arXiv:2203.08335v4 [physics.acc-ph] 22 Sep 2022

Snowmass21 Accelerator Modeling Community White Paper

by the Beam and Accelerator Modeling Interest Group (BAMIG)[*]

Authors (alphabetical): S. Biedron[13], L. Brouwer[1], D.L. Bruhwiler[7], N. M. Cook[7], A. L. Edelen[6], D. Filippetto[1], C.-K. Huang[9], A. Huebl[1], T. Katsouleas[15], N. Kuklev[4], R. Lehe[1], S. Lund[12], C. Messe[1], W. Mori[10], C.-K. Ng[6], D. Perez[9], P. Piot[4,5], J. Qiang[1], R. Roussel[6], D. Sagan[2], A. Sahai[11], A. Scheinker[9], E. Stern[14], M. Thévenet[16], F. Tsung[10], J.-L. Vay[1], D. Winklehner[8], and H. Zhang[3]

[1]*Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*
[2]*Cornell University, Ithaca, NY 14853, USA*
[3]*Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, USA*
[4]*Argonne National Laboratory, Lemont, IL 60439, USA*
[5]*Northern Illinois University, DeKalb, IL 60115, USA*
[6]*SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*
[7]*RadiaSoft LLC, Boulder, CO 80301, USA*
[8]*Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*
[9]*Los Alamos National Laboratory, Los Alamos, NM 87545, USA*
[10]*University of California at Los Angeles, Los Angeles, CA 90095, USA*
[11]*University of Colorado Denver, Denver, CO 80204, USA*
[12]*Michigan State University, East Lansing, MI 48824, USA*
[13]*University of New Mexico, Albuquerque, NM 87106, USA*
[14]*Fermi National Accelerator Laboratory, Batavia, IL 60563, USA*
[15]*University of Connecticut, Storrs, CT 06269, USA*
[16]*Deutsches Elektronen-Synchrotron DESY, 22607 Hamburg, Germany*

September 26, 2022

[*]AccBeamModelSnowmass21@lbl.gov

Contents

Executive summary

Computer modeling is essential to beam and accelerator physics research, as well as to the design, commissioning and operation of particle accelerators. Somewhat surprisingly, despite accelerator physics being a field with extreme levels of coordination and long-range planning for the research, design, construction, and operation of its largest accelerator complexes, e.g., at CERN or at Fermilab, the development of beam and accelerator physics codes has often been largely uncoordinated. This comes at a great cost, is not desirable and may not be tenable.

Accelerator simulation is a large, complex topic, and much time and effort has been spent in developing a large collection of simulation software that cover an expanding range of intertwined physics topics. The complexity of the overall endeavor has risen sharply in the last decade with the impetus to adapt the algoritms and codes to rapidly changing computing hardware and software environments, and the additional task of having to reimagine the algorithms and recast them for quantum computing. This is compounded by the need to infuse a rapidly evolving set of AI/ML technologies, which represent tremendous opportunities but can also be very disruptive.

After a summary of relevant comments and recommendations from various reports over the last ten years, this *community* paper examines the modeling needs in accelerator physics, from the modeling of single beams and individual accelerator elements, to the realization of virtual twins that replicate all the complexity to model a particle accelerator complex as accurately as possible. A discussion follows on cutting-edge and emerging computing opportunities, such as advanced algorithms, AI/ML and quantum computing, computational needs in hardware, software performance, portability and scalability, and needs for scalable I/O and in-situ analysis. Considerations of reliability, long-term sustainability, user support and training are covered next, followed by an overview of the benefits of ecosystems with integrated workflows based on standardized input and output, and with integrated frameworks and data repositories developed as a community. The last section highlights how the community can work more collaboratively and efficiently through the development of consortia and centers, and via collaboration with industry.

The following high-level recommendations are provided to synthesize the summary of recommendations for the topics discussed in the paper and listed afterward:

1. Develop a comprehensive portfolio of particle accelerator and beam physics modeling tools in support of achieving Accelerator and Beam Physics Thrust Grand Challenges on intensity, quality, control, and prediction.

2. Develop software infrastructure to enable end-to-end virtual accelerator modeling and corresponding virtual twins of particle accelerators.

3. Develop advanced algorithms and methods including AI/ML modalities and quantum computing technologies.

4. Develop efficient and scalable software frameworks and associated tools to effectively leverage next generation high-performance and high-throughput computing hardware.

5. Develop sustainable and reliable code maintenance practices, community benchmarking capabilities, and training opportunities to foster the cooperative application of accelerator software.

6. Foster an open community that spans academia, national labs and industry to (a) develop software ecosystems, libraries, frameworks and standards, (b) curate data repositories, and (c) establish dedicated centers and distributed consortia with open governance models.

Summary of recommendations (extended version)

The Beam and Accelerator Modeling Interest Group (BAMIG) proposes the following recommendations to the Snowmass21 conveners:

1. **Recommendation on Modeling needs:** Support the development of a comprehensive portfolio of particle accelerator and beam physics modeling tools for all types of particle accelerators (e.g., RF-based, plasma-based, structured-based wakefield, plasmonic), accelerator components (e.g., materials, superconducting magnets, structured plasmas), and which target the Accelerator and Beam Physics Thrust Grand Challenges on intensity, quality, control, and prediction.

    (a) **Subrecommendation on RF-based acceleration:** *Support the development of modeling tools and methods that target the Accelerator and Beam Physics Thrust Grand Challenges (intensity, quality, control, and prediction), which will require modeling of collective effects with improved fidelity on long time scales, improving computational speed to allow for statistical ensembles and design optimization, and improved integration with realistic magnet and RF modeling.*

    (b) **Subrecommendation on plasma-based wakefield acceleration:** *Support the development of modeling tools and methods that will enable start-to-end simulations that predict the full 6-D (+ spin) evolution of beams in a PBA-based linear collider, from their creation to their final focusing at the interaction point, and include all the intermediate phases of acceleration, transport, manipulations, collisions, etc.*

    (c) **Subrecommendation on structure-based wakefield acceleration:** *Support the development of efficient and accurate algorithms capable of modeling the beam interaction with its wakefield over long interaction lengths in structures with arbitrary geometries and constitutive parameters.*

    (d) **Subrecommendation on PetaVolts per meter plasmonics:** *Support the development of a new quantum-kinetic approach to model large-amplitude plasmons sustained by oscillations of ultra-dense conduction-band free electron Fermi gas. Modeling the dynamics of ionic lattice and the energy band structure under the influence of PetaVolts per meter plasmonic fields of relativistic, nonlinear plasmons is critical to understand the effect on materials.*

    (e) **Subrecommendation on materials modeling for accelerator design:** *Support the development of automated scale-bridging methods that can autonomously parameterize higher-scale models from large numbers of lower-scale calculations, so as to enable predictive materials studies over a broad space of materials and conditions.*

    (f) **Subrecommendation on structured plasmas:** *Support the development and integration of fluid and kinetic codes to meet the modeling demands for a new class of structured plasma devices coupling macroscopic plasma properties with strict requirements on kinetic interactions.*

    (g) **Subrecommendation on superconducting magnets:** *Support the development of novel mixed finite element formulations and algorithms, and their implementation in open source software tailored for superconducting magnet design.*

2. **Recommendation on the next frontier: ultraprecise, ultrafast virtual twins of particle accelerators:** Support the development of accelerator modeling software that

4

orchestrate interdisciplinary set of tools with standardized data representations to enable end-to-end virtual accelerator modeling and virtual twins of particle accelerators, which combine first-principle models together with machine learning-based surrogate models, for tunability from maximum precision for accurate and realistic accelerator design to maximum speed for online particle accelerator tuning.

(a) **Subrecommendation on interdisciplinary simulations:** *Support interdisciplinary simulations especially efforts to establish standards that would ease the sharing of information and data across codes as well as standards for interfacing codes.*

(b) **Subrecommendation on end-to-end Virtual Accelerators (EVA):** *Support the development of software that are capable of end-to-end virtual accelerator (EVA) modeling (Grand Challenge of Accelerator and Beam Physics) that incorporate all components (including both conventional and AAC sections) and all pertinent physical effects.*

(c) **Subrecommendation on virtual twins of particle accelerators:** *Support the development of virtual twins of particle accelerators, which combine high-performance computing, first-principle models together with machine learning-based surrogate models, with tunability from maximum precision for accurate and realistic accelerator design to maximum speed for online particle accelerator tuning.*

3. **Recommendation on cutting-edge and emerging computing opportunities:** Support the research and development on cutting-edge and emerging computing opportunities, including advanced algorithms, AI/ML methods, quantum computing algorithms for beam and accelerator physics, as well as on the development of storage ring quantum computers.

(a) **Subrecommendation on Advanced algorithms:** *Support research on algorithms, from refining the understanding of the properties and bottlenecks of existing algorithms to the elaboration of novel algorithms that exhibit better properties, remove the bottlenecks, and improve the speed and accuracy of accelerator modeling.*

(b) **Subrecommendation on Artificial intelligence, machine learning, and differentiable simulations:** *Support the development of ML modeling techniques and their integration into accelerator simulation and control systems, with an emphasis on fast-executing (up to real-time) and differentiable models, continual learning and adaptive ML for time-varying systems and distribution shifts, uncertainty quantification to assess confidence of model predictions, and physics-informed methods to enable broader model generalization to new conditions and reduced reliance on large training data sets.*

(c) **Subrecommendation on Quantum computing:** *Support quantum computing algorithm and code development for accelerator modeling, feasibility study on quantum computing implementation in accelerator modeling, and quantum computing education in accelerator community.*

4. **Recommendation on computational needs:** Support the development of increasingly powerful and specialized High-Performance Computing (HPC) and High-Throughput Computing (HTC) capabilities for accelerator modeling, and maintenance of software to run efficiently on these hardware (e.g., port of codes to GPUs) with efficient and scalable I/O, post-processing and in situ data analysis solutions, which will be needed to support ab initio modeling at increasing fidelity, training of surrogate models and AI/ML guided designs.

(a) **Subrecommendation on hardware: CPU/GPU time, memory, archive:** *Support the development of increasingly powerful and specialized High-Performance Computing*

(HPC) and High-Throughput Computing (HTC) capabilities for accelerator modeling, which will be needed to support ab initio modeling at increasing fidelity, interactive and parallel data analysis, training of surrogate models and AI/ML guided designs.

(b) **Subrecommendation on software performance, portability and scalability:** *Foster the development and maintenance of codes that run efficiently on the latest hardware (e.g., add support for GPUs/FPGAs) by using maintainable single-source, portable solutions, and that are scalable on leadership-scale supercomputers with multiple levels of parallelization and support for effective dynamic load balancing.*

(c) **Subrecommendation on scalable I/O and in-situ analysis:** *Support the development and maintenance of efficient and scalable I/O, post-processing, in situ data analysis, and data sharing solutions in particle accelerator codes. Coordinate on scientific data documentation, standardization, development of interactive and reproducible analysis workflows and foster data reuse.*

5. **Recommendation on sustainability, reliability, user support, training:** Provide sufficient resources for code maintenance, automated testing, benchmarking, documentation and code reviews. Convene a community effort to identify topics and teaching teams to deliver academic classes designed to foster sharing and cooperation, to be taught at the U.S. Particle accelerator school.

(a) **Subrecommendation on code robustness, validation & verification, benchmarking, reproducibility:** *Establish and maintain open review, automated testing, validation and benchmark procedures for modeling tools and ensure reproducibility by tracking all changes in a documented and openly accessible manner.*

(b) **Subrecommendation on usability, user support and maintenance:** *Consider maintenance of code reviews for improvements, documentation, installation and testing to be central to the mission of usable scientific software: establish open feedback and support channels, perform regular releases with change logs, use permissive open source licensing whenever possible and cover all scientific functionality with automated tests.*

(c) **Subrecommendation on training and education:** *Convene a community effort to identify topics and teaching teams to deliver academic classes designed to foster sharing and cooperation. The classes should be taught in the US Particle Accelerator School with course materials and linked tutorials/extensions regularly maintained and publicly posted.*

6. **Recommendation on community ecosystems & data repositories:** Organize the beam and accelerator modeling tools and community through the development of (a) ecosystems of codes, libraries and frameworks that are interoperable via open community data standards, (b) open access data repositories for reuse and community surrogate model training, (c) dedicated Centers and distributed consortia with open community governance models and dedicated personnel to engage in cross-organization and -industry development, standardization, application and evaluation of accelerator and beam modeling software and data.

(a) **Subrecommendation on loose integration: Integrated workflows:** *Foster the adoption and continued development of open community (meta-)data standards and their implementation in modeling tools and data acquisition for seamless integration of the community accelerator modeling tools into multiphysics workflows.*

(b) **Subrecommendation on tighter integration: Integrated frameworks:** *Establish open, contributable, modular libraries and integrated frameworks for central computing,*

modeling and analysis tasks that foster the sharing of common functionalities between applications, using open licenses and best practices/policies.

(c) **Subrecommendation on data repositories:** *Establish open access data repositories and foster publishing of modeled and measured accelerator & beam data to allow re-use (e.g. beam transport to applications), model training (e.g. AI/ML), preservation, recasting and reinterpretation.*

(d) **Subrecommendation on centers & consortia, collaborations with industry:** *Organize the beam and accelerator modeling community through the development of dedicated Centers and distributed consortia. Dedicate resources to adopt open community governance models and dedicate personnel to engage in cross-organization and -industry cooperation.*

1 Introduction

For particle accelerators—among the most complex and largest high-tech devices ever built—computational tools are indispensable. Computer simulations are critical to the design, commissioning, operation, and upgrading of accelerator facilities which cost many millions to billions of dollars. It is thus widely recognized that the importance of accelerators to society and the high cost of new accelerator facilities demand that the most advanced and sophisticated high-performance computing (HPC) tools be brought to bear on modeling activities in accelerator science and technology [1–6].[]

Accelerator simulation is a large, complex topic, and much time and effort has been spent in developing a large collection of simulation software, many developed by a single accelerator physicist and some by interdisciplinary collaborations of computational accelerator physicists, computer scientists, applied mathematicians and software engineers. Some address a single physics topic while others involve multiphysics, interdisciplinary frameworks, or workflows. Some software has grown very large. Some include very sophisticated algorithms that are pushing the state-of-the-art in accelerator modeling, and sometimes in applications outside of accelerator physics (e.g., astrophysics). While some accelerator modeling research or design can be done on a laptop or workstation, others need the full power of the largest supercomputers. For a significant fraction of these activities, approximations, idealizations or others means to reduce the computational needs are necessary to fit the simulations within the available computer memory and runtime, eventually compromising accuracy and fidelity. In addition to pursue its longstanding tradition in research and development of novel algorithms, the accelerator modeling community is investigating emerging opportunities based on machine learning and quantum computing.

Despite accelerator physics being a field with extreme levels of coordination and long-range planning for the research, design, construction, and operation of its largest accelerator complexes, e.g., at CERN or at Fermilab, the development of beam and accelerator physics codes has often been largely uncoordinated. This comes at a great cost, is not desirable and may not be tenable. Due to developers retiring or moving on to other projects, numerous simulation programs have been completely abandoned or are seldom used. This has resulted in a collection of codes that are not interoperable, use different I/O formats and quite often duplicate some physics functionalities using the exact same underlying algorithms. Frequently there is a huge impediment to maintaining these programs due to poorly-written code and lack of documentation. Additionally, many of the programs that are available tend to be "rigid". That is, it is generally difficult to modify a program to simulate something it is not designed to simulate *a priori*. Adding a new type of lattice element that a particle can be tracked through is one such example [5]. Abandoned simulation programs represent a huge cost [7], not only in terms of time and money spent in developing a program, but also in terms of researchers leveraging existing technology. Indeed, a researcher who wants to simulate something that existing programs are unable to, will, due to time and monetary constraints, generally not be able to fully develop a comprehensive simulation program from scratch as compared to what could have been done if existing software could be leveraged. As simulation programs become more complex due to the ever-increasing demands placed upon machine performance, the situation will become worse if not addressed.

After a summary of relevant comments and recommendations from various reports over the last ten years, this paper examines the modeling needs in accelerator physics, from the modeling of single beams and individual accelerator elements, to the realization of virtual twins that replicate all the complexity to model a particle accelerator complex as accurately as possible. We then discuss cutting-edge and emerging computing opportunities, such as advanced algorithms, AI/ML and quantum computing, computational needs in hardware, software performance, portability and scalability, and needs for scalable I/O and in-situ analysis. Considerations of reliability, long-term

sustainability, user support and training are considered next, before discussing the benefits of ecosystems with integrated workflows based on standardized input and output, and with integrated frameworks and data repositories developed as a community. Last, we highlight how the community can work more collaboratively and efficiently through the development of consortia and centers, and via collaboration with industry. A recommendation is proposed at the beginning of each section and subsection.

2   Previous Reports and Recommendations

This section summarizes relevant comments and recommendations from various reports over the last ten years.

2012 Office of HEP Accelerator R&D Task Force report

Software should help researchers optimize operating regimes and reduce the overall risk that underlies all modern accelerator design. Much of the current software has not taken advantage of the many computer improvements that have been developed in the last few decades. Such progress includes vastly increased processor speed, exploding memory capabilities, disk storage growth and cloud computing. There is not enough overall commercial demand for such high performance accelerator design software to have confidence this problem will be solved without US government intervention.
...
Accelerators across the board also need advanced simulation studies, and long-term support for code development and maintenance is therefore needed.

— Office of HEP Accelerator R&D Task Force report, 2012

2014 Particle Physics Project Prioritization Panel (P5) report

The present practice is to handle much of the computing within individual projects. Rapidly evolving computer architectures and increasing data volumes require effective crosscutting solutions that are being developed in other science disciplines and in industry. Mechanisms are needed for the continued maintenance and development of major software frameworks and tools for particle physics and long-term data and software preservation, as well as investments to exploit next-generation hardware and computing models. Close collaboration of national laboratories and universities across the research areas will be needed to take advantage of industrial developments and to avoid duplication.
Recommendation 29: Strengthen the global cooperation among laboratories and universities to address computing and scientific software needs, and provide efficient training in next-generation hardware and data-science software relevant to particle physics. Investigate models for the development and maintenance of major software within and across research areas, including long-term data and software preservation.
...
Computing in particle physics continues to evolve based on needs and opportunities. For example, the use of high-performance computing, combined with new algorithms, is advancing full 3-D simulations at realistic beam intensities of nearly all types of accelerators. This will enable "virtual prototyping" of accelerator components on a larger scale than is currently possible.

— Report of the Particle Physics Project Prioritization Panel (P5), 2014

2015 HEPAP report

Recommendation 3. Support a collaborative framework among laboratories and universities that assures sufficient support in beam simulations and in beam instrumentation to address beam and particle stability including strong space charge forces.
...
With Scenario B funding, an ambitious computational accelerator science program could be initiated to develop new algorithms, techniques, and generic simulation code with the goal of end-to-end simulations of complex accelerators that will guide the design, and improve the operations, of future accelerators of all types. Advancing the capabilities of accelerator simulation codes to capitalize on the drive toward exascale computing would have large benefits in improving accelerator design and performance. New computational algorithms coupled with the latest computer architectures are likely to reduce execution times for many classes of simulation code by several orders of magnitude, thereby making practical end-to-end simulations of complex accelerator systems. Such capabilities will enable cost-effective optimization of wakefield accelerators, as well as near-real-time simulations of large operational machines such as megawatt proton accelerators or a very high-energy proton-proton collider. In the near term, advanced simulation tools will maximize the productivity of R&D for all future accelerators.
...
One area in which there has been some recent movement towards a nationally unified effort is in accelerator-related computation. Effort in the area has been boosted by funding from the SciDAC (Scientific Discovery through Advanced Computing) program jointly funded by ASCR (DOE Office of Advanced Scientific Computational Research) and HEP. One of the outgrowths of this effort is the CAMPA (Consortium for Advanced Modeling of Particle Accelerators) initiative from LBNL, SLAC, and Fermilab to establish a national program in advanced modeling of accelerators. There are, however, still many isolated simulation efforts within the program.
...
It is likely that there will be significant developments in accelerating technologies, both conventional (NCRF, SRF) and advanced (DWFA, PWFA, LWFA) technologies, in the coming decades. ... Accelerator physics and simulation support in these areas are crucial for making progress.
...
Computer simulations play an indispensable role in all accelerator areas. Currently, there are many simulation programs used for accelerator physics. There is, however, very little coordination and cooperation among the developers of these codes. Moreover there is very little effort currently being made to make these codes generally available to the accelerator community and to support the users of these codes. The CAMPA framework is an exception, and such activities should be encouraged.

The direction of development in computer technologies makes it mandatory that the accelerator simulation codes (as well as all other HEP-related codes) adapt to modern computer architectures. High performance computers are another resource that HEP has not yet sufficiently exploited. The effort to coordinate such advanced computational activities for HEP is taking place within the Forum for Computing Excellence (FCE). Accelerator simulation effort in the direction of advanced computing should also be an integral part of the FCE, as are the other areas of HEP computation. An overall goal of this coordinated effort is to maintain and update mainline accelerator computer codes to take ad-vantage of the most modern computer architectures.

Advances in simulations, as well as in computational capabilities, raise the exciting possibility of making a coherent set of comprehensive numerical tools available to enable virtual prototyping of accelerator components as well as virtual end-to-end accelerator modeling of beam dynamics. It should be possible to construct real-time simulations to support accelerator operations and experiments, allowing more rapid and detailed progress to be made in under-standing accelerator performance.

Simulation efforts are vital for new accelerator development and supporting experimental accelerator R&D studies. Such coherent efforts could be tailored after the successful LARP model that identified mutual study goals for assuring success of a given project (HL-LHC in the case of LARP) and supported collaboration among various university and laboratory partners.

— HEPAP report, 2015

3   Modeling needs

Recommendation: *Support the development of a comprehensive portfolio of particle accelerator and beam physics modeling tools for all types of particle accelerators (e.g., RF-based, plasma-based, structured-based wakefield, plasmonic), accelerator components (e.g., materials, superconducting magnets, structured plasmas), and which target the Accelerator and Beam Physics Thrust Grand Challenges on intensity, quality, control, and prediction.*

3.1   RF-based acceleration

contributed by J. Qiang, C. Mitchell, C-K. Ng

Recommendation: *Support the development of modeling tools and methods that target the Accelerator and Beam Physics Thrust Grand Challenges (intensity, quality, control, and prediction), which will require modeling of collective effects with improved fidelity on long time scales, improving computational speed to allow for statistical ensembles and design optimization, and improved integration with realistic magnet and RF modeling.*

In an RF accelerator, a train of charged particles is accelerated to very high energy (TeVs) by RF cavities, and confined by magnetic elements for high energy physics applications. For such a high intensity charged particle beam, the collective effects from charged particle interactions among themselves and from the other beams play an important role in limiting the beam quality and accelerator performance. These collective effects include space-charge, intrabeam scattering, coherent synchrotron radiation, short-range and long-range wakefields, beam-beam, electron cloud and electron cooling effects.

