, hardware platform and training data used. Here we will investigate ways to develop a tool that can intelligently select best computational resources and ML frameworks based on analysis and modeling of the application behaviors.
- **Agility:** We will explore ways to efficiently choose and utilize computational resources that may be available at large data centers, on the grid or cloud. Workflow management systems such as the ATLAS PanDA/iDDS [3] workload management system and other similar tools will be studied for efficient work-dispatching in real time.

Since most of the ML workloads will likely be deployed on a remote system, issues related to cyber security such as user authentication will also need to be addressed. To accomplish this, we will develop an integrated software and hardware test bed based on the Institutional Cluster resources available at Brookhaven National Laboratory (BNL). The success of the proposed research will pave the way for the development of a scalable distributed ML service to provide large-scale processing resources to ML applications across the BNL science programs and beyond.

2 Science Drivers and Current State of Practice

Machine Learning has played an increasingly important role in large-scale data analysis and simulations at the experimental facilities such as the Large Hadron Collider at CERN [2, 3] and National Synchrotron Light Source II [4]. For the next-generation experiments such as the Electron Ion Collider (EIC) at BNL, it is anticipated that ML will play an even more important role than currently plausible given the rapid advances in this field. From the EIC detector design to the accurate reconstruction of parton distribution functions, ML can significantly accelerate the pace of scientific research in the EIC era. The upcoming Deep Underground Neutrino Experiment is also investing heavily in ML-assisted simulations and data analysis, such as track reconstruction and signal processing [5]. Many recent ML-related efforts in High Energy Physics were reported at the recent international conference on Computing in High Energy Physics [6], where challenges and progress related to distributed machine learning and deep learning were discussed. One common message is that distributed training on GPUs can greatly reduce the training time without significant loss of accuracy. However, different ML applications tend to customize their own optimization strategies, owing to the proliferation of ML software and frameworks available. These frameworks include scikit-learn, TensorFlow, PyTorch, and Keras, with new software and frameworks being released at a fast rate. These manual one-off optimizations are labor-intensive and hard to scale to the vast landscape of the potential ML workflows in the High Energy Physics, Nuclear Physics and other scientific communities.

Here the challenges associated with the manual optimization approach are further explained. The ML community sees many new topologies published each year, with an increasing number of contributions from the domain scientists. These topologies typically build on existing ML frameworks [7] such as TensorFlow, PyTorch etc. However, achieving high performance on each new topology remains challenging, as each requires some level of deep changes within each framework to improve the training performance for each target hardware backend, which can be CPU, GPU, or TPU-like accelerators. This problem is compounded by the need for large-scale training and optimization in data-intensive applications. While several distributed machine learning frameworks exist, such as Apache Hadoop, Spark, Hovorod etc., their suitability for data-intensive scientific workflows varies case by case. And it often falls on the scientists to choose and fine tune these tools manually. Such diversity and complexities of the ML software ecosystem limit scientists' ability to efficiently apply ML in their workflows, hindering the overall scientific progress. For example, the ATLAS FastCaloGAN [2] requires 100 GPU-days to do just one training pass on the full detector. Coupled with the expected order of magnitude higher computational demand of the next-generation physics experiments such as High-Luminosity LHC, DUNE and EIC, the need for a portable and scalable ML service environment is even more urgent.

While ML services aiming to abstract hardware details from users exist for commercial clouds, such as Microsoft Azure ML service, such a service environment is lacking for the more traditional HPC setting. One of the challenges is in the interaction between the end users and the remote systems due to issues such as user authentication. Recently a new class of serverless computing such as Function as a Service (FaaS) has gained significant attention of the scientific community as a way for the users to write light-weight functions to be executed in a cloud computing-like environment. For example, funcX [8] is a high-performance FaaS system designed to orchestrate scientific workloads across heterogeneous computing resources, from laptops and local servers to campus clusters, cloud computers, and even supercomputing facilities. Users first set up funcX endpoints, a persistent service launched on the remote compute nodes that serves as a gateway for routing functions to and executing them on the nodes. The authentication is enabled through Globus to ease the administration overhead. Then, users register the workload Python function to be run on the endpoints with funcX Client, which returns a universally unique identifier (UUID), a 128-bit label used for information in computer systems. By specifying the UUIDs representing the workload and the target endpoint, funcX can then run the Python function remotely and asynchronously and retrieve the result. To our knowledge, such a service does not exist for ML workloads running on HPC centers. But similar funcationalities and services may be just what our proposed ML service framework needs.

