Defining a Cyberinfrastructure for Plasma Science and Space Weather Simulations

EROMA ABEYSINGHE, Indiana University, USA
CAMERON W. SMITH, Rensselaer Polytechnic Institute, USA
LIANG WANG, Princeton University, USA
JAMES JUNO, Princeton Plasma Physics Laboratory, USA
MATTHEW MCCALL, Rensselaer Polytechnic Institute, USA
MARLON PIERCE, Indiana University, USA
MARK S. SHEPHARD, Rensselaer Polytechnic Institute, USA

This paper presents a novel cyberinfrastructure designed to facilitate advanced plasma and space weather simulations across various high performance computing platforms through a user friendly web-based interface. The interface offers a comprehensive, open source solution with workflow control that enables the modeling of plasma systems from terrestrial magnetospheres to laboratory experiments in weakly collisional or collisionless plasma regimes. The discussion covers the supported use cases that cater to a diverse user community, an overview of the cyberinfrastructure software components, and community building efforts.

CCS Concepts: • **Information systems → Computing platforms**; • **Computing methodologies → Simulation tools**.

Additional Key Words and Phrases: space weather, plasma science, workflows, science gateway

ACM Reference Format:
Eroma Abeysinghe, Cameron W. Smith, Liang Wang, James Juno, Matthew McCall, Marlon Pierce, and Mark S. Shephard. 2023. Defining a Cyberinfrastructure for Plasma Science and Space Weather Simulations. In *Practice and Experience in Advanced Research Computing, July 23–27, 2023, Portland, OR.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3569951.3597592

1 INTRODUCTION

Physicists, scientists, and engineers studying plasma science and space weather rely on computational methods and tools to gain a deeper understanding of physical processes and to aid the design and evaluation of physical experiments. Unfortunately, in this field, there is a significant barrier to overcome when using powerful computational software at the scale required. Software users need working knowledge of one or more programming languages, the underlying numerical and mathematical models, and the computing environment they are running in, which can range from their laptop to a supercomputer at a national laboratory. To significantly reduce the users' effort to produce their first research result or educational outcome, a cyberinfrastructure (CI) for plasma physics and space weather simulations is defined. It employs permissively licensed open source software that is easy to access through web-based interfaces and uses multiple computing platforms to ensure high availability.

PEARC23, July 23–27, 2023, Portland, OR

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/3569951.3597592

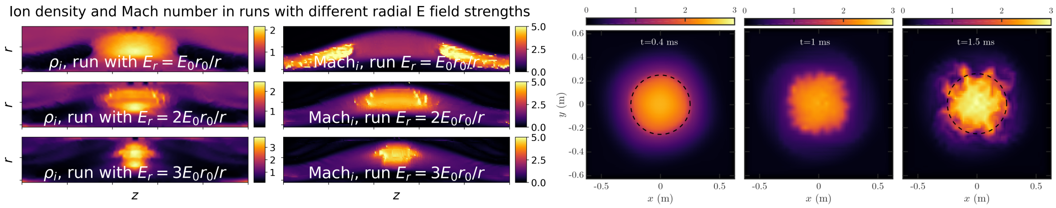

Fig. 1. *Left*: Numerical verification of the Centrifugal Mirror Fusion Experiment (CMFX) using the Gkeyll Mirror App. Centrifugal forces from spinning plasma direct and confine plasma towards mid-plane. The simple geometry of CMFX provides low-cost path to breakthrough fusion energy devices [9]. *Right*. Turbulence saturation in the Large Plasma Device (LAPD) in a simulation using the Gkeyll LAPD App [18].

The key computational framework in the CI is Gkeyll [1]; an open source software framework for high performance simulations for plasma science and space weather with the flexibility to deploy on diverse high performance computing (HPC) architectures [10]. The Plasma Science Virtual Laboratory (Plasma VLab) [2] is the gateway that drives the Gkeyll workflow executions on a wide range of HPC systems. The Plasma VLab is deployed with hosted gateway platform SciGaP [15], using Apache Airavata middleware [13] as the underlying middleware that provides gateways.

The following sections discuss the use cases for the CI, Section 2; how the Plasma VLab supports those users via a web-based interface, Section 3; specific sequences of CI operations that support the needed use cases, Section 4; efforts to leverage the CI for education and outreach, Section 5; and, lastly, a summary of future work.

