



PIConGPU: Predictive Simulations of Laser-Particle Accelerators with Manycore Hardware

Dipl.-Phys. Axel Hübl

Born on: 27th February 1989 in Karl-Marx-Stadt, now Chemnitz

Dissertation

to achieve the academic degree

Doctor rerum naturalium (Dr. rer. nat.)

Referees

Prof. Dr. Ulrich Schramm

Prof. Dr. Thomas E. Cowan

Prof. Dr. Paul Gibbon

Supervisor

Dr. Michael Bussmann

Submitted on: 25th April 2019

Defended on: 25th June 2019



If not indicated otherwise, all materials published in this dissertation by Axel Huebl (orcid:0000-0003-1943-7141) are licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.[§] Permissions beyond the scope of this license may be available by contacting a.huebl@hzdr.de .

Last update: 3rd July 2019 (Post-Defense)
doi:10.5281/zenodo.3266820

[§]<https://creativecommons.org/licenses/by-sa/4.0/>

Für Lotti, meine Eltern, Großeltern und meinen Bruder.

Abstract

The presented thesis establishes simulations on modern massively parallel computing hardware to investigate relativistic laser-driven plasmas. The latter are of special interest as they may provide a compact source for energetic ion beams. Computer simulations provide valuable insight into ultrafast plasma processes, evolving in the ultrahigh intensity ($I_0 \gg 10^{18} \text{ W/cm}^2$) focus of the ultrashort ($\tau_0 = 30 - 500 \text{ fs}$) laser pulses driving the interaction. Such simulations require high numerical resolution and full geometric treatment for reliable predictions, which can only be addressed with high-performance computing. The open source particle-in-cell code PICongPU, which is developed in the framework of this thesis, answers these demands, providing speed and scalability to run on the world's largest supercomputers. PICongPU is designed with a modular and extensible implementation, allowing to compute on current and upcoming hardware from a single code base. Furthermore, challenges arising for generated data rates, reaching 1 PByte per simulation, are resolved with scalable data reduction techniques and novel workflows, such as interactive simulations.

Numerical studies are performed on two novel targets for laser-proton acceleration with near-critical and mass-limited properties. A micrometer-scale spherical target is explored with realistic temporal laser contrast, providing an interpretation for experimental results collected at the PW-class laser system PHELIX ($\tau_0 = 500 \text{ fs}$ pulse length). In this study, 3D modeling with the GPU supercomputer Titan enabled the identification of pre-expansion to near-critical target conditions, which uncovers a regime of volumetric laser-electron interaction generating a highly directed proton beam. Furthermore, a novel cryogenic hydrogen jet target is researched in close collaboration to experiments at the laser system DRACO ($\tau_0 = 30 \text{ fs}$). This target system provides a unique setup for the isolated investigation of multi-species effects and their influence on the generated ion energy distribution. A novel analytical model provides a link between characteristic modulations in the ion energy spectra and ensemble properties of the microscopic electron distribution. In view of a potential experimental realization, parametric scans are performed confirming the feasibility of the proposed setup.

Publications by the Author

Important parts of this thesis are published in or submitted to scientific journals and were presented at international conferences. Individual chapters announce the accompanying publications on which wording and figure overlaps might exist.

- [H1] A Debus, R Pausch, A Huebl, K Steiniger, R Widera, T Cowan, U Schramm, and M Bussmann. "Breaking the dephasing and depletion limits of laser-wakefield acceleration". In: *Submitt. to Phys. Rev. X* (2019).
- [H2] A Huebl, M Rehwald, L Obst-Huebl, T Ziegler, M Garten, R Widera, K Zeil, TE Cowan, M Bussmann, U Schramm, and T Kluge. "Spectral Control via Multi-Species Effects in PW-Class Laser-Ion Acceleration". In: *Submitted* (2019). arXiv: 1903.06428.
- [H3] P Hilz, TM Ostermayr, A Huebl, V Bagnoud, B Borm, M Bussmann, M Gallei, J Gebhard, D Haffa, J Hartmann, T Kluge, FH Lindner, P Neumayr, CG Schaefer, U Schramm, PG Thirolf, TF Rösch, F Wagner, B Zielbauer, and J Schreiber. "Isolated proton bunch acceleration by a petawatt laser pulse". In: *Nat. Commun.* 9.1 (2018), p. 423. DOI: 10.1038/s41467-017-02663-1.
- [H4] L Obst-Huebl, T Ziegler, FE Brack, J Branco, M Bussmann, TE Cowan, CB Curry, F Fiuza, M Garten, M Gauthier, S Göde, SH Glenzer, A Huebl, A Irman, JB Kim, T Kluge, SD Kraft, F Kroll, J Metzkes-Ng, R Pausch, I Prencipe, M Rehwald, C Roedel, HP Schlenvoigt, U Schramm, and K Zeil. "All-optical structuring of laser-driven proton beam profiles". In: *Nat. Commun.* 9.1 (2018), p. 5292. DOI: 10.1038/s41467-018-07756-z.
- [H5] R Pausch, A Debus, A Huebl, U Schramm, K Steiniger, R Widera, and M Bussmann. "Quantitatively consistent computation of coherent and incoherent radiation in particle-in-cell codes—A general form factor formalism for macro-particles". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 909.4 (2018), pp. 419–422. DOI: 10.1016/j.nima.2018.02.020.
- [H6] JP Couperus, R Pausch, A Köhler, O Zarini, JM Krämer, M Garten, A Huebl, R Gebhardt, U Helbig, S Bock, K Zeil, A Debus, M Bussmann, U Schramm, and A Irman. "Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator". In: *Nat. Commun.* 8.1 (2017), p. 487. DOI: 10.1038/s41467-017-00592-7.

- [H7] C Fortmann-Grote, AA Andreev, K Appel, J Branco, R Briggs, M Bussmann, A Buzmakov, M Garten, A Grund, A Huebl, Z Jurek, ND Loh, M Nakatsutsumi, L Samoylova, R Santra, EA Schneidmiller, A Sharma, K Steiniger, S Yakubov, CH Yoon, MV Yurkov, U Zastra, B Ziaja-Motyka, and AP Mancuso. "Simulations of ultrafast x-ray laser experiments". In: *Proc. SPIE 10237, Adv. X-ray Free. Lasers Instrum. IV*. Vol. 10237. 2017, 102370S. DOI: 10.1117/12.2270552.
- [H8] A Huebl, R Widera, F Schmitt, A Matthes, N Podhorszki, JY Choi, S Klasky, and M Bussmann. "On the scalability of data reduction techniques in current and upcoming HPC systems from an application perspective". In: *Lect. Notes Comput. Sci.* 10524.4 (2017), pp. 15–29. DOI: 10.1007/978-3-319-67630-2_2.
- [H9] A Matthes, R Widera, E Zenker, B Worpitz, A Huebl, and M Bussmann. "Tuning and optimization for a variety of many-core architectures without changing a single line of implementation code using the alpaka library". In: *Lect. Notes Comput. Sci.* Vol. 10524. 2017, pp. 496–514. DOI: 10.1007/978-3-319-67630-2_36.
- [H10] L Obst, S Göde, M Rehwald, FE Brack, J Branco, S Bock, M Bussmann, TE Cowan, CB Curry, F Fiuza, M Gauthier, R Gebhardt, U Helbig, A Huebl, U Hübner, A Irman, L Kazak, JB Kim, T Kluge, S Kraft, M Loeser, J Metzkes, R Mishra, C Rödel, HP Schlenvoigt, M Siebold, J Tiggesbäumker, S Wolter, T Ziegler, U Schramm, SH Glenzer, and K Zeil. "Efficient laser-driven proton acceleration from cylindrical and planar cryogenic hydrogen jets". In: *Sci. Rep.* 7.1 (2017). DOI: 10.1038/s41598-017-10589-3.
- [H11] R Pausch, M Bussmann, A Huebl, U Schramm, K Steiniger, R Widera, and A Debus. "Identifying the linear phase of the relativistic Kelvin-Helmholtz instability and measuring its growth rate via radiation". In: *Phys. Rev. E* 96.1 (2017), p. 013316. DOI: 10.1103/PhysRevE.96.013316.
- [H12] U Schramm, M Bussmann, A Irman, M Siebold, K Zeil, D Albach, C Bernert, S Bock, F Brack, J Branco, JP Couperus, TE Cowan, A Debus, C Eisenmann, M Garten, R Gebhardt, S Grams, U Helbig, A Huebl, T Kluge, A Köhler, JM Krämer, S Kraft, F Kroll, M Kuntzsch, U Lehnert, M Loeser, J Metzkes, P Michel, L Obst, R Pausch, M Rehwald, R Sauerbrey, HP Schlenvoigt, K Steiniger, and O Zarini. "First results with the novel petawatt laser acceleration facility in Dresden". In: *J. Phys. Conf. Ser.* 874.1 (2017), p. 012028. DOI: 10.1088/1742-6596/874/1/012028.
- [H13] C Fortmann-Grote, AA Andreev, R Briggs, M Bussmann, A Buzmakov, M Garten, A Grund, A Huebl, S Hauff, A Joy, Z Jurek, ND Loh, T Rüter, L Samoylova, R Santra, EA Schneidmiller, A Sharma, M Wing, S Yakubov, CH Yoon, MV Yurkov, B Ziaja, and AP Mancuso. "SIMEX: Simulation of Experiments at Advanced Light Sources". In: *11th NOBUGS Conf.* Copenhagen, 2016, pp. 29–34. DOI: 10.17199/NOBUGS2016.21.
- [H14] A Matthes, A Huebl, R Widera, S Grottel, S Gumhold, and M Bussmann. "In situ, steerable, hardware-independent and data-structure agnostic visualization with ISAAC". In: *Supercomput. Front. Innov.* 3.4 (2016). DOI: 10.14529/jsfi160403.