3 Proposed Approach

We envision the scalable distributed ML service framework will be able to provide end-to-end services including data ingestion, data preparation, model buidling and training and model deployment, as illustrated in Figure 1. To realize this, as mentioned in Section 1, challenges related to portability, scalability, agility and cyber security need to be addressed. We propose to set up an integrated software and hardware test bed through the Institutional Cluster resources at BNL, where different CPU and GPU architectures are available for the study of portability, interconnected compute nodes are available for the scalability investigation, and we will work with them to develop and test user authentication methods for the ML service.

There are three key components to the ML service framework: user interface, intermediate provisioning and backend execution, which we discuss in more detail below:

- **User Interface (UI).** To make it easy for the potential users to use the ML service, an intu-

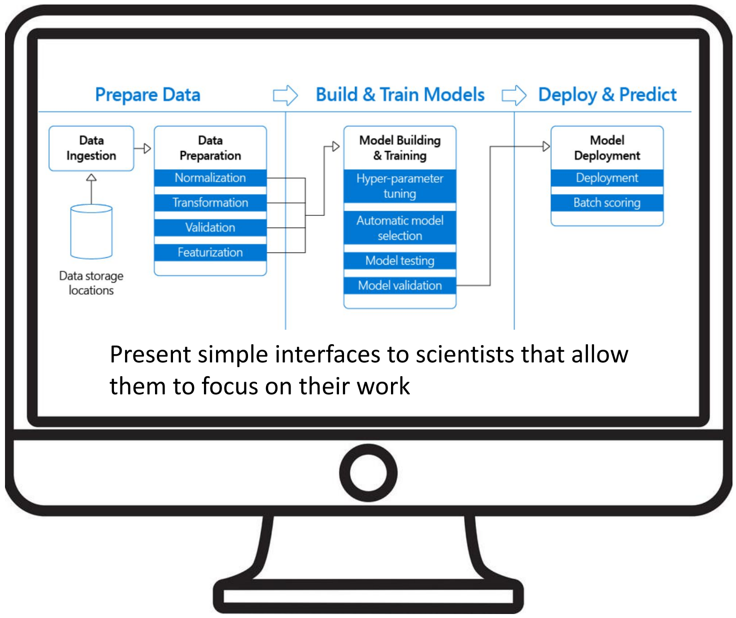

Figure 1: Illustration of the proposed ML service framework.

itive and fast user interface is essential. This is where users can register the ML workloads/functions that need to be processed. There are several candidates for the user interface: text-based, graphics-based and web-based. The text-based UI can be a set of scripts that the users modify. The graphics-based UI will likely not work well in this situation and will not be considered a top choice. Web-based UI such as Jupyter Notebook has been widely adopted in scientific data analysis, in particular with applications written in Python, and we will explore its functionalities in support of the ML service framework. The current JupyterHub service provided by the BNL Scientific Data and Computing Center (SDCC) may make it a suitable choice for the ML service we are proposing here. It is also being adopted at other supercomputing centers such as NERSC, and will make it possible for our ML service framework to be available on other data centers.