A number of computational methods have been developed to model these collective effects in an RF accelerator. The particle-in-cell method has been used to simulate space-charge, beam-beam, and electron cloud effects on massive parallel computers [8–13]. The modeling of coherent synchrotron radiation effects was reviewed in a recent publication [14], where the recent advance has led to multi-dimensional and self-consistent simulation capability for the first time. A Monte-Carlo method has been used to simulate the intrabeam scattering effect [15]. A Langevin method was developed to self-consistently simulate the intrabeam scattering effect by solving the Fokker-Planck equations with the Landau/Rosenbluth potential [16]. The fast multipole method was developed to simulate both the space-charge and the electron cooling effects [17, 18].

Even though significant progress has been made in modeling RF accelerators, self-consistent modeling of the above collective effects remains challenging, especially for understanding dynamics on a long time scale. This includes long-term simulation of space-charge and beam-beam effects, self-consistent first-principles modeling of intrabeam scattering and electron cooling, electron cloud and coherent synchrotron radiation effects. For example, one key goal involves the accurate prediction of halo formation and low-level beam loss at high intensity [19]. Fast advanced computational methods are needed to improve both the speed and the accuracy of modeling these effects. The computing time required must be sufficiently short to enable the large ensembles of runs needed for design optimization. Parallel programming paradigms that can make use of the latest computer hardware such as multi-node GPUs are needed to further improve the speed of the simulation. Studies are needed to understand numerical artifacts such as numerical noise associated with long-term simulations. A final area lies in improved magnet modeling and improved integration of RF and magnet models with existing tracking tools. Methods such as surface methods that can include realistic external field effects are needed for improving the fidelity of RF accelerator designs.

The manufacture of accelerator structures and components is a major cost of an RF accelerator. Virtual prototyping of accelerator components has been a key process in the design and optimization of accelerators. Therefore, the ability to virtually prototype with HPC tools to create designs that work "out of the box" will substantially reduce the R&D and operational costs of these accelerator components by eliminating the delicate, labor-intensive modifications employed in previous practices. Virtual prototyping requires multi-physics modeling capabilities for determining RF parameters such as shunt impedance, wakefield effects and higher-order-mode (HOM) damping, temperature distribution and thermal heat load, as well as mechanical stress and shape deformation [20].

To provide the required luminosity for a linear collider, it is critical to reduce emittance dilution in the main linac from machine tolerances such as cavity misalignments and imperfections [21]. For example, in a superconducting linac, misalignments can arise from individual cavity misalignments and changes in the properties of coupled HOMs in a cryomodule with misaligned and deformed cavities. The RF parameters and fields of the HOMs can be evaluated for random distributions of cavity offsets (in a cryomodule) and cavity deformations along the full linac. The fields are then used for beam emittance dilution evaluation. A statistical analysis [22] using the constraints form realistic fabrication and component placement tolerances will be facilitated by the HPC capabilities of advanced simulation codes.

3.2   Plasma-based wakefield acceleration

contributed by J.-L. Vay, W. Mori, F. Tsung

Recommendation: *Support the development of modeling tools and methods that will enable start-to-end simulations that predict the full 6-D (+ spin) evolution of beams in a PBA-based linear collider, from their creation to their final focusing at the interaction point, and include all the intermediate phases of acceleration, transport, manipulations, collisions, etc.*

This section summarizes the modeling needs for plasma-based accelerators (PBAs). A more detailed description is given in [6], Section 2.1.

Several aspects of plasma-based accelerators (PBA) are particularly challenging to model. These include detailed kinetic modeling of the wake excitation; trapping of background plasma electrons in the wakefields; ion motion for a nonlinear beam loading scenario; the need to model the processes in three dimensions; the propagation of ultra-low emittance and energy spread beams in the plasma over meter distances (many time steps), the long (ps to ns) evolution of the perturbed plasma; the disparity in time and space scales of wake and driver evolution; and the large number of 3D simulations required to study the tolerances to nonideal effects.

The most widely used numerical tool to study plasma accelerators is the particle-in-cell (PIC) algorithm [23]. However, ab initio simulations of particle accelerators with PIC codes are limited by the number of time steps that are needed to resolve the beam propagation through the accelerator and by the number of grid cells that are sometimes needed to cover the wide disparity of spatial scales. Several methods to reduce the computational cost, while maintaining physical fidelity, have been developed. *Reduced dimensionality:* When the beam and structure are nearly azimuthally symmetric, a two-dimensional (r-z) representation with a truncated series of azimuthal modes may be used [24–27]. *Quasistatic approximation:* For PBA, one of the most successful methods to speed up simulations is the *quasistatic* approximation [28–33], which relies on the separation of time scales to make approximations and decouple the (long time scale) driver evolution and the (short time scale) plasma wave excitation, providing orders of magnitude speedup over standard PIC. *Boosted frame:* An alternative to performing a quasi-static approximation to handle the disparity of scales

is to shrink the range of scales by using the *Lorentz-boosted frame method* [34]. With this method, the ultrahigh relativistic velocity of the driver is taken advantage of by using a frame of reference that is moving close to the speed of light, and in which the range of space and time scales of the key physics parameters is reduced by orders of magnitude, lowering the number of time steps—thus speeding up the simulations—by the same factor.

While enormous progress has been made in the algorithms and codes to model PBAs, the modeling of a chain of tens to thousands of PBA stages for a multi-TeV collider is extremely challenging, needing further developments [3, 6]. It is also essential to integrate physics models beyond single-stage plasma physics in the current numerical tools, including ionization and recombination, coupling with conventional beamlines, production of secondary particles, spin polarization, QED physics at the interaction point, collisions, and in some cases plasma hydrodynamics for long-term plasma/gas evolution and accurate plasma/gas profiles. Modeling PBA is currently challenging because many beam-loading scenarios are still being considered. These choices may differ between accelerating electrons and positrons, and the best numerical choices for each scenario are not the same.

3.3 Structure-based wakefield acceleration

contributed by P. Piot

Recommendation: *Support the development of efficient and accurate algorithms capable of modeling the beam interaction with its wakefield over long interaction lengths in structures with arbitrary geometries and constitutive parameters.*

Structure-based wakefield acceleration (SWFA) relies on high-charge "drive" bunches [O(10–100 nC)] passing through slow-wave structures (SWSs) to excite electromagnetic wakefields. SWFAs can be configured in either two-beam acceleration (TBA) or collinear wakefield acceleration (CWA). The produced wakefields can be directly used to accelerate a delayed "main" bunch (CWA) or be out-coupled and guided to an optimized accelerating structure that accelerates the main bunch (TBA) in a parallel beamline. CWA offers a simpler configuration where both the drive and main bunches are transported along the same beamline; the TBA scheme decouples the drive and bunch beam dynamics at the expense of increased complexity (e.g., two parallel beamlines are required). The beam dynamics associated with the simultaneous transport of the accelerating main bunch and decelerating drive bunch is one of the major challenges in CWA.

Significant achievements in TBA research include accelerating gradients in excess of ~ 300 MV/m, wakefield-based power generation of ~ 0.5 GW, [35] and demonstration of staged acceleration [36]. Similarly, recent progress along the CWA includes the beam-based demonstration of 300-MV/m gradient in a THz structure [37], the generation of record transformer ratio of > 5 via improved longitudinal beam shaping [38, 39], and significant advances in the theoretical understanding of beam stability [40].

These developments have been possible due to consistent progress in software capable of modelling the beam-wave interactions in complex accelerating and power-extracting and transmission structures (PETSs). Modeling of the SWFA includes radiofrequency (RF) design of complex structures where the RF properties (e.g., R/Q, operating frequency, scattering parameters) are optimized using a frequency-domain solver for eigenmode or scattering parameters (e.g., such as OMEGA3P and S3P [41]). Likewise, time-domain simulations are performed to understand the coupling (wakefield generation and acceleration) of an electron bunch with the mode supported by the structures. These investigations are generally performed with computer programs such as T3P [41] and WARPX [42].

In unison with the roadmap elaborated in Ref. [43] the demonstration of SWFA as a mature

13

technology for a linear collider will require development to (*i*) perform integrated experiments in dedicated test facilities, (*ii*) explore high-efficiency schemes, and (*iii*) push accelerating gradient. Such developments can only be guided with high-fidelity simulations. Developing integrated experiments demands software capable of simulating the long-term beam dynamics within meter-scale structures. The target for high-efficiency acceleration relies on a fine control of the phase-space distribution of the drive and possibly main beams [44]. Understanding the precise evolution of these phase-space distributions over long interaction distances is critical to understanding the onset of the beam break-up (BBU) instability. Simulations over long integration time are also needed for TBA where the generation of GW-level power involves the deceleration of a train of bunches (with 100's ps separation) in a structure. Finally, pushing the accelerating gradient requires optimization of complex electromagnetic structures operated in a transient mode. Implementing boundaries with complex geometries frequency-dependent electromagnetic and material properties is critical to the development of exotic structures and understanding of their dynamical response to intense fields [37]; see also Section 3.5.

3.4 PetaVolts per meter plasmonics and Plasmonic acceleration

contributed by Aakash Sahai along with T. Katsouleas, D. Filippetto
and other contributors of the PV/m plasmonics Snowmass paper

Recommendation: *Support the development of a new quantum-kinetic approach to model large-amplitude plasmons sustained by oscillations of ultra-dense conduction-band free electron Fermi gas. Modeling the dynamics of ionic lattice and the energy band structure under the influence of PetaVolts per meter plasmonic fields of relativistic, nonlinear plasmons is critical to understand the effect on materials.*

PetaVolts per meter (PV/m) plasmonics effort [45–48] focuses on access to extreme electromagnetic fields using plasmonics. Plasmonics is based upon the unique properties of condensed matter that stem from the structure of the ionic lattice which inevitably gives rise to electronic energy bands. The electrons that inherently occupy the conduction band in conductive materials are free to move across the entire lattice in response to external excitation. Plasmonic modes are collective oscillations of the Fermi gas that is constituted by the conduction band electrons [49].

Below we briefly summarize the PV/m initiative and its modeling needs. Further details are available in the above referred work as well as the independent PV/m Plasmonic Snowmass paper [50].

The scope of PV/m plasmonics is limited to mechanisms where there exists a strongly correlated ionic lattice and corresponding electronic energy bands over relevant timescales. Under these conditions, quantum mechanical effects dominate. Specifically, our effort focuses on acceleration and focusing gradients as high as PetaVolts per meter [45–48].

PV/m plasmonics initiative [45–48] introduces a novel class of non-perturbative plasmonic modes that are strongly electrostatic. These novel modes are excited by intense ultra-short charged particle beams. Optical photons cannot directly excite these strongly electrostatic modes as their energy is significantly smaller than metallic density plasmons. Moreover, high amplitude optical plasmons are technologically infeasible due to the large pre-pulse or pedestal energy in high-intensity optical pulses which ablate and disrupt the ionic lattice. This turns any sample into solid-state plasma which are out of the scope of the PV/m plasmonic effort.

The electrostatic nature of plasmons makes it possible to access unprecedented PetaVolts per meter electromagnetic fields. Strong excitation of the Fermi electron gas drives the oscillations to

the wavebreaking limit. Because metals have an equilibrium free electron density n_0, as high as 10^{24}cm^{-3}, the coherence or "wavebreaking" limit [51] of EM fields is,

$$\text{E}_{\text{plasmon}} = 0.1\sqrt{\text{n}_0(10^{24}\text{cm}^{-3})}\text{PVm}^{-1}.$$

A comprehensive approach to modeling of relativistic and highly nonlinear plasmonic modes in conductive materials is required to appropriately incorporate the multi-scale and multi-physics nature of the underlying processes. Being strongly electrostatic, nonlinear plasmonic modes significantly differ from the conventional optical plasmons and do not lend themselves to being modeled using purely electromagnetic codes. Optical plasmons are conventionally modeled using the Finite-Difference-Time-Domain (FDTD) method where the perturbative electron oscillations are simply approximated using constitutive parameters. The FDTD approach with constitutive parameters cannot be used when the collective oscillations become non-perturbative.

Our work [45, 46, 48] has adopted the kinetic approach along with Particle-In-Cell (PIC) computational modeling of collective oscillations of the free electrons Fermi gas to account for the nonlinear characteristics of strongly electrostatic plasmons. It is noted that PIC methodology already utilizes the FDTD solver for calculation of electromagnetic fields but utilizes charge and current densities based upon particle tracking. This approach utilizes the collisionless nature of relativistic oscillations of the Fermi gas and does not assume any constitutive parameters as part of the initial conditions. However, initialization and self-consistent evolution of the electron density implicitly accounts for most of the constitutive parameters. Moreover, as relativistic oscillations of the free electron Fermi gas have been experimentally observed to go beyond the Ohm's law, evolution of these oscillations still remains unaccounted for by collisionless modeling approach. Specifically, conductivity, which is a critical constitutive parameter based upon electron-ion collisions, is not properly understood in the relativistic plasmonic regime. Therefore, heating processes resulting from relativistic plasmonic modes are not fully incorporated in the kinetic modeling approach.

In consideration of this immense promise of PV/m plasmonics using nanomaterials, we call for the support of the modeling community to engage with the our effort and help understand several new challenges that are not part of existing modeling tools.

Below we outline a few of the key challenges identified through our ongoing efforts and bring out the unique requirements our modeling effort:

- Modeling the effect of relativistic energy gain of free electron Fermi gas in the the ionic lattice potential.
- Understanding the effects of the energy density of plasmonic modes on the energy-band structure.
- Accounting for self-fields within the ultra-dense electron beams where particle-to-particle interaction may play a relevant role.
- Incorporating effects from atomistic modeling within the kinetic approach.
- Devising an approach to handle collisions to determine the longer term effects of electrostatic plasmonic fields and account effects related to conductivity.

3.5 Materials modeling for accelerator design

contributed by Danny Perez

Recommendation: *Support the development of automated scale-bridging methods that can autonomously parameterize higher-scale models from large numbers of lower-scale calculations, so as to enable predictive materials studies over a broad space of materials and conditions.*

The quest for ever-higher acceleration gradients leads to stringent demands on the ability of materials to maintain high performance in the harsh environments typical of high-gradient accelerator cavities. Of special interest is the understanding of the origin and properties of RF breakdown precursors, with the goal of eventually using this understanding to develop breakdown-mitigation strategies. These simulations are typically conducted at multiple scales, ranging from first-principle quantum calculations using methods such as density functional theory (DFT). Such calculations allow one to accurately estimate elastic, thermodynamic, and electronic properties of different alloys. However, the steep scaling of these methods (cubic in the number of electrons) strongly limits possible simulation sizes and times. In order to obtain nanostructural information, such as defect nucleation, propagation, and reactions, classical molecular dynamics (MD) simulations are instead employed. These simulations are typically extremely scalable ($O(N_{\mathrm{atoms}}$ or $O(N_{\mathrm{atoms}} \log N_{\mathrm{atoms}}$). A key challenge with moving from DFT to MD is that electrons are typically integrated out, which implies that electronic physics (e.g., the coupling of electrons with the electric and magnetic fields in the cavity) has to be re-introduced "by hand", which often requires *ad hoc* assumptions and is typically extremely time-consuming. While MD simulations are orders of magnitude cheaper than their DFT counterparts, simulations times ($< \mu$s) and sizes ($\sim 10^9$ atoms) remain extremely limited compared to engineering scales. This requires the introduction of yet larger-scale models that operate at the continuum level, including for example crystal plasticity or continuum heat and mass transport codes. While these continuum codes are often extremely scalable, accurately parameterizing the models with materials-specific physics is also extremely tedious and requires a significant level of understanding of the relevant factors, which is often not unambiguously available. We note that the systematic and automated parameterization of continuum models from lower-scale simulations is often not possible, except for "simple" properties such as elastic constants.

The Materials Genome Initiative has pioneered the use of large numbers of quantum simulations (typically at the DFT level) to explore the space of possible material in search of solutions with optimized performance characteristics. These approaches have proved very successful at exploring the space of functional materials (such as materials for batteries), and extensions to structural materials are also possible when suitable figures of merit (i.e., computable from computationally-affordable DFT calculations) can be defined. First applications to accelerator materials have recently been carried out for dilute binary Cu alloys, which identified new candidates for breakdown-tolerant materials, as well as rationalized the observed performance of certain alloys like Cu/Ag. Extending this exploration to more complex multi-component materials will prevent exhaustive enumeration and is instead likely to require more sophisticated machine-learning-guided active-learning approaches. Similarly, including complex microstructural effects in the figure of merit (e.g., to also take into account the microstructure of the materials) is likely to require a combination of state-of-the-art computational methods with data-driven ML algorithms.

Further, the development of robust and automated workflows, such as those described in section 8.1 ,for the estimation and upscaling of materials properties to inform higher-scale models from lower-scale simulations would be extremely valuable, as it would enable the development of predictive materials models in a practical amount of human effort, compensating with a more aggressive use of high-performance computing resources.

3.6   Structured plasmas

contributed by Nathan Cook

Recommendation: *Support the development and integration of fluid and kinetic codes to meet the modeling demands for a new class of structured plasma devices coupling macroscopic plasma properties with strict requirements on kinetic interactions.*

This section summarizes the modeling needs for structured plasmas. A more detailed description is given in [3], Section 2.4.

Structured plasmas, systems for which plasma density, temperature, composition, and ionization state are tailored to enhance an interaction between a beam or laser with that plasma system, are increasingly important for future accelerators. Controllable plasma channels are essential to increasing peak energy and quality across all types of PBAs; they permit the guiding of intense lasers over long distances for laser-driven schemes [52–55], control of plasma density over meter-scale distances as required for beam-driven schemes [56–58], and the realization of plasma columns for novel positron acceleration schemes [59, 60].Structured plasmas have also found application as flexible focusing elements. Discharge capillaries can produce orders-of-magnitude larger magnetic field gradients than traditional electromagnets [61], subsequently enabling the compact staging of plasma accelerators [62]. Alternative approaches include passive plasma lenses, consisting of a narrow plasma jet outflow generated by laser pre-ionization [63, 64], which can provide comparable focusing at high densities. Discharge plasmas have been employed as tunable dechirpers to remove correlated energy spreads from GeV-scale electron beams [65]. Finally, a class of proposed non-destructive diagnostics with high spatiotemporal resolution will rely on controllable plasma densities in concert with electron and laser beams [66].

Although specific needs vary with device type, scale, and application, structured plasmas impose unique modeling requirements as compared with simulations of other accelerator systems. Capillary discharge dynamics require the characterization of a discharge current and corresponding plasma transport properties, including electrical resistivity and thermal conductivity. Magnetohydrodynamic (MHD) codes are well suited to capturing the basic physics of these systems, while maintaining larger timesteps and reduced resolution requirements from kinetic approaches. Significant progress has been made in demonstrating their agreement with 1D analytical models and experimental results for waveguides and active plasma lenses [67, 68]. The coupling of such MHD codes with laser envelope models [69] have contributed to record-breaking LPA acceleration of electrons to 8 GeV in 20 cm capillary [70, 71]. Similarly, active plasma lens studies have benefited from the application of MHD to reproduce species-dependent nonlinearities in current flow and magnetic field [72–74]. MHD codes have also been used to support a new class if structured plasmas predicted on the use of an extended laser focus, known as hydrodynamic optical-field-ionized channels (OFI or HOFI) [75, 76]; Initial investigations into these systems have been supported by 1D hydrodynamic simulations [75, 77]. Despite these advances, MHD simulations face challenges in accurately capturing dynamics at early timescales, where low ionization rates and large temperature gradients impose constraints on fundamental assumptions of local thermal equilibrium (LTE), and may introduce numerical aberrations resulting from poor convergence of equation of state methods and inaccurate transport models [78]. Overcoming these limitations can require prohibitive increases in simulation runtime. More work is needed to improve the speed, reliability, and accessibility of MHD codes for modeling these systems.

Pre-ionized plasma sources, such as those used for PWFA stages and passive lenses, present additional modeling challenges. Neutral gas dynamics necessitate hydrodynamic simulation on μs-timescales to generate proper initial conditions, while the ionization laser requires propagation of an intense electromagnetic fields on fs-timescales. The resulting plasma does not constitute a local thermal equilibrium, therefore the LTE dynamics implemented by most MHD codes is insufficient to capture the ionization and heating dynamics. Moreover, vacuum-plasma interfaces are

of interest for matching incident drive beams [79], and multi-species plasmas have been employed for high-brightness injection scheme [80]. In these cases, kinetic codes, for example PIC techniques, may be coupled with hydrodynamic codes to obtain reasonable approximations, or alternatively, first principles models may be employed [63, 64, 81–83]. Similar challenges are faced in understanding the longterm evolution of these systems following laser or beam interactions. More work is needed to improve the interfacing between fluid and kinetic modeling tools to address the multi-physics nature of these systems across disparate spatiotemporal regimes.

3.7 Superconducting magnets

contributed by Lucas Brouwer and Christian Messe

> **Recommendation:** *Support the development of novel mixed finite element formulations and algorithms, and their implementation in open source software tailored for superconducting magnet design.*

Advanced modeling tools are currently utilized across the full range of US Magnet Development Program (US-MDP) research activities [84], enabling the design of improved conductors, magnets, and diagnostics. A diverse and challenging set of new modeling tools are required to continue this effort and ultimately improve design time, cost, and performance of future superconducting accelerator magnets. The new developments include: (1) simulation of conductor and cable, (2) advanced modeling of interfaces and other potential sources of training in stress-managed designs, (3) modeling of LTS/HTS hybrid magnets, and (4) radiation environment thermal effects on magnets [85].

The material laws for superconducting materials are highly non-linear, which implies that a lot of iterations must be made per simulated timestep to achieve a numerically sufficient convergence. To reduce the high computational cost of these material laws cause, a novel finite element formulation called **h-φ** is being discussed within the HTS modeling community [86, 87]. By using the scalar magnetic potential ϕ rather than the vector potential **a** that is traditionally used to model the magnetic field in air and vacuum domains. Since the computational effort is roughly proportional to the number of unknowns to the power of two, this means that reducing the magnetic potential form a vector to a scalar can reduce the computational effort up to a factor of nine.