- [H15] E Zenker, R Widera, A Huebl, G Juckeland, A Knüpfer, WE Nagel, and M Bussmann. "Performance-portable many-core plasma simulations: Porting PIconGPU to OpenPower and beyond". In: *Lect. Notes Comput. Sci.* 9945.4 (2016), pp. 293–301. DOI: 10.1007/978-3-319-46079-6_21.
- [H16] E Zenker, B Worpitz, R Widera, A Huebl, G Juckeland, A Knupfer, WE Nagel, and M Bussmann. "Alpaka - An abstraction library for parallel kernel acceleration". In: *Proc. - 2016 IEEE 30th Int. Parallel Distrib. Process. Symp. IPDPS 2016* 60.4 (2016), pp. 631–640. DOI: 10.1109/IPDPSW.2016.50.
- [H17] A Huebl, D Pugmire, F Schmitt, R Pausch, and M Bussmann. "Visualizing the radiation of the Kelvin-Helmholtz instability". In: *IEEE Trans. Plasma Sci.* 42.10 (2014), pp. 2638–2639. DOI: 10.1109/TPS.2014.2327392.
- [H18] R Pausch, H Burau, M Bussmann, J Couperus, TE Cowan, A Debus, A Huebl, A Irman, A Köhler, U Schramm, K Steiniger, and R Widera. "Computing angularly-resolved far field emission spectra in particle-in-cell codes using GPUs". In: *Proceeding of IPAC2014*. Vol. MOPRI069. 2014, pp. 761–764. DOI: 10.18429/JACoW-IPAC2014-MOPRI069.
- [H19] M Bussmann, H Burau, TE Cowan, A Debus, A Huebl, G Juckeland, T Kluge, WE Nagel, R Pausch, F Schmitt, U Schramm, J Schuchart, and R Widera. "Radiative signatures of the relativistic Kelvin-Helmholtz instability". In: *SC'13 Proc. Int. Conf. High Perform. Comput. Networking, Storage Anal.* Vol. 60. 4. 2013, pp. 1–12. DOI: 10.1145/2503210.2504564.

Scientific Software by the Author

Over the time of this thesis, the author published and contributed significantly to several scientific software projects in the domain of high-performance computing and plasma physics. A selection of the most popular scientific applications and libraries which are either directly authored or created under the author's supervision and mentorship are listed here.

- **PICongPU** - *Particle-in-Cell Simulations for the Exascale Era* (28 releases)
Role: core developer, maintainer and scientific lead with Michael Bussmann
Developers as of 0.4.3: A. Huebl, R. Widera, B. Worpitz, R. Pausch, H. Burau, M. Garten, S. Starke, A. Grund, A. Debus, A. Matthes, S. Bastrakov, K. Steiniger, I. Göthel, S. Rudat, J. Kelling, M. Bussmann
 - ▶ 0.1.0 doi:10.5281/zenodo.17827 (2015)
 - ▶ 0.2.0 doi:10.5281/zenodo.168390 (2016)
 - ▶ 0.3.0 doi:10.5281/zenodo.810041 (2017)
 - ▶ 0.4.0 doi:10.5281/zenodo.1467038 (2018)
- **openPMD-api** - *C++ & Python API for File Format Agnostic Scientific I/O* (14 releases)
Role: core developer, maintainer and scientific lead
Developers as of 0.8.0: F. Koller, F. Poeschel, A. Huebl
Background: advised F. Koller in undergraduate research project, now continued in F. Pöschel's Master's thesis (2019)
 - ▶ 0.1.1 doi:10.14278/rodare.28 (2018)
 - ▶ 0.8.0 doi:10.14278/rodare.119 (2019)
- **Alpaka** - *Abstraction Library for Parallel Kernel Acceleration* (8 releases)
Role: project mentor, developer
Developers as of 0.3.5: B. Worpitz, A. Matthes, E. Zenker, A. Huebl, R. Widera
Background: co-advised B. Worpitz' Master's thesis (2015), doi:10.5281/zenodo.49768
 - ▶ 0.1.0 (2016)
 - ▶ 0.2.0 (2017)
 - ▶ 0.3.0 (2018)

- **CUDA support in Cling** - *Interactive C++ Interpreter*
Role: supervisor
Developers: S. Ehrig
Background: conceptualized and co-advised S. Ehrig's Diploma thesis (2018),
Oral presentations at ROOT Users' Workshop (2018),
doi:10.5281/zenodo.1412256 and doi:10.5281/zenodo.1412537
- **ISAAC** - *In Situ Animation of Accelerated Computations* (10 releases)
Role: project mentor, contributor
Developers as of 1.4.0: A. Matthes, F. Meyer, E. Zenker, A. Huebl
Background: co-advised A. Matthes' Diploma thesis (2016), doi:10.5281/zenodo.55329
 - ▶ 1.0.0 (2016)
 - ▶ 1.2.0 (2017)
 - ▶ 1.3.0 (2017)
 - ▶ 1.4.0 (2018)
- **ParaTAXIS** - *PARAllel Tracer for Arbitrary X-Ray Interaction and Scattering*
Role: project mentor
Developers: A. Grund
Background: co-advised A. Grund's Diploma thesis (2016), doi:10.5281/zenodo.1164654
- **mallocMC** - *Memory Allocator for Many Core Architectures* (8 releases)
Role: project mentor, maintainer
Developers: C.H.J. Eckert, R. Widera, A. Huebl
Background: co-advised C.H.J. Eckert's GroÙer Beleg (2014), doi:10.5281/zenodo.34461
 - ▶ 2.1.0 (2015)
 - ▶ 2.2.0 (2015)
 - ▶ 2.3.0 (2018)

Technical Documents by the Author

- **openPMD** - *Open Standard for Particle-Mesh Data Files* (3 releases)

The author of this thesis is the original author and continued scientific lead of the openPMD standard, an open community standard for particle-mesh data files.

- ▶ 1.0.0 doi:10.5281/zenodo.33624 (2015)
- ▶ 1.0.1 doi:10.5281/zenodo.1069534 (2017)
- ▶ 1.1.0 doi:10.5281/zenodo.1167843 (2018)

- **PICongPU Manual** (290 pages as of 0.4.3)

The author initiated the document and contributed significantly to its content.

<https://picongpu.readthedocs.io>

Contents

1. Introduction	1
1.1. Motivation	1
1.2. Content of this Thesis	3
2. Theoretical Background	5
2.1. Laser-Plasma Interaction	5
2.1.1. Single Electron	6
2.1.2. Plasma	9
2.2. Ion Acceleration	12
2.2.1. Collisionless Absorption Mechanisms	12
2.2.2. Plasma-Rear Expansion into Vacuum	15
2.2.3. Energy Conversion Model	18
2.2.4. Mass-Limited Targets	19
2.2.5. Ion Energy Spectra and Acceleration Mechanisms	21
2.3. Electro-Magnetic Particle-in-Cell Simulations	24
2.3.1. Governing Set of Equations	24
2.3.2. Discretization	25
2.3.3. Explicit Temporal Evolution	27
2.3.4. Reduced Dimensionality	30
2.3.5. Resolution and Computational Estimates	31
2.3.6. Extended Fundamental Processes	34
2.3.7. Alternative Numerical Methods	35
2.4. Basics of Scientific Computing	36
3. Exascale-Era Simulations with PICongPU	39
3.1. Scope and Applications	39
3.2. Community Code and Automation	41
3.3. Zero-Overhead Abstractions	42
3.3.1. Algorithmic Agility	42
3.3.2. Supercomputing Systems	43

3.3.3.	Scalability	45
3.3.4.	HPC Software Stack	47
3.3.5.	Abstract Accelerator Programming Model	49
3.3.6.	Data Parallelism	53
3.3.7.	Task Parallelism	54
3.4.	IO and Data Reduction	56
3.4.1.	Theoretical Background	56
3.4.2.	openPMD	61
3.4.3.	SIMEX Platform	63
3.5.	Interactive Simulations	64
3.5.1.	In Situ Coupling	64
3.5.2.	Interactive Control and Visualization	66
3.5.3.	Cling CUDA	68
3.6.	Conclusions and Perspectives	71
4.	Volumetric PW-Level Ion Acceleration on the Micrometer-Scale	73
4.1.	Assessment of Realistic Interaction Conditions	75
4.2.	Petabyte-Scale 3D3V Modeling	78
4.2.1.	Temporal Evolution	79
4.2.2.	Ponderomotive Focusing of Electrons	81
4.2.3.	Directed Ion Acceleration	83
4.2.4.	Discussion	85
4.3.	Summary and Perspectives	88
5.	Multi-Species Effects in PW-Class Laser-Ion Acceleration	91
5.1.	Unique Target Properties	92
5.2.	Target Geometry	93
5.3.	Systematic Characterization of Multi-Species Effects	96
5.3.1.	Analytical Model	97
5.3.2.	Systematic Target Preparation	100
5.3.3.	Simulation Setup	100
5.3.4.	Influence of Deuterium Concentration	102
5.3.5.	Phase Space Dynamics	104
5.3.6.	Scaling with Laser Pulse Amplitude	105
5.3.7.	Ion Emission Characteristics	107
5.4.	Summary and Outlook	109
6.	Conclusions and Perspectives	111
A.	Appendix	i
A.1.	Template Meta-Programming Examples	i
A.1.1.	Generalized Particle Shapes and Current Deposition	i
A.1.2.	Particle Species and Compile-Time Code Generation Switch	ii
A.2.	Phase Space Evolution	iii
A.3.	Measuring Prompt Electrons in a Simulation	iv

Glossary

vii

Bibliography

xi

Contents

1. Introduction

1.1. Motivation

Less than 20 years ago, experiments with a new generation of ultrahigh intensity lasers demonstrated that ions can be accelerated to multiple MeV per nucleon within a few micrometers [1–4]. Turning the material in the focus of these laser pulses into a plasma, accelerating fields reach amplitudes in the order of TV/m, four orders of magnitude higher than achievable with conventional accelerator technology [1, 5, 6]. Over the following years, experiments reported maximum ion energies reaching 94 MeV [7–10]. Contrary to conventional accelerators with nanosecond-long bunches, ion bunches from these laser-plasma accelerators can be as short as the driving femtosecond scale laser pulses ($\tau_0 \approx 30 - 500$ fs) and still provide a significant amount of charge [11].