- **Intermediate Provisioning.** After the user requests are received through the UI, a server, likely located locally at the BNL data center, will perform the provisioning and orchestration of the workflows to be deployed to the remote computing units – which we will call the backend of the ML service. The intermediate server will handle the Data Ingestion and Data Preparation stages of the workflow as shown in Figure 1. The data source could be on the user's local computer, in the cloud or remotely in the target data center. Depending on the workflows and the data source, the data may or may not need to be preprocessed and transferred over to the remote backend. For this component, the key research direction is to investigate and develop a tool that is adaptive to the various scenarios of the ML workflows. For simple, computational inexpensive, workflows, the ML tasks may be performed directly on the intermediate server (with the user's permission). To enable this, performance analysis of the workflows will need to be implemented. Here input on requirements and types of workflows from the applications is crucial, which we will survey at the beginning of the project, and periodically throughout the project to keep up with the development of the ML user community.

- **Backend Execution.** The ML service backend will perform the heavy-lifting of the framework where computation-intensive tasks such as ML model training and hyperparameter optimization will take place. The portability, scalabity and agility aspects of the ML service framework mostly pertain to this component.

  - *Portability*. Given that the software environments and hardware architectures may vary between different backends (i.e. different data centers), the most portable solution is to dispatch the workflows in a containerized environment with portable ML frameworks. The challenge is, how do we turn simple user requests, likely in just a few high-level functions, into an executable that can be run on the compute nodes, likely with GPUs? Here is where the FaaS paradigm will be investigated, to see if user requests can be executed efficiently and portably on different architectures. The current approach, which we call "direct optimization", requires deep changes within each framework to improve the training performance for each hardware backend (CPUs, GPUs, FPGAs, ASICs). These hardware diverge in terms of memory organization and, compute functional units. CPU, GPU, and TPU-like accelerators require different on-chip memory architectures and compute primitives. This divergence must be addressed when generating optimized code.
  - *Scalability*. For distributed ML training, existing frameworks such as Apache Spark, Hovorod and MPI_Learn may be employed. However, the suitability and scalability of each framework may depend on the target platform and the applications. Fast convergence of ML algorithms on large HPC machines need better communication model, synchronization effort and selection of parameters for the algorithms. This requires scanning huge search space to find the best solution. Currently, no tools are available that can help the HEP physicists to optimize the ML algorithms for the target computing resources. We will develop a set of tools that can analyze the performance characteristics of the applications and model the outcome on different platforms, to help develop the scalable execution of the machine learning workloads across a wide spectrum of computing resources.
  - *Agility*. Existing distributed workflow management tools such as PanDA [9] and Parsl [10] will be studied to see if and how they can be used to facilitate fast and flexible executation of the ML workflows on the backend. Since these workflow systems are designed for more traditional computing workflows, theirs features and gaps for ML workflows will be investigated. An Application Programming Interface (API) will need to be defined to facilitate the interaction of the workflow management system with the user ML applications.

Throughout the project, we will use ATLAS and DUNE ML applications as driving examples, such as the ATLAS FastCaloGAN training mentioned previously. Performance benchmarks (in terms of time to solution, coding efforts and other metrics) will be documented periodically to track the progress of the project. We note that this is not a pure software development project, as there are still many unknowns as to how such an ML service framework should be designed and developed. The focus of the project is to investigate different technologies available to us, and find the best path towards a scalable and distributed machine learning service framework through an application-framework codesign approach. The outcome of the project will be the evaluation of different technologies as mentioned above along with a set of prototype tools developed that can be used to guide the design and development of a production-level ML service framework in the future.

4 Milestones, Deliverables and Success Criteria

Year 1: Testbed set up. Investigate single-node portability solutions with containerization and portable ML frameworks. Initial study of the workflow management systems to help define APIs.

Year 2: Investigate scalability solutions, and develop performance analysis and modeling tools for distributed ML.

Year 3: Investigate ways to integrate portability and scalability components with real-time workload management.

Deliverables: Evaluation results of different technologies relevant to the project; software tools for scalable distributed machine learning services;

Success Criteria: Working prototype to be deployed on the BNL IC testbed [1] At least 10x performance improvement for the FastCaloGAN benchmark use case.