One challenge in sufficiently implementing this novel formulation into a finite element code lies in the hybrid use of node and edge-based elements, the way boundary conditions are imposed and how efficient implementations must be designed. Another challenge lies in efficiently solving the system matrices which are non-symmetric, non-positive-definite and ill-conditioned.

Regarding the **h-φ** formulation, some of the most pressuring questions discussed in the HTS modeling community are:

- imposing current boundary conditions in a user-friendly fashion[88, 89]
- discretization of tape structures (thin shells) in 2D and 3D space [90]
- quench prediction (maxwell-thermal coupling)
- delamination (maxwell-thermal-mechanical coupling)

Most of the published formulations have been implemented on top of commercial toolboxes such as COMSOL and are not yet available to the general public. Moreover, it is in the very nature of closed source codes to limit the access the developer has to the data structure and the knowledge of underlying algorithms. The desire of having full control over the data structure motivated us to develop a custom codebase that is tailored to the needs of HTS modeling. Having defined the

project goals below, we decided to develop a new finite element framework from scratch to achieve them:

- Support **h-a** and **h-ϕ** formulations for 2D and 3D, as well as thin shells, both with first and higher order elements.
- Support multiphysics, specifically thermal and mechanical coupling, as well as current sharing.
- Have a text-based user interface that is tailored to the needs of HTS magnet and cable modeling.
- Use popular open-source data formats, such as HDF5, GMSH, and Exodus II (ParaView)
- Link against modern sparse linear algebra solvers such as STRUMPACK, PETSc, PARDISO and MUMPS.
- Run in parallel using the MPI standard.
- Be readable, extendable and maintainable.

4   To the next frontier: ultraprecise, ultrafast virtual twins of particle accelerators

contributed by Jean-Luc Vay, Axel Huebl, Rémi Lehe, Chengkun Huang, Nikita Kuklev, Ji Qiang, David Sagan

Recommendation: *Support the development of accelerator modeling software that orchestrate interdisciplinary set of tools with standardized data representations to enable end-to-end virtual accelerator modeling and virtual twins of particle accelerators, which combine first-principle models together with machine learning-based surrogate models, for tunability from maximum precision for accurate and realistic accelerator design to maximum speed for online particle accelerator tuning.*

The development of future particle accelerators requires predictive modeling tools that will range from (i) very detailed, full physics and dimensionality, first-principle kinetic simulation tools that are needed for detailed runs for physics studies (typically based on Monte-Carlo and Particle-In-Cell methods) [91], to (ii) ultrafast simulation tools that are needed for ensemble runs for design studies (using a combination of, e.g., reduced physics, low dimensionality, low resolution, artificial intelligence (AI)/machine learning (ML) surrogates [92, 93]).

4.1   Interdisciplinary simulations

Recommendation: *Support interdisciplinary simulations especially efforts to establish standards that would ease the sharing of information and data across codes as well as standards for interfacing codes.*

Interdisciplinary simulations are important in a number of areas. In vacuo particle tracking coupled with particle/matter interactions is an example of a growing need. One application is in simulating the radiation induced by "dark current" electrons in accelerating cavities. This radiation may cause damage to cavities which leads to shortened lifetimes of the devices and a radiation safety hazard for the surrounding environment. Dark current induced problems have been observed at many facilities such as the CEBAF[94], LCLS-II[95], ANL, etc. Sufficient shielding is required to properly contain the radiation which in turn requires a good understanding and prediction of radiation levels through simulations. Another example is the modeling of positron production in a

target from the impact of high-energy electron beams accelerated through a linac injector. These simulations require accurate calculations using electromagnetic RF codes for accelerator structures and beam dynamics codes for particle transport in a beamline to characterize the beam profile before it hits the accelerator enclosure or the target.

Particle/matter simulation codes exist. Examples include Geant4[96], FLUKA[97], and MARS[98] which have traditionally been used for detector simulation in HEP experiments. However, since these codes and accelerator simulation codes have all been developed without common standards, interfacing them is a laborious task. A seamless simulation requires the proper transfer of field and particle data from accelerator to radiation codes. Communication in a standardized format such as openPMD[99], which has been adopted in some accelerator codes, would help ensure efficient and error-free field and particle data transfers[100]. Another issue with an integrated simulation is in matching of the geometry of the vacuum chamber surface. The surface geometry in accelerator simulations is generally poorly defined if at all. The most comprehensive simulations define the surface using a finite element mesh generated from a CAD model (e.g., both WARP [8] and OPAL [13, 101] have some limited capacity for this). In contrast, radiation codes generally employ a faceted representation of the CAD model boundary. A converter for mapping finite element curved surfaces to faceted divisions on an interface boundary is required to accurately determine the location of a particle crossing from one computational domain to another. Much time and effort would be saved if the surface geometry descriptions were standardized so that a single converter module could be used in multiple codes.

Increasingly, accelerator simulation tools are also incorporating more micro-physics models to better describe the complex interplay of the various physics phenomena. One particular example being the emission modeling of a high-brightness electron photocathode gun. A photocathode gun provides a high-brightness electron source for the downstream accelerator beamline where the beam brightness can only be degraded, not improved. Thus, it is essential to understand the cathode emission characteristics and the method to control the beam quality in the gun environment through validated simulations. While Monte Carlo photo-electron emission simulations have been widely employed in studying photocathode performance for dedicated experiments, its potential in integrated simulations has only been explored recently [102]. For such purposes, a tight integration of the micro-physics models into existing gun simulations can be achieved via the best practices and standardization as discussed below.

4.2   End-to-end Virtual Accelerators (EVA)

Recommendation: *Support the development of software that are capable of end-to-end virtual accelerator (EVA) modeling (Grand Challenge of Accelerator and Beam Physics) that incorporate all components (including both conventional and AAC sections) and all pertinent physical effects.*

The realization of software that are capable of end-to-end virtual accelerator (EVA) modeling has been identified as a Grand Challenge of Accelerator and Beam Physics [103].

Thorough modeling of a full particle accelerator in individual parts (injector, magnets, beam dynamics, etc.) is already an important aspect of particle accelerator development; one, without which no project nowadays can proceed to the building stage. For example, start-to-end simulations have been used in x-ray light source accelerator designs [104–106]. However, simulating the system in many small units poses two significant challenges: (i) The various codes used for the individual parts are often from different areas, lack the state-of-the-art high-fidelity models for extremely bright and intense beams, written in different languages (C/C++, Fortran, Python, etc.), and use different

standards for data I/O (particles, fields, etc.), slowing down the design and simulation process; (ii) Without multiphysics couplings, subtle, but important effects (collective effects, halo, coherent synchrotron radiation, etc.) might not be considered in the design or without sufficient accuracy, ultimately leading to issues during commissioning, the need to modify the machine (shimming, additional steering, shielding, etc.), and longer downtimes due to added radiation. Furthermore, opportunities for more efficient working points that can be achieved only from system optimization of the entire accelerator may be missed entirely.

Thus, there is a clear need for full physics 6-D computer simulations of the entire accelerator system that incorporate all components (including both conventional and AAC sections), all pertinent physical effects, and that execute fast and reliably. Such EVAs need to be validated with experiments at both the component and system levels. Furthermore, these EVAs should be able to leverage modern computing infrastructure like HPC clusters and GPU computing, and fully integrate AI/ML tools to maximize efficiency for practical applications. The development of such tools requires continuous advances in fundamental beam theory and applied mathematics, improvements in mathematical formulations and algorithms, and their optimized implementation on the latest computer architectures.

4.3 Virtual twins of particle accelerators

Recommendation: *Support the development of virtual twins of particle accelerators, which combine high-performance computing, first-principle models together with machine learning-based surrogate models, with tunability from maximum precision for accurate and realistic accelerator design to maximum speed for online particle accelerator tuning.*

Ultimately, the availability of multiphysics, end-to-end simulation capabilities will lead to the realization of virtual twins of particle accelerators. These will enable the design and optimization of future accelerators at scale on supercomputers, with unprecedented levels of accuracy and speed. This capability will dramatically increase the breadth of parameter space that can be explored, thereby enabling the design of particle accelerators that have not been possible before. Ultrafast and ultraprecise tunability will let users simulate as needed: from maximum precision for accurate and realistic accelerator design to maximum speed for online particle accelerator tuning. Modern ML optimization frameworks can be used to take advantage of this tunability via multi-fidelity algorithms, allowing for faster scalable design refinement. Automated uncertainty quantification can be used to provide hints as to where the computational budget should be applied for best overall simulation quality.

5 Cutting-edge and emerging computing opportunities

Recommendation: *Support the research and development on cutting-edge and emerging computing opportunities, including advanced algorithms, AI/ML methods, quantum computing algorithms for beam and accelerator physics, as well as on the development of storage ring quantum computers.*

Just as it is essential to continually invest in fundamental studies of accelerator technologies and beam science, it is also essential to do the same for the software and algorithms that enable such studies using computers. This includes the continuation of the study and development of novel algorithms and optimization strategies, as well as leveraging AI/ML and exploring algorithm that will take advantage of upcoming quantum computers. This also includes the development of

modeling tools to study and design quantum computers based on storage rings.

5.1 Advanced algorithms

contributed by Jean-Luc Vay, Ji Qiang

> **Recommendation:** *Support research on algorithms, from refining the understanding of the properties and bottlenecks of existing algorithms to the elaboration of novel algorithms that exhibit better properties, remove the bottlenecks, and improve the speed and accuracy of accelerator modeling.*

Computational beam and accelerator physics has spurred the development of many algorithmic advances [33, 34, 107–164] that have boosted the computational capabilities tremendously, in some instances by orders of magnitude at a time. Yet, the modeling of particles accelerators remains very demanding, and even extremely challenging in the case of design of future plasma-based or structure-based particle wakefield HEP colliders. It is thus essential to pursue the research on algorithms, from refining the understanding of the properties and bottlenecks of existing algorithms to the elaboration of novel algorithms that exhibit better properties, remove the bottlenecks, and improve the speed and accuracy of accelerator modeling. The discovery of new and better algorithms for beam and accelerator modeling is a fundamental topic that is poised to provide big boosts to the other fundamental topics that are the discovery of new and better accelerator technologies and concepts, and the discoveries of new laws of natures with colliders that they enable.

This also includes the adoption of algorithms that have proven to be very effective into other fields, such as adaptive mesh refinement [120], which enables to zoom in on regions that need higher resolution (e.g., sharp edge of high-intensity beam) while zooming out to simulate larger regions (e.g., beam with pipe to include image charges and halo effects). This is discussed in more detail in [6].

5.2 Artificial intelligence, machine learning, and differentiable simulations

contributed by Auralee Edelen, Remi Lehe, Nikita Kuklev, Alexander Scheinker, Ryan Roussel

> **Recommendation:** *Support the development of ML modeling techniques and their integration into accelerator simulation and control systems, with an emphasis on fast-executing (up to real-time) and differentiable models, continual learning and adaptive ML for time-varying systems and distribution shifts, uncertainty quantification to assess confidence of model predictions, and physics-informed methods to enable broader model generalization to new conditions and reduced reliance on large training data sets.*

Machine learning (ML) and artificial intelligence (AI) have revolutionized many computational tasks in recent years, from protein folding prediction to fusion energy control. Speeding up computationally-intensive accelerator simulations and aiding optimization of particle accelerators are two key areas where ML & AI have also been making substantial contributions. Many ML/AI tools are now technically mature, production-ready, and can be used as part of standard computational toolkits for particle accelerators. Notable examples so far include fast-to-execute accelerator models that can be used online as "digital twins", as well as in design and online optimization using Bayesian optimization and other methods, whereby a model that is learned on-the-fly is used to guide sampling of the parameter space. These methods, and tools supporting their use, are beginning to be integrated into dedicated simulation workflows and into control systems for use online that

share common standards for interoperability. For example, the open source package LUME [165] has been used to deploy online ML and simulation models for a variety of systems, and xopt [166] has been used to generate datasets (both offline and online) for training ML models.

So far, most studies in AI/ML for accelerators have focused on proof-of-principle tests of new methods under narrow operating conditions, rather than integration into dedicated use online where they would be tested under a wide variety of conditions and face additional challenges in robustness. Looking forward, there is great potential to now expand upon the fundamental methods algorithmically, combine the strengths of different methods in novel ways (for example, combining neural network models or differentiable simulations with uncertainty quantification and Bayesian optimization), and address the challenges of dedicated deployment that often do not arise in R&D testing of each new algorithm (e.g. robust uncertainty quantification and adaptation to account for changing conditions over time or exploration of new regions of parameter space, extension to very high dimensional systems with hundreds of freely-tunable accelerator settings, etc.). We highlight some current and future directions below.

Fast-executing, adaptive accelerator models Physics-based accelerator models can very accurately predict expected beam phase space evolution. These models can be calibrated offline by hand or online based on feedback from diagnostics, such as using energy spread spectra to tune a model to predict the beam's longitudinal phase space at FACET [167]. The major limitation of physics-based simulation models is that they are computationally expensive and therefore cannot give real-time (e.g. at rates of 120 Hz to MHz) predictions in the control room. For physics-based simulations that include the major nonlinear collective beam effects, prediction latencies can be on the order of tens-of-minutes to hours depending on the simulation fidelity. In contrast, trained ML models do have the potential to provide real-time predictions, in some cases providing orders of magnitude improvements to execution speed [168]. Many studies in ML-based modeling of accelerator systems have relied on supervised learning on bulk inputs/outputs and fairly standard neural network architectures, such as multi-layer perceptrons [168–172] and convolutional neural network (CNN) based encoder-decoder architectures [173–176], as well as Gaussian process models [177, 178]. Common inputs and outputs include accelerator settings such as magnet currents, cavity gradients, and initial conditions such as laser parameters; common outputs include scalar beam parameters such as the emittance and bunch length, and 2D projections of the beam phase space. When trained on large simulation datasets of input-output variables, these models have been demonstrated to have as much as a million times speedup in execution speed with reasonable prediction accuracy [168]. Training on measured data (either from scratch or by adapting a model that is pre-trained on simulations) also enables differences between the as-built machine and the simulation to be represented.

Looking forward, there is a need to further develop methods for making ML-based accelerator models reliable under shifting input-output data distributions. For example, this can include both time-varying changes in inputs from an operational accelerator and deliberate changes in operating conditions or simulation inputs that go beyond the statistical distribution of the training data. Methods to flag when model performance is degraded, as well as methods to adapt models to new conditions are needed. Having accurate uncertainty estimates in addition to model predictions is critical for applications in accelerators, especially in cases where decisions about control actions, analysis, or design are made based on model predictions. Uncertainty quantification methods, such as the use of Bayesian neural networks[179] and ensembling [180], have been used in accelerators to highlight cases where model uncertainty is high. However, additional work is needed to ensure robustness of uncertainty estimates under changing conditions. In general, further work incorporating

state-of-the-art methods from continual online learning [181] and transfer learning [182] could be used to maintain model accuracy over time and aid re-use of models in different settings. This could also help reduce reliance on large training data sets when transferring models between specific accelerators (for example, between injectors with similar designs). Operation in tandem with local feedback algorithms can also help retain model accuracy over time as conditions change. Sample-efficiency is a key consideration for accelerator problems in which the data collection is very time-consuming or detailed beam measurements would interrupt operations. By incorporating physics constraints within ML frameworks, such as developing beam dynamics models based on Taylor maps, it is possible to greatly reduce the amount of data needed for re-training models [183], as discussed in more detail below. Further work is also needed in developing advanced methods for cases where online re-training is not feasible or practical. Some avenues toward addressing this challenge without re-training include the use of adaptive feedback / local optimization to make the overall ML system more robust to unknown and time-varying system changes and inputs. Some work towards the challenge of ML without re-training is utilizing adaptive feedback within the ML framework by adjusting ML model inputs or internal parameters with gradient-free adaptive feedback algorithms [176, 184]. In the context of accelerator tuning, it can also be accomplished by combining ML system models for control with local adaptive feedback or optimization algorithms [185].

Physics-informed and physics-guided ML modeling   Recently there has been interest in incorporating physics information more directly into ML models. This is expected to improve sample efficiency and accurate generalization to unseen regions of parameter space. Such "physics-informed" modeling spans a wide variety of techniques that use physics information more or less directly (see [186] for a review). In the most broad sense of "physics informed", training on very broad, high-fidelity simulation data sets introduces observational biases that allowing ML models to learn functions, vector fields and operators that reflect the physical structure of the data [186]. Numerous examples exist in the literature where training on broad accelerator simulation data sets enables enough of the underlying physics to be represented in a sufficiently generalizable form that the learned models can reliably interpolate to unseen configurations or provide qualitative estimates of beam behavior even outside of the training distribution [168, 172], for both scalar predictions and beam phase space projections. A recent example highlights that this can also be extended to unobserved components of the beam parameter space [176]. In that work, a convolutional neural network (CNN)-based encoder-decoder trained and then utilized in an unsupervised adaptive approach to predict the 15 unique 2D projections of a beam's 6D phase space for an unknown and time-varying input beam distribution. The approach in [176] is to adaptively tune the low-dimensional latent space of the encoder-decoder to match a measured 2D (longitudinal phase space) projection of an unknown 6D phase space distribution and in doing so accurately predict the other unmeasured 2D projections of the 6D distribution.

More direct methods of making "physics-informed" models include adding physics-based constraints into the learning process [186, 187], or structuring the learning algorithm in a way that mathematically resembles the physics problem (e.g., replacing one component of a Hamiltonian update with a neural network [188]). In accelerator physics this avenue has been less explored but holds significant potential. In one accelerator example, the Hessian from a physics-based or learned model is used to inject expected correlations into a Gaussian process kernel [189], which results in more sample-efficient learning and optimization. Another accelerator example uses a differentiable formulation of beam dynamics in a ring, implemented in the tensorflow machine learning library, to enable learning of free parameters using traditional ML workflows [183]. For

more efficient predictions, Galerkin-based or symplectic reduced order modeling (ROM) may present new opportunities for nonlinear beam dynamics study beyond the theoretical formulation.

Differentiable simulators   While there is a long history of using automatic differentiation in accelerator physics for higher-order calculations in rings [190], accelerator physics simulation codes do not readily support arbitrary computation of gradients. This necessitates the use of numerical estimation of gradients when using gradient-based optimization methods. Having simulation codes that inherently support automatic differentiation would enable gradient-based optimization to be more readily used with accelerator physics simulations. Tying together simulations with gradient-based learning algorithms also has significant promise; for example, this could aid gradient-based learning of free parameters, as was done for accelerators with a differentiable lie map formulation in [183] and for a differential hysteresis model in [191]. This is similar to examples using adaptive feedback that approximates gradient descent to optimize free parameters in physics models [167]. Gradient-based learning of external ML algorithms (such as learning a control policy in reinforcement learning) can also be aided by a differentiable forward model, as was done for example in [171] but with a differentiable learned ML model rather than a differentiable simulator. Tying together standard simulation tools with ML components (e.g., using a learned model for one machine section or one computation component) would benefit substantially from having end-to-end differentiability. Expanding this concept to detailed differentiable simulations of non-linear collective effects (e.g., with particle-in-cell codes) could enable new capabilities in optimizing particle accelerator simulations and using them in conjunction with AI/ML. One example of integrating optimization with a differentiable physics model is given in [191], where an analytic model of hysteresis is made differentiable and used in conjunction with measured data to efficiently learn an accurate model of a magnet response that includes hysteresis; this is then used for hysteresis-aware Bayesian optimization that can more precisely optimize the resultant beam.

ML/AI frameworks have highly-optimized implementations of automatic-differentiation and other low level computational tools for running efficiently on CPU/GPU. Major libraries (PyTorch, TensorFlow, JAX, Julia) provide low level hooks for these routines. Some frameworks can be used to provide gradients for existing codes written in C or Fortran. This presents two promising avenues toward making differentiable physics simulators: (1) re-writing or translating existing low-level accelerator physics codes into formats compatible with automatic differentiation and (2) wrapping sections of existing code with external software to provide derivatives. This also opens a path toward using efficient AI/ML learning algorithms and frameworks directly with physics simulations.

Model-based optimization   The development and use of advanced algorithms for accelerator design and optimization aims to efficiently solve challenging optimization problems in high dimensional parametric spaces. The so called "curse of dimensionality", where optimization difficulty scales exponentially with the number of free parameters, limits the number of optimization problems that can be done, especially when running high-fidelity accelerator simulations is computationally expensive. To make a problem solvable with a limited computational budget, either assumptions are made to simplify the simulation itself or the free parameter space must be heavily restricted. However, if novel techniques that substantially improve the efficiency of simulated optimization can be implemented, previously unsolvable optimization problems in accelerator physics can be realistically tackled.

Traditional design optimization in accelerators has relied on computationally intensive algorithms such as genetic algorithms (e.g. NSGA-II [192]) and particle swarm optimization, which require many simulation samples to converge and do not leverage any learned representations of the system.

ML-based optimization methods that attempt to learn a model on-the-fly during optimization or use an existing model can improve sample efficiency and convergence speed by enabling more judicious choices of input variables to examine.

In cases where gradient information is not available, known as "Black Box" problems, model based optimization (such as Bayesian optimization) techniques can be used to maximize optimization efficiency. These algorithms have already been shown to perform well in certain experimental [177, 189] and simulated [193, 194] contexts. However, there remains substantial room for further innovation and improvement in sample efficiency, scaling, and tailoring to specific applications. For example, multi-fidelity Bayesian optimization [195], where evaluations of the optimization objective at a lower fidelity are used as a proxy for high fidelity simulations, is naturally well-suited to simulated optimization of accelerator based problems where multiple fidelities are straightforward to interpret. These algorithms could be used to help balance computational resources and required fidelity automatically during search refinement. Furthermore, if models used for these algorithms also provide an associated confidence metric in their predictions, they can be used in information-based optimization to either speed up optimization in high dimensional parameter spaces based on model confidence regions [196] or improve mutual information gain during nested optimization procedures [197].

Model based algorithms also enable the use of transfer learning based-techniques, where previously created models of similar systems are used to inform algorithms applied to novel optimization problems and aid sample-efficiency. For example, using information about correlated input parameters from a physics simulation to inform the internal model in Bayesian optimization of an operating accelerator [177, 189] can also be considered a form of kernel-based transfer learning. Physics-informed models or differentiable simulators also have great potential for use in sample-efficient, model-based optimization across different systems. They can be used in conjunction with gradient-based optimization to learn system model parameters on-the-fly with limited data and exploit this information in control. For example, in [191], the parameters of a general differentiable hysteresis model for magnets is learned from beam measurements and is used in conjunction with Bayesian optimization for hysteresis-aware fine-tuning.