2 PLASMA SCIENCE AND SPACE WEATHER

Gkeyll is written in a combination of C/C++ kernels and LuaJIT architecture, and fully embraces continuous integration testing through Travis CI. Post-processing is performed using Python and the VTK family of tools. At present, Gkeyll has three major solver packages implemented: the multi-fluid multi-moment model, the Vlasov-Maxwell (VM) model, and the gyrokinetic (GK) model. These models support a range of different physics at different computational costs, allowing flexible deployment of Gkeyll based on the desired physics fidelity of the problem of interest.

Laboratory plasma experiments. Terrestrial laboratory plasma experiments are a rich testbed in plasma sciences, spanning a huge array of applications in plasma physics, from fusion energy sciences to astrophysical plasmas. The diversity of models within Gkeyll has been extensively deployed on a variety of these lab plasma experiments, from the Large Plasma Device (LAPD) at UCLA [18] to fusion experiments such as the Centrifugal Mirror Experiment (CMFX) — see Fig 1. Gkeyll's flexibility to tackle the range of physics observed in these experiments allows for both predictive and interpretive modeling of experimental campaigns. Gkeyll can be deployed for a variety of devices through the Plasma VLab, and ongoing work is being done to extend support to the full spectrum of geometries, boundary conditions, and synthetic diagnostics needed for completely general lab plasma support.

Space weather: multi-fluid ten-moment model. A particularly unique application of Gkeyll is the development and deployment of the multi-fluid, multi-moment model for space weather applications [8, 12, 17, 20, 22]. The key advantage of these new models is that they can accurately capture crucial kinetic physics entirely missing from conventional fluid-based models. This breakthrough represents a significant step forward in our ability to model plasma systems where kinetic effects can play a crucial role in many plasma phenomena. The multi-fluid five-moment and ten-moment

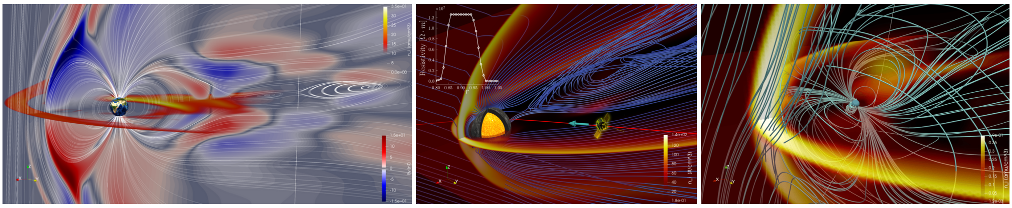

Fig. 2. Gkeyll simulations of magnetospheres for Earth, Mercury, and Uranus. The simulation of Earth demonstrates the dynamic generation of flux ropes, which could potentially result in geomagnetic substorms and other space weather events [22]. Additionally, the simulation for Mercury showcases its large conducting mantle layer [8], while the simulation for Uranus highlights its dipole orientation towards the sun [12].

models in Gkeyll have been thoroughly benchmarked and applied to a variety of plasma physics studies.

Results demonstrate that the multi-fluid ten-moment model can effectively simulate collisionless magnetic reconnection, which typically requires a more expensive kinetic treatment [14, 21]. The multi-fluid model has also been tested against data from NASA's Magnetospheric Multiscale mission, which observed magnetic reconnection in the Earth's magnetosphere [19]. More realistic, large-scale systems, such as the global magnetosphere of Earth [22], Mercury [8], Uranus [12] (see Fig. 2 for demonstration), and other bodies such as Ganymede (a moon of Jupiter) [20] and the Asteroid 16 Psyche [17] have been studied.

The middle panel of Fig. 2 shows Mercury's 3D global magnetosphere configuration. Magnetospheric characteristics such as the bow shock, magnetosheath, magnetopause, and magnetotail are clearly captured by this new model. MESSENGER's second flyby trajectory is plotted in red, pointing from night/dusk side to day/dawn side and near Mercury's equatorial plane. These types of studies have enabled Gkeyll to resolve for the first time the detailed dynamics of electrons and ions in a global context and to make sense of local observations of space weather phenomena such as substorms and the associated ion and electron heating. In the future, these types of studies hold the promise of better understanding and predictive capability, for example, of geomagnetically induced currents.

3 VIRTUAL LABORATORY

Towards significantly reducing the effort needed for users to produce their first research result or educational outcome, the CI defines tailored web-based interfaces that provide end-to-end control over plasma science and space weather simulation workflows. These web-based interfaces exist within the Plasma VLab. The VLab is the main access point for user engagements that will build and grow the community around the Gkeyll framework. This requires supporting a range of users, from advanced researchers with Gkeyll expertise to students following guided lesson plans. Accordingly, all the complexities of installing and maintaining Gkeyll across various HPC resources will be hidden from users. Likewise, new users to VLab are not required to have their own HPC compute and storage allocations to begin running simple analyses. Once they are well matured in their research and have secured the appropriate allocations, they can leverage them directly through VLab.