5 Project Team

Our multidisciplinary team will consist of scientific ML users/developers (led by Viren/Wenaus) and HPC/ML scientists and engineers (led by Lin), with collaborators from EIC (Aschenauer) to help define the application requirements. One postdoc at CSI will work with Lin to investigate and develop the ML service framework, with 1/2 of a postdoc at NPPS working with Viren and Wenaus to test and validate the developed tools. We have long-standing fruitful collaborations in other projects [11–13], and will make a great team. Our intimate understanding of the data-intensive ML application needs (Viren, Wenaus), and expertise in performance portable HPC (Lin) and distributed computing (Viren, Wenaus) will help ensure the project's success.

6 Impact and Return on Investment

Our research will lead to the successful development of a hardware and application agnostic distributed ML service at scale. This will be an enabler for expanding scientific creativity in AI/ML applications towards models of a scale and complexity otherwise impractical, or with prohibitive development/training time requirements. The project aligns with and is an enabler for BNL's strategic growth plan to use AI/ML to create solutions to experiment-driven computing challenges across the Laboratory, leveraging BNL and wider computing resources to do so.

The development and applications of ML technologies have been the focus of several recent DOE funding calls, including the *Artificial Intelligence, Machine Learning and Data Analytics Codesign* call from ASCR in 2019, and the *Data, Artificial Intelligence, and Machine Learning at DOE Scientific User Facilities* call sponsored by several DOE offices in 2020 and the recent *Integrated Computational and Data Infrastructure for Scientific Discovery* call from ASCR. These funding opportunities typically give awards ranging from $500,000/year to $2,000,000/year. This can also be a good topic for Computational Partnership programs such as SciDAC. Our research will place us at a strong position to participate in future similar funding calls.

[1] While the eventual goal of the ML service is to have it deployed at different computing centers, realizing it will require collaboration and support of the target computing centers. We have a close relationship with SDCC at BNL, making it more realistic to have the service deployed at BNL IC for the LDRD.

References

[1] R. Stevens, V. Taylor, J. Nichols, A. B. Maccabe, K. Yelick, and D. Brown, "Ai for science,"

[2] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, "Machine learning at the energy and intensity frontiers of particle physics," *Nature*, vol. 560, no. 7716, pp. 41–48, 2018.

[3] D. Bourilkov, "Machine and deep learning applications in particle physics," *International Journal of Modern Physics A*, vol. 34, p. 1930019, Dec 2019.

[4] S. I. Campbell, D. B. Allan, A. M. Barbour, D. Olds, M. S. Rakitin, R. Smith, and S. B. Wilkins, "Outlook for artificial intelligence and machine learning at the NSLS-II," *Machine Learning: Science and Technology*, vol. 2, p. 013001, mar 2021.

[5] H. Yu, M. Bishai, W. Gu, M. Lin, X. Qian, Y. Ren, A. Scarpelli, B. Viren, H. Wei, H. Yu, *et al.*, "Augmented signal processing in liquid argon time projection chambers with a deep neural network," *Journal of Instrumentation*, vol. 16, no. 01, p. P01036, 2021.

[6] "International conference on computing in high energy physics." https://indico.cern.ch/event/948465/contributions/4323968/.

[7] G. Nguyen, S. Dlugolinsky, M. Bobák, V. Tran, Á. López García, I. Heredia, P. Malík, and L. Hluchý, "Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey," *Artificial Intelligence Review*, vol. 52, no. 1, pp. 77–124, 2019.

[8] "Funcx." https://funcx.readthedocs.io/en/latest/reference.html.

[9] K. De, A. Klimentov, T. Maeno, P. Nilsson, D. Oleynik, S. Panitkin, A. Petrosyan, J. Schovancova, A. Vaniachine, and T. Wenaus, "The future of PanDA in ATLAS distributed computing," *Journal of Physics: Conference Series*, vol. 664, p. 062035, dec 2015.