5.3 Quantum computing

contributed by He Zhang, Ji Qiang

Recommendation: *Support quantum computing algorithm and code development for accelerator modeling, feasibility study on quantum computing implementation in accelerator modeling, and quantum computing education in accelerator community.*

The start-to-end simulation of an accelerator using real beams with billions or more particles remains challenging and expensive even with state-of-the-art exascale machines. The development of quantum computers, which can potentially provide an exponential improvement in efficiency for some classes of simulations, may bring new opportunities to enhance the particle accelerator community's simulation abilities.

The technique for building quantum computer has entered the Noisy Intermediate Scale Quantum (NISQ) era [198]. Current state-of-the-art quantum computers have above 50 qubits and quantum supremacy has been demonstrated on some specific problems [199, 200]. Quantum computing is currently available to the public through cloud services provided by some commercial companies such as IBM [201], D-Wave [202], Amazon [203] and Microsoft [204]. Meanwhile, studies on quantum algorithms have also been booming in the past few years. One hot topic is quantum machine learning. Machine learning can be used to optimize the performance of accelerators in design,

commissioning and operation. Some time-consuming simulations can also be replaced by well-trained learning systems. Quantum computing can speed up many learning algorithms, including SVM [205, 206], Bayesian networks [207], convolutional neural networks [208], Bayesian deep learning [209], reinforcement learning [210, 211], etc.

Another topic that is relevant for accelerator modeling is solving linear systems [212–216] as well as ordinary differential equations (ODEs) and partial differential equations (PDEs) [217–221]. Solving Poisson's equation and the Vlasov equations — which are often used in the simulation of collective effects such as space charge and beam-beam interactions — with quantum computing is being explored [222–224]. From the perspective of programming, quantum computer simulators, as code developing-, debugging-, and testing platforms, are available for almost all mainstream programming languages, e.g. C/C++, Python, Java, Matlab, etc. There also exist some languages specifically designed for quantum computing [225]. All these provide the community with the fundamental blocks to build simulation tools for beam and accelerator physics.

To make a problem suitable for quantum simulations, it has to be described by unitary operators so that the system, which may be classical, is mathematically equivalent to a quantum system. Additionally, in quantum computing, all variables are stored in quantum states and reading out the results accurately can be time intensive, which should be carefully managed. Many problems in accelerator modeling and beam dynamic simulations meet the above criteria. The motion of a charged particle, even under collective effects, can often be described as Hamiltonians, the evolution of which can be carried out by unitary operators. Recently, a quantum Schrodinger approach was used to solve the Poisson-Vlasov equations to simulate space-charge effects in high intensity beams [226]. Although particles are used to model beams, in many cases we care more about the macroscopic property of the beam, e.g., emittance, momentum spread, bunch size, luminosity, etc., rather than the individual particle and reading out the results of all the particles can be avoided.

However, most previous work on quantum algorithms focused on the realization of an algorithm with quantum circuits, but applying the algorithm to solve a practical scientific problem is seldom discussed [227]. Clearly there is a gap between the development of the algorithms and their implementation. To implement the quantum computing technique in accelerator modeling, we need to keep abreast with the latest developments and get involved in algorithm design and program development. We expect, at least in the near future, a quantum computer will not replace a classical computer but will probably work together with one. It is thus probably better for now to focus on how to make new quantum simulation tools that collaborate with the existing pool of accelerator modeling programs. Ideally a protocol will be invented through which the quantum packages could be called by the classical programs. Also needed are innovative analyzing tools to process the simulation results before the reading-out. To achieve these tasks, the accelerator physics problems that will benefit from quantum computing need to be identified and for each of them the proper mathematical model will need to be established, which may be different from the conventional classical one. For this, contributions from experts in both accelerator physics and quantum computing are required.

5.4 Storage ring quantum computers

contributed by Sandra Biedron, Jean-Luc Vay

> **Recommendation:** *Support research and development of storage ring quantum computers.*

In addition to benefit from quantum computing in the future, storage rings can be used as quantum computing devices themselves [228], with the advantages, over Paul [229, 230] or Penning

traps [231], of longer coherence time, the ability to store a significantly higher number of ions and the possibility of storing multiple ion chains [228]. A storage ring quantum computer (SRQC) could potentially store thousands of ions to form an Ion Coulomb Crystal in a closed trajectory. For SRQC, we anticipate unique challenges associated with the unprecedented large number of ions in the ICC. For example, on timing the laser pulses to efficiently cool the ICC beam to its equilibrium energy configuration, from which ion manipulations can be initiated. The solutions to the equilibrium positions of the ICC can be solved numerically for an arbitrary N, but they quickly become lengthy and impractical as N becomes large, from hours to days. A solution that employs AI/ML techniques was described recently that brings the prediction to a few seconds [232].

SRQC presents a particularly interesting case, where the design of the storage ring can benefit from high-performance computing and AI/ML, to produce quantum computing devices that will open unprecedented capabilities in computing with application to many fields, including beam and accelerator physics and designs themselves.

6   Computational needs

Recommendation: *Support the development of increasingly powerful and specialized High-Performance Computing (HPC) and High-Throughput Computing (HTC) capabilities for accelerator modeling, and maintenance of software to run efficiently on these hardware (e.g., port of codes to GPUs) with efficient and scalable I/O, post-processing and in situ data analysis solutions, which will be needed to support ab initio modeling at increasing fidelity, training of surrogate models and AI/ML guided designs.*

6.1   Hardware: CPU/GPU time, memory, archive

contributed by Jean-Luc Vay, Axel Huebl

Recommendation: *Support the development of increasingly powerful and specialized High-Performance Computing (HPC) and High-Throughput Computing (HTC) capabilities for accelerator modeling, which will be needed to support ab initio modeling at increasing fidelity, interactive and parallel data analysis, training of surrogate models and AI/ML guided designs.*

Many activities in accelerator modeling research and design both push the envelope in High-Performance Computing (HPC) as well as High-Throughput Computing (HTC). Typical HPC problems are spearheaded by plasma-based accelerator modeling, where simulations are limited by the available computational resources, whether in spatial grid resolutions, number of macroparticles, number of time steps, even if using the most powerful supercomputers available at the time of writing [233]. Emerging techniques like machine-guided optimization of ensemble studies and AI/ML workflows are often HTC workflows, where many parallel runs need to be coordinated. Likewise, parallel data analysis tasks are often interactive HTC workflows that furthermore can require interactivity and elasticity of the allocated compute resources (e.g., Jupyter Notebooks). Over time as computational resources available at leadership-scale computing facilities and individual institutions grow, some prior HPC workflows become HTC workflows as exploration phases evolve into design phases.

Some accelerator software have been or are being ported to GPUs and are being ready to run efficiently on upcoming Exascale supercomputers [234, 235]. It is expected that alongside GPUs,

additional specialized compute hardware such as Field Programmable Gate Arrays (FPGA), reduced instruction set (RISC) multicore-CPUs, tensor-processing units (TPUs) and specialized neural engines will be part of near-term computing architectures.

While the mentioned improvements on GPUs will increase the fidelity and accuracy of the simulations, the bar of fidelity and accuracy to aim at is being raised at the same time. The needs in computational hardware are expected to increase in the foreseeable future, whether to simulate up to every real particle, including in the halo, of beams that can be composite of many bunches, to account with high fidelity of collective effects in complex geometries or to simulate beams for very long times in conventional, plasma-based or structure-based wakefield colliders. While there is an increasing amount of efforts aimed at speeding up simulations with surrogate models, the training of these models will rely on simulations, in addition to experimental data. Producing the large volume of data from simulations will require tremendous needs in computational power, memory and archiving capabilities.

6.2 Software performance, portability and scalability

contributed by Axel Huebl, David Sagan

Recommendation: *Foster the development and maintenance of codes that run efficiently on the latest hardware (e.g., add support for GPUs/FPGAs) by using maintainable single-source, portable solutions, and that are scalable on leadership-scale supercomputers with multiple levels of parallelization and support for effective dynamic load balancing.*

6.2.1 Performance

Performance increase for modeling software is generally desirable, even outside of large-scale HPC runs, and improves overall scientific productivity when studying particle accelerators at all scales. For instance, most optimization and machine-learning workflows are GPU-accelerated as well to increase computational throughput for studies, while ultraprecise numerical schemes can often use GPU-accelerated linear algebra and FFT routines; thus GPU acceleration also benefits small model sizes.

Developing software for multiple GPU-platforms in parallel to existing CPU architectures is an undertaking that requires the redesign of many existing algorithms and code-bases. Programming models and languages evolve rapidly and the need to port algorithms to significantly more fine-grained parallelism is a good incentive to rewrite and adopt modern software practices and popular programming languages.

6.2.2 Portability

Portability of implementations is important as many scientists work with three major platforms (Linux, macOS, Windows), although the former two (Unix-like) variants are certainly dominating with HPC developers and machines. An additional challenge arises with the diversity in GPU and other compute hardware, which ideally should be addressed with *performance-portable* software design. In such a design, an algorithm is ideally implemented only in a single, abstract form and then specialized/optimized for the specific hardware needs at compilation-time.

Performance portability relies on using and contributing to standard, industry-supported programming languages to allow such a single-source approach. Performance portability layers (e.g., Kokkos [236], Alpaka [237], Raja [238]) are currently implemented predominantly as modern C++

libraries. The Julia language natively supports coroutines, multi-threading, distributed and GPU computing. Other languages lacking significantly in release time, compiler variety and robustness. It is to be expected that performance portability architecture will evolve into standardized programming language components, such as ISO C++, in coming years. HEP modeling needs can provide important input to their standardization by openly sharing critical community use cases and codes and interacting directly with computer scientists for development.

Common numerical and data management libraries (e.g., AMReX [239], CoPA [240]) build on such performance portability layers and only as the next steps, domain-specific libraries and applications are implemented to benefit from all aforementioned developments.

6.2.3   Parallel scalability

Parallel scalability on leadership-scale machines continues to be only achievable with additional parallelization over multiple coupled compute nodes (each powered with CPUs or GPUs). Methods that (a) split the computational domain into local sub-problems, and (b) ideally overlap communications (usually of data in guard cell regions) between collaborating nodes with computation within each node, are essential for such capability simulation runs. Furthermore, at large-scale, load balancing of particle-based methods as well as refinement of fidelity needs (e.g., adaptive mesh-refinement) of the simulation domain become important to make optimal use of available resources.

Looking at the rapidly evolving HPC landscape, many themes in advancing accelerator and beam modeling software are common, e.g., Poisson solvers, robust particle pushers, need to model hybrid particle accelerators combining conventional and advanced accelerator elements, etc. It is thus evident that collaborative efforts between conventional and advanced accelerator development can benefit the community, relying for instance on similar or compatible low-level parallelization layers, but also coordination when implementing and sharing domain-specific numerics can ensure that scalable solution can be readily extended for research needs.

6.3   Scalable I/O and in-situ analysis

contributed by Axel Huebl, Jean-Luc Vay

Recommendation: *Support the development and maintenance of efficient and scalable I/O, post-processing, in situ data analysis, and data sharing solutions in particle accelerator codes. Coordinate on scientific data documentation, standardization, development of interactive and reproducible analysis workflows and foster data reuse.*

While modern programming models, software design and advanced algorithms ensure that simulations can execute efficiently, one further challenge emerges to scientific productivity on HPC hardware: When comparing the computational peak performance evolution to the parallel system storage bandwidth of HPC machines as in Table 1, the gap between potential simulation fidelity that can be computed and data that can be stored for analysis widens with every new machine generation.

Following this trend, modeling applications that run at a system's capability and follow the traditional workflow of simulation + post-processing will spend more of their time in large-scale data input and output, which eventually can dominate the time-to-solution. Driven by this trend, several steps can be taken by simulation developers.

As a first step, adopting scalable, parallel I/O libraries with adequate data layouts [241] can extent the applicability to truly make optimal use of parallel filesystems [242]. Implementing these

System Specs	Titan	Summit	Frontier
Peak Performance	27 PFLOP	200 PFLOP	>1.5 EFLOP
Storage Bandwidth	1 TB/s	2.5 TB/s	2-4x Summit
Ratio	27:1	80:1	150-300:1

Table 1: Evolution of leadership-scale machines at the Oak Ridge Leadership Computing Facility with respect to peak performance and storage bandwidth, according to https://www.olcf.ornl.gov/frontier/.

libraries in both data producing as well as data consuming codes opens another possibility to standardize on and share the community solutions, as demonstrated in the openPMD project [99].

As a second step and long-term solution, data analysis and processing need to be transitioned to in situ-processing, running online with the modeling simulation and avoiding traditional, file-based long-term storage for raw data. This approach essentially moves the design of data analysis to the setup and design phase of a simulation that, similar to an experimental setup, needs to plan the locations, acceptance and dynamic ranges of "virtual diagnostics" before even starting a high-fidelity modeling run. Virtual diagnostics can reduce data sizes and thus data throughput needs by orders of magnitudes, e.g., by calculating beam momenta, phase space histograms, spectra and even 3D videos as a simulation runs instead of storing realistically trillions of particles in regular intervals to disk for later analysis. While in-situ data processing and analysis are not new, and have even been mainstream in some accelerator codes or frameworks (e.g., Warp [126], PIConGPU [243]), the discrepancy between peak performance and storage bandwidth reported in Table 1 triggers the development of community libraries that enable a more systematic description of new diagnostics in an in-situ approach, ensuring high efficiency and scalability.

A challenge arises in the rapid setup of in situ processing pipelines, since typical approaches of implementing a virtual detector in the simulation code itself can be more complicated compared to scripted (serial) analysis workflows on files. Innovative software design can potentially alleviate this challenge, with approaches such as streaming data with close-to-file-like analysis syntax [244] or embedding of interpreters into running simulations to execute scripts on in-memory data. Continued research in this area and close collaboration with computer science and engineering teams will be essential to ensure that flexible, science-case specific analysis can be designed by accelerator and beam physicists.

7   Sustainability, reliability, user support, training

contributed by Axel Huebl, Rémi Lehe, David Sagan, Christopher Mayes, Jean-Luc Vay, Kiersten Ruisard, Steven Lund

Recommendation: *Provide sufficient resources for code maintenance, automated testing, benchmarking, documentation and code reviews. Convene a community effort to identify topics and teaching teams to deliver academic classes designed to foster sharing and cooperation, to be taught at the U.S. Particle accelerator school.*

An important aspect of the development of scientific modeling tools is to ensure that the corresponding software is robust, easy to use and can be extended for new science cases. This is made more difficult by the fact that these tools are oftentimes constantly evolving, with new capabilities being continuously added to meet ever changing computational needs. Here we mention a few important practices [245] that facilitate this process.

Accelerator simulation is a large, complex topic, and much time and effort has been spent in developing simulation software. Nevertheless, in the field of accelerator physics, simulation code development has often been a haphazard affair. Due to developers retiring or moving on to other projects, numerous simulation programs have been completely abandoned or are seldom used.

As simulation programs become more complex due to the ever-increasing demands placed upon machine performance, the situation will become worse if not addressed. Such demands include the accelerator and beam physics Grand Challenges that have been identified recently by the community [91, 103]

- Increasing beam intensities by orders of magnitude.
- Increasing the beam phase-space density by orders of magnitude, towards the quantum degeneracy limit.
- Complete and highly accurate start-to-end "virtual particle accelerators" simulations.
- Fast and accurate multi-objective optimization methods to speed up the design process.

Accelerator and beam modeling software development should allow for extensive testing of new functionality while preserving demonstrated capabilities on previously validated scenarios[245]. Performance and interoperability must be constantly improved to increase understanding (multi-physics problems) and optimization (machine learning). One must also ensure a transfer of knowledge over generations of scientists in form of formal education and easy accessibility to the tools.

Addressing these Grand Challenges will require a community effort to coordinate and modernize the current set of modeling tools, with capabilities that extend far beyond what the current software can do, including interdisciplinary simulations and advanced models for virtual prototyping of complete accelerator systems.

7.1 Code robustness, validation & verification, benchmarking, reproducibility

Recommendation: *Establish and maintain open review, automated testing, validation and benchmark procedures for modeling tools and ensure reproducibility by tracking all changes in a documented and openly accessible manner.*
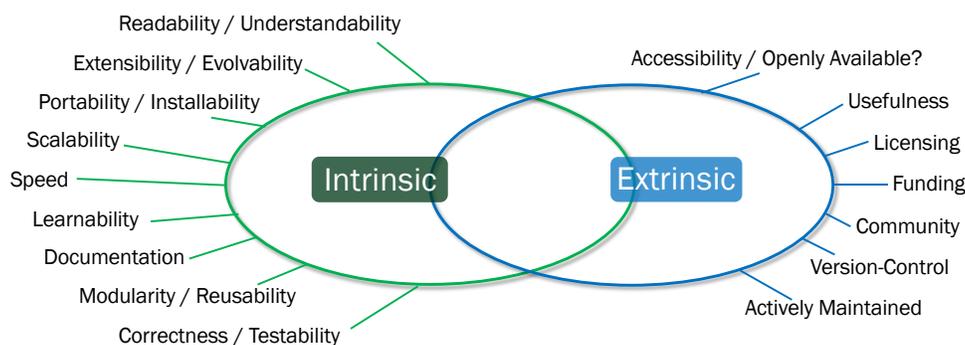
Figure 1: There are a number of aspects that make software sustainable. Broadly, they can be grouped into the "intrinsic" characteristics of the software itself and the "extrinsic" environment in which the software is developed and used.

There are several aspects that must be addressed to enable the development of the quality software that will be needed for the machines of tomorrow. One facet can be put under the rubric of "software sustainability" which can be defined as [246]:

"the capacity of the software to endure. In other words, sustainability means that the software will continue to be available in the future, on new platforms, meeting new needs."

There are many aspects to software sustainability, as illustrated in Fig 1. Broadly, these aspects can be grouped into the "intrinsic" characteristics of the software itself and the "extrinsic" environment in which the software is developed and used. Software sustainability has been studied academically and there is even a Software Sustainability Institute [247], which promotes "the advancement of software in research by cultivating better, more sustainable, research software to enable world-class research".

As mentioned above, many software packages developed for simulating accelerators, while showing excellent results for their specific application and their era, are not "sustainable" in the long run. However, software sustainability is extremely important given the limited resources that the accelerator community has for code development in conjunction with the even more limited resources for maintaining codes. To meet future needs, it is imperative that there is a community wide effort to promote sustainable practices.

Because modeling tools are being continuously improved and extended by developers, there is always a risk that a given change to the source code may introduce bugs and produce erroneous results in some specific cases. It is therefore paramount that the developers track the changes to the source code (with a version control tool such as `git`), and that they simultaneously maintain a comprehensive suite of tests that the code should always successfully pass. In the case of scientific modeling codes, these tests can consist of a number of physical setups for which there are well-known theoretical predictions. It can thus be checked that the results of the code conform to these predictions. In order for these tests to be most effective, they need to be run *automatically* whenever the source code is modified – a process known as *continuous integration*.

In addition to validating against theoretical predictions, it is also critical to benchmark against experimental observations. The relevance of code predictions to accelerator performance and output may be limited by a number of factors, such as:

- results from experiments that are not fully characterized (e.g., magnet misalignments)
- initial beam distribution is unknown or incompletely known (e.g., beam phase-space projections from experiments may not allow full reconstruction of beam 6-D phase-space accurately)
- diagnostics or simulations with insufficient range or resolution to fully characterize relevant parameters (e.g., phase space diagnostics that are sensitive to halo / beam loss levels [19])
- missing or incompletely described physics in simulations

Proper benchmarking of codes against theory and experiments is time-consuming and requires dedicated resources. In general, there are few resources dedicated to code benchmarking at experimental facilities and in simulation teams. As the complexity of codes and the computers that they run on increase, it is essential to ensure enough resources for proper benchmarkings, including dedicated beam time and test stands for well-characterized experimental results [19].

In addition, the development of large-scale scientific tools nowadays involve a whole community of developers and users rather than a single, isolated developer. Therefore, the robustness and quality of the software is generally greatly improved by the adoption of collaborative development platforms. These online platforms allow the developers to easily and efficiently review each other's changes to the source code, before these changes are incorporated in the mainline version of the code. They also provide a central venue for users to raise issues and questions, and generally communicate their needs with the developers.

Finally, another aspect of a code's robustness is its interface with the user. As the code evolves, the interface with the user often needs to be modified so as to enable functionalities or generalize existing ones. It is not uncommon that these changes break some of the users' scientific workflows. Thus, it is important for the developers to regularly publish releases of the source code, where these changes are clearly documented. The different releases of a given software can be made citable and available through an online archival service, so that users can roll back to a previous release if needed and scrutinize changes in a reproducible manner.

7.2 Usability, user support and maintenance

Recommendation: *Consider maintenance of code reviews for improvements, documentation, installation and testing to be central to the mission of usable scientific software: establish open feedback and support channels, perform regular releases with change logs, use permissive open source licensing whenever possible and cover all scientific functionality with automated tests.*

7.2.1 Usability

In order to maximize the impact of a given software tool on its scientific community, it is also key that this tool be user-friendly. For example, a common obstacle to the adoption of some scientific tools is their complex installation procedure, especially if many dependencies are involved. The developers can drastically simplify the installation procedure of a given software tool by making it available through modern package managers [248–250]. These package managers can automatically download and compile the source code as well as its set of dependencies. They can also ensure that compatible versions of the dependencies are being used. Importantly, some of these package managers can allow users to have the exact same installation workflow on large-scale HPC resources as on their desktop workstation.

Another key aspect of the user experience is the code documentation. For scientific modeling tools, the documentation usually includes sections on how to install and use the code, as well as a description of algorithm being used along with relevant references. Since, again, scientific tools are often continuously evolving, it can be beneficial that the documentation be generated automatically, whenever the source code is modified or whenever a new release is published. Just as the modeling tools themselves, documentation should be easily accessible for the whole community and easy to improve from developers and users alike.

7.2.2 User support

In order for scientific codes to continuously improve and meet the needs of the community, it is important to have a communication channel for users to provide feedback [245]. This feedback may include potential bug reports, questions on installation issues, and inquiries on typical usage. Again, many tools exist, including issue trackers (e.g., GitHub issues [251]), chat rooms (e.g., Gitter [252] or Slack [253]), mailing list, forums, stack-overflow, etc. These different tools have different purposes and target different segments of the user pool. The aim is thus certainly not for a developer to embrace all of them simultaneously, but rather to select one or two channels that are best suited for their particular code. These communication channels should also be clearly identified in the documentation.