Despite its potential to substantially reduce the size of particle accelerators, the extreme compactness of a few micrometers acceleration length and femtosecond scale acceleration dynamics is also a great challenge for providing reliable particle beams. The ultrafast coupling of a driving laser pulse to the electrons in the few-micrometer laser focus of a plasma target is highly non-linear as well as the desired energy transfer from electrons to ions. As target materials are usually opaque, direct observation of the experimental target conditions in the important few tens of femtoseconds of highest laser intensity is highly demanding for diagnostics and not yet available. Consequently, observed generated particle beams and other forms of target radiation are temporally integrated with little direct evidence of which microscopic condition varied under macroscopic changes in experiments.

Simulations are an essential instrument in order to fill this gap with insight, linking experimental observables and the complex plasma evolution. While analytical models for many-particle systems have to estimate laser-particle acceleration under simplified conditions, computer models provide an ab initio description from the fundamental processes and initial conditions alone. Consequently, first principle simulations can investigate the non-equilibrium dynamics of many-particle plasmas in full within a controlled environment. All the more, both experimental and simulation results must match effective mathematical models for conclusive evidence. Yet still, predictive capabilities arise not from a single

simulation and a well-matched mathematical model, but require parametric scans over an ensemble of uncertainties of typical initial conditions, just as evidence gathered in experiments.

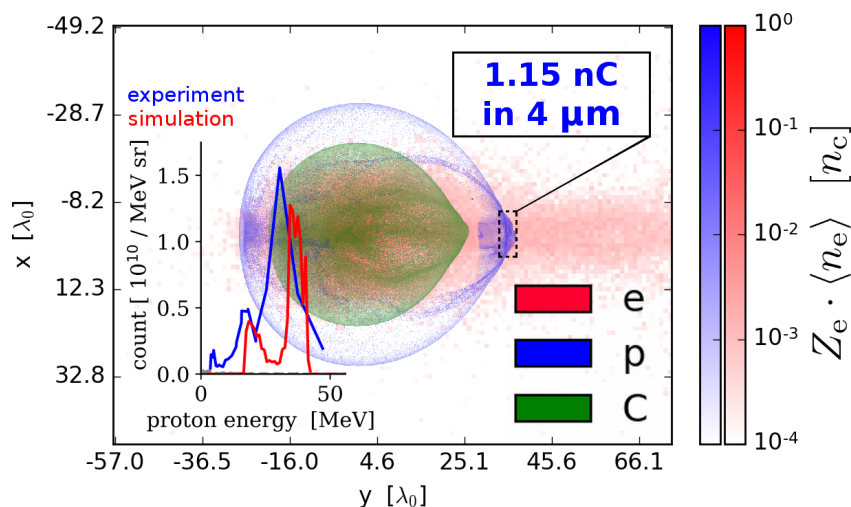


Figure 1.1.: Ultrashort, highly directed proton bunch from laser-ion acceleration with a spherical target. The inset compares the predicted, narrow-width ion spectra from simulation and experiment [11]. Results are presented in chapter 4 of this thesis.

Such simulations with proper geometry and high resolution for even the fastest occurring plasma processes are computationally demanding. As a matter of fact, for adequate predictions such as the simulation shown in Fig. 1.1, our community already uses the largest supercomputers in the world. Computer science methods are integrated deeply in today's research and single-person projects relying on computational frameworks developed a decade ago are unsuited to utilize current and upcoming machines. At the same time, simulations describing electro-magnetic plasmas need to be extended for effects obtainable with upcoming laser systems, as well as previously neglected fundamental processes. Such additional models, e.g. to describe the plasma interaction with an X-ray probe pulse, increase the computational requirements as well.

With modern, capable simulation frameworks, the laser-plasma community can explore favorable new regimes with targets such as the spherical, near-wavelength sized sphere shown in Fig. 1.1, which demands proper geometric modeling. Correlating macroscopic observables from experiments with certain microscopic processes in the plasma can provide control over the ultrafast acceleration process. Such control is essential if one wants to design stable laser-plasma accelerators and further increase their provided beam energy, as both properties are indispensable for future applications such as oncology [12–16], laboratory astrophysics [17, 18], proton radiography [19, 20], warm-dense matter and fusion research [21–23], and nuclear physics [24–27].

1.2. Content of this Thesis

In this thesis, laser-plasma acceleration with an emphasis on novel ion accelerators is studied with simulations and theory. For that, simulations required for current and upcoming supercomputers are developed, recent experimental campaigns are systematically modeled and explained, and analytical theory is derived.

PICongGPU is introduced as a modern, scalable, accelerated simulation framework for laser-plasma physics. It is developed into a state suitable for production-quality runs on state-of-the-art supercomputers driven by manycore accelerators as well as legacy architectures with a single, maintainable code base. On the basis of fast turn-around simulations with these machines, high-resolution and full-geometry studies are possible for systematic understanding of laser-plasma problems. For example, the influence of experimental uncertainties are modeled and explored in extensive scans instead of single simulations. Both software and workflows are conceived, implemented, documented, tested, designed to be scrutable and collaborative, which in some cases even allows to interact live with a study subject. Apart from that, PICongGPU provides flexible numerical methods for agile problem-specific modeling. With high computational capability, the gathering of scientific insight and understanding from simulation-generated data is addressed and novel workflows and techniques to handle large three-dimensional output are presented (chapter 3).

PICongGPU's capabilities are applied and verified against experimental results from near-critical and mass-limited targets such as a micrometer-scale sphere levitating in a Paultrap (chapter 4) and a cryogenic jet target (chapter 5). Both targets provide unique properties for the progress of laser-ion acceleration, such as isolated arrangements and optimal access for diagnostics (Paultrap) or suitability for high-repetition rate experiments and reduced material complexity (cryogenic jet), with favorable density and geometry for an enhanced absorption of laser energy. Results and variations in macroscopic observables are explained under realistic experimental conditions such as a finite temporal laser contrast and significant variations of geometry from ab initio simulations and connected with the underlying microscopic physics.

Target materials of different constituents often complicate the identification of acceleration mechanisms in laser-ion acceleration and few analytical models exist for the influence of multi-ion components on the generated particle beams. Hence, for interpretation of simulations and experiments with multiple species, an analytical model correlating microscopic plasma physics with the resulting spectral modulations of the generated ion beams is derived (chapter 5). Therein, emphasis is set again on predictable and sensitive parameters observable in the integrated beam properties accessible in experiments. A potential experimental realization is proposed and its applicability is confirmed with theory and extensive parameter scans from PICongGPU simulations.

2. Theoretical Background

In this chapter, theoretical models and fundamental effects relevant for laser-ion acceleration are presented. Lasers used in this research field are of ultrashort duration (usually tens of femtoseconds to half a picosecond) and ultrahigh intensity ($10^{18} - 10^{22} \text{ W/cm}^2$), ionizing matter nearly instantaneously which justifies its treatment as a plasma. In order to understand the relativistic processes driven by such immense intensities from first principles, the conversion of energy from fs-scale laser pulses to electrons is discussed first. Furthermore, important collective plasma properties and expansion dynamics are introduced. The latter leads to the widely studied ion acceleration mechanism of target normal sheath acceleration (TNSA). Scaling laws determining the maximum energy of thereby generated ion beams are introduced and applied to favorable target conditions. Most of the laser-plasma processes presented herein are highly non-linear and determined by ultrashort-scale, many-body dynamics. Ab initio descriptions require simulations as an essential instrument for systematic investigations. Hence, the particle-in-cell (PIC) algorithm and its applicability for laser-plasma simulations are discussed.

2.1. Laser-Plasma Interaction

Before discussing the collective interaction of high-power lasers and electro-magnetic plasmas, it is worth considering the case of a single electron under the influence of an electro-magnetic wave. Relativistic effects and the typical normalization of units are introduced first. The fundamentals of collisionless, hot plasmas considered in this thesis as well as their interaction with electro-magnetic waves are described.

2.1.1. Single Electron

Plane Wave




Figure 2.1.: Electro-magnetic wave with linear polarization of the electric field in x , propagating in y .

Figure 2.1 depicts a sketch of an electro-magnetic wave with a wavelength λ_0 and an amplitude E_0 , which propagates freely along the y axis. In such a wave's electric field, an electron at $y = 0$ oscillates with its central frequency ω_0 under the Lorentz force, as

$$\omega_0 = 2\pi c/\lambda_0 \quad (2.1)$$

$$F_{\perp} = q_e E_0 \cos(\omega_0 t + \varphi_0), \quad (2.2)$$

where c is the speed of light in vacuum, q_e is the electron charge and the phase φ_0 is set to zero. Due to the oscillating nature of the electric field, the electron's lateral velocity v_{\perp} will turn its sign twice per laser period.

As the electron oscillates transversely, the magnetic component of the electro-magnetic wave excites an oscillation of the electron in longitudinal direction

$$\begin{aligned} F_{\parallel} &= q_e v_{\perp} B_0 \cos(\omega_0 t) \\ &= -q_e v_{\perp} \frac{E_0}{c} \cos(\omega_0 t). \end{aligned} \quad (2.3)$$

The influence of the magnetic component of the exciting wave reaches the order of the electric component as soon as relativistic velocities are approached. Due to the inertia of an electron oscillating with lateral velocity v_{\perp} , the change of sign in its velocity occurs with a temporal delay to the zero-crossing of the electro-magnetic field and thus the oscillation in y occurs with a frequency of $2\omega_0$ [28].