[10] "Parsl project." http://parsl-project.org/.

[11] H. Yu, Z. Dong, K. Knoepfel, M. Lin, B. Viren, and K. Yu, "Evaluation of portable acceleration solutions for lartpc simulation using wire-cell toolkit," *arXiv preprint arXiv:2104.08265*, 2021.

[12] H. Yu, M. Bishai, W. Gu, M. Lin, X. Qian, Y. Ren, A. Scarpelli, B. Viren, H. Wei, H. Yu, *et al.*, "Augmented signal processing in liquid argon time projection chambers with a deep neural network," *Journal of Instrumentation*, vol. 16, no. 01, p. P01036, 2021.

[13] C. Leggett, Z. Dong, H. Gray, M. Lin, V. R. Pascuzzi, and K. Yu, "Porting hep parameterized calorimeter simulation code to gpus," *Frontiers in Big Data*, vol. 4, p. 32, 2021.

Torre Wenaus, Brookhaven National Laboratory

Email: wenaus@gmail.com

Education and Training

Ph.D. HEP, Leptonic Decays of the Z (L3), Massachusetts Institute of Technology, 1990

B.Sc. Hon Physics, University of Toronto, 1984

Recent Research and Professional Experience

2019-present: EIC Users Group Software Working Group Co-Convener

2019-present: Nuclear and Particle Physics Software Group Leader, BNL Physics Dept

2018-present: US ATLAS HL-LHC Computing Co-Manager

2017-2018: ATLAS Computing Coordinator

2016-2017: Deputy ATLAS Computing Coordinator

2015-present: Senior Physicist, Brookhaven National Laboratory

2014-2017: HEP Software Foundation startup team co-leader

2014-2016: ATLAS Distributed Computing Co-Coordinator

2010-2015: Software and Computing Manager, U.S. ATLAS Operations Program

2005-present: PanDA distributed production and analysis project co-leader and developer

Recent Honors and Professional Service

Computing in High Energy Physics Conference International Advisory Committee Member 2019-2022

Invited Chair of Jefferson Lab 12 GeV Program Software Review, November 2018, November 2021

Chair of sPHENIX Software & Computing Review, June 2018

Chair of Jefferson Lab 12 GeV Program Software Reviews in 2017, 2015, 2013 and 2012

GeantV review co-chair, October 2016

Session Convener, CHEP 2016, October 2016

Computing Session Convener, APS DPF Meeting, August 2015

Keynote talk at the New York Scientific Data Summit conference, August 2015

BNL Science & Technology Award (BNL's highest employee award), June 2015

DOE HEP advisory panel on HEP computing invited panelist, 2013-2015

Recent Selected Publications

W. Guan et al, Towards an Intelligent Data Delivery Service, EPJ Web Conf. 245 (2020) 04015

P. Svirin et al, BigPanDA: PanDA Workload Management System and its Applications beyond ATLAS, EPJ Web Conf. 214 (2019) 03050

S. Campana et al, An ATLAS distributed computing architecture for HL-LHC, J.Phys.Conf.Ser. 1085 (2018) 3, 032029

MEIFENG LIN

Computational Science Initiative
Brookhaven National Laboratory
Upton, NY 11973-5000

office: (631) 344 4379
fax: (631) 344 5751
email: mlin@bnl.gov

Education

Ph.D. Theoretical Particle Physics, Columbia University, 2007
B.S. Physics, Peking University, Beijing, China, 2001