The HPC resources available through VLab range from institute-based systems, to traditional static ACCESS resources, to dynamically scalable cloud systems such as Jetstream2 [11]. Currently, the Princeton University Stellar cluster is available for Princeton users, and for educational and classroom purposes, Jetstream2 fits well. National users are supported primarily on the ACCESS

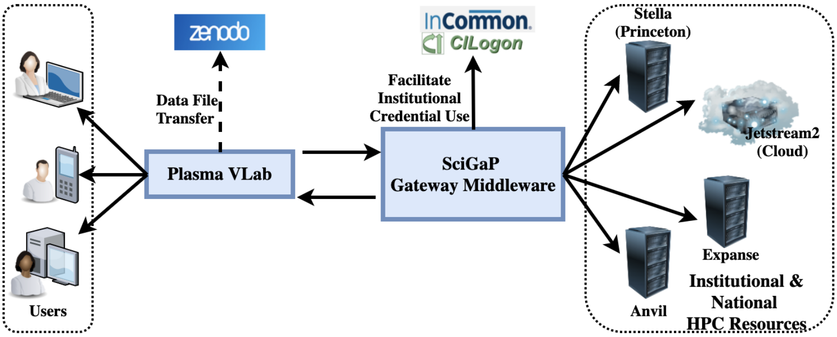

Fig. 3. Users interact with the Plasma VLab. The VLab communicates with Zenodo for cloud-based data storage and SciGap (1) for federated login (CILogon) and (2) HPC experiment execution.

San Diego Super Computer Center's Expanse [4] and Purdue University's Anvil [3] resources. Both are traditional multi-petaflop HPC resources with both CPUs and GPUs.

The Plasma VLab is developed using the Python-based Django framework [6] and accesses HPC resources via SciGaP. The VLab provides users with a secured [7, 16] environment to run their simulations on a range of HPCs, and enables sharing their work with other VLab users or transfering to Zenodo [5] for archival access (i.e., publications)(Fig. 3). The VLab is also equipped with user data storage, a workspace for users to manage their simulation experiments, and functionality to create user groups to share their work. A complete set of instructions on using these features, tutorials on running simulations, and information on the Gkeyll framework will be provided through the Wagtail content management system.

4 SCIENCE WORKFLOWS

Scientific workflows provide researchers and educators with pre-defined sequences of tasks and associated control logic that automates each step across single or multiple computational resources. Fig. 4 is a typical Gkeyll workflow. The VLab aims to support this workflow by providing an intuitive and versatile interface.

VLab applications are defined with a front-end, user-facing component, and a back-end component that transforms and passes front-end user inputs to executables installed on the target HPC system. Specifically for Gkeyll, the front-end provides customized input fields presented to the user in VLab for each application. Those inputs then drive simple variable substitution or text block replacement mechanisms in Lua scripting files that ultimately define the complete set of inputs Gkeyll needs for execution.

The left half of the workflow depicts the user's logic to select from an existing VLab application ('App'). When a pre-defined application does not exist, an advanced user with sufficient Gkeyll and VLab administrator knowledge may define new applications. Once an application is selected (blue shaded box) the user then interacts with the web-based front end (green shaded boxes) to define the required input parameters, execute the simulation, and monitor the results.

For example, a space weather scientist using VLab would pick a planet of interest. Baseline information is automatically filled in, and the user modifies parameters of interest for their problem. The set of inputs is then checked for internal consistency and errors, and the user is provided with an approximation of the initial state that they can review before executing the simulation. The simulation is then launched and maintained (e.g., checkpoint/restart as needed) on an appropriate computing resource. Results are passed through plotting and analysis scripts to inform user

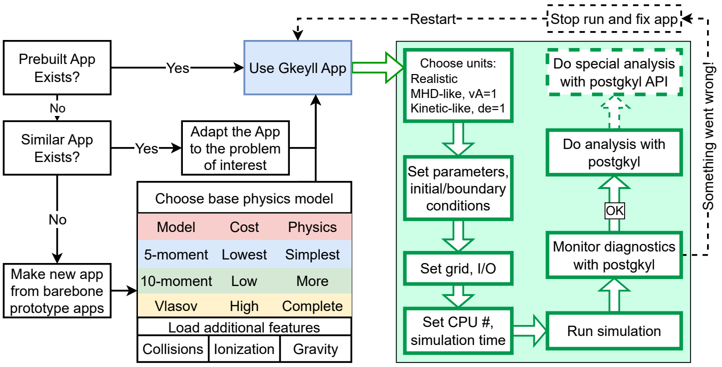

Fig. 4. Typical Gkeyll science workflow. The green shaded boxes depict the sequence of steps after a Gkeyll application is selected in VLab.

modifications to the simulation. Lastly, upon completion, simulation (meta) data is transferred and archived in Zenodo for sharing, publication, and further study.