One defines the ratio between velocity v and speed of light c as $\beta \equiv v/c$ and the Lorentz factor as $\gamma \equiv (1 - \beta^2)^{-1/2}$. In the following, an index e stands for electron quantities, e.g. m_e denotes the electron rest mass. An initially resting electron acquires a momentum of $2m_e c$

or more ($\beta_e \geq 0.4$), within a half-wave of the electric field

$$\beta_e \gamma_e = \frac{E_0 q_e}{m_e c} \int_{-\pi/2\omega_0}^{\pi/2\omega_0} \cos(\omega_0 t) dt \quad (2.4)$$

$$= \frac{E_0 q_e}{m_e c} \frac{2}{\omega_0}, \quad (2.5)$$

if E_0 is large enough to fulfill

$$\frac{E_0 q_e}{m_e c \omega_0} \equiv a_0 \geq 1. \quad (2.6)$$

The normalized electric field amplitude a_0 is the so-called laser strength parameter. The relativistic regime is reached for electric field amplitudes exceeding $a_0 \gtrsim 1$ while relativistic effects are negligible for smaller a_0 .

Nowadays, ultrashort pulse ultrahigh-intensity lasers nominally reach an intensity up to $I_0 = 10^{22} \text{ W/cm}^2$ in a focal spot size of a few microns [29–31], which corresponds to an a_0 of

$$a_0 = \frac{q_e \lambda_0}{\pi m_e c} \sqrt{\frac{I_0}{2 \epsilon_0 c^3}} \quad (2.7)$$

$$\approx 68$$

for linear polarization and 800 nm central wavelength [32]. The symbol ϵ_0 is the dielectric constant. Comparing this to the onset of the relativistic regime, an amplitude equivalent to $a_0 = 1$ corresponds to an intensity of $I_0 = 2.1 \cdot 10^{18} \text{ W/cm}^2$ at 800 nm wavelength.

One might wonder under which conditions relativistic effects become relevant for ions. Replacing the electron mass in Eq. (2.6) with the proton mass, a corresponding relativistic electric field is reached for $a_0 = 1836$. This in turn translates quadratically into a required laser intensity of

$$I_{0,p+} \geq 7.2 \cdot 10^{24} \text{ W/cm}^2, \quad (2.8)$$

which is nearly three orders of magnitude higher than currently available technology. Laser intensity scales linearly with required laser power, which in turn determines the size and costs of such laser systems.

The electron momentum $\mathbf{p} = \gamma m_e \mathbf{v}$ in the fast oscillating term from Eq. (2.2) can be described in terms of the introduced laser strength parameter a_0 in the temporal average $\langle \dots \rangle$ over a laser cycle

$$\frac{p_{\perp}}{m_e c} = a_0 \cdot \sin(\omega_0 t), \quad (2.9)$$

$$\frac{\langle p_{\perp}^2 \rangle}{m_e^2 c^2} = \frac{a_0^2}{2\pi} \int_0^{2\pi} \sin^2(\varphi) d\varphi$$

$$= \frac{a_0^2}{2}, \quad (2.10)$$

and an additional average longitudinal drift [33]

$$\frac{\langle p_{||} \rangle}{m_e c} = \frac{a_0^2}{4}. \quad (2.11)$$

Separating the electron gamma factor into a fast oscillating (||) and a slowly changing (s) component allows to estimate its cycle-average as [34]

$$\langle \gamma_e \rangle \approx \sqrt{1 + \frac{\langle p_s^2 \rangle}{m_e^2 c^2} + \frac{\langle p_{\perp}^2 \rangle}{m_e^2 c^2}} \quad (2.12)$$

$$\approx \sqrt{1 + \frac{\langle p_s^2 \rangle}{m_e^2 c^2} + \frac{a_0^2}{2}}. \quad (2.13)$$

Inhomogeneous Field

In applications of ultrahigh-intensity laser pulses, the electro-magnetic field in the laser focus is both temporally and spatially limited, e.g. in the form of a finite beam with lateral Gaussian envelope. An electron which is initially positioned at the point of highest intensity in the focus will be displaced over a half-wave of the laser pulse and experiences a lower restoring force at the new position, due to the inhomogeneity of the field. Consequently, the electron is not restored to its initial position and drifts away from the region of maximum intensity with each period of the laser pulse. The time-averaged, non-linear force resulting from the spatial intensity envelope of a laser pulse is the so-called ponderomotive force [34], equal to the time-averaged change of the slowly varying component of the electron motion in Eq. (2.13)

$$F_{\text{pond}} = -\frac{dp_s}{dt} = -m_e c^2 \cdot \nabla \langle \gamma_e \rangle. \quad (2.14)$$

In the sub-relativistic case ($a_0 \ll 1$), the ponderomotive force depends on the gradient of the intensity [34]

$$F_{\text{pond}} = -\frac{m_e c^2}{4} \cdot \nabla \langle (E [a_0])^2 \rangle. \quad (2.15)$$

Although the gradient of the averaged laser pulse envelope provides a first qualitative description, the resulting motion of a relativistic electron depends on its time-dependent relativistic gamma factor $\gamma_e(t)$ and numerical treatment of its motion is necessary for most scenarios.

As an outlook, the ponderomotive force is of central importance as the driver for radiation pressure acceleration (RPA) as well as non-linear plasma waves in so-called laser wakefield acceleration (LWFA) of electrons [32]. Computational research in laser-electron acceleration was performed during the time of this thesis [35–37], but results are excluded here for sake of the concentration on laser-ion acceleration. Yet, chapter 4 presents an alternative application of the ponderomotive force, demonstrating ponderomotive focusing of electrons in a favorable regime of laser-ion acceleration.

2.1.2. Plasma

For the relativistic intensities discussed in this work, the considered fields are orders of magnitude higher than typical material ionization thresholds ($I > 10^{12-14}$ W/cm² [38]). The collectively ionized matter can therefore be described as a plasma with electron density n_e . An electrically neutral ion-electron plasma in equilibrium is called a quasi-neutral plasma [39]. One can define the plasma frequency for oscillations of the electron density as [39]

$$\omega_{p,e} := \sqrt{\frac{n_e q_e^2}{\epsilon_0 \langle \gamma_e \rangle m_e}}. \quad (2.16)$$

Disturbances of the system, such as a shift of electrons by a strong laser pulse, will cause a relaxation in form of oscillations against the ion background with this frequency. The definitions of collective plasma properties are usually approximations since assumptions such as negligible ion motion and effective average $\langle \gamma_e \rangle$ are incorporated.

A dense enough plasma can strongly screen an externally driven excitation with its electrons, i.e. an electro-magnetic wave with oscillation frequency ω_0 . One defines the critical density n_c at the point of equivalence between the plasma frequency and the frequency of the wave $\omega_{p,e} = \omega_0$ as

$$n_c := \langle \gamma_e \rangle m_e \epsilon_0 \left(\frac{\omega_0}{q_e} \right)^2. \quad (2.17)$$

For near-infrared electro-magnetic waves of $\lambda_0 = 800$ nm, the critical density corresponds to $n_e = 1.74 \cdot 10^{21}$ cm⁻³. Materials which, when ionized, act as overcritical plasmas ($n_e > n_c$) for this wavelength are, e.g. solid materials such as metals (Cu $n_e \approx 1400 n_c$) or plastics (PMMA $n_e \approx 230 n_c$), liquids (water $n_e \approx 192 n_c$), and even cryogenic or strongly compressed gases ($n_e \approx 1 - 30 n_c$) [40-48].

While electro-magnetic waves can propagate in undercritical plasmas ($n_e < n_c$), overcritical plasmas damp the field amplitude exponentially with depth from the surface. The collisionless skin depth δ characterizes the exponential scale-length of the evanescent wave from the critical surface layer [34]

$$\delta = c/\omega_{p,e}. \quad (2.18)$$

For many relevant cases of laser-matter interaction, an equilibrium electron temperature $\langle T_e \rangle$ can be introduced equivalent to the thermodynamic concept. For sufficiently high mean kinetic energy and low enough density, the temperature can be identified as the mean of the Boltzmann electron energy distribution

$$\frac{dN}{dE} = \frac{N_e}{k_B \langle T_e \rangle} \exp\left(-\frac{E}{k_B \langle T_e \rangle}\right). \quad (2.19)$$

The plasmas considered in this thesis are so-called hot plasmas, with temperature or (non-thermal) average kinetic energy of tens of keV and more. Plasma oscillations can

propagate in a thermal plasma due to the finite group velocity $v_g = d\omega/dk$ in the dispersion relation [39]

$$\omega^2 = \omega_{p,e}^2 + \frac{3}{2}k^2 \frac{2k_B \langle T_e \rangle}{m_e}. \quad (2.20)$$

The so-called Debye length is introduced as [39]

$$\lambda_D := \sqrt{\frac{\epsilon_0 k_B \langle T_e \rangle}{q_e^2 n_e}}. \quad (2.21)$$

Here as well, the ion motion is neglected. Within a Debye length, an exciting electric potential drops exponentially due to charge screening as $\Phi(r) = \Phi_0 \exp(-r/\lambda_D)$ assuming the potential's time-dependence is negligible, e.g. if a periodic excitation of wavelength λ_0 is large compared to λ_D . Shielding of an external potential in a plasma is a kinetic process and depends on further conditions to justify the collective description above. Additionally, enough particles N_D must contribute to the so-called Debye sphere

$$N_D = n_e 4/3\pi \lambda_D^3 \gg 1, \quad (2.22)$$

otherwise structural effects, e.g. the potential of an ion lattice, cannot be neglected or the number of screening charge carriers is too small for a collective description. However, for plasmas considered in this thesis N_D is sufficiently large. The Debye sphere N_D is identical up to a numeric constant to the so-called plasma parameter [34]

$$\Lambda := \frac{12\pi n_e}{Z_i} \lambda_D^3. \quad (2.23)$$

Z_i is the number of free electrons per ion. The term $\ln \Lambda$ is called the Coulomb logarithm, with which one defines the electron-ion collision frequency ν_{e-i} via [34]

$$\nu_{e-i} \approx 2.91 \cdot 10^{-12} \cdot Z_i n_e \langle T_e \rangle^{-3/2} \ln \Lambda \quad (2.24)$$

$$\propto \frac{n_e}{\langle T_e \rangle^{3/2}} \ln \left(\frac{\langle T_e \rangle^3}{n_e} \right). \quad (2.25)$$

In the above formula, frequency, lengths and density are in SI units and the temperature is in eV. The dynamics of a plasma is dominated by plasma oscillations over electron-ion collisions for $\nu_{e-i} \omega_{p,e}^{-1} \ll 1$ [49, 50], justifying collisionless treatment.