Professional Experience

07/2019 - present, Group Leader, High Performance Computing, Brookhaven National Laboratory
11/2018 - 05/2019, Acting Group Leader, Quantum Computing, Brookhaven National Laboratory
10/2018 - present, Computational Scientist, Brookhaven National Laboratory
05/2017 - 05/2020, Adjunct Associate Professor, Stony Brook University
10/2016 - 09/2018, Associate Computational Scientist, Brookhaven National Laboratory
11/2013 - 09/2016, Assistant Computational Scientist, Brookhaven National Laboratory
03/2013 - 09/2013, Assistant Computational Scientist, Argonne National Laboratory
11/2012 - 03/2013, Research Scientist, Boston University
10/2009 - 09/2012, Postdoctoral Research Associate, Yale University
09/2007 - 09/2009, Postdoctoral Research Associate, Massachusetts Institute of Technology

Select Publications

1. H. Yu, Z. Dong, K. Knoepfel, M. Lin, B. Viren, and K. Yu, "Evaluation of portable accelera-tion solutions for lartpc simulation using wire-cell toolkit,"arXiv preprint arXiv:2104.08265,2021.
2. H. Yu, M. Bishai, W. Gu, M. Lin, X. Qian, Y. Ren, A. Scarpelli, B. Viren, H. Wei, H. Yu, et al., "Augmented signal processing in liquid argon time projection chambers with a deepneural network,"Journal of Instrumentation, vol. 16, no. 01, p. P01036, 2021.
3. C. Leggett, Z. Dong, H. Gray, M. Lin, V. R. Pascuzzi, and K. Yu, "Porting hep parameterized calorimeter simulation code to gpus," Frontiers in Big Data, vol. 4, p. 32, 2021
4. Z. Dong, Y.-L.L. Fang, X. Huang, H. Yan, S. Ha, W. Xu, Y-S. Chu, S.I. Campbell, ..., M. Lin, *High-performance multi-mode ptychography reconstruction on distributed GPUs*, 2018 New York Scientific Data Summit (NYSDS), 1-5
5. P.A. Boyle, M.A. Clark, C. DeTar, M. Lin, V. Rana and A. Vaquero, *Performance Portability Strategies for Grid C++ Expression Templates*, Proceedings of the 35th International Symposium on Lattice Field Theory (Lattice 2017)

Select Professional Activities

Technical Paper Program Committee, SC2021
Technical Paper Program Committee, PASC2021
Reviewer, International Conference on Parallel Processing (ICPP20), August 17-20, 2020
Reviewer, 13th International Workshop on OpenMP (IWOMP 2017), 2017

Brett Viren (bv@bnl.gov)

Education

Ph.D. Department of Physics and Astronomy, Stony Brook University, 1993–2000.

B.S. College of Creative Studies (Physics), University of California at Santa Barbara, 1988–1992.

Professional Experience

Physicist Physics Department, BNL, (2010-present).

Advanced Technology Architect Physics Department, BNL, (2006-2010).

Advanced Technology Engineer Physics Department, BNL, (2003-2006).

Postdoc Physics Department, BNL (2000-2003).

Research Assistant Department of Physics and Astronomy, Stony Brook University (1994-2000).

Select Software Projects

- Wire-Cell Toolkit: Multi-threaded, component based toolkit for dynamic construction of applications driven by configuration. Provides current state of the art components for LArTPC noise and signal simulation, digital noise filtering and signal processing and 3D imaging from 2D tomographic views.
- PTMP: ProtoDUNE Trigger Message Passing system, a distributed, asynchronous message passing system with included application to provide self-triggering the ProtoDUNE LArTPC detector.
- ZIO: General purpose toolkit for developing asynchronous, self-agregating distributed systems. Applied to Wire-Cell Toolkit it provides multi-producer/multi-consumer asynchronous HDF5 I/O and expected to succeed PTMP for DUNE Far Detector DAQ.
- LARF: A boundary-element method solver for 3D drift and weighting fields required for simulating large LArTPC and similar detectors.
- GeGeDe Constructive solid geometry authoring system to develop configurable parameterized detector descriptions are used to generate geometry for ROOT and Geant4 and with experimental support for FreeCAD.
- Daya Bay Offline Software: primary on geometry, simulation, file I/O, Gaudi framework adaptation, build system, documentation.

