Letter of Interest submitted to Snowmass 2021

Computational modeling needs of plasma-based accelerators towards future colliders

August 31, 2020

W. B. Mori[†1], C. B. Schroeder[‡2], C. Benedetti[2], M. Bussmann[3], N. M. Cook[4], A. Debus[3], B. Z. Djordjevic[5], E. Esarey[2], R. Fonseca[6], A. Huebl[2], R. Lehe[2], L. O. Silva[6], K. Steiniger[3], M. Thévenet[7], A. G. R. Thomas[8], J.-L. Vay[2], H. Vincenti[9], S. Webb[4], Xinlu Xu[10]

[1] *Department of Physics and Astronomy, University of California, Los Angeles, California 90095, USA*
[2] *Lawrence Berkeley National Laboratory, Berkeley, California 94720 USA*
[3] *Helmholtz-Zentrum Dresden-Rossendorf, 01328 Dresden, Germany*
[4] *RadiaSoft LLC, Boulder, Colorado 80304, USA*
[5] *Lawrence Livermore National Laboratory, Livermore, California 94550, USA*
[6] *Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal*
[7] *Deutsches Elektronen Synchrotron, Notkestraße 85, 22607 Hamburg, Germany*
[8] *Center for Ultrafast Optical Science, University of Michigan, Ann Arbor, Michigan 48109, USA*
[9] *LIDYL, CEA, CNRS, Université Paris-Saclay, CEA Saclay, 91 191 Gif-sur-Yvette, France*
[10] *SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA*

† mori@physics.ucla.edu
‡ cbschroeder@lbl.gov

Plasma-based acceleration (PBA) has the potential to provide ultra-high gradient structures, enabling compact accelerators. In PBA, plasma wakefields may be driven by charged particle beams, referred to as a plasma wakefield accelerator (PWFA), or high-power, short-pulse lasers, referred to as a laser wakefield accelerator (LWFA). Acceleration gradients of tens of GeV/m have been experimentally demonstrated [1,2]. These plasma accelerators can be the building blocks for a future linear collider [3,4]. The FACET-II [5] and BELLA [6] facilities are aimed at demonstrating that high quality electron/positron beams can be produced and controllably accelerated in cm to meter long PWFA or LWFA stages. In this Letter of Interest, we describe the short-term modeling challenges for experiments at the FACET-II and BELLA facilities, and the longer term needs to provide end-to-end modeling and design of a PBA-based collider.

Rapid experimental progress in PBA has been aided by a more complete theoretical understanding of the beam- and laser-plasma interactions and the emergence of numerical tools to simulate the experimental configurations. When modeling PBA, it is necessary to model the self-consistent generation and evolution of the wakefield, the driver evolution over pump depletion distances, the injection and capture of witness beams of electrons/positrons, and the acceleration and the beam loading of the wakefield by the witness beams. High-fidelity modeling also requires accurate knowledge of the plasma formation. Furthermore, in end-to-end modeling of a PBA-based collider, it is necessary to model how accelerator sections are staged, how the beams are transported from one stage to the next, how the witness beam is focused to the interaction point, and the disruption that occurs when the beams collide. In some schemes, such as the afterburner of an existing collider (perhaps based on LWFA) or the use of a compressed proton beam from the LHC, then studying PBA stages that are more than tens of meters rather than a meter long is also of interest.

Within one PBA stage, there are many similarities between modeling PWFAs and LWFAs. A key difference is that, when modeling an LWFA, the smallest spatial scale that needs to be resolved is often the laser wavelength and not the wakefield wavelength (typically of order of the plasma wavelength). There is often a difference of more than two orders of magnitude between the plasma and laser wavelengths. In addition, there can be differences between the beam loading scenarios and the spot size of the witness bunch. However, there is much in common and the current methods used to model PBA are applied to both PWFA and LWFA. These currently involve particle-in-cell (PIC) methods, i.e., solving the Maxwell-Vlasov system with a Lagrangian representation of the plasma using macro-particles.

The relevant spatial and temporal scales, as well as the necessary physics that needs to be included when modeling PBA, are reasonably well known. For the evolution of the driver, the computational cost is determined by the ratio of plasma length to the driver depletion length scales, as well as to the betatron wavelength in PWFA and the Rayleigh length in LWFA. For the witness beam, it is the spot size and bunch length compared to the size of the wakefield and the betatron wavelength to acceleration length (which could be over many stages). In addition to fully-explicit PIC in three dimensions (3D), several methods have been developed to reduce the computational cost of modeling PBA, including fluid/moment models, ponderomotive guiding center (PGC) and envelope solvers for the laser, quasi-static PIC (with PGC for LWFA), and solving in Lorentz boosted frames that reduce the scale disparity by transforming to a frame where the scales are comparable. In addition, there are axisymmetric (r-z) codes that exploit near cylindrical symmetry to reduce the computation expense by expanding the fields into azimuthal modes and truncating the expansion. There are a variety of different codes in the advanced accelerator community that use each of these methods. An effort to develop the capability for real-time steering of experiments is on-going, and this should be possible through the use of hardware accelerated, reduced and/or surrogate models. Determination of PBA-based collider tolerances to non-ideal effects will require a large number of 3D simulations, and reduced and/or surrogate models will also be needed for these studies.