Table 2.1 exemplifies three density ranges of typical laser-solid interaction with ultrahigh intensity, ultrashort-pulse conditions. The free electron density is an upper estimate for perfect absorption and high ionization rates [52]. As visible from the ratio of collision frequency over plasma frequency in the $\nu_{e-i} \omega_{p,e}^{-1}$ column, collisional treatment of the plasma is important for regimes of high densities and low average kinetic energies. Examples are low-intensity pre-pulse, multi-picosecond pre-expansion dynamics and high density, bulk-dominated mechanisms such as electron transport and resistance calculations in thick, high-Z plasmas [50]. As a last note, ion-ion collision rates are significantly lower as the

$n_e [n_c]$	a_0	$\langle T_e \rangle [\text{keV}]$	$\nu_{e-i} \omega_{p,e}^{-1}$	example distribution	example material
30	16	0.5	1.4	pre-plasma bulk prompt	cryogenic H ⁺
		10	$1.6 \cdot 10^{-2}$		
		2600	$1.0 \cdot 10^{-5}$		
230	16	0.5	3.9	pre-plasma bulk prompt	C ⁶⁺ in PMMA
		10	$4.6 \cdot 10^{-2}$		
		2600	$2.9 \cdot 10^{-5}$		
1400	16	0.5	10	pre-plasma bulk prompt	Cu ²⁹⁺
		10	$1.1 \cdot 10^{-1}$		
		2600	$7.2 \cdot 10^{-5}$		
30	0.1	1.3	$3.5 \cdot 10^{-1}$	bulk	cryogenic H ⁺
230	0.1	1.3	1.0	bulk	C ⁶⁺
1400	0.1	1.3	2.5	bulk	Cu ²⁹⁺

Table 2.1.: Influence of electron-ion collisions in laser-plasmas, exemplified for three density ranges of solids. The temperature of the thermalized plasma bulk depends heavily on target thickness and absorption and is given for an order of magnitude estimate. Prompt average kinetic energy for laser-accelerated electrons is taken from Ref. [51] Eq. (9). Material example for 230 n_c is polymethyl methacrylate (PMMA). Blue highlighted configurations need collisional treatment. The critical density corresponds to a central laser wavelength of $\lambda_0 = 800 \text{ nm}$.

differential scattering cross-section decreases with mass of the participating particles [53]. Electron-electron collisions become relevant for high-density plasmas with few keV such as the bulk Cu²⁹⁺ case in Tab. 2.1 [54].

In thermal plasmas, electro-static, longitudinal density oscillation can be excited, analogous to acoustic waves in a gas. An ion acoustic wave is a form of energy transport with regions of density compression and rarefaction [39]. Ion acoustic waves propagate with the (local) speed-of-sound defined as [39]

$$c_{s,i} := \sqrt{\frac{\langle \gamma_e \rangle Z_i k_B \langle T_e \rangle + \langle \gamma_i \rangle k_B \langle T_i \rangle}{m_i}} \quad (2.26)$$

$$\approx \sqrt{\frac{\langle \gamma_e \rangle Z_i k_B \langle T_e \rangle}{m_i}}. \quad (2.27)$$

m_i is the ion mass. The average electron velocity in such a wave is large compared to the ion motion and overall drift of the wave, causing a quasi-equilibrium of the electrons and an overall quasi-neutral energy transport [39]. An important case for laser-ion acceleration is discussed in the following section 2.2.2, describing ion acoustic waves in form of pure rarefactions at density gradients [55].

2.2. Ion Acceleration

Figure 2.2 provides an overview over the laser-target interaction and summarizes introduced plasma quantities of the last section. The available laser intensity provided by state-of-the-art laser systems is too low to directly accelerate ions to high energies, as addressed in section 2.1.1. Instead, laser pulse energy is primarily transferred to the electrons of a plasma which then couple further to ions. A laser pulse that is incident on an over-critical plasma target is partly reflected at the critical surface n_c and as the electro-magnetic wave interacts with this surface up to the skin depth δ , energy is absorbed by electrons. The absorption of energy itself, in the simplest case as in section 2.1.1, changes the local average γ_e of front surface electrons, which in turn increases the critical density from Eq. (2.17) in a non-linear manner. Therefore, the overall process of laser-electron coupling is highly non-linear and depends on the specific laser-target conditions, like the material density, the density gradient, the local electron velocity, the surface roughness, etc..




Figure 2.2.: Overview of laser absorption at the target surface into electrons. The interaction depth of the laser pulse depends on the surface density gradient and density, which determines the penetration depth over the skin depth δ .

2.2.1. Collisionless Absorption Mechanisms

In the following, important absorption mechanisms for relativistic intensities of $a_0 \gg 1$ shall be summarized. The most significant mechanism for high a_0 is $\vec{v} \times \vec{B}$ acceleration from the surface skin depth [56, 57]. For a linear polarized plane wave, the electrons are laterally accelerated in each half-wave of the electric field as in Eq. (2.2), changing twice the sign of the perpendicular velocity per period. In the first quarter-wave, the $\vec{v} \times \vec{B}$ -term in Eq. (2.3) will drag electrons from the surface into the vacuum. In the next quarter-wave, the magnetic field is reversed and the $\vec{v} \times \vec{B}$ -term pushes the electrons in forward direction. In

the second half-wave, the perpendicular velocity is reversed and a further pull-push follows.

In the low-intensity case, this surface oscillation will excite plasma oscillations as charges from layers beyond the skin depth try to compensate the excitation in the form of plasma oscillations. However, in the case of ultrahigh laser intensity, forward-pushed electrons will escape the influence of the laser pulse as they are accelerated rapidly into the target to regions deeper than the skin depth [58]. Figure 2.3 displays a typical longitudinal electron phase space with two laser-accelerated electron bunches per laser period and laser amplitude $a_0 = 8$. In the presented temporal snapshot, the laser is about to push the next electron bunch behind the critical density layer at the target front ($y \approx 3.1 \mu\text{m}$).




Figure 2.3.: Longitudinal electron phase space for a near-critical flat hydrogen target with $n_e = 30 n_c$, thickness $d_0 = 2 \mu\text{m}$ (dashed lines are the initial target surfaces) and a laser pulse with wavelength $\lambda_0 = 800 \text{ nm}$, an amplitude $a_0 = 8$ and a pulse length $\tau_0^{\text{FWHM},l} = 30 \text{ fs}$ (same conditions as in chapter 5). On the right, a momentum histogram of the integrated phase space on the left shows the underlying distribution of forward-accelerated electrons. Models in this section try to approximate such distributions with an average kinetic energy $\langle T_e \rangle$ for varying laser-target conditions. The snapshot is taken at the time of the laser peak intensity on target.

Wilks et al. [57] estimated this energy transfer mechanism with the ponderomotive potential

$$T_e^{\text{Wilks}} = m_e c^2 \cdot \left(\sqrt{1 + a_0^2/2} - 1 \right). \quad (2.28)$$

Albeit, this expression can only serve as an estimate of the average free electron quiver energy, see Eq. (2.10), and overestimates the energy transfer to collectively responding electrons at a critical surface. Therefore, the ensemble averaging over a laser pulse period must be performed properly over all γ_e instead of the momentum [51]

$$\langle T_e \rangle_t = m_e c^2 \cdot \left(\left\langle \sqrt{1 + p_y^2} \right\rangle_t - 1 \right). \quad (2.29)$$

A generalized solution of this total, time-averaged electron energy is presented at the end of the section. For the moment, further microscopic absorption effects are discussed.

An important energy transfer mechanism is known as anharmonic resonance absorption [59]. The amplitude of the laser field decays from the target critical surface in an evanescent wave over the plasma skin depth in Eq. (2.18). Hence, electron layers from different target depths are excited as a (laser-)driven harmonic oscillator with varying excitation. Each layer of thickness d corresponds to another eigenfrequency $\omega_{\text{layer}}(d, \chi)$, as the electro-static potential at its location varies with distance to the surface and displacement χ , which lowers its density and increases γ_e in ω_{pe} . The eigenfrequency is large compared to the laser frequency for small excitations and goes to zero for large oscillation amplitudes χ_0

$$\omega_{\text{layer}} = \frac{\pi}{4} \cdot \sqrt{\frac{\omega_{pe}^2 d}{\chi_0}}. \quad (2.30)$$

Resonance with high energy gain occurs when a layer's eigenfrequency equals the exciting frequency. Anharmonic resonance absorption is responsible for electrons accelerated to many times the mean quiver energy in the high-energy tail of the electron distribution. Contrary to linearly polarized laser pulses described so far, the $\vec{v} \times \vec{B}$ force in circular polarization does not oscillate normal to the target surface. Hence, electrons are steadily accelerated by the intensity envelope of the laser pulse. As such, the plasma surface is not harmonically excited and hence anharmonic resonance absorption is strongly suppressed.