Select Recent Publications

- "Three-dimensional Imaging for Large LArTPCs": Qian, Zhang, Viren, Diwan, JINST, **13** #05, 2018.
- "Ionization electron signal processing in single phase LArTPCs. Part I. Algorithm Description and quantitative evaluation with MicroBooNE simulation": MicroBooNE collaboration, JINST **13** #07, 2017

1. EQUIPMENT (Reference: DOE Order 413.2C for guidance on equipment restrictions)

Will LDRD funding be used to purchase equipment? Y/N_Y_

If "Yes," provide cost and description of equipment
Year 1 - $10,000 to purchase a GPU server to be used as the provisioning server.

Year 2 - $

Year 3 - $

Description:

2. <u>HUMAN SUBJECTS</u> (Reference: DOE Order 443.1)

Are human subjects involved from BNL or a collaborating institution? Human Subjects is defined as "A living individual from whom an investigator obtains either (1) data about that individual through intervention or interaction with the individual, or (2) identifiable, private information about that individual".

If **yes**, attach copy of the current Institutional Review Board Approval and Informed Consent Form from BNL and/or collaborating institution. Y/N N

3. <u>VERTEBRATE ANIMALS</u>

Are live, vertebrate animals involved? Y/N N

If **yes**, attach copy of approval from BNL's Institutional Animal Care and Use Committee. Y/N ______

4. <u>NEPA REVIEW</u>

Are the activities proposed similar to those now carried out in the Department/Division which have been previously reviewed for potential environmental impacts and compliance with federal, state, local rules and regulations, and BNL's Environment, Safety, and Health Standards? (Therefore, if funded, proposed activities would require no additional environmental evaluation.) Y/N Y

If **no**, has a NEPA review been completed in accordance with the National Environmental Policy Act (NEPA) and Cultural Resources Evaluations Subject Area and the results documented? Y/N ______

(**Note:** If a NEPA review has not been completed, submit a copy of the work proposal to the BNL NEPA Coordinator for review. No work may commence until the review is completed and documented.)

5. <u>ES&H CONSIDERATIONS</u>

Does the proposal provide sufficient funding for appropriate decommissioning of the research space when the experiment is complete? Y/N Y

Is there an available waste disposal path for project wastes throughout the course of the experiment? Y/N Y

Is funding available to properly dispose of project wastes throughout the course of the experiment? Y/N Y

Are biohazards involved in the proposed work? If yes, attach a current copy of approval from the Institutional Biosafety Committee. Y/N N

Can the proposed work be carried out within the existing safety envelope of the facility (Facility Use Agreement, Nuclear Facility Authorization Agreement, Accelerator Safety Envelope, etc.) in which it will be performed? Y/N Y

If **no**, attach a statement indicating what has to be done and how modifications will be funded to prepare the facility to accept the work.

6. TYPE OF WORK Select Basic, Applied or Development Basic

7. ALIGNMENT WITH THE LABORATORY PRIORITIES

This proposal supports **4. Artificial Intelligence and Data Science**, **5. High Energy Physics and 1. Nuclear Physics.** In particular, this proposal addresses the “Migration to Operation” challenge in the priority area of **Discovery Science Driven by Human-AI-Facility Integration**.

8. POTENTIAL FUTURE FUNDING

DOE ASCR, DOE HEP, DOE NP and potentially DOE BES/BER.
We anticipate that funding opportunities in this area will arise frequently in the next few years given DOE's emphasis on scientific machine learning and artificial intelligence. For us to be competitive, we will likely be in a good position to respond to such funding opportunities in Year 2 and beyond, with the expectation that we will have some innovations that can be competitive then.