The PBA community has been a leader in driving innovation in the PIC method and is well positioned to take the next steps. Areas for future numerical algorithmic development include higher-order (including pseudo-spectral) Maxwell solvers, (adaptive) mesh refinement, and higher-order and/or more accurate particle pushers that ensure accurate spatial resolution and accurate modeling of the emittance evolution of the witness bunches. To mitigate grid heating over millions of betatron periods in a staged PBA collider simulation, symplectic PIC algorithms may be necessary. Areas for future physics algorithmic development include physics modules for accurately including radiation reaction in both classical and quantum limits, and photon production (including single quanta events), electron-positron pair creation, collisions, non-LTE atomic physics, modeling the evolution of spin-polarization in the witness beams, as well as detailed modeling of experimental diagnostics. Another area for R&D is modeling the long-term evolution of the plasma wakefield, including modeling how heat is deposited in and extracted from the plasma sources. Considering a collider is driven at an average power in the 100 MW range at multi-kHz levels, it will be necessary to have high-fidelity modeling of the source performance. While there are already separate efforts nationally and internationally on all of these areas, the formation of an international consortium with funding mechanisms to support close collaborations within it should be established.

The community should work together so that physics modules and algorithms used within existing codes are widely available. For example, as more physics packages are included, it is important that the details of these algorithms be widely accessible to ensure confidence in the results. This will not only advance progress, it will ensure that computing resources are well-spent. PBA simulations on leadership-scale supercomputers are routinely performed and are energy-intensive, so it is essential that these are accurate. The community should also work to ensure that the source code of modules and production codes are freely available through open source or other mechanisms that permit innovation.

Modeling PBA requires effective use of supercomputer architectures. Currently the PIC algorithm using the FDTD method has been scaled to more than 10^6 CPU cores and more than 10^8 GPU cores, using both strong and weak scaling. Some quasi-static PIC codes have been scaled to more than 10^5 cores. These methods have also achieved >30% of peak speed on some of these scaling studies. However, ensuring this scaling and effective use of peak speed on physics studies will also require effective dynamic load balancing. The time to solution is determined by the critical path of the application, e.g., the node with the highest workload. Hybrid methods that use on-node acceleration and multi-node decomposition are the de facto default and further levels of parallelism are to be expected. Dynamic load balancing will be ever more challenging as more physics algorithms are included, more interoperable modules are integrated together, mesh refinement is utilized, and computing devices are operating at dynamic (boosted) frequencies. The advent of many core architectures, which include SIMD units with evolving widths and hierarchies of memory, presents an additional challenge to effectively using leadership-class high-performance supercomputing facilities. It is doubtful that compilers alone will achieve effective performance. New data structures and algorithms, which provide a set of parameters that can then be tuned by a compiler or through experimentation, will be desirable. General algorithms that are not specific to one architecture will also be desirable. Furthermore, accelerated computing architectures are increasingly incorporating specialized cores (e.g., tensor cores) for which optimizations of individual algorithms will be needed. Generation, analysis, and movement of large amounts of data will present a larger challenge than already exists today (>PBytes). A detailed and accurate assessment of expected scientific “workflows” regarding data I/O and data lifecycle management should be conducted, which includes in situ solutions.

Key to the development of PBA technology for collider applications will be close collaboration in the accelerator modeling community, as well as coordination between national laboratories and university groups. If a collider is to be built using PBA, then several generations of new researchers will need to be trained. Universities play an essential role in the training of future scientists, who will provide the new concepts and ideas that are needed to advance the field. Developing educational tools using state-of-the-art research codes and training of students should be a high priority.

References

1. Blumenfeld et al., Nature, 445, 741 (2007).
2. A. J. Gonsalves et al., Phys. Rev. Lett. 122, 084801 (2019).
3. S. Gessner et al., Letter of Interest submitted to Snowmass 2021.
4. C. B. Schroeder et al., Letter of Interest submitted to Snowmass 2021.
5. https://facet.slac.stanford.edu
6. https://bella.lbl.gov