Many further absorption mechanisms exist for laser-plasma interaction, yet have less influence on the electron energy density relevant for ion acceleration from solid targets with the intensities and interaction time scales discussed in this thesis. Depending on incidence angle and steepness of the plasma gradient, the plasma oscillation can be resonantly driven by the laser pulse in the critical surface layer. This so-called resonance absorption can contribute significantly for non-relativistic intensities with [60]

$$T_e^{\text{res.abs.}} [\text{keV}] \approx 1938 \cdot (a_0^2 \cdot T_e^{\text{at } n_c})^{1/3}. \quad (2.31)$$

$T_e^{\text{at } n_c}$ is the temperature of the background plasma at the critical layer in keV. As other authors already speculated [61], the electron temperature in resonance absorption scales similarly to an empirical model by Beg et al. [62]

$$T_e^{\text{Beg}} \propto a_0^{2/3} \quad (2.32)$$

for intensities between $10^{16} - 10^{19} \text{ W/cm}^2$. Hence, resonance absorption might be the dominating mechanism for the regime reported by Beg et al., yet underestimates observations in the relativistic regime [51]. Furthermore, an oblique incidence angle of the laser pulse gives rise to so-called vacuum heating or Brunel heating, which accelerates electrons via the electric field of the laser pulse into the target [63, 64]. This mechanism can generate electrons with average energies up to [64]

$$T_e^{\text{Brunel}} \propto (a_0^2 \lambda_0^2)^{1/3-1/2}, \quad (2.33)$$

under ideal conditions. Simulations and theory found that vacuum heating is more efficient than resonance absorption for steep plasma gradients [63, 64]. The exact transition under

ultrahigh laser intensities was currently revisited with experiments and simulations in Ref. [65].

In Kluge et al. [51], a general scaling law for laser pulse interaction with solids based on a general Lorentz invariant electron distribution ansatz is derived with predictions for both non-relativistic and relativistic intensities and without the assumption of a specific absorption mechanism. While this model provides no prediction about the number of accelerated electrons, it has been found to robustly describe the average electron kinetic energy of prompt, laser-accelerated electrons with experiments and simulations

$$T_e^{\text{Kluge}} = m_e c^2 \cdot \left(\frac{\pi}{2K(-a_0^2)} - 1 \right). \quad (2.34)$$

$K(-a_0^2)$ is the complete elliptical integral of the first kind and for $a_0 \gg 1$ one can approximate to

$$T_e^{\text{Kluge}} \approx m_e c^2 \cdot \left(\frac{\pi a_0}{\ln 16 + 2 \ln a_0} - 1 \right). \quad (2.35)$$

As this paper also introduced, the term "heating" in the case of the $\vec{v} \times \vec{B}$ force in the relativistic regime should be interpreted with some caution, since the $\vec{v} \times \vec{B}$ push excites bunches of highly forward-directed electrons into the target. Due to their thin spatial extent and long relaxation time with respect to the background plasma, such electron bunches are not in thermodynamic equilibrium and thus cannot be described by a thermodynamic temperature [51, 66].

2.2.2. Plasma-Rear Expansion into Vacuum




Figure 2.4.: Thermodynamic expansion of a hot plasma into vacuum.

The expansion dynamics of hot, collisionless plasma slabs has been studied for laser-heated plasmas as early as the 1970s. Fluid-models for thermal, sharp rear-gradient plasmas were discussed, e.g. in Refs. [67, 68], under the topic of plasma rarefaction waves. Depending on developing density gradients, the expansion front can exceed the local ion speed of sound which leads to rarefaction shocks.

All models discussed here aim to describe the following situation depicted in Fig. 2.4a. A hot plasma density slab "relaxes" into vacuum causing a density decompression wave of

electrons and ions. For a thermal plasma with electrons following a Boltzmann distribution as in Eq. (2.19), the electrons expand from the front in the order of the Debye lengths (Debye sheath), as shown in Fig. 2.4b. The resulting charge separation causes an electro-static field, from which some high-energy electrons can escape while most remain in proximity of the ion surface, forming an electron sheath.

The heavier ion background reacts to that ambipolar field and expands in an exponential manner. A so-called self-similar solution for the plasma density over time is found with a self-similar length determined by the ion speed of sound c_s , if one assumes quasi-neutrality $n_e \approx \sum_i Z_i n_i$ in the expanding plasma

$$n_e/Z_i \approx n_i = n_{i,0} \exp(-y/c_s t - 1). \quad (2.36)$$

The expansion is called self-similar, since the overall shape of electrons and ions is nearly identical besides the described thermal variation at its front.

Besides this ambipolar description of the overall expansion shape, early plasma-fluid models [67–69] further studied the expansion dynamics of multi-modal electron distributions, as laser-heated plasmas do not thermalize before the start of the expansion. Only part of the electrons absorb the incoming laser light, leading to a "hot" distribution in the laser-accelerated ions and a "cold"(er) distribution in the plasma background bulk electrons.

After the invention of (optical) chirped pulse amplification (CPA) in 1985 [70], the technical implementation of ultrahigh intensity lasers became possible. Deploying early PW-class lasers, experiments by Hattchett et al. (2000) [1] and Wilks et al. (2001) [6] observed the acceleration of ions from solid targets irradiated with $3 \cdot 10^{20} \text{ W/cm}^2$ near-infrared laser pulses and attributed the acceleration mechanism to the aforementioned vacuum expansion, coining the term target normal sheath acceleration (TNSA) due to the target-normal nature of the surface charge separation fields.

Revisiting earlier plasma expansion works in order to formulate scaling laws for this newly found acceleration mechanism, a popular model was published in 2003 by Mora [71]. As it turns out, the assumption of overall quasi-neutrality and neglecting the initial sheath expansion phase, in which the expansion scale is still smaller than the Debye length, leads to inconsistent solutions. Although the self-similar expansion model describes the overall shape of the ion density well, the electro-static field at the ion front is not self-consistent with the predicted velocity in the self-similar (s-s) picture. Interpolating from a numerical solution of a Lagrangian solver, the self-consistent Poisson-solution for the ion front over time was deduced as

$$E_{s-s} = \frac{2E_0}{\omega_{pi} t}, \quad (2.37)$$

$$E_{\text{Mora}} \approx \frac{2E_0}{\sqrt{2e + \omega_{pi}^2 t^2}}, \quad (2.38)$$

$$E_0 = \sqrt{n_{e,0} k_B T_e / \epsilon_0}. \quad (2.39)$$

With this, the ion front velocity and position can be calculated for $\omega_{p,i}t \gg 1$ as

$$y_{\text{front}}(t) = c_s \cdot [2 \ln(\omega_{p,i}t) + \ln 2 - 1], \quad (2.40)$$

$$v_{\text{front}}(t) = c_s t \cdot [2 \ln(\omega_{p,i}t) + \ln 2 - 3]. \quad (2.41)$$

Notably, since the whole expansion process is driven by plasma electrons, the term $\omega_{p,i} = (n_{e,0} Z_i^2 q_e^2 m_i^{-1} \epsilon_0^{-1})^{1/2}$ is a function of density of the electron population [51]. Accordingly, the energy of the ion front, which corresponds to the fastest ions in TNSA, follows as

$$K_{\text{max}}(\langle T_e \rangle, \tau_{\text{acc}}) \cong 2Z_i k_B \langle T_e \rangle [\ln(2\tau_{\text{acc}})]^2. \quad (2.42)$$

Even with Mora's correction of the electric field at the ion front, the final energy K_{max} for $\tau_{\text{acc}} \rightarrow \infty$ still diverges. As already noted in his paper, there are several physical arguments why it must converge, most importantly because of the finite laser energy applied to the system. The acceleration time, on which energy transfer from laser pulse to electrons, rear expansion, and energy transfer from electrons to ions, is limited by the finite laser pulse length t_{laser} driving the interaction. A typical time scale for the overall process has been reported empirically as [72]

$$\tau_{\text{acc}} = t_{\text{acc}} \cdot Z_i (m_e/m_i)^{0.5} \cdot \omega_{p,e} \quad (2.43)$$

$$t_{\text{acc}} \approx 1.3 \cdot t_{\text{laser}}. \quad (2.44)$$

In Mora's derivation of K_{max} in Eq. (2.42), the reachable ion cutoff energy depends on the temperature $\langle T_e \rangle$ of an isothermal, one-temperature electron distribution. Later publications identified the density and average kinetic energy of pre-thermal laser-accelerated electrons as an appropriate generalization of the model [51, 73]. Intriguingly, the temporal average of such electrons still resembles the characteristics of a Boltzmann distribution for thick enough targets [66].

With regards to the normalization of the acceleration time in Eq. (2.43), the underlying dependency of the ion cutoff energy in Eq. (2.42) is transitively a function of (prompt) electron density

$$K_{\text{max}}(\langle T_e \rangle, \tau_{\text{acc}}) = K_{\text{max}}(\langle T_e \rangle, n_e, t_{\text{acc}}). \quad (2.45)$$

One can validate this observation with the initial assumptions of the TNSA model. An expanding electron sheath on the order of the Debye length sets up an electro-static field in which ions are accelerated. Following Gauß's law, this electric field is a function of electron charge density $\rho_e = q_e n_e$ as well as the average kinetic electron energy in Eq. (2.21). The former contributes linear to the field strength and the latter determines the spatial expansion scale length of the electrons $\lambda_D(\langle T_e \rangle^{1/2})$. Consequently, in order to increase ion cutoff energies one must increase the *average electron energy density*. Increasing the average kinetic energy, as the compact formula in Eq. (2.42) suggests, is a suitable approach as long as the number of laser-accelerated electrons does not drop disproportionately. Alternatively, one can investigate methods to increase the number of laser-accelerated

electrons [74], as discussed in section 2.2.4.

As a final note, later works [75, 76] analyze further effects on the plasma expansion, such as the self-consistent cooling of the driving average kinetic energy $\langle T_e \rangle$ during the expansion, as well as thin targets and large rear pre-expansions ($L \gg \lambda_D$). These models for average kinetic electron energy scaling with high laser intensity in TNSA are not purely analytically solvable and the complex in-target evolution of the accelerated electron density is only properly accessible through simulations (section 2.3 and chapter 3). In order to draw connections to scaling laws described in this section, details on measuring the prompt in-target electron distribution in simulations are described in appendix A.3.

2.2.3. Energy Conversion Model

A second class of models for ion acceleration describes ion acceleration based on the transfer of energy. Popular is the model by Schreiber et al. [77], including general parameters such as a finite pulse length driving the expansion and a lateral spread of transported electrons in thick targets.