5 EDUCATION AND OUTREACH

Building education and research communities is one of the most important aspects of the CI. Critical to this is the VLab's ability to support a large number of users with minimum effort. Educational and outreach events will focus on ensuring the VLab meets the communities' needs and growing the user community by providing required features and assistance. In order to attract users and contributors, efforts will focus on participating in summer school programs, webinar series, and conferences. Engagements that emphasize diversity and equity will be a high priority, such as visits to undergraduate institutes serving minority communities. Outreach that drives interactions with similar communities, such as PlasmaPy, and those that promote software best practices such as DOE IDEAS will also be prioritized.

6 FUTURE WORK

The current work is part of a five-year plan. Efforts started with the deployment of the Gkeyll code and its integration with VLab to run simulations on HPC systems. Currently, work is being done on advanced workflow requirement gathering, defining VLab software lifecycle management practices, and outreach targeting summer school programs. Workflows will then be enabled in VLab while efforts to build the user and contributor communities around VLab and Gkeyll continue.

ACKNOWLEDGMENTS

The work is supported by multiple NSF (National Science Foundation) awards. CSSI (Cyberinfrastructure for Sustained Scientific Innovation) program #2209471, Advanced Cyberinfrastructure Coordination Ecosystem: Services Support (ACCESS) program (2138259, 2138286, 2138307, 2137603, and 2138296), SciGaP hosting platform (#1339774). Related research is also conducted under the Laboratory Directed Research and Development (LDRD) at Princeton Plasma Physics Laboratory, operated by Princeton University for the U.S. Department of Energy under Prime Contract No. DE-AC02-09CH11466.

REFERENCES

[1] 2023. The GKEYLL 2.0 code. https://gkyl.readthedocs.io/
[2] 2023. *Plasma Science VLab*. https://vlab.plasmascience.scigap.org/
[3] 2023. *Purdue Anvil*. https://www.rcac.purdue.edu/compute/anvil
[4] 2023. *SDSC Expanse*. https://www.sdsc.edu/services/hpc/expanse/index.html

[5] 2023. *Zenodo open repository*. https://zenodo.org/

[6] Marcus Christie, Suresh Marru, Eroma Abeysinghe, Dimuthu Upeksha, Sudhakar Pamidighantam, Stephen Paul Adithela, Eldho Mathulla, Aarushi Bisht, Shivam Rastogi, and Marlon Pierce. 2020. An extensible django-based web portal for apache airavata. In *Practice and Experience in Advanced Research Computing*. 160–167.

[7] Marcus A Christie, Anuj Bhandar, Supun Nakandala, Suresh Marru, Eroma Abeysinghe, Sudhakar Pamidighantam, and Marlon E Pierce. 2017. Using keycloak for gateway authentication and authorization. (2017).

[8] Chuanfei Dong, Liang Wang, Ammar Hakim, Amitava Bhattacharjee, James A. Slavin, Gina A. DiBraccio, and Kai Germaschewski. 2019. Global Ten-Moment Multifluid Simulations of the Solar Wind Interaction with Mercury: From the Planetary Conducting Core to the Dynamic Magnetosphere. *Geophysical Research Letters* 46, 21 (Nov. 2019), 11584–11596. https://doi.org/10.1029/2019gl083180

[9] Gkeyll Team. 2023. Theory, Modeling, and Validation for a Range of Innovative Fusion Concepts. In *2023 ARPA-E Energy Innovation Summit*. Advanced Research Projects Agency - Energy U.S. Department of Energy, National Harbor, MD.

[10] Ammar Hakim and James Juno. 2020. Alias-Free, Matrix-Free, and Quadrature-Free Discontinuous Galerkin Algorithms for (Plasma) Kinetic Equations. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press, Article 73, 15 pages.