It can also be helpful for these communications channels to be openly accessible and easily searchable, so as to avoid that multiple users report the same issues. In this regard, some of the

conversations in chat rooms may sometimes have to be transferred to an issue tracker, or summarized in the documentation.

Open user support channels establish significant synergies not only across related accelerator modeling projects, but with the computational physics community as a whole. For instance, documented computational and numerical issues, testing approaches, chosen conventions and implemented automation frameworks add significantly to the effectivity of the community, in terms of development speed, issue triage, installation, solution/work-around propagation, communication with high-performance computing centers and adoption of best-practices, to name a few.

7.2.3 Maintenance

Version control and software releases - It is now widely accepted, in the scientific community, that any significant software development project should be managed with a version control system (e.g., `git` [254]). However, a less common practice is that of regularly publishing new releases of the code.

Code releases are important because scientific codes are constantly evolving and improving, and, as part of this process, the code interface with the user may change, bugs may be fixed, and new bugs may occasionally be introduced. It is therefore not uncommon for a new version of the code to suddenly break a workflow that a user previously relied on. Well-documented releases (e.g., with a `CHANGELOG` document) help the user understand how the code has evolved with each release, and how to update their workflow to adapt to these changes. Archiving each release (e.g., with tools such as Zenodo [255]) may also allow the user to roll back to a previous version of the code, if a new version is altogether incompatible with a prior workflow. Also here, readily available services can speed up the generation of `CHANGELOG` documents and publication with package managers through so-called automated deployments.

Access and licensing - over the last decade, it emerged that scrutinizable implementations and low-barrier contributions to accelerator modeling software are highly desirable for scientific productivity. Consequently, a significant part of the community considers open source development [256–258] and permissive open source licensing [259, 260] the preferred way for software access, contribution, teaching, collaboration and reuse [245, 256, 261, 262]. For a few modules where such a model is not possible, e.g., due to export control constrains, other community sharing methods such as limited access through a memorandum of understanding (MoU) exist.

Automated testing - In order to minimize the risk of new bugs, it is important to verify that the code still works as expected, whenever the source code is modified. In the case of scientific codes, this can be done for instance by comparing the results of the code against known solutions, for a number of analytically-tractable problems. This process has long been done by hand, but experience shows that it is time-consuming and may in practice be often omitted by the developers.

Instead, these tests can be done automatically and systematically, whenever the source code is modified. Automated tests usually also results in faster code development, since the different contributors to the code can be rapidly assured that their changes do not break previous functionalities. A number of tools that are free to open source software allow to easily setup those automated tests, including GitHub Actions [263], Azure Pipelines [264] and Travis-CI [265]. Automated test on exotic hardware or HPC platforms is intricate, but addressed for instance by the E4S [266] effort within the Exascale Computing Project.

7.3 Training and education

Recommendation: *Convene a community effort to identify topics and teaching teams to deliver academic classes designed to foster sharing and cooperation. The classes should be taught in the US Particle Accelerator School with course materials and linked tutorials/extensions regularly maintained and publicly posted.*

Training and education are key to the long term viability of positions put forth in this white paper. University classes cover the core of essential programming skills, numerical algorithms, AI/ML, and some software development and maintenance tools. Traditionally workers in the field rely heavily on online documentation and training materials as well as help from colleagues. Much benefit is derived from the larger plasma and EM modeling communities. Accelerator science and engineering also relies on specialized training provided by the US Particle Accelerator School (USPAS). The USPAS convenes two, intensive-format sessions per year to provide graduate-level training in topics that are not practical to teach regularly at universities. Recommendations for enhancing the USPAS are covered in a separate Education, Outreach, and Diversity white paper [267]. Specific to accelerator modeling, USPAS offerings have been primarily centered on code demonstrations and applications in various accelerator dynamics, design, and applications classes. Recent classes have covered optimization and ML, python/matlab applications, and levels of self-consistent modeling in beams and plasmas. It would benefit the community if more offerings were regularly available in the USPAS. The USPAS classes are also an excellent opportunity to foster community cooperation via team teaching efforts and to help ingrain productive collaborative work approaches in the younger generation. Focused topical collaborations centered on classes needed could also be used to generate needed tutorial materials as well as disseminate and archive useful code tools while demonstrating the advantages of effective use in classes. Tutorials and archives of materials from the courses would improve modeling efforts and encourage increased productivity via sharing.

8 Toward community ecosystems & data repositories

contributed by Jean-Luc Vay, Axel Huebl, Rémi Lehe, Warren B. Mori, Frank Tsung, David Bruhwiler, Ji Qiang, David Sagan

Recommendation: *Organize the beam and accelerator modeling tools and community through the development of (a) ecosystems of codes, libraries and frameworks that are interoperable via open community data standards, (b) open access data repositories for reuse and community surrogate model training, (c) dedicated Centers and distributed consortia with open community governance models and dedicated personnel to engage in cross-organization and -industry development, standardization, application and evaluation of accelerator and beam modeling software and data.*

8.1 Loose integration: Integrated workflows

Recommendation: *Foster the adoption and continued development of open community (meta-)data standards and their implementation in modeling tools and data acquisition for seamless integration of the community accelerator modeling tools into multiphysics workflows.*

Integrated workflows are often needed to answer complex contemporary research questions, such as the start-to-end description of a hybrid accelerator with conventional and plasma elements.
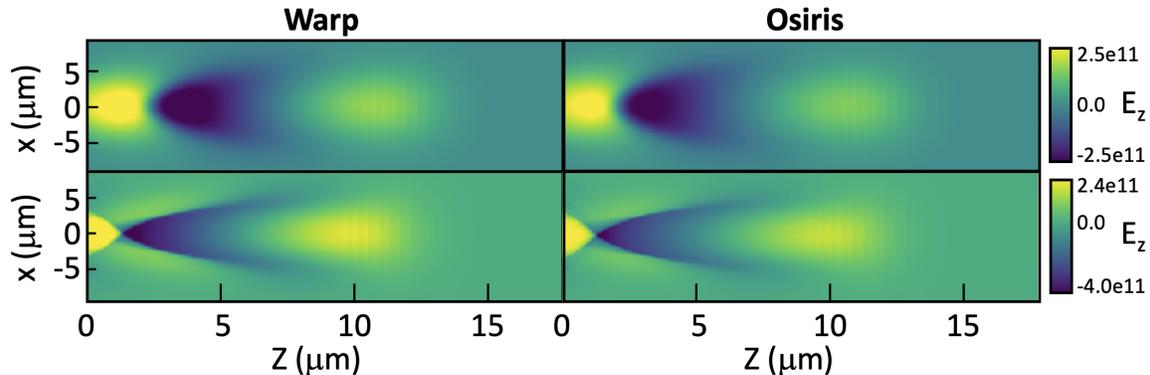
Figure 2: Longitudinal electric field (in V/m) in a laser-driven plasma acceleration stage at two times (top: $t \approx 300\text{fs}$, bottom: $t \approx 600\text{fs}$) along the laser propagation from 2-D PIC simulations with: (left) Warp; (right) Osiris. Plots are based on rendering from the openPMD-viewer.

Furthermore, all optimization, AI/ML training as well as exploratory studies benefit from establishing and maintenance of workflows. These workflows might (a) link several codes together to solve a larger problem or solve a problem with more physical processes, (b) benchmark and validate the various codes against each other and known solutions, and (c) establish reproducibility for new scientific results.

8.1.1 I/O standardization

At the moment, many simulation codes are developed independently with little coordination between various groups. In order to establish workflows that span multiple applications, a typical approach so far is to rely on file-based data exchange and code-specific high-level wrapping of input options. Along these lines, the Consortium for Advanced Modeling of Particle Accelerators (CAMPA [268]) supports (among other activities) the development of standardized input and output formats through the openPMD [99] and PICMI [269] projects.

The Open Standard for Particle-Mesh Data (openPMD) is a meta-data standard for research data, adding scientific self-description on top of modern, portable, scalable file-formats such as ADIOS [270] and HDF5 [271]. openPMD achieves this by collaboratively defining an open standard document that evolves in versions, similar to software. Community members then implement openPMD data handling in their simulation codes and analysis types either directly against supported file formats or use an open source library layer, openPMD-api [272].

The standard input format for Particle-In-Cell codes (PICMI) strives to unify the usage of accelerator modeling codes by standardization on an API for simulation design. Conceptually, PICMI defines a common, explicit PIC modeling interface (API) for a wide range of numerical and initialization options. A PICMI script can be run with multiple codes, given that both implement the requested features, and simplifies comparability of physics cases tremendously. Code "backend" choices can be picked based on performance or numerical needs, without breaking context for the user and relearning the simulation design for each specific application.

As part of the CAMPA efforts to foster compatibility and coordination between particle accelerator modeling applications, a code-validation workflow between the Warp [126] and Osiris [273] code has been carried out. In the study on the two codes in Figure 2, both independent code bases modeled the same physical setup of an LWFA (with a laser driving a wake in a plasma column) and used standardized openPMD output and the same, community-developed post-processing tool

(openPMD-viewer) to analyze the data. A one-to-one comparison to machine precision between the two codes was beyond the scope of the study because: (a) while using the same physical parameters to initialize the simulations with each code, some secondary parameters were left unspecified such as, e.g., the exact phase of the laser oscillations, (b) some numerical aspects are different such as, e.g., the method of initialization of the laser in the boosted frame. Nonetheless, it is rewarding to observe on Figure 2 that the two codes predict very similar wake evolutions and dephasing of the laser as it propagates through the plasma column.

Another series of comparisons is underway that also incorporates a common Python input script using PICMI, and including more codes, in the workflow, to validate the workflow concept for further adoption and development.

Another example is the integrated electromagnetics and beam dynamics simulation on high performance computers using the ACE3P and the IMPACT codes [274]. In this case, the design of the RF cavity using the ACE3P and the beam dynamics simulation using the IMPACT were integrated into an optimization workflow for final beam quality optimization.

8.2 Tighter integration: Integrated frameworks

Recommendation: *Establish open, contributable, modular libraries and integrated frameworks for central computing, modeling and analysis tasks that foster the sharing of common functionalities between applications, using open licenses and best practices/policies.*

Integrated developments of frameworks could implement more tightly coupled workflows on a library level instead of relying on high-level application coupling. As presented at the beginning of the section in the concept of an modular software ecosystem in [6], modular, cross-cutting software designs can provide reusability, efficient development investments, and compatibility. Built-in features of lower-level software such as GPU-support, multi-node parallelism, translation to AI/ML frameworks and multi-physics support can, if evolved together, be inherited by a wide community. In particular, common or compatible in-memory data structures have not been attempted in accelerator and beam modeling so far, but would provide significant performance and software maintenance benefits for tightly-coupled, co-developed physics modules for integrated modeling frameworks that address grand challenges.

Several community examples exist to demonstrate modular development. For instance, the codes WarpX [233], HiPACE++ [275] and PIConGPU [243] as well as several analysis frameworks share the development of the I/O library openPMD-api [272], which provides access to high-performant, portable low-level I/O libraries and a standardized metadata schema [99]. WarpX, FBPIC implement PICMI assisted by a shared, central interface library, with implementation underway in other codes (e.g., OSIRIS). WarpX, HiPACE++ [276] and related codes share the data structures from AMReX [239] and parts of their code bases.

8.2.1 Community development

The adoption of a more coordinated and collaborative approach is also driven by the need to transition a large body of software from CPUs to GPUs, which is more disruptive than most past transitions, including the one from serial to parallel codes. Furthermore, this latest transition is anticipated to be one of many transitions to come to adapt to the rapidly evolving computer hardware.

More realisms, better numerics and advanced algorithms are continuously being added to many accelerator codes, and code developers have developed best practices for adding new modules while maintaining the performance needed to run on state-of-the-art HPC centers around the world. Thus,

instead of porting solely to GPUs, accelerator modeling teams already teamed up with computer scientists and embrace libraries that abstract computer hardware specific details for performance and sustainability. Establishing more community standards, funding reliable software dependencies and reusable physics modules are the logical steps to speed up development of code bases further.

An efficient mean for guiding community development is the collaborative setting of and adherence to common policies. Rather than starting from scratch, one can adopt an existing libraries and codes and build a beam and accelerator modeling ecosystem on it. As a blueprint, one could build on community development concepts of the Interoperable Design of Extreme-scale Application Software (IDEAS) project [277] and its Extreme-scale Scientific Software Development Kit (xSDK) [278].

The main goal of the IDEAS project is to "help move scientific software development toward an approach of building new applications as a composition of reusable, robust, and scalable software components and libraries, using the best available practices and tools."[277], while the xSDK project is being developed to "provide a coordinated infrastructure for independent mathematical libraries to support the productive and efficient development of high-quality applications" [278]:

> Rapid, efficient production of high-quality, sustainable extreme-scale scientific applications is best accomplished using a rich ecosystem of state-of-the art reusable libraries, tools, lightweight frameworks, and defined software methodologies, developed by a community of scientists who are striving to identify, adapt, and adopt best practices in software engineering. The vision of the xSDK is to provide infrastructure for and interoperability of a collection of related and complementary software elements—developed by diverse, independent teams throughout the high-performance computing (HPC) community—that provide the building blocks, tools, models, processes, and related artifacts for rapid and efficient development of high-quality applications.

Although xSDK targets "extreme-scale scientific applications" of relevance to exascale super-computing, its derived community policies apply well to all scales of computing, from laptops to clusters or Cloud computing. Hence, the paradigm and tools are readily applicable to the full set of modeling needs of the particle accelerator community.

Individual packages in such an ecosystem are composable libraries, frameworks, and domain components developed by individual groups in the community. Each package publishes technical design documents, documentation (reference, tutorials, how-to guides), API definitions, and a concrete implementation [245]. In order to solve a specific physics case, a typical application uses functionalities from several packages, often in an innovative way.

There are many advantages to adopting policies and tools such as the IDEAS/xSDK:

- Community policies have been established over years by teams of specialists in scientific software development and computing, and they include best practices in software development.
- For reusable components, the policies require the use of *permissive open source licenses* ("Non-critical optional dependencies can use any OSI-approved license.") that allow reuse of source and binary code, including for commercial and proprietary applications, fostering collaborations across laboratories, academia and industrial partners.
- A wide-ranging set of open source, interoperable tools from a variety of backgrounds (including numerical solvers, e.g., Hypre, multi-parameter optimizer, e.g., LibEnsemble, and mesh-refinement frameworks, e.g., AMReX) can be combined and used as foundation of accelerator and beam physics toolkits and codes.
- A wide community of developers that can help improve software capability and sustainability across projects.
- Partial overlap in functionality is not problematic and in fact improves the diversity of the whole.

- Time of development and maintenance of domain-specific software is greatly reduced, as the domain scientists can concentrate on the domain-specific functionalities while lower-level, cross-cutting, numerical packages are maintained by dedicated specialists.
- Portability across platforms (CPUs, GPUs, etc.) of low-level libraries ensures portability of the software that is built upon them. Building new tools on these low-level libraries greatly reduces the overall burden on the community for porting codes to new platforms.
- The domain-specific packages that are built, following established policies are portable across platforms and interoperable, and can be further combined to form larger toolkits and codes.
- Since packages are reused across software and duplication is minimized, bugs and inefficiencies are spotted earlier.
- Many codes, libraries and packages exist in the community that can be reused as building blocks, and progressively modernized and blended in a module that is portable and runs efficiently on modern CPUs and GPUs. Hence, not all functionalities have to be rewritten from scratch at once, offering a progressive path from the current status to an ecosystem of accelerator modeling tools.

It is to be emphasized that the goal is not to propose that only one code be developed to study, e.g., RF accelerators or plasma-based accelerators. On the contrary, the goal is to provide a coordinated infrastructure that enables inclusive and collaborative development of modeling tools for the community, and for these tools, to follow modern practices for scientific software: be portable, efficient, robust and leading to consistently accurate and reproducible results.

8.3 Data repositories

Recommendation: *Establish open access data repositories and foster publishing of modeled and measured accelerator & beam data to allow re-use (e.g. beam transport to applications), model training (e.g. AI/ML), preservation, recasting and reinterpretation.*

In current community practice, the reporting and archival of many modeling results are often limited to prose and figures in papers. As accelerator systems and support for experimental campaigns become more complex, it is desirable to establish reuse of modeling results, especially if simulations themselves made use of significant HPC resources.

Similar to needs in HEP detector data [279], it is desirable to systematically categorize and share accelerator data, such as beam distributions, source information, field maps, among others. The benefits of such archives are manifold: sharing data enables preservation, recasting, reinterpretation [279] and meta-analysis of published results; archives establish training sets for AI/ML workflows, calibration and benchmark cases for new models and theories, beam transport from one model (and code) into another, design of hybrid particle accelerators combining conventional and plasma elements/sources, accurate input for modeling at interaction points, etc. The standardization of data and meta-data (see previous section) for machine-readability, findability, accessibility and reuse would maximize the productivity when sharing data over such repositories.

Following Open Science practices [280, 281], such data archives and repositories should be established in an open access manner [282], similar to existing data repositories in particle physics experiments [283, 284]. Lastly, open science compatible data policies need to be advanced on the national, laboratory and university level. For instance, as of today data sharing in the U.S. Department of Energy laboratories is bound to first assert (often individually) copyright for all but raw data, which can produce significant workflow overheads for scientists and entry burdens to contribute to community data repositories. The policies for universities vary with individual states.

Clear advise also needs to be established for corner-cases such as publishing interactive notebooks and code snippets, which traditionally fall under software copyright and publishing policies, with data sets in a single and time-effective authorization and license.

8.4 Centers & consortia, collaborations with industry

Recommendation: *Organize the beam and accelerator modeling community through the development of dedicated Centers and distributed consortia. Dedicate resources to adopt open community governance models and dedicate personnel to engage in cross-organization and -industry cooperation.*

8.4.1 Centers & consortia

It quickly becomes clear that, in order to achieve what is described and proposed in the earlier sections, a coherent and consolidated effort is needed. This is best achieved in the form of dedicated Centers for Accelerator and Beam Physics Modeling [285]. Other areas of computer science have already embraced this. New colleges for computing are established at universities to consolidate the dispersed computing efforts of the various departments (e.g., MIT's Schwarzman College of Computing [286]), and new centers for Quantum Computing [287, 288] have been built. Exascale computing has been embraced through the Exascale Computing Project [289]. The US Department of Energy (DOE) has founded SciDAC [290] to accelerate progress in scientific computing across the different programs supported by DOE: Advanced Scientific Computing Research, Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, High-Energy Physics, and Nuclear Physics.

The respective communities have benefited strongly from these new centers and the partnerships across disciplines.

Accelerator and Beam Physics Modeling would no doubt benefit similarly. The centers can be at a given location or distributed geographically and among institutions across laboratories, academia and industrial partners. They would bring together domain scientists (computational accelerator and beam physicists), applied mathematicians, computer scientists and software engineers with collaborations across the full landscape of accelerator modeling. In addition, some of the computer science centers mentioned above are already supporting accelerator modeling efforts on which the Centers for Accelerator and Beam Physics Modeling could build.

Depending on the overall size, the centers could enable part or all of the following:

- Community development and maintenance of codes using industry-standard quality processes by dedicated, specialized teams [245].
- Collect libraries for field solvers, particle trackers, and other modules.
- Provide a modular community ecosystem for multiphysics particle accelerator modeling and design [291].
- Standardize input scripts, output data, lattice description and start-to-end workflows [100].
- Provide compatibility layers to use the same libraries and modules in a number of programming languages.
- Development and maintenance of End-to-end Virtual Accelerators (EVA) [91].
- User support, high-quality and detailed documentations, online tutorials, and training.
- Easy-to-use, standardized, user interfaces for preparation and analysis of simulations.
- Automated tools for ensemble simulations for optimization with builtin AI/ML support.

- Suite of test problems with well-characterized solutions for benchmarking, quality assurance and regression testing.
- Development, analysis and efficient implementation of novel algorithms and numerical methods (e.g., high-order solvers, symplectic multiparticle tracking, Fast Multipole Methods, adaptive mesh refinement).
- Providing a space to meet (physically or virtually) for the integration of developments from contributors into larger codes, such as PhD projects from external groups, organizing development hackathons, knowledge-transfer, and onboarding.
- Developing and organizing workshops for developers and users of codes alike. Inviting national and international speakers/developers (travel/hosting funds).
- Interacting with existing schools, by developing and maintaining state-of-the-art educational resources (e.g., tutorials, lectures) on codes.
- Exploration of novel use of machine learning for accelerator modeling, and, further in the future, of quantum computing [92].

Multiple Centers can be organized through a Consortium (e.g., CAMPA [268]). Except for special restrictions such as export control, it would be desirable for the software developed by the Center to be open source, enabling crosschecking, testing and contribution by the community at large, beyond the participants to the Center(s) [256].

8.4.2 Collaborations with industry

Effective long-term collaboration between national laboratories, academia and industry will lead to important benefits for the entire HEP community. Labs and universities will have access to better software with lower lifecycle costs [292]. Companies will be strengthened by knowledge transfer from labs and universities. Computational scientists will be better able to concentrate on core competencies, without spending time on user interface design, ease of use, cloud computing, etc. Society will reap the benefits of better science, more innovation, and stronger businesses. State-of-the-art simulation codes will become readily available to students. Training time and associated costs will be reduced, as new team members will become productive more quickly. This will contribute to equity, diversity and inclusion (EDI), as barriers to entry are removed for scientists in developing countries and for those at US institutions with less federal funding and no direct access to code developers.

In order to facilitate the necessary collaborations and to provide the entire community with confidence that the software will be widely available and adequately supported over decades, an open source license is required for the industry software and can be very helpful for the entire software ecosystem [256]. This imposes an open source business model on the corresponding businesses, at least with regard to this specific activity. The software design objectives must include seamless integration with legacy codes, low barrier to entry for new users, easily moving between GUI and command-line modes, cataloging of provenance to aid reproducibility, and simplified collaboration through multi-modal sharing.

9 Outlook

Computer simulations will continue to be essential to particle accelerator research, design and operation. Its relative importance is even expected to grow, thanks to improvements in algorithms, computer hardware, and new opportunities in machine learning and quantum computing. These will enable accelerator modeling capabilities that include more physics that is integrated self-consistently

to model accelerators with ever increasing fidelity and accuracy, toward the ultimate realization of the grand challenge of virtual twins of particle accelerators. A more collaborative and coordinated approach that enables the development of community ecosystems, adopting best practices in software developments and maintenance, is needed to meet the challenge in a realistic budget envelope and timeframe.