Schreiber's description is based on a radially confined electron surface charge at the rear. The total energy of a laser pulse $E_L = P_L \tau_L$ is converted with efficiency $0 < \eta < 1$ into prompt electrons. The maximum possible energy an ion species i can gain for infinitely long acceleration is given by

$$E_{i,\infty} = q_i 2 m_e c^2 \sqrt{\frac{\eta P_L}{P_e}}. \quad (2.46)$$

The laser power P_L is normalized by the relativistic power unit $P_e = m_e c^3 / r_e = 8.71$ GW with r_e as the classical electron radius. Conversion efficiency is assumed to scale as in Refs. [74, 78]

$$\eta \approx 1.2 \cdot 10^{-15} \cdot I_L^{3/4}. \quad (2.47)$$

And the laser intensity I_L on target is related to the laser power P_L , assuming a diffraction limited beam propagation

$$I_L = \frac{P_L}{\pi r_L^2} \quad (2.48)$$

where r_L is the spot size of the laser beam.

The finite length of the laser pulse is considered in the assumption that the acceleration time is limited in the integration of the equation of motion. The maximum energy can be approximated by [79]

$$E_{\max} \approx E_{i,\infty} \tanh^2 \left(\frac{\tau_L}{2\tau_0} \right), \quad (2.49)$$

in which τ_0 is the reference time approximating the time an ion will be accelerated by the

electron surface charge,

$$\tau_0 = \frac{R}{\sqrt{2E_{i,\infty}/m_i}}, \quad (2.50)$$

$$R = r_L + d \tan \theta. \quad (2.51)$$

Here, d denotes the target thickness and θ the half-angle of the electron propagation cone, accounting for electron divergence in realistic, non-1D geometries.

2.2.4. Mass-Limited Targets




Figure 2.5.: Prompt (red) and cold (green) electron currents in laser-ion acceleration targets of varied size (interacting volume).

During the interaction of an intense laser pulse with the flat surface of a thick foil, electrons that absorb part of the incoming laser pulse energy are accelerated into target regions beyond the critical density [5, 64, 80]. As prompt electron currents spread both longitudinally and transversely out of the main interaction volume of the laser pulse, cold, thermalized "bulk" return currents continuously compensate the laser-accelerated electron current due to Lenz's law. Hence, since cold electrons from the depth of the target steadily stream into the laser interaction area, the laser pulse effectively deposits energy in a wide volume (or mass) of the target, while only a part of those electrons contribute to the acceleration of ions at the rear. Prompt electrons that reach the target rear surface with high enough energy will propagate into the vacuum, yet a significant part will experience strong fields from the beginning sheath expansion leading to re-attraction towards the target (see again Fig. 2.3 for electrons with $p_y < 0$).

As shown in the sketch in Fig. 2.5, limiting the thickness of the target can exploit the fact that recirculating electrons with average velocity $\langle v_e \rangle \propto \langle T_e \rangle^{0.5}$ that return to the interaction volume of the laser focus before the end of the laser pulse experience additional laser heating, resulting in an increased average electron energy [80–83]. A thin-enough target with thickness d or long-enough laser pulse τ_L is required for repetitive absorption of laser energy into the same electrons $\tau_L \cdot \langle v_e \rangle > 2(d + \lambda_D)$.

The resulting increase in average electron energy density leads to significant enhancement of ion energy, in line with the models of the previous sections [84]. Experimental campaigns were able to reproduce this energy enhancement as long as non-optimal laser contrast

conditions (pre-pulses, ASE level, etc.) did not destroy the target before the interaction with the main pulse [79]. In studies with plasma mirrors [85–91], an on-shot technique improving the temporal laser contrast, thickness scans down to several nm target thickness observed an optimum between the complex interplay of target pre-expansion and the onset of relativistic transparency [92] and volumetric heating [88, 92–95] for

$$d \lesssim \delta \propto \sqrt{\frac{\langle \gamma_e \rangle}{n_e}}, \quad (2.52)$$

see Eq. (2.18) for the skin depth δ .

The latter two effects can be exploited for enhanced absorption of laser pulse energy with targets of near-critical density. In overcritical plasma regions, the exciting laser field is exponentially damped. For targets below the relativistic critical density n_c , the laser pulse interacts with less electron density than at its iso-surface of maximum penetration. Consequently, the highest absorption into electrons occurs at the critical surface for the relativistic $\mathbf{v} \times \mathbf{B}$ force. If the overall target is only moderately overcritical and the skin depth in Eq. (2.18) is long, the laser pulse can penetrate deeper into the critical surface, causing volumetric heating of the target electrons up to the onset of target transparency, see Eq. (2.52).

As shown in the third case of Figure 2.5, lateral limitation of the target volume leads to the confinement of electrons that would otherwise escape the interaction area transversely, leading to an increased electron circulation and reheating and potentially a spatially confined Debye sheath which results in a higher energy density at the target rear surface. This effect is present in experiments with gas cluster targets, water droplets as well as cropped foils and will be revisited in chapter 4 [5, 96, 97]. Electrons that would laterally escape the interaction region may experience a combination of reflection at the lateral surface and Brunel heating (see section 2.2.1). An intuitive classification of laterally limited targets with width w in a typical laser pulse focal spot with width w_0 , compared to the lateral drift distance $t_{\text{acc}} v_{\perp}$ during the acceleration time and pulse duration τ_0 is given in Tab. 2.2 [83]. It is worth adding that only in case of sufficient laser intensity does condition IV lead to a swift removal of target electrons and a subsequent Coulomb-explosion of ions [98, 99].

In summary, an ideal target from an analytical point of view would be limited both in thickness and laterally, to benefit from the effect of electron recirculation. Such mass-limitation enables re-heating of electrons and limits energy dissipation in transverse direction. Furthermore, the target material density should be near-critical, either already in its cold state or due to laser pulse energy conversion into relativistic electrons (which raises the density threshold for n_c in Eq. (2.17)), for optimal energy absorption. The production and alignment of such targets at a high repetition rate in experiments remains challenging, not least due to the undesired influence of target mounting structures on the acceleration [97]. Chapters 4 and 5 of this thesis present ion acceleration results from novel target types that answer some of those technical issues and feature the favorable properties mentioned above.

	condition	w [μm]	electron dynamics	effect
I	$w > t_{\text{acc}} v_{\perp}$	> 200	infinite lateral drift	cold bulk return current (flat foil)
II	$t_{\text{acc}} v_{\perp} > w > \tau_0 v_{\perp}$	$70 - 200$	drift to edges and back to focus region during ion acceleration	increased "hot" density
III	$\tau_0 v_{\perp} > w > w_0$	$w_0 - 70$	drift to edges and back to focus region during laser pulse interaction	increased "hot" density and average energy
IV	$w_0 \geq w$	$< w_0$	lateral and longitudinal electro-static fields approximately equal	potentially Coulomb-explosion

Table 2.2.: Simplified classification of regimes for laterally limited targets for ion acceleration [83]. Symbols denote the target width w , laser focal spot w_0 , pulse duration τ_0 and lateral drift distance $t_{\text{acc}} v_{\perp}$. This table assumes overcritical targets and ultrashort laser pulses. As shown in chapter 4 of this thesis, the electron dynamics may change significantly with target conditions and longer laser pulses.

2.2.5. Ion Energy Spectra and Acceleration Mechanisms

In laser-ion acceleration, energy spectra are a central observable for experiments. Since the main interaction between ultrashort laser pulse and target occurs on time scales of femtoseconds in mostly opaque materials, plasma processes which are microscopically influencing the acceleration process are not directly accessible through diagnostics in experiments. Even important properties in the focus of an ultrahigh intensity laser pulse require estimates, e.g. radial intensity distribution and peak intensity. Complex probing diagnostics such as shadowgraphy [100], interferometry [101], reflectometry [102, 103] and phase contrast imaging are used today, but upcoming experiments at XFEL facilities (e.g. HIBEF) might be the first providing an experimental insight for fundamental observables like the charge state distribution [104–107]. On the other hand, generated ion beams can be recorded with less complex detectors, providing a macroscopic and time-integrated measurement of the accelerated charge, spectral shape and cutoff energy.

Figure 2.6 correlates the initial position of an ion in a foil target ($3 \mu\text{m} < y_0 < 5 \mu\text{m}$) with its final energy after target normal sheath acceleration (TNSA). From an abstract point of view, one could interpret this plot with its clear relation between proton position in the target rear ($y_0 \geq 4.9 \mu\text{m}$) and reachable energy as "TNSA potential" for the overall acceleration. Exemplified on the right of the figure for the integrated charge distribution, TNSA generates an exponential ion energy spectrum [71]. The high energy tail of the energy distribution is generated from surface ions (red box) undergoing the strongest acceleration in the ambipolar fields of the Debye sheath. Under unmatched conditions, the maximum ion (cutoff) energy increases proportionally to the laser strength parameter a_0 [6, 27, 71]. Yet, TNSA cutoff energies were recently found to scale favorably with $\tau_0^2 a_0^2$ when matching laser pulse energy and duration τ_0 to the target thickness [51, 77, 79].

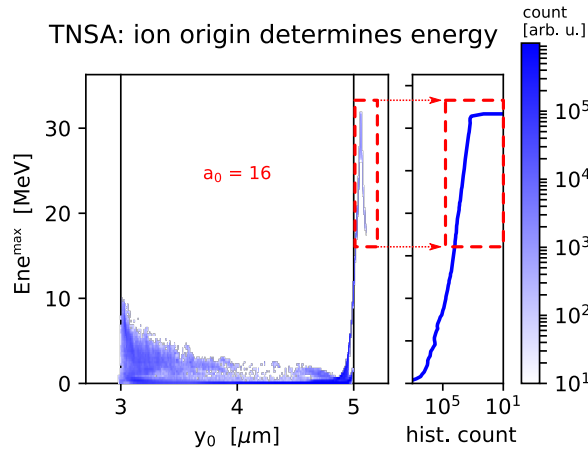


Figure 2.6.: Correlation of a proton's initial in-target position y_0 with its maximum energy after acceleration for typical TNSA conditions. Particles selected for this plot are from a narrow angular emission region in forward direction ($\pm 2^\circ$ from the laser pulse propagation axis). The target material is solid hydrogen in flat geometry at $n_e = 30 n_c$, with a thickness of $2 \mu\text{m}$ and an exponential pre-plasma scale length of 20 nm . The laser pulse strength in focus is $a_0 = 16$ and the pulse length is $\tau_0 = 30 \text{ fs}$. The presented data is generated with a 2D3V simulation from PICongPU (chapter 3) for a scenario studied in chapter 5.