9. BUDGET JUSTIFICATION

We request to have 0.2 FTE PI time each year for Lin, 0.1 FTE PI time each year for Viren. Wenaus will devote 10% of his time on this project with other related funding. 1 CSI PD will work on the main research and development tasks, while ½ NPP PD will work on supplying application use cases and testing, both expected to start on June 1, 2021. Computer purchase ($5000 each) and relocation expenses ($5000 each) for the PD new hires in the first year. A GPU server purchase (at a cost of ~$10,000) in the first year is requested to serve as the intermediate provisioning host. Funds are requested for computing time on the BNL Institutional Cluster for approximately 1 node*year. Moderate travel (if allowed by DOE and federal guidelines) to disseminate results at conferences and workshops.

10. NAME OF SUGGESTED BNL REVIEWERS

- Shinjae Yoo (CSI)
- Alexei Klimentov (NPP)

APPROVALS (NPP)

Business Operations Manager

Susan M. Pankowski

Susan M. Pankowski

Department Chair/Division Manager

Hong Ma

Hong Ma

Associate Laboratory Director for
Nuclear and Particle Physics

Haiyan Gao

APPROVALS (CSI)

Business Operations Manager	Grace Giuffre
Department Chair/Division Manager	Frank Alexander
Computational Science Initiative Director	Kerstin Kleese

Distributed Machine Learning
(FY22 Call TYPE A)
T. Wenaus / B. Viren / M. Lin

Resource Category	DESCRIPTION		FY22		FY23		FY24		FY25	
050	Salary - Scientific		7,457		21,965		24,087		17,477	
051	Salary - Research Assoc		20,038		51,634		53,222		37,788	
050	Salary - Professional		0		0		0		0	
050	Salary -Technical		0		0		0		0	
050	Salary - Management & Admin.		0		0		0		0	
	Total FTEs		**0.60**		**1.80**		**1.80**		**1.26**	
TOTAL SALARY/WAGE & FRINGE			**27,494**		**73,599**		**77,309**		**55,264**	
201	RECHARGE LABOR - CSI		61,134		202,517		202,517		139,229	
TOTAL PURCHASED LABOR			**61,134**		**202,517**		**202,517**		**139,229**	
645	Computing Recharge Services		1,000		2,062		2,126		2,192	
TOTAL OTHER			**1,000**		**2,062**		**2,126**		**2,192**	
290	Domestic Travel		2,000		5,000		5,000		5,000	
various	Material Purchases		10,000		2,000		2,000		1,000	
TOTAL MSTC			**12,000**		**7,000**		**7,000**		**6,000**	
170	Relocation Expense		10,000							
271	Communications									
TOTAL COM/MISC			**10,000**		**0**		**0**		**0**	
212	Service Contracts/Publications				1,000		1,000		1,000	
			0		0		0		0	
TOTAL SPECIAL PURCHASES			**0**		**1,000**		**1,000**		**1,000**	
312/314	Capital Equip - Low Value		10,000		0		0		0	
TOTAL EQUIPMENT			**10,000**		**0**		**0**		**0**	
250	Electric		0		0		0		0	
TOTAL ELECTRIC			**0**		**0**		**0**		**0**	
480	Space									
TOTAL SPACE			**0**		**0**		**0**		**0**	
TOTAL DIRECT COSTS			**121,628**		**286,178**		**289,952**		**203,685**	
251	Electric Distributed (Electric Power Burden)		220		589		618		442	
700/701/481	Organizational Burden		3,052		8,169		8,581		6,134	
TOTAL ORGANIZATIONAL BURDEN			**3,272**		**8,758**		**9,200**		**6,576**	
745	Procurement (Material Handling)		1,540		560		560		490	
710	G&A Burden		0		0		0		0	
720	Common Support		45,560		125,503		127,288		89,249	
722	Safeguards & Security Assess		0		0		0		0	
746	Adjs to Procurement Burden									
TOTAL LABORATORY BURDEN			**47,100**		**126,063**		**127,848**		**89,739**	
705	LDRD Burden		0		0		0		0	
TOTAL PROGRAM COSTS			**172,000**		**421,000**		**427,000**		**300,000**	**1,320,000**