References

[1] *Accelerating Discovery: A Strategic Plan for Accelerator R&D in the U.S., Report of the High Energy Physics Advisory Panel (HEPAP) Accelerator Research and Development Subpanel, 2015.* https://science.osti.gov/-/media/hep/hepap/pdf/Reports/Accelerator_RD_Subpanel_Report.pdf.

[2] *Advanced Accelerator Concepts Research Roadmap Workshop Report.* 2016. URL: https://science.energy.gov/~/media/hep/pdf/accelerator-rd-stewardship/Advanced_Accelerator_Development_Strategy_Report.pdf.

[3] ALEGRO. "Towards an Advanced Linear International Collider". In: (Jan. 2019). URL: http://arxiv.org/abs/1901.10370.

[4] S. Nagaitsev et al. "Accelerator and Beam Physics Research Goals and Opportunities". In: (Jan. 2021). URL: https://arxiv.org/abs/2101.04107.

[5] D. Sagan et al. "Simulations of future particle accelerators: issues and mitigations". In: *Journal of Instrumentation* 16.10 (Oct. 2021), T10002. DOI: 10.1088/1748-0221/16/10/t10002. URL: https://doi.org/10.1088/1748-0221/16/10/t10002.

[6] J.-L. Vay et al. "Modeling of advanced accelerator concepts". In: *Journal of Instrumentation* 16.10 (Oct. 2021), T10003. DOI: 10.1088/1748-0221/16/10/t10003. URL: https://doi.org/10.1088/1748-0221/16/10/t10003.

[7] Carole Goble. "Better Software, Better Research". In: *IEEE Internet Computing* 18.5 (2014), pp. 4–8. DOI: 10.1109/MIC.2014.88.

[8] A. Fiedman et al. "3D particle simulation of beams using the warp code". In: *Part. Accel.* 37-38 (1992), pp. 131–139.

[9] Ji Qiang et al. "An Object-Oriented Parallel Particle-in-Cell Code for Beam Dynamics Simulation in Linear Accelerators". In: *Journal of Computational Physics* 163.2 (2000), pp. 434–451. ISSN: 0021-9991. DOI: https://doi.org/10.1006/jcph.2000.6570.

[10] J. Amundson et al. "Synergia: An accelerator modeling tool with 3-D space charge". In: *Journal of Computational Physics* 211.1 (2006), pp. 229–248. ISSN: 0021-9991. DOI: https://doi.org/10.1016/j.jcp.2005.05.024.

[11] Ji Qiang et al. "Three-dimensional quasistatic model for high brightness beam dynamics simulation". In: *Physical Review Special Topics - Accelerators and Beams* 9.4 (Apr. 2006), p. 044204. ISSN: 1098-4402. DOI: 10.1103/PhysRevSTAB.9.044204. URL: https://link.aps.org/doi/10.1103/PhysRevSTAB.9.044204.

[12] M.A. Furman et al. "Self-consistent 3D modeling of electron cloud dynamics and beam response". In: *Proceedings of the IEEE Particle Accelerator Conference.* 2007. ISBN: 1424409179. DOI: 10.1109/PAC.2007.4441093.

[13] Andreas Adelmann et al. "OPAL a Versatile Tool for Charged Particle Accelerator Simulations". In: *arXiv:1905.06654 [physics]* (May 2019). arXiv: 1905.06654. URL: http://arxiv.org/abs/1905.06654 (visited on 05/17/2019).

[14] C. Mayes. "Computational approaches to Coherent Synchrotron Radiation in two and three dimensions". In: *Journal of Instrumentation* 16 (2021), P10010.

[15] P. Yu, Y. Wang, and W. Huang. "Lattice and beam dynamics for the pulse mode of the laser-electron storage ring for a Compton x-ray source". In: *Phys. Rev. ST Accel. Beams* 12 (2009), p. 061301.

[16] J. Qiang, R. D. Ryne, and S. Habib. "Self-consistent Langevin simulation of Coulomb collisions in charged-particle beams". In: *Proceedings of the IEEE/ACM SC2000 Conference*. Dallas Texas, USA, Nov. 4–10, 2000.

[17] He Zhang and Martin Berz. "The fast multipole method in the differential algebra framework". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 645.1 (2011). The Eighth International Conference on Charged Particle Optics, pp. 338–344. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2011.01.053.

[18] A. Al Marzouk et al. "Efficient algorithm for high fidelity collisional charged particle beam dynamics". In: *Phys. Rev. ST Accel. Beams* 24 (2021), p. 074601.

[19] Kiersten Ruisard, Alexander Aleksandrov, and Sarah Cousineau. "Loss prediction through modeling of high dynamic range beam distributions". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/AF/SNOWMASS21-AF1_AF2_Kiersten_Ruisard-145.pdf.

[20] L. Xiao et al. "Advances in Multiphysics Modeling for Parallel Finite-Element Code Suite ACE3P". In: *IEEE J. Multiscale and Multiphysics Comp. Tech.* 4 (2019), pp. 298–306. DOI: 10.1109/JMMCT.2019.2954946.

[21] V. Akcelik et al. "SRF Cavity Imperfection Studies Using Advanced Shape Uncertainty Quantification Tools". In: *Proceedings of LINAC08*. Victoria, BC, Canada, Sept. 29–Oct. 3, 2008.

[22] A. Lunin et al. "Statistical analysis of the eigenmode spectrum in the SRF cavities with mechanical imperfections". In: *Proceedings of 13th Int. Computational Accelerator Physics Conf.* Key West, FL, USA, Sept. 20–24, 2018.

[23] C K Birdsall and A B Langdon. *Plasma Physics Via Computer Simulation*. Adam-Hilger, 1991, Xxvi+479 Pp. ISBN: 0 07 005371 5.

[24] B B Godfrey and MISSION RESEARCH CORP ALBUQUERQUE NM. *The IPROP Three-Dimensional Beam Propagation Code*. Defense Technical Information Center, 1985. URL: https://books.google.com/books?id=hos_OAAACAAJ.

[25] A F Lifschitz et al. "Particle-in-Cell modelling of laser plasma interaction using Fourier decomposition". In: *Journal of Computational Physics* 228.5 (2009), pp. 1803–1814. ISSN: 0021-9991. DOI: http://dx.doi.org/10.1016/j.jcp.2008.11.017. URL: http://www.sciencedirect.com/science/article/pii/S0021999108005950.

[26] C. Benedetti et al. "Efficient modeling of laser-plasma accelerators with INF&RNO". In: *AIP conference proceedings* 1299.1 (2010), pp. 250–255.

[27] A. Davidson et al. "Implementation of a hybrid particle code with a PIC description in r–z and a gridless description in ϕ into OSIRIS". In: *Journal of Computational Physics* 281 (2015), p. 1063.

[28] P Sprangle, E Esarey, and A Ting. "Nonlinear-Theory Of Intense Laser-Plasma Interactions". In: *Physical Review Letters* 64.17 (Apr. 1990), pp. 2011–2014. ISSN: 0031-9007.

[29] P. Mora and T. M. Antonsen. "Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas". In: *Physics of Plasmas* (1997), p. 217.

[30] C. K. Huang et al. "QUICKPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas". In: *Journal of Computational Physics* 217 (2006), p. 658.

[31] W. An ad V. K. Decyk, W. B. Mori, and T. M. Atonsen. "An improved iteration loop for thre three dimensionnal quasi-static particle-in-cell algorithm: QuickPIC". In: *Jour. Comp. Phys.* 250 (2013), pp. 165–177.

[32] T Mehrling et al. "HiPACE: a quasi-static particle-in-cell code". In: *Plasma Physics and Controlled Fusion* 56.8 (July 2014), p. 084012. DOI: 10.1088/0741-3335/56/8/084012. URL: https://doi.org/10.1088/0741-3335/56/8/084012.

[33] F. Li et al. "A quasi-static particle-in-cell algorithm based on an azimuthal Fourier decomposition for highly efficient simulations of plasma-based acceleration: QPAD". In: *Computer Physics Communications* 261 (2021), p. 107784.

[34] J.-L. Vay. "Noninvariance Of Space- And Time-Scale Ranges Under A Lorentz Transformation And The Implications For The Study Of Relativistic Interactions". In: *Physical Review Letters* 98.13 (2007), pp. 1–4. ISSN: 0031-9007.

[35] J. Picard and al. *Generation of 1 GW of 11.7 GHz Power using a Metamaterial-based Power Extractor for Structure-based Wakefield Acceleration.* Presented at the APS April Meeting 2021, Volume 66, Number 5. 2021. URL: https://meetings.aps.org/Meeting/APR21/Session/T08.1.

[36] C. Jing et al. "Electron acceleration through two successive electron beam driven wakefield acceleration stages". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 898 (2018), pp. 72–76. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2018.05.007. URL: https://www.sciencedirect.com/science/article/pii/S0168900218305928.

[37] B. D. O'Shea et al. "Observation of acceleration and deceleration in gigaelectron-volt-per-metre gradient dielectric wakefield accelerators". In: *Nature Communications* 7.1 (Sept. 2016), p. 12763. ISSN: 2041-1723. DOI: 10.1038/ncomms12763. URL: https://doi.org/10.1038/ncomms12763.

[38] Q. Gao et al. "Observation of High Transformer Ratio of Shaped Bunch Generated by an Emittance-Exchange Beam Line". In: *Phys. Rev. Lett.* 120 (11 Mar. 2018), p. 114801. DOI: 10.1103/PhysRevLett.120.114801. URL: https://link.aps.org/doi/10.1103/PhysRevLett.120.114801.

[39] R. Roussel et al. "Single-Shot Characterization of High Transformer Ratio Wakefields in Nonlinear Plasma Acceleration". In: *Phys. Rev. Lett.* 124 (4 Jan. 2020), p. 044802. DOI: 10.1103/PhysRevLett.124.044802. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.044802.

[40] S. S. Baturin and A. Zholents. "Stability condition for the drive bunch in a collinear wakefield accelerator". In: *Phys. Rev. Accel. Beams* 21 (3 Mar. 2018), p. 031301. DOI: 10.1103/PhysRevAccelBeams.21.031301. URL: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.21.031301.

[41] *ACE3P*. URL: https://portal.slac.stanford.edu/sites/ard_public/acd/Pages/Default.aspx.

[42] *WarpX*. URL: https://github.com/ECP-WarpX/WarpX.

[43] *Advanced Accelerator Development Strategy Report: DOE Advanced Accelerator Concepts Research Roadmap Workshop*. English. Tech. rep. Sponsor Org.: USDOE Office of Science (SC), High Energy Physics (HEP). United States, 2016. DOI: 10.2172/1358081. URL: https://doi.org/10.2172/1358081.

[44] S. Nagaitsev et al. *Accelerator and Beam Physics Research Goals and Opportunities*. 2021. arXiv: 2101.04107 [physics.acc-ph].

[45] A. A. Sahai. "Nanomaterials Based Nanoplasmonic Accelerators and Light-Sources Driven by Particle- Beams". In: *IEEE Access* 9 (2021), pp. 54831–54839. DOI: 10.1109/ACCESS.2021.3070798. URL: https://ieeexplore.ieee.org/document/9395099.

[46] A. A. Sahai. "Emergence of TeraVolts per meter plasmonics using relativistic surface plasmonic modes". In: *Proc. Vol. 11797, Plasmonics: Design, Materials, Fabrication, Characterization, and Applications XIX*. Vol. 117972A. SPIE, 2021. DOI: 10.1117/12.2596637. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11797/2596637/Emergence-of-TeraVolts-per-meter-plasmonics-using-relativistic-surface-plasmonic/10.1117/12.2596637.short?SSO=1.

[47] A. A. Sahai. *anostructure nanoplasmonic accelerator, high-energy photon source, and related methods*. 2021. URL: https://patentimages.storage.googleapis.com/55/0a/6c/7b757efb9e9547/WO2021216424A1.pdf.

[48] A. A. Sahai. "Plasmonic nano-focusing of particle beams by surface crunch-in plasmons excited in tapered tubes". In: *Proc. SPIE 11999, Ultrafast Phenomena and Nanophotonics XXVI*. Vol. 1199903. SPIE, 2022. DOI: 10.1117/12.2605722. URL: https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11999/2605722/Plasmonic-nano-focusing-of-particle-beams-by-surface-crunch-in/10.1117/12.2605722.short.

[49] D. Bohm and D. Pines. "A Collective Description of-Electron Interactions: III. Coulomb Interactions in a Degenerate Electron Gas". In: *Phys. Rev.* 92 (1953), p. 609. DOI: 10.1109/ACCESS.2021.3070798. URL: https://ieeexplore.ieee.org/document/9395099.

[50] A. Sahai and collaborators. *PetaVolts per meter Plasmonics: Snowmass21 White Paper*. 2022. URL: https://arxiv.org/abs/2203.11623.

[51] J. M. Dawson. "Nonlinear Electron Oscillations in a Cold Plasma". In: *Phys. Rev. (1959)* 113 (1959), p. 383. DOI: 10.1109/ACCESS.2021.3070798. URL: https://ieeexplore.ieee.org/document/9395099.

[52] C. G. R. Geddes et al. "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding". In: *Nature* 431.7008 (Sept. 30, 2004), pp. 538–541. ISSN: 0028-0836. URL: http://dx.doi.org/10.1038/nature02900.

[53] W. P. Leemans et al. "GeV electron beams from a centimetre-scale accelerator". In: *Nat Phys* 2.10 (Oct. 2006), pp. 696–699. ISSN: 1745-2473. URL: http://dx.doi.org/10.1038/nphys418.

[54] T. P. A. Ibbotson et al. "Laser-wakefield acceleration of electron beams in a low density plasma channel". In: *Phys. Rev. ST Accel. Beams* 13 (3 Mar. 2010), p. 031301. DOI: `10.1103/PhysRevSTAB.13.031301`. URL: `http://link.aps.org/doi/10.1103/PhysRevSTAB.13.031301`.

[55] W. P. Leemans et al. "Multi-GeV Electron Beams from Capillary-Discharge-Guided Sub-petawatt Laser Pulses in the Self-Trapping Regime". In: *Phys. Rev. Lett.* 113 (24 Dec. 2014), p. 245002. DOI: `10.1103/PhysRevLett.113.245002`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.113.245002`.

[56] B. E. Blue et al. "Plasma-Wakefield Acceleration of an Intense Positron Beam". In: *Phys. Rev. Lett.* 90 (21 May 2003), p. 214801. DOI: `10.1103/PhysRevLett.90.214801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.90.214801`.

[57] Ian Blumenfeld et al. "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator". In: *Nature* 445 (Feb. 15, 2007), 741 EP -. URL: `http://dx.doi.org/10.1038/nature05538`.

[58] M. Litos et al. "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator". In: *Nature* 515 (Nov. 5, 2014), 92 EP -. URL: `http://dx.doi.org/10.1038/nature13882`.

[59] Spencer Gessner et al. "Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator". In: *Nature Communications* 7 (June 2, 2016), 11785 EP -. URL: `http://dx.doi.org/10.1038/ncomms11785`.

[60] S. Diederichs et al. "Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns". In: *Phys. Rev. Accel. Beams* 22 (8 Aug. 2019), p. 081301. DOI: `10.1103/PhysRevAccelBeams.22.081301`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.081301`.

[61] J. van Tilborg et al. "Active Plasma Lensing for Relativistic Laser-Plasma-Accelerated Electron Beams". In: *Phys. Rev. Lett.* 115 (18 Oct. 2015), p. 184802. DOI: `10.1103/PhysRevLett.115.184802`. URL: `http://link.aps.org/doi/10.1103/PhysRevLett.115.184802`.

[62] S. Steinke et al. "Multistage coupling of independent laser-plasma accelerators". In: *Nature* 530.7589 (Feb. 11, 2016), pp. 190–193. ISSN: 0028-0836. URL: `http://dx.doi.org/10.1038/nature16525`.

[63] R. Lehe et al. "Laser-plasma lens for laser-wakefield accelerators". In: *Phys. Rev. ST Accel. Beams* 17 (12 Dec. 2014), p. 121301. DOI: `10.1103/PhysRevSTAB.17.121301`. URL: `https://link.aps.org/doi/10.1103/PhysRevSTAB.17.121301`.

[64] C. E. Doss et al. "Laser-ionized, beam-driven, underdense, passive thin plasma lens". In: *Phys. Rev. Accel. Beams* 22 (11 Nov. 2019), p. 111001. DOI: `10.1103/PhysRevAccelBeams.22.111001`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.22.111001`.

[65] R. D'Arcy et al. "Tunable Plasma-Based Energy Dechirper". In: *Phys. Rev. Lett.* 122 (3 Jan. 2019), p. 034801. DOI: `10.1103/PhysRevLett.122.034801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.122.034801`.

[66] Paul Scherkl et al. *Plasma-photonic spatiotemporal synchronization of relativistic electron and laser beams*. 2019. arXiv: `1908.09263 [physics.plasm-ph]`.

[67] G. A. Bagdasarov et al. "Laser beam coupling with capillary discharge plasma for laser wakefield acceleration applications". In: *Physics of Plasmas* 24.8 (Aug. 7, 2017), p. 083109. ISSN: 1070-664X. DOI: `10.1063/1.4997606`. URL: `https://doi.org/10.1063/1.4997606`.

[68] C. V. Pieronek et al. "Laser-heated capillary discharge waveguides as tunable structures for laser-plasma acceleration". In: *Physics of Plasmas* 27.9 (2020), p. 093101. DOI: `10.1063/5.0014961`. eprint: `https://doi.org/10.1063/5.0014961`. URL: `https://doi.org/10.1063/5.0014961`.

[69] C Benedetti et al. "An accurate and efficient laser-envelope solver for the modeling of laser-plasma accelerators". In: *Plasma Physics and Controlled Fusion* 60.1 (Oct. 2017), p. 014002. DOI: `10.1088/1361-6587/aa8977`. URL: `https://doi.org/10.1088/1361-6587/aa8977`.

[70] G. A. Bagdasarov et al. "Creation of an axially uniform plasma channel in a laser-assisted capillary discharge". In: *Physics of Plasmas* 28.5 (2021), p. 053104. DOI: `10.1063/5.0046428`. eprint: `https://doi.org/10.1063/5.0046428`. URL: `https://doi.org/10.1063/5.0046428`.

[71] A. J. Gonsalves et al. "Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide". In: *Phys. Rev. Lett.* 122 (8 Feb. 2019), p. 084801. DOI: `10.1103/PhysRevLett.122.084801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.122.084801`.

[72] J. van Tilborg et al. "Nonuniform discharge currents in active plasma lenses". In: *Phys. Rev. Accel. Beams* 20 (3 Mar. 2017), p. 032803. DOI: `10.1103/PhysRevAccelBeams.20.032803`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.032803`.

[73] C. A. Lindstrøm et al. "Emittance Preservation in an Aberration-Free Active Plasma Lens". In: *Phys. Rev. Lett.* 121 (19 Nov. 2018), p. 194801. DOI: `10.1103/PhysRevLett.121.194801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.121.194801`.

[74] N M Cook et al. "Modeling of capillary discharge plasmas for wakefield acceleration and beam transport". In: *Journal of Physics: Conference Series* 1596 (Sept. 2020), p. 012063. DOI: `10.1088/1742-6596/1596/1/012063`. URL: `https://doi.org/10.1088/1742-6596/1596/1/012063`.

[75] R. J. Shalloo et al. "Hydrodynamic optical-field-ionized plasma channels". In: *Phys. Rev. E* 97 (5 May 2018), p. 053203. DOI: `10.1103/PhysRevE.97.053203`. URL: `https://link.aps.org/doi/10.1103/PhysRevE.97.053203`.

[76] B. Miao et al. "Optical Guiding in Meter-Scale Plasma Waveguides". In: *Phys. Rev. Lett.* 125 (7 Aug. 2020), p. 074801. DOI: `10.1103/PhysRevLett.125.074801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.125.074801`.

[77] A. Picksley et al. "Guiding of high-intensity laser pulses in 100-mm-long hydrodynamic optical-field-ionized plasma channels". In: *Phys. Rev. Accel. Beams* 23 (8 Aug. 2020), p. 081303. DOI: `10.1103/PhysRevAccelBeams.23.081303`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.081303`.

[78] A. Diaw et al. "Impact of electron transport models on capillary discharge plasmas". In: *Physics of Plasmas* 29.6 (2022), p. 063101. DOI: `10.1063/5.0091809`. eprint: `https://doi.org/10.1063/5.0091809`. URL: `https://doi.org/10.1063/5.0091809`.

[79] X. L. Xu et al. "Physics of Phase Space Matching for Staging Plasma and Traditional Accelerator Components Using Longitudinally Tailored Plasma Profiles". In: *Phys. Rev. Lett.* 116 (12 Mar. 2016), p. 124801. DOI: `10.1103/PhysRevLett.116.124801`. URL: `https://link.aps.org/doi/10.1103/PhysRevLett.116.124801`.

[80] A. Deng et al. "Generation and acceleration of electron bunches from a plasma photocathode". In: *Nature Physics* 15.11 (2019), pp. 1156–1160. ISSN: 1745-2481. DOI: `10.1038/s41567-019-0610-9`. URL: `https://doi.org/10.1038/s41567-019-0610-9`.

[81] Y. Xi et al. "Hybrid modeling of relativistic underdense plasma photocathode injectors". In: *Phys. Rev. ST Accel. Beams* 16 (3 Mar. 2013), p. 031303. DOI: 10.1103/PhysRevSTAB.16.031303. URL: https://link.aps.org/doi/10.1103/PhysRevSTAB.16.031303.

[82] G. G. Manahan et al. "Advanced schemes for underdense plasma photocathode wakefield accelerators: pathways towards ultrahigh brightness electron beams". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377.2151 (Aug. 12, 2019), p. 20180182. DOI: 10.1098/rsta.2018.0182. URL: https://doi.org/10.1098/rsta.2018.0182.

[83] S. Gessner and the AWAKE Collaboration. *Evolution of a plasma column measured through modulation of a high-energy proton beam.* 2020. arXiv: 2006.09991 [physics.acc-ph].

[84] S. Gourlay et al. "The US Magnet Development Program Plan". In: (2016). URL: https://osti.gov/scitech/servlets/purl/1306334.

[85] D. Arbelaez et al. "Numerical Modeling for Superconducting Accelerator Magnets". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF0-AF7_AF0-027.pdf.

[86] Julien Dular, Christophe Geuzaine, and Benot Vanderheyden. "Finite-Element formulations for systems with high-temperature superconductors". In: *IEEE Transactions on Applied Superconductivity* 30.3 (2019), pp. 1–13. ISSN: 1051-8223. DOI: 10.1109/tasc.2019.2935429. URL: https://doi.org/10.1109/tasc.2019.2935429.