[11] David Y Hancock, Jeremy Fischer, John Michael Lowe, Winona Snapp-Childs, Marlon Pierce, Suresh Marru, J Eric Coulter, Matthew Vaughn, Brian Beck, Nirav Merchant, et al. 2021. Jetstream2: Accelerating cloud computing via Jetstream. In *Practice and Experience in Advanced Research Computing*. 1–8.

[12] S. Jarmak, E. Leonard, A. Akins, E. Dahl, D.R. Cremons, S. Cofield, A. Curtis, C. Dong, E.T. Dunham, B. Journaux, D. Murakami, W. Ng, M. Piquette, A. Pradeepkumar Girija, K. Rink, L. Schurmeier, N. Stein, N. Tallarida, M. Telus, L. Lowes, C. Budney, and K.L. Mitchell. 2020. QUEST: A New Frontiers Uranus orbiter mission concept study. *Acta Astronautica* 170 (May 2020), 6–26. https://doi.org/10.1016/j.actaastro.2020.01.030

[13] Suresh Marru, Lahiru Gunathilake, Chathura Herath, Patanachai Tangchaisin, Marlon Pierce, Chris Mattmann, Raminder Singh, Thilina Gunarathne, Eran Chinthaka, Ross Gardler, et al. 2011. Apache airavata: a framework for distributed applications and computational workflows. In *Proceedings of the 2011 ACM workshop on Gateway computing environments*. ACM, 21–28.

[14] Jonathan Ng, Ammar Hakim, A. Bhattacharjee, Adam Stanier, and W. Daughton. 2017. Simulations of anti-parallel reconnection using a nonlocal heat flux closure. *Physics of Plasmas* 24, 8 (2017), 082112. https://doi.org/10.1063/1.4993195

[15] Marlon Pierce, Suresh Marru, Eroma Abeysinghe, Sudhakar Pamidighantam, Marcus Christie, and Dimuthu Wannipurage. 2018. Supporting science gateways using apache airavata and scigap services. In *Proceedings of the Practice and Experience on Advanced Research Computing*. 1–4.

[16] Isuru Ranawaka, Suresh Marru, Juleen Graham, Aarushi Bisht, Jim Basney, Terry Fleury, Jeff Gaynor, Dimuthu Wannipurage, Marcus Christie, Alexandru Mahmoud, et al. 2020. Custos: Security middleware for science gateways. In *Practice and Experience in Advanced Research Computing*. 278–284.

[17] Elisabeth Rulke, Liang Wang, and Chuanfei Dong. 2021. Three-dimensional, ten-moment multifluid simulation of the solar wind interaction with Asteroid 16 Psyche. In *AGU Fall Meeting Abstracts*, Vol. 2021. Article SM53C-08, SM53C-08 pages.

[18] E. L. Shi, G. W. Hammett, T. Stoltzfus-Dueck, and A. Hakim. 2017. Gyrokinetic continuum simulation of turbulence in a straight open-field-line plasma. *Journal of Plasma Physics* 83, 3 (May 2017). https://doi.org/10.1017/s002237781700037x

[19] Jason M. TenBarge, J. Ng, James Juno, Liang Wang, Ammar H. Hakim, A. Bhattacharjee, Ng Jonathan, James Juno, Liang Wang, Ammar H. Hakim, and Amitava Battacharjee. 2019. An Extended MHD Study of the 16 October 2015 MMS Diffusion Region Crossing. *Journal of Geophysical Research: Space Physics* (2019), 1–14. https://doi.org/10.1029/2019JA026731

[20] Liang Wang, Kai Germaschewski, Ammar Hakim, Chuanfei Dong, Joachim Raeder, and Amitava Bhattacharjee. 2018. Electron Physics in 3-D Two-Fluid 10-Moment Modeling of Ganymede's Magnetosphere. *Journal of Geophysical Research: Space Physics* 123, 4 (April 2018), 2815–2830. https://doi.org/10.1002/2017ja024761

[21] Liang Wang, Ammar H. Hakim, A. Bhattacharjee, and K. Germaschewski. 2015. Comparison of multi-fluid moment models with particle-in-cell simulations of collisionless magnetic reconnection. *Physics of Plasmas* 22, 1 (Jan. 2015), 012108. https://doi.org/10.1063/1.4906063

[22] Liang Wang, Ammar H. Hakim, Jonathan Ng, Chuanfei Dong, and Kai Germaschewski. 2020. Exact and locally implicit source term solvers for multifluid-Maxwell systems. *J. Comput. Phys.* 415 (Aug. 2020), 109510. https://doi.org/10.1016/j.jcp.2020.109510

Received 21 April 2023