[87] Julien Dular et al. "On the stability of mixed finite-element formulations for high-temperature superconductors". In: *IEEE Transactions on Applied Superconductivity* 31.6 (2021), pp. 1–12. ISSN: 1051-8223. DOI: 10.1109/tasc.2021.3098724. eprint: 2106.00313. URL: https://doi.org/10.1109/tasc.2021.3098724.

[88] Alexandre Arsenault, Frdric Sirois, and Francesco Grilli. "Implementation of the h-ϕ formulation in COMSOL Multiphysics for simulating the magnetization of bulk superconductors and comparison with the h-formulation". In: *IEEE Transactions on Applied Superconductivity* 31.2 (2021), pp. 1–11. ISSN: 1051-8223. DOI: 10.1109/tasc.2020.3033998. eprint: 2006.13784. URL: https://doi.org/10.1109/tasc.2020.3033998.

[89] Alexandre Arsenault, Bruno de Sousa Alves, and Frdric Sirois. "COMSOL Implementation of the h-ϕ-formulation with thin cuts for modeling superconductors with transport currents". In: *IEEE Transactions on Applied Superconductivity* 31.6 (2021), pp. 1–9. ISSN: 1051-8223. DOI: 10.1109/tasc.2021.3097245. URL: https://doi.org/10.1109/tasc.2021.3097245.

[90] Bruno de Sousa Alves et al. "Thin-shell approach for modeling superconducting tapes in the h-ϕ finite-element formulation". In: *arXiv* (2021). h-phi thin. eprint: 2108.08828.

[91] J.-L. Vay et al. "End-to-End Virtual Accelerators (EVA)". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF0-AF1_AF0_Vay-067.pdf.

[92] R. Lehe et al. "Machine learning and surrogates models for simulation-based optimization of accelerator design". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF3-AF1_AF6_Lehe-075.pdf.

[93] D. Winklehner and A. Adelmann. "Application of Machine Learning to Particle Accelerator Simulations". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF3_CompF0-AF1_AF0_Winklehner-108.pdf.

[94] C. Hovater et al. "Operation of the CEBAF 100 MV Cryomodules". In: *Proc. of Linear Accelerator Conference (LINAC'16), East Lansing, MI, USA, 25-30 September 2016* (East Lansing, MI, USA). Linear Accelerator Conference 28. Geneva, Switzerland: JACoW, May 2017, pp. 65–67. ISBN: 978-3-95450-169-4. DOI: `https://doi.org/10.18429/JACoW-LINAC2016-MOOP11`. URL: `http://jacow.org/linac2016/papers/moop11.pdf`.

[95] J. R. Lewandowski et al. "RF Gun Dark Current Suppression with a Transverse Deflecting Cavity at LCLS". In: *Proc. 37th Int. Free Electron Laser Conf. (FEL'15)* (Daejeon, Korea). https://doi.org/10.18429/JACoW-FEL2015-WEP001. JACoW Publishing, Aug. 2015, pp. 583–586. DOI: `doi:10.18429/JACoW-FEL2015-WEP001`.

[96] J. Allison et al. "Recent developments in Geant4". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225.

[97] Alfredo Ferrari et al. "FLUKA: A multi-particle transport code (Program version 2005)". In: (Oct. 2005). DOI: `10.2172/877507`.

[98] Nikolai V. Mokhov and Catherine C. James. "The MARS Code System User's Guide Version 15(2016)". In: (Feb. 2017). DOI: `10.2172/1462233`.

[99] Axel Huebl et al. "openPMD: A meta data standard for particle and mesh based data". In: (2015). DOI: `10.5281/zenodo.591699`. URL: `https://doi.org/10.5281/zenodo.591699`.

[100] A. Huebl et al. "Develop/integrate data standards and start-to-end workflows". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF7-AF1_AF0_Huebl-079.pdf`.

[101] Daniel Winklehner et al. "Realistic simulations of a cyclotron spiral inflector within a particle-in-cell framework". In: *Physical Review Accelerators and Beams* 20.12 (Dec. 2017), p. 124201. DOI: `10.1103/PhysRevAccelBeams.20.124201`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.20.124201` (visited on 10/15/2018).

[102] C.-K. Huang et al. "Physics-based high-fidelity modeling of high brightness beam injectors". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/AF/SNOWMASS21-AF7_AF1-CompF2_CompF0_Huang-183.pdf`.

[103] S. Nagaitsev et al. "Accelerator and Beam Physics: Grand Challenges and Research Opportunities". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/AF/SNOWMASS21-AF1_AF7_S_Nagaitsev-056.pdf`.

[104] J Qiang et al. "Start-to-end simulation of x-ray radiation of a next generation light source using the real number of electrons". In: *Physical Review Special Topics-Accelerators and Beams* 17.3 (2014), p. 030701.

[105] Ji Qiang. "Start-to-End Beam Dynamics Optimization of X-Ray FEL Light Source Accelerators". In: *Proceedings of North American PAC2016*. Chicago, IL, USA, Oct. 2016, pp. 838–842.

[106] Ji Qiang. "X-ray FEL linear accelerator design via start-to-end global optimization". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (2022), p. 166294.

[107] J.-L. Vay and C. Deutsch. "A three-dimensional electromagnetic particle-in-cell code to simulate heavy ion beam propagation in the reaction chamber". In: *Fusion Engineering and Design* 32-33.1-4 (1996). ISSN: 09203796. DOI: `10.1016/S0920-3796(96)00502-9`.

[108] J.L. Vay and C. Deutsch. "Charge compensated ion beam propagation in a reactor sized chamber". In: *Physics of Plasmas* 5.4 (1998). ISSN: 1070664X. DOI: `10.1063/1.872648`.

[109] Ji Qiang and Salman Habib. "Second-order stochastic leapfrog algorithm for multiplicative noise Brownian motion". In: *Physical Review E* 62.5 (Nov. 2000), p. 7430. ISSN: 1063651X. DOI: `10.1103/PhysRevE.62.7430`. URL: `https://journals.aps.org/pre/abstract/10.1103/PhysRevE.62.7430`.

[110] J.-L. Vay. "An Extended Fdtd Scheme For The Wave Equation: Application To Multiscale Electromagnetic Simulation". In: *Journal of Computational Physics* 167.1 (Feb. 2001), pp. 72–98. ISSN: 0021-9991.

[111] J. Qiang and R. D. Ryne. "Parallel 3D Poisson Solver for a Charged Beam in a Conducting Pipe". In: *Comp. Phys. Comm.* 138 (2001), p. 18.

[112] Ji Qiang, Miguel A. Furman, and Robert D. Ryne. "Strong-strong beam-beam simulation using a Green function approach". In: *Physical Review Special Topics - Accelerators and Beams* 5.10 (Oct. 2002), pp. 60–66. ISSN: 10984402. DOI: `10.1103/PHYSREVSTAB.5.104402/FIGURES/9/MEDIUM`. URL: `https://journals.aps.org/prab/abstract/10.1103/PhysRevSTAB.5.104402`.

[113] J.-L. Vay. "Asymmetric Perfectly Matched Layer For The Absorption Of Waves". In: *Journal of Computational Physics* 183.2 (Dec. 2002), pp. 367–399. ISSN: 0021-9991. DOI: `10.1006/Jcph.2002.7175`.

[114] J.-L. Vay et al. "Mesh refinement for particle-in-cell plasma simulations: Applications to and benefits for heavy ion fusion". In: *Laser and Particle Beams* 20.4 (2002). ISSN: 02630346. DOI: `10.1017/S0263034602204139`.

[115] Ji Qiang and Robert L. Gluckstern. "Three-dimensional Poisson solver for a charged beam with large aspect ratio in a conducting pipe". In: *Computer Physics Communications* 160.2 (July 2004), pp. 120–128. ISSN: 0010-4655. DOI: `10.1016/J.CPC.2004.03.002`.

[116] J.-L. Vay, J.-C. Adam, and A Heron. "Asymmetric Pml For The Absorption Of Waves. Application To Mesh Refinement In Electromagnetic Particle-In-Cell Plasma Simulations". In: *Computer Physics Communications* 164.1-3 (Dec. 2004), pp. 171–177. ISSN: 0010-4655. DOI: `10.1016/J.Cpc.2004.06.026`.

[117] J.-L. Vay et al. "Application Of Adaptive Mesh Refinement To Particle-In-Cell Simulations Of Plasmas And Beams". In: *Physics Of Plasmas* 11.5 (May 2004), pp. 2928–2934. ISSN: 1070-664X. DOI: `10.1063/1.1689669`.

[118] J. Qiang, D. Todd, and D. Leitner. "A 3D model for ion beam formation and transport simulation". In: *Comp. Phys. Comm.* 175 (2006), p. 416.

[119] V. K. Decyk. "UPIC: A framework for massively parallel particle-in-cell codes". In: *Computer Physics Communications* 177 (2007), p. 95.

[120] J L Vay. "Simulation Of Beams Or Plasmas Crossing At Relativistic Velocity". In: *Physics Of Plasmas* 15.5 (May 2008), p. 56701. DOI: `10.1063/1.2837054`.

[121] J.-L. Vay et al. "Application Of The Reduction Of Scale Range In A Lorentz Boosted Frame To The Numerical Simulation Of Particle Acceleration Devices". In: *Proc. Particle Accelerator Conference*. Vancouver, Canada, 2009. URL: `https://accelconf.web.cern.ch/PAC2009/papers/tu1pbi04.pdf`.

[122] R H Cohen et al. "An Implicit "Drift-Lorentz{"} Mover For Plasma And Beam Simulations". In: *Nuclear Instruments & Methods In Physics Research Section A-Accelerators Spectrometers Detectors And Associated Equipment* 606.1-2 (July 2009), pp. 53–55. ISSN: 0168-9002. DOI: `10.1016/J.Nima.2009.03.083`.

[123] J L Vay et al. "Numerical Methods For Instability Mitigation In The Modeling Of Laser Wakefield Accelerators In A Lorentz-Boosted Frame". In: *Journal of Computational Physics* 230.15 (July 2011), pp. 5908–5929. DOI: `10.1016/J.Jcp.2011.04.003`.

[124] J -L. Vay et al. "Modeling Of 10 Gev-1 Tev Laser-Plasma Accelerators Using Lorentz Boosted Simulations". In: *Physics Of Plasmas* 18.12 (Dec. 2011). ISSN: 1070-664X. DOI: `10.1063/1.3663841`.

[125] Jl Vay et al. "Effects Of Hyperbolic Rotation In Minkowski Space On The Modeling Of Plasma Accelerators In A Lorentz Boosted Frame". In: *Physics Of Plasmas* 18.3 (Mar. 2011), p. 30701. DOI: `10.1063/1.3559483`.

[126] J.-L. Vay et al. "Novel methods in the particle-in-cell accelerator code-framework warp". In: *Computational Science and Discovery* 5.1 (2012), 014019 (20 pp.) ISSN: 1749-4680.

[127] Brendan B Godfrey and Jean-Luc Vay. "Numerical stability of relativistic beam multidimensional {PIC} simulations employing the Esirkepov algorithm". In: *Journal of Computational Physics* 248.0 (2013), pp. 33–46. ISSN: 0021-9991. DOI: `http://dx.doi.org/10.1016/j.jcp.2013.04.006`. URL: `http://www.sciencedirect.com/science/article/pii/S0021999113002556`.

[128] Jean Luc Vay, Irving Haber, and Brendan B. Godfrey. "A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas". In: *Journal of Computational Physics* 243 (2013), pp. 260–268.

[129] Brendan B Godfrey, Jean-Luc Vay, and Irving Haber. "Numerical stability analysis of the pseudo-spectral analytical time-domain {PIC} algorithm". In: *Journal of Computational Physics* 258.0 (2014), pp. 689–704. ISSN: 0021-9991. DOI: `http://dx.doi.org/10.1016/j.jcp.2013.10.053`. URL: `http://www.sciencedirect.com/science/article/pii/S0021999113007298`.

[130] Brendan B Godfrey and Jean-Luc Vay. "Suppressing the numerical Cherenkov instability in {FDTD} {PIC} codes". In: *Journal of Computational Physics* 267.0 (2014), pp. 1–6. ISSN: 0021-9991. DOI: `http://dx.doi.org/10.1016/j.jcp.2014.02.022`. URL: `http://www.sciencedirect.com/science/article/pii/S0021999114001429`.

[131] Brendan B. Godfrey, Jean Luc Vay, and Irving Haber. "Numerical stability improvements for the pseudospectral EM PIC algorithm". In: *IEEE Transactions on Plasma Science* 42.5 (2014), pp. 1339–1344.

[132] Brendan B. Godfrey and Jean Luc Vay. "Improved numerical Cherenkov instability suppression in the generalized PSTD PIC algorithm". In: *Computer Physics Communications* 196 (2015), pp. 221–225.

[133] Remi Lehe et al. "Elimination of numerical Cherenkov instability in flowing-plasma particle-in-cell simulations by using Galilean coordinates". In: *Physical Review E* 94.5 (Nov. 2016), p. 053305. ISSN: 2470-0045. DOI: `10.1103/PhysRevE.94.053305`. URL: `https://link.aps.org/doi/10.1103/PhysRevE.94.053305`.

[134]   M. Kirchen et al. "Stable discrete representation of relativistically drifting plasmas". In: *Physics of Plasmas* 23.10 (Oct. 2016), p. 100704. ISSN: 1070-664X. DOI: 10.1063/1.4964770. URL: http://aip.scitation.org/doi/10.1063/1.4964770.

[135]   H. Vincenti and J.-L. Vay. "Detailed analysis of the effects of stencil spatial variations with arbitrary high-order finite-difference Maxwell solver". In: *Computer Physics Communications* 200 (Mar. 2016), pp. 147–167. ISSN: 00104655. DOI: 10.1016/j.cpc.2015.11.009. URL: https://apps.webofknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=1&SID=1CanLFIHrQ5v8O7cxqV&page=1&doc=2.

[136]   J. Qiang. "Efficient Three-Dimensional Poisson Solvers in Open Rectangular Conducting Pipe". In: *Comp. Phys. Comm.* 203 (2016), p. 122.

[137]   G. Blaclard et al. "Pseudospectral Maxwell solvers for an accurate modeling of Doppler harmonic generation on plasma mirrors with particle-in-cell codes". In: *Physical Review E* 96.3 (Sept. 2017), p. 033305. ISSN: 2470-0045. DOI: 10.1103/PhysRevE.96.033305. URL: https://link.aps.org/doi/10.1103/PhysRevE.96.033305.

[138]   S. Jalas et al. "Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods". In: *Physics of Plasmas* 24.3 (Mar. 2017), p. 033115. ISSN: 1070-664X. DOI: 10.1063/1.4978569. URL: http://aip.scitation.org/doi/10.1063/1.4978569.

[139]   Ji Qiang. "Symplectic multiparticle tracking model for self-consistent space-charge simulation". In: *Physical Review Accelerators and Beams* 20.1 (Jan. 2017), p. 014203. ISSN: 24699888. DOI: 10.1103/PhysRevAccelBeams.20.014203.

[140]   J. Qiang. "A Fast Numerical Integrator for Relativistic Charged Particle Tracking". In: *Instruments & Methods in Physics Research A* 885 (2017), p. 55.

[141]   J.-L. Vay et al. "Warp-X: A new exascale computing platform for beam–plasma simulations". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* (Jan. 2018). ISSN: 0168-9002. DOI: 10.1016/J.NIMA.2018.01.035. URL: https://www.sciencedirect.com/science/article/pii/S0168900218300524.

[142]   J. L. Vay et al. "Toward plasma wakefield simulations at exascale". In: *2018 IEEE Advanced Accelerator Concepts Workshop, ACC 2018 - Proceedings*. Institute of Electrical and Electronics Engineers Inc., Mar. 2019. ISBN: 9781538677216. DOI: 10.1109/AAC.2018.8659392.

[143]   H. Vincenti and J.-L. Vay. "Ultrahigh-order Maxwell solver with extreme scalability for electromagnetic PIC simulations of plasmas". In: *Computer Physics Communications* (2018). ISSN: 00104655. DOI: 10.1016/j.cpc.2018.03.018.

[144]   J. Qiang. "Symplectic particle-in-cell model for space-charge beam dynamics simulation". In: *Phys. Rev. ST Accel. Beams* 21 (2018), p. 054201.

[145]   Haithem Kallala, Jean Luc Vay, and Henri Vincenti. "A generalized massively parallel ultra-high order FFT-based Maxwell solver". In: *Computer Physics Communications* 244 (Nov. 2019), pp. 25–34. ISSN: 00104655. DOI: 10.1016/j.cpc.2019.07.009.

[146]   Olga Shapoval, Jean-Luc Vay, and Henri Vincenti. "Two-step perfectly matched layer for arbitrary-order pseudo-spectral analytical time-domain methods". In: *Computer Physics Communications* 235 (Feb. 2019), pp. 102–110. ISSN: 0010-4655. DOI: 10.1016/J.CPC.2018.09.015. URL: https://www.sciencedirect.com/science/article/pii/S0010465518303291?via%3Dihub.

[147] J. Qiang. "Fast 3D Poisson solvers in elliptical conducting pipe for space-charge simulation". In: *Phys. Rev. ST Accel. Beams* 22 (2019), p. 104601.

[148] Manuel Kirchen et al. "Scalable spectral solver in Galilean coordinates for eliminating the numerical Cherenkov instability in particle-in-cell simulations of streaming plasmas". In: *Physical Review E* 102.1 (July 2020), p. 13202. ISSN: 24700053. DOI: 10.1103/PhysRevE.102. 013202. URL: https://journals.aps.org/pre/abstract/10.1103/PhysRevE.102.013202.

[149] Olga Shapoval et al. "Overcoming timestep limitations in boosted-frame particle-in-cell simulations of plasma-based acceleration". In: *Physical Review E* 104.5 (Nov. 2021), p. 055311. ISSN: 2470-0045. DOI: 10.1103/PHYSREVE.104.055311/FIGURES/6/MEDIUM. URL: https://journals.aps.org/pre/abstract/10.1103/PhysRevE.104.055311.

[150] Edoardo Zoni et al. *A Hybrid Nodal-Staggered Pseudo-Spectral Electromagnetic Particle-In-Cell Method with Finite-Order Centering* https://arxiv.org/abs/2106.12919v1. 2021.

[151] Remi Lehe et al. *Absorption of charged particles in Perfectly-Matched-Layers by optimal damping of the deposited current* https://arxiv.org/abs/2201.09084. 2022.

[152] V. K. Decyk annd C. D. Norton. "UCLA Parallel PIC Framework". In: *Computer Physics Communications* 164 (2004), p. 80.

[153] Samuel F Martins et al. "Numerical Simulations Of Laser Wakefield Accelerators In Optimal Lorentz Frames". In: *Computer Physics Communications* 181.5 (May 2010), pp. 869–875. ISSN: 0010-4655. DOI: 10.1016/J.Cpc.2009.12.023.

[154] V. K. Decyk and T. V. Singh. "Adaptable Particle-in-Cell algorithms for graphical processing units". In: *Computer Physics Communications* 182 (2011), p. 641.

[155] Xinlu Xu et al. "Numerical instability due to relativistic plasma drift in EM-PIC simulations". In: *Computer Physics Communications* 184.11 (2013), pp. 2503–2514. ISSN: 0010-4655. DOI: http://dx.doi.org/10.1016/j.cpc.2013.07.003. URL: http://www.sciencedirect.com/science/article/pii/S0010465513002312.

[156] V. K. Decyk and T. V. Singh. "Particle-in-Cell algorithms for emerging computer architectures". In: *Computer Physics Communications* 185 (2014), p. 708.

[157] Peicheng Yu et al. "Elimination of the numerical Cerenkov instability for spectral EM-PIC codes". In: *Computer Physics Communications* 192 (July 2015), pp. 32–47. ISSN: 00104655. DOI: 10.1016/j.cpc.2015.02.018. URL: https://apps.webofknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=2&SID=1CanLFIHrQ5v8O7cxqV&page=1&doc=3.

[158] Peicheng Yu et al. "Mitigation of numerical Cerenkov radiation and instability using a hybrid finite difference-FFT Maxwell solver and a local charge conserving current deposit". In: *Computer Physics Communications* 197 (Dec. 2015), pp. 144–152. ISSN: 00104655. DOI: 10.1016/j.cpc.2015.08.026. URL: https://apps.webofknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=2&SID=1CanLFIHrQ5v8O7cxqV&page=1&doc=2.

[159] P. Yu et al. "Enabling Lorentz boosted frame particle-in-cell simulations of laser wakefield acceleration in quasi-3D geometry". In: *Journal of Computational Physics* 316 (2016), p. 747.

[160] Fei Li et al. "Controlling the numerical Cerenkov instability in PIC simulations using a customized finite difference Maxwell solver and a local FFT based current correction". In: *Computer Physics Communications* 214 (May 2017), pp. 6–17. ISSN: 0010-4655. DOI: 10.1016/J.CPC.2017.01.001.

[161] F. Li et al. "A new field solver for modeling of relativistic particle-laser interactions using the particle-in-cell algorithmhttps://arxiv.org/abs/2004.03754". In: (2020).

[162] F. Li et al. "Accurately simulating nine-dimensional phase space of relativistic particles in strong fields". In: *Jour. Comp. Phys.* 438 (2021), p. 110367.

[163] X. L. Xu et al. "On nnumerical errors to the fields surrounding a relativistically moving particle in PIC codes". In: *Jour. Comp. Phys.* 413 (2020), p. 109451.

[164] M. Vranic et al. "Classical radiation reaction in particle-in-cell simulations". In: *Comp. Phys. Comm.* 204 (2016), pp. 141–151.

[165] C. E. Mayes et al. "Lightsource unified modeling environment (LUME), a start-to-end simulation ecosystem". In: *Proc. of IPAC.* 2021, THPAB217.

[166] Christopher Mayes. *xopt: Simulation optimization, based on DEAP.* https://github.com/ChristopherMayes/xo

[167] Alexander Scheinker and Spencer Gessner. "Adaptive method for electron bunch profile prediction". In: *Physical Review Special Topics-Accelerators and Beams* 18.10 (2015), p. 102801. DOI: https://doi.org/10.1103/PhysRevSTAB.18.102801.

[168] Auralee Edelen et al. "Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems". In: *Phys. Rev. Accel. Beams* 23 (4 Apr. 2020), p. 044601. DOI: 10.1103/PhysRevAccelBeams.23.044601. URL: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.044601.

[169] A. L. Edelen et al. "Neural Networks for Modeling and Control of Particle Accelerators". In: *IEEE Transactions on Nuclear Science* 63.2 (2016), pp. 878–897. DOI: 10.1109/TNS.2016.2543203.

[170] C. Emma et al. "Machine learning-based longitudinal phase space prediction of particle accelerators". In: *Phys. Rev. Accel. Beams* 21 (11 Nov. 2018), p. 112802. DOI: 10.1103/PhysRevAccelBeams.21.112802. URL: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.21.112802.

[171] A. Edelen et al. "Using Neural Network Control Policies for Rapid Switching Between Beam Parameters in a free electron laser". In: *NeurIPS 2017.* https://dl4physicalsciences.github.io/files/nips_dlps_2017_16.pdf. Long Beach, CA, 2017.

[172] A. Edelen et al. "Machine Learning Models for Optimization and Control of X-ray Free Electron Lasers". In: *NeurIPS 2019.* https://ml4physicalsciences.github.io/2019/files/NeurIPS_ML4PS_2019_90.pdf. Vancouver, Canada, 2019.

[173] AL Edelen et al. "First steps toward incorporating image based diagnostics into particle accelerator control systems using convolutional neural networks". In: *arXiv preprint arXiv:1612.05662* (2016).

[174] Lipi Gupta et al. "Improving surrogate model accuracy for the LCLS-II injector frontend using convolutional neural networks and transfer learning". In: *Machine Learning: Science and Technology* 2.4 (Oct. 2021), p. 045025. DOI: 10.1088/2632-2153/ac27ff. URL: https://doi.org/10.1088/2632-2153/ac27ff.

[175] J. Zhu et al. "High-Fidelity Prediction of Megapixel Longitudinal Phase-Space Images of Electron Beams Using Encoder-Decoder Neural Networks". In: *Phys. Rev. Applied* 16 (2 Aug. 2021), p. 024005. DOI: 10.1103/PhysRevApplied.16.024005. URL: https://link.aps.org/doi/10.1103/PhysRevApplied.16.024005.

[176] Alexander Scheinker. "Adaptive machine learning for time-varying systems: low dimensional latent space tuning". In: *Journal of Instrumentation* 16.10 (2021), P10008. URL: `https://doi.org/10.1088/1748-0221/16/10/P10008`.

[177] Joseph Duris et al. "Bayesian optimization of a free-electron laser". In: *Physical review letters* 124.12 (2020), p. 124801.

[178] J. Ögren, C. Gohil, and D. Schulte. "Surrogate modeling of the CLIC final-focus system using artificial neural networks". In: *Journal of Instrumentation* 16.05 (May 2021), P05012. DOI: `10.1088/1748-0221/16/05/p05012`. URL: `https://doi.org/10.1088/1748-0221/16/05/p05012`.

[179] Aashwin Ananda Mishra et al. "Uncertainty quantification for deep learning in particle accelerator applications". In: *Physical Review Accelerators and Beams* 24.11 (2021), p. 114601.

[180] Owen Convery et al. "Uncertainty quantification for virtual diagnostic of particle accelerators". In: *Phys. Rev. Accel. Beams* 24 (7 July 2021), p. 074602. DOI: `10.1103/PhysRevAccelBeams.24.074602`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.074602`.

[181] Anusha Nagabandi, Chelsea Finn, and Sergey Levine. "Deep online learning via meta-learning: Continual adaptation for model-based RL". In: *arXiv preprint arXiv:1812.07671* (2018).

[182] Sinno Jialin Pan and Qiang Yang. "A survey on transfer learning". In: *IEEE Transactions on knowledge and data engineering* 22.10 (2009), pp. 1345–1359.

[183] Andrei Ivanov and Ilya Agapov. "Physics-based deep neural networks for beam dynamics in charged particle accelerators". In: *Physical Review Accelerators and Beams* 23.7 (July 2020), p. 074601. ISSN: 2469-9888. DOI: `10.1103/PhysRevAccelBeams.23.074601`. URL: `https://doi.org/10.1103/PhysRevAccelBeams.23.074601%20https://link.aps.org/doi/10.1103/PhysRevAccelBeams.23.074601`.

[184] Alexander Scheinker et al. "An adaptive approach to machine learning for compact particle accelerators". In: *Scientific reports* 11.1 (2021), pp. 1–11. URL: `https://doi.org/10.1038/s41598-021-98785-0`.

[185] Alexander Scheinker et al. "Demonstration of model-independent control of the longitudinal phase space of electron beams in the linac-coherent light source with femtosecond resolution". In: *Physical review letters* 121.4 (2018), p. 044801. URL: `https://doi.org/10.1103/PhysRevLett.121.044801`.

[186] George Em Karniadakis et al. "Physics-informed machine learning". In: *Nature Reviews Physics* 3.6 (2021), pp. 422–440.

[187] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations". In: *Journal of Computational physics* 378 (2019), pp. 686–707.

[188] Alvaro Sanchez-Gonzalez et al. "Hamiltonian Graph Networks with ODE Integrators". In: (Sept. 2019). arXiv: `1909.12790 [cs.LG]`.

[189] Adi Hanuka et al. "Physics model-informed Gaussian process for online optimization of particle accelerators". In: *Physical Review Accelerators and Beams* 24.7 (2021), p. 072802.

[190] M. Berz. "Differential Algebraic Description of Beam Dynamics to Very High Orders". In: *Part. Accel.* 24 (1989), pp. 109–124.

[191] R Roussel et al. "Differentiable Preisach Modeling for Characterization and Optimization of Accelerator Systems with Hysteresis". In: *arXiv preprint arXiv:2202.07747* (2022).

[192] K. Deb et al. "A fast and elitist multiobjective genetic algorithm: NSGA-II". In: *IEEE Transactions on Evolutionary Computation* 6.2 (2002), pp. 182–197. DOI: `10.1109/4235.996017`.

[193] Ryan Roussel, Adi Hanuka, and Auralee Edelen. "Multiobjective Bayesian optimization for online accelerator tuning". In: *Phys. Rev. Accel. Beams* 24 (6 June 2021), p. 062801. DOI: `10.1103/PhysRevAccelBeams.24.062801`. URL: `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.062801`.

[194] Ryan Roussel et al. "Turn-key constrained parameter space exploration for particle accelerators using Bayesian active learning." In: *Nature Communications* 12 (1 2021). URL: `https://doi.org/10.1038/s41467-021-25757-3`.

[195] Jian Wu et al. "Practical Multi-fidelity Bayesian Optimization for Hyperparameter Tuning". In: *Uncertainty Quantification in Artificial Intelligence*. 2019.

[196] David Eriksson et al. "Scalable Global Optimization via Local Bayesian Optimization". In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019. URL: `https://proceedings.neurips.cc/paper/2019/file/6c990b7aca7bc7058f5e98ea909e924b-Paper.pdf`.

[197] Willie Neiswanger, Ke Alexander Wang, and Stefano Ermon. "Bayesian Algorithm Execution: Estimating Computable Properties of Black-box Functions Using Mutual Information". In: *Proceedings of the 38th International Conference on Machine Learning*. Ed. by Marina Meila and Tong Zhang. Vol. 139. Proceedings of Machine Learning Research. PMLR, 18–24 Jul 2021, pp. 8005–8015. URL: `https://proceedings.mlr.press/v139/neiswanger21a.html`.

[198] John Preskill. "Quantum computing in the NISQ era and beyond". In: *Quantum* 2 (2018), p. 79.

[199] Frank Arute et al. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779 (2019), pp. 505–510.

[200] Han-Sen Zhong et al. "Quantum computational advantage using photons". In: *Science* 370.6523 (2020), pp. 1460–1463.

[201] *IBM Quantum Computing*. `https://www.ibm.com/quantum-computing/`. 2020 (accessed March 08, 2022).

[202] *Leap²*. `https://www.dwavesys.com/`. 2020.

[203] *Amazon Quantum Solutions Lab*. `https://aws.amazon.com/quantum-solutions-lab/`. 2020.

[204] *Azure Quantum*. `https://azure.microsoft.com/en-us/services/quantum/`. 2020.

[205] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. "Quantum support vector machine for big data classification". In: *Physical review letters* 113.13 (2014), p. 130503.

[206] Rupak Chatterjee and Ting Yu. "Generalized coherent states, reproducing kernels, and quantum support vector machines". In: *arXiv preprint arXiv:1612.03713* (2016).

[207] Guang Hao Low, Theodore J Yoder, and Isaac L Chuang. "Quantum inference on Bayesian networks". In: *Physical Review A* 89.6 (2014), p. 062315.

[208] Iris Cong, Soonwon Choi, and Mikhail D Lukin. "Quantum convolutional neural networks". In: *Nature Physics* 15.12 (2019), pp. 1273–1278.

[209] Zhikuan Zhao et al. "Bayesian deep learning on a quantum computer". In: *Quantum Machine Intelligence* 1.1 (2019), pp. 41–51.

[210] Vedran Dunjko, Jacob M Taylor, and Hans J Briegel. "Quantum-enhanced machine learning". In: *Physical review letters* 117.13 (2016), p. 130501.

[211] Vedran Dunjko, Jacob M Taylor, and Hans J Briegel. "Advances in quantum reinforcement learning". In: *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE. 2017, pp. 282–287.

[212] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. "Quantum algorithm for linear systems of equations". In: *Physical review letters* 103.15 (2009), p. 150502.

[213] B David Clader, Bryan C Jacobs, and Chad R Sprouse. "Preconditioned quantum linear system algorithm". In: *Physical review letters* 110.25 (2013), p. 250504.

[214] Andrew M Childs, Robin Kothari, and Rolando D Somma. "Quantum algorithm for systems of linear equations with exponentially improved dependence on precision". In: *SIAM Journal on Computing* 46.6 (2017), pp. 1920–1950.

[215] Carlos Bravo-Prieto et al. "Variational quantum linear solver: A hybrid algorithm for linear systems". In: *arXiv preprint arXiv:1909.05820* (2019).

[216] Yonghae Lee, Jaewoo Joo, and Soojoon Lee. "Hybrid quantum linear equation algorithm and its experimental test on ibm quantum experience". In: *Scientific reports* 9.1 (2019), p. 4778.

[217] Sarah K Leyton and Tobias J Osborne. "A quantum algorithm to solve nonlinear differential equations". In: *arXiv preprint arXiv:0812.4423* (2008).

[218] Dominic W Berry. "High-order quantum algorithm for solving linear differential equations". In: *Journal of Physics A: Mathematical and Theoretical* 47.10 (2014), p. 105301.

[219] Juan Miguel Arrazola et al. "Quantum algorithm for nonhomogeneous linear partial differential equations". In: *Physical Review A* 100.3 (2019), p. 032306.

[220] Andrew M Childs, Jin-Peng Liu, and Aaron Ostrander. "High-precision quantum algorithms for partial differential equations". In: *arXiv preprint arXiv:2002.07868* (2020).

[221] Tao Xin et al. "Quantum algorithm for solving linear differential equations: Theory and experiment". In: *Physical Review A* 101.3 (2020), p. 032307.

[222] Yudong Cao et al. "Quantum algorithm and circuit design solving the Poisson equation". In: *New Journal of Physics* 15.1 (2013), p. 013021.

[223] Shengbin Wang et al. "Quantum Fast Poisson Solver: the algorithm and modular circuit design". In: *arXiv preprint arXiv:1910.09756* (2019).

[224] Alexander Engel, Graeme Smith, and Scott E Parker. "Quantum algorithm for the Vlasov equation". In: *Physical Review A* 100.6 (2019), p. 062315.

[225] *Q# and the Quantum Development Kit*. https://www.microsoft.com/en-us/quantum/development-kit. 2020 (accessed March 08, 2022).

[226] J. Qiang. "Simulation of space-charge effects using a quantum Schrodinger approach". In: *Phys. Rev. ST Accel. Beams* 25 (2022), p. 034602.

[227] Ilya Y Dodin and Edward A Startsev. "On applications of quantum computing to plasma simulations". In: *arXiv preprint arXiv:2005.14369* (2020).

[228] K. A. Brown and T. Roser. "Towards storage rings as quantum computers". In: *Physical Review Accelerators and Beams* 23.5 (May 2020), p. 054701. ISSN: 24699888. DOI: 10.1103/PHYSREVACCELBEAMS.23.054701/FIGURES/4/MEDIUM. URL: https://journals.aps.org/prab/abstract/10.1103/PhysRevAccelBeams.23.054701.

[229] Wolfgang Paul. "Electromagnetic traps for charged and neutral particles". In: *Reviews of Modern Physics* 62.3 (July 1990), p. 531. ISSN: 00346861. DOI: 10.1103/RevModPhys.62.531. URL: https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.62.531.

[230] D F V James. *Applied Physics B Lasers and Optics Quantum dynamics of cold trapped ions with application to quantum computation.* Tech. rep. 1998, pp. 181–190.

[231] Lowell S. Brown and Gerald Gabrielse. "Geonium theory: Physics of a single electron or ion in a Penning trap". In: *Reviews of Modern Physics* 58.1 (Jan. 1986), p. 233. ISSN: 00346861. DOI: 10.1103/RevModPhys.58.233. URL: https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.58.233.

[232] Bohong Huang et al. "Artificial Intelligence-Assisted Design and Virtual Diagnostic for the Initial Condition of a Storage-Ring-Based Quantum Information System". In: *IEEE Access* 10 (2022), pp. 14350–14358. ISSN: 21693536. DOI: 10.1109/ACCESS.2022.3147727.

[233] J. L. Vay et al. "Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs". In: *Physics of Plasmas* 28.2 (Feb. 2021), p. 23105. ISSN: 10897674. DOI: 10.1063/5.0028512. URL: https://doi.org/10.1063/5.0028512.

[234] A. Myers et al. "Porting WarpX to GPU-accelerated platforms". In: *Parallel Computing* 108 (Dec. 2021), p. 102833. ISSN: 0167-8191. DOI: 10.1016/J.PARCO.2021.102833.

[235] *Synergia.* URL: http://web.fnal.gov/sites/synergia.

[236] *Kokkos · GitHub.* URL: https://github.com/kokkos.

[237] Erik Zenker et al. "Alpaka - An Abstraction Library for Parallel Kernel Acceleration". In: *Proceedings - 2016 IEEE 30th International Parallel and Distributed Processing Symposium, IPDPS 2016* (Feb. 2016), pp. 631–640. DOI: 10.1109/ipdpsw.2016.50. URL: https://arxiv.org/abs/1602.08477v1.

[238] *GitHub - LLNL/RAJA: RAJA Performance Portability Layer (C++).* URL: https://github.com/LLNL/RAJA.

[239] *AMReX.* URL: https://amrex-codes.github.io/.

[240] The Copa Team. *Copa.* URL: https://github.com/ECP-copa.

[241] Lipeng Wan et al. "Improving I/O Performance for Exascale Applications through Online Data Layout Reorganization". In: *accepted in IEEE Transactions on Parallel and Distributed Systems* (2021). URL: https://arxiv.org/abs/2107.07108.

[242] Axel Huebl et al. "On the Scalability of Data Reduction Techniques in Current and Upcoming HPC Systems from an Application Perspective". In: *High Performance Computing.* Ed. by Julian M. Kunkel et al. Cham: Springer International Publishing, 2017, pp. 15–29. ISBN: 978-3-319-67630-2.

[243] M. Bussmann et al. "Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability". In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis.* SC '13. Denver, Colorado: ACM, 2013, 5:1–5:12. ISBN: 978-1-4503-2378-9. DOI: 10.1145/2503210.2504564. URL: http://doi.acm.org/10.1145/2503210.2504564.

[244] Franz Poeschel et al. "Transitioning from file-based HPC workflows to streaming data pipelines with openPMD and ADIOS2". In: *submitted* (2021).

[245] R. Lehe et al. "Embracing modern software tools and user-friendly practices, when distributing scientific codes". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF0_Lehe-076.pdf.

[246] Daniel S. Katz. *Scientific Software Challenges and Community Responses*. Talk given at RTI International on 7 December 2015, discussing 12 scientific software challenges and how the scientific software community is responding to them. 2015. URL: `https://www.slideshare.net/danielskatz/scientific-software-challenges-and-community-responses`.

[247] Software Sustainability Institute. The University of Edinburgh. URL: `https://www.software.ac.uk/`.

[248] *Spack*. `https://spack.io/`.

[249] *Conda*. `https://docs.conda.io/en/latest/`.

[250] *The Phython package Installer*. `https://pip.pypa.io/en/stable/`.

[251] *GitHub Issues*. `https://guides.github.com/features/issues/`.

[252] *Gitter*. `https://gitter.im/`.

[253] *Slack*. `https://slack.com/`.

[254] *Git*. `https://git-scm.com/`.

[255] *Zenodo*. `https://zenodo.org/`.

[256] A. Huebl et al. "Aspiration for Open Science in Accelerator & Beam Physics Modeling ". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF7-AF1_AF0_Huebl-081.pdf`.

[257] Matthew Feickert et al. "Long Term Reproducibility and Sustainability of Scientific Software". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF7_CompF0_Matthew_Feickert-107.pdf`.

[258] Daniel S. Katz et al. *Software Sustainability & High Energy Physics*. Oct. 2020. DOI: `10.5281/zenodo.4095837`. URL: `https://arxiv.org/abs/2010.05102`.

[259] Open Source Initiative. *The Open Source Definition*. 2004. URL: `https://opensource.org/docs/osd`.

[260] Free Software Foundation. *The Free Software Definition*. 1990. URL: `https://www.gnu.org/philosophy/free-sw.en.html`.

[261] R. Schmitz, T. Eichlersmith, and A. Huebl. "Barriers to Entry in Physics Computing". In: *Snowmass21 LOI* (2020). URL: `https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF0_CompF0_Ryan_Schmitz-104.pdf`.

[262] Xiaoli Chen et al. "Open is not enough". In: *Nature Physics* 15.2 (Nov. 2018), pp. 113–119. DOI: `10.1038/s41567-018-0342-2`.

[263] *GitHub actions*. `https://github.com/features/actions`.

[264] *Azure pipelines*. `https://azure.microsoft.com/en-us/services/devops/pipelines/`.

[265] *Travis CI*. `https://travis-ci.org/`.

[266] *E4S*. `https://e4s-project.github.io/`.

[267] M. Bai et al. *Snowmass AF1 White Paper: Strategies in Education, Outreach, and Inclusion to Enhance the US Workforce in Accelerator Science and Engineering*. 2022.

[268] *CAMPA: Consortium for Advanced Modeling of Particle Accelerators*. URL: `http://campa.lbl.gov`.

[269] *PICMI*. URL: `https://github.com/picmi-standard`.

[270] William F. Godoy et al. "ADIOS 2: The Adaptable Input Output System. A framework for high-performance data management". In: *SoftwareX* 12 (2020), p. 100561. ISSN: 2352-7110. DOI: https://doi.org/10.1016/j.softx.2020.100561. URL: https://www.sciencedirect.com/science/article/pii/S2352711019302560.

[271] The HDF Group. *Hierarchical data format version 5*. URL: http://www.hdfgroup.org/HDF5.

[272] Axel Huebl et al. "openPMD-api: C++ & Python API for Scientific I/O with openPMD". In: (2018). DOI: 10.14278/rodare.27. URL: https://github.com/openPMD/openPMD-api.

[273] R. A. Fonseca. "No Title". In: *Lec. Notes In Comp. Sci.* 2329 (2002), p. 342.

[274] D.A. Bizzozero et al. "Multi-objective optimization with an integrated electromagnetics and beam dynamics workflow". In: *Nucl. Instrum. Methods Phys. Res. A* 1020 (2021), p. 165844.

[275] Severin Diederichs et al. *HiPACE++: a portable, 3D quasi-static Particle-in-Cell code*. 2022. eprint: arXiv:2109.10277. URL: https://arxiv.org/abs/2109.10277.

[276] Maxence Thévenet et al. *HiPACE++: Highly efficient Plasma Accelerator Emulation, quasistatic particle-in-cell code*. 2021. URL: https://github.com/Hi-PACE/hipace.

[277] *IDEAS: Interoperable Design of Extreme-scale Application Software*. URL: https://ideas-productivity.org.

[278] *xSDK: Extreme-scale Scientific Software Development Kit*. URL: http://xsdk.info.

[279] Bierlich, Christian and Buckley, Andy Cranmer, Kyle and others. "Data and Analysis Preservation, Recasting, and Reinterpretation". In: *Snowmass21 Whitepaper to TF07 and CompF7* (2022).

[280] United Nations Educational, Scientific and Cultural Organization (UNESCO). *Open Science Movement*. URL: http://www.unesco.org/new/en/communication-and-information/portals-and-platforms/goap/open-science-movement/.

[281] FOSTER Plus (European Union funded open project in Horizon 2020 and beyond). *Open Science Definition*. URL: https://www.fosteropenscience.eu/foster-taxonomy/open-science-definition.

[282] United Nations Educational, Scientific and Cultural Organization (UNESCO). *What is Open Access?* URL: https://en.unesco.org/open-access/what-open-access.

[283] *CERN Data Portal*. URL: http://opendata.cern.ch.

[284] *FAIR Principles*. URL: https://www.go-fair.org/fair-principles/.

[285] Jean-Luc Vay et al. "Center(s) for Accelerator and Beam Physics Modeling". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF0-AF1_AF0_Vay-069.pdf.

[286] *MIT Schwarzman College of Computing*. MIT Schwarzman College of Computing. URL: https://computing.mit.edu/ (visited on 06/23/2021).

[287] *QuICS*. URL: https://quics.umd.edu/ (visited on 06/29/2021).

[288] Chris Fisher. *IBM — Quantum Computing*. URL: https://www.ibm.com/quantum-computing//(visited on 06/29/2021).

[289] *Exascale Computing Project*. Exascale Computing Project. URL: https://www.exascaleproject.org/ (visited on 06/23/2021).

[290] *Scientific Discovery through Advanced Computing*. URL: https://www.scidac.gov/ (visited on 06/23/2021).

[291]   J.-L. Vay et al. "A modular community ecosystem for multiphysics particle accelerator modeling and design". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CompF/SNOWMASS21-CompF2_CompF0-AF1_AF0_Vay-070.pdf.

[292]   David Bruhwiler et al. "Collaboration between industry and the HEP community". In: *Snowmass21 LOI* (2020). URL: https://www.snowmass21.org/docs/files/summaries/CommF/SNOWMASS21-CommF1_CommF0-AF0_AF1_Bruhwiler-066.pdf.