Besides TNSA, additional acceleration mechanisms are predicted from simulations and theory [27], the confirmation of which being the subject of a number of experiments [5, 6, 108–111]. Light sail radiation pressure acceleration (LS RPA) assumes sufficiently high laser pulse intensity to push a very thin, overcritical electron sheath with ions following suit in phase with a laser wavefront. Expected ion spectra are mono-energetic and the peak energy scales proportionally with the maximum time pushing the critical layer $\tau_{\text{push},0}$ and the laser intensity, as $\tau_{\text{push},0}^2 a_0^2$ [112]. In thicker targets, ions at the front surface might experience hole boring radiation pressure acceleration (HB RPA), which pushes electrons and ions from the critical front layer promptly into the target [113, 114]. Such ions propagate through the target and may excite further acceleration mechanisms such as collisionless shocks [115–119], in case the ion front's velocity exceeds the local speed of sound in the target. Background ions, which are potentially reflected under such favorable conditions, are accelerated to twice the shock velocity. As ions finally reach the target rear-side, subsequent energy enhancement through TNSA might occur [81, 120]. For sufficiently small targets and high enough laser intensity [98], an ultrashort laser pulse is able to swiftly remove all electrons from a target leading to a Coulomb-explosion of remaining ions [5, 99, 121]. Coulomb-explosion energy spectra often have a suppressed low-energy tail compared to TNSA [122, 123], yet may be modified by further mechanisms during the expansion [124]. Additional acceleration mechanisms are observed or predicted for special laser-target conditions [27, 125, 126].

Solving the inverse problem of reconstructing an acceleration mechanism from an measured spectral ion distribution, requires acknowledging that the latter is an integrated emission signal projected far from its source in the target. For example, surface contamination layers on the target front and rear present a spatially limited source volume for protons. An observed peaked, quasi-monoenergetic ion spectrum could for instance be attributed to a non-TNSA acceleration mechanism or could equally be the result of spatially

confining the proton source (compare to Fig. 2.6)[83, 127, 128].¹ Furthermore, variations in absorption at the target surface with finite temporal laser pulse contrast and in-target electron transport, both potentially modified by, e.g. plasma filamentation instabilities [104, 129], influence the acceleration process and the resulting ion spectra. Even target mount structures [97] and mass-limitation (see section 2.2.4), can modify the accelerating fields and might mimic signatures of alternative acceleration mechanisms [43, 46]. Consequently, attributing spectral signatures to a certain aspect of the acceleration is highly ambiguous.

As will be discussed in chapter 4, more than one acceleration mechanism or respective geometric effect can occur simultaneously under realistic experimental conditions. Chapter 5 presents how the choice of target material modifies the ion energy spectra and even generates narrow bandwidth beams within the standard TNSA description. Due to such complexities posed in the interpretation of contemporary laser-ion experiments, simulations that are able to connect the microscopic plasma conditions with measurable macroscopic signatures are introduced in the following section 2.3 and researched in chapter 3 of this thesis.

¹Target contamination layers are present on the target front and rear surface and accelerated low-energy protons might still contribute to a final spectrum.

2.3. Electro-Magnetic Particle-in-Cell Simulations

2.3.1. Governing Set of Equations

The particle-in-cell (PIC) method is a popular numerical scheme to solve N-body problems with long-range potentials, such as plasmas. Based on first principles of classical electromagnetism, relativistic plasmas can be modeled ab initio and including kinetic effects. The governing system of equations is the full set of Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{1}{\epsilon_0} \sum_s \rho_s, \quad (2.53)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (2.54)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (2.55)$$

$$\nabla \times \mathbf{B} = \mu_0 \left(\sum_s \mathbf{J}_s + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right), \quad (2.56)$$

for multiple particle species s . $\mathbf{E}(t)$ represents the electric field, $\mathbf{B}(t)$ the magnetic field, ρ_s the charge density and $\mathbf{J}_s(t)$ the current density, μ_0 is the vacuum permeability and ϵ_0 is the dielectric constant.

A many-particle system in non-equilibrium can be described via the distribution function f_s for particle species s and fulfills the collisionless Boltzmann-equation

$$\frac{\partial f_s}{\partial t} + \frac{\mathbf{p}}{m_s} \cdot \nabla f_s + \mathbf{F} \cdot \frac{\partial f_s}{\partial \mathbf{p}} = \left(\frac{\partial f_s}{\partial t} \right)_{\text{coll}} \equiv 0. \quad (2.57)$$

\mathbf{F} is the force field, \mathbf{p} the momentum, and m_s the mass of a species s . The first term in Eq. (2.57) is the change of the distribution function with time, the second term describes spatial particle transport with momentum, the third term the coupling to external fields with resulting change in momentum and the right-hand-side can be added to describe collisions between particles.

For collisionless, electro-magnetic plasmas the equations of motion are determined by the Lorentz force

$$\frac{d}{dt} \mathbf{P}_s(t) = q_s [\mathbf{E}(\mathbf{X}_s(t), t) + \mathbf{V}_s(t) \times \mathbf{B}(\mathbf{X}_s(t), t)], \quad (2.58)$$

$$\frac{d}{dt} \mathbf{X}_s(t) = \mathbf{V}_s(t). \quad (2.59)$$

$\mathbf{X}_s = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)_s$ and $\mathbf{P}_s = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N)_s$ are vectors for the ensemble of N particles in species s . For an individual particle at index i , the momentum \mathbf{p}_i is the relativistic expression $\mathbf{p}_i = \gamma_i m_s \mathbf{v}_i$, and accordingly $\mathbf{V}_s = (\mathbf{v}_1, \mathbf{v}_2, \dots)_s$. In Cartesian coordinates, a single particle position is just $\mathbf{x}_i = (x, y, z)^T$.

It follows the Vlasov-Maxwell-Equation

$$\partial_t f_s(\mathbf{x}, \mathbf{p}, t) + \frac{\mathbf{p}}{m_s} \cdot \nabla_{\mathbf{x}} f_s(\mathbf{x}, \mathbf{p}, t) + \frac{q_s}{m_s} [\mathbf{E}(\mathbf{x}, t) + \mathbf{v} \times \mathbf{B}(\mathbf{x}, t)] \cdot \nabla_{\mathbf{v}} f_s(\mathbf{x}, \mathbf{p}, t) = 0 \quad (2.60)$$

with q_s/m_s as the charge to mass-ratio of a species.

2.3.2. Discretization




Figure 2.7.: Lagrangian particles (orange, green) distributed over an Eulerian mesh for the electro-magnetic field (red, blue). Discrete positions of the Eulerian \mathbf{E} and \mathbf{B} fields are staggered via the Yee scheme [130].

In the numerical discretization of 3D3V electro-magnetic PIC, fields are described in an Eulerian frame of reference and particles in a Lagrangian frame of reference. The Eulerian description focuses on specific locations and observes the change of fields over time while the Lagrangian specification follows individual kinetic markers in space and time.

Position \mathbf{X}_s and momentum \mathbf{P}_s are interpreted as vectors of generalized marker coordinates approximating the distribution function f_s . In order to scale to the represented physical system, each marker is assigned an individual weight $w_s = (w_1, w_2, \dots, w_N)_s$. The temporal evolution of the real distribution function is simulated by advancing the markers over time. The evolution of each marker i represents w_i equivalent particle trajectories in order to save computational time.

Instead of directly solving the six dimensions in space and momentum of the distribution function $f_s(\mathbf{x}, \mathbf{p}, t)$ over time, the problem is reduced to three spatial dimensions (3D) as sketched in Fig. 2.7. Fields such as $\mathbf{E}(\mathbf{x}, t)$ and $\mathbf{B}(\mathbf{x}, t)$ are discretized on a regular mesh and updated over time, described via an Eulerian frame of reference. An element of the mesh is usually referred to as cell. Particle markers are distributed over the cells at arbitrary in-cell positions. In order to interpolate between both domains for force and particle flux calculations, markers can be evaluated at discrete positions of the mesh with a density assignment function of order m , approximating the represented charge density $\rho_s(\mathbf{x}, t)$.

A typical choice for marker assignment functions are B-splines $S^m(x)$ [131]

$$S^0(x) = \begin{cases} 1 & \text{if } 0 \leq x < 1 \\ 0 & \text{else} \end{cases}, \quad (2.61)$$

$$S^m(x) = (S^{m-1} \star S^0)(x). \quad (2.62)$$

Here, x is the normalized coordinate in a 1D unit cell of length Δx . B-splines are not only easy to implement but additionally offer the advantages of symmetry, straight-forward extensibility to higher orders, and a compact explicit form. Table 2.3 summarizes notations for assignment functions in popular PIC literature [131, 132].

order m	density notation ("cloud" [131])	shape notation (assignment function S^m [132])
0	nearest-grid-point (NGP)	piecewise constant shape
1	cloud-in-cell (CIC)	piecewise linear shape
2	triangular-shaped cloud (TSC)	piecewise quadratic shape (PQS)
3	piecewise quadratic cloud shape (PQS)	piecewise cubic shape (PCS)
4		piecewise quartic shape (P4S)
5		piecewise quintic shape (P5S)

Table 2.3.: Commonly used naming schemes for particle B-splines in PIC codes, depending on the polynomial order m of the spline function. The differentiation of some authors between assignment function S^m and evaluated density "cloud" causes naming conflicts.

Assignment functions are also referred to as particle shape and represent the pre-integrated, normalized density function ("cloud") for the evaluation at cell corners [131]. For example, the zero-order shape $S^0(x)$ will contribute 100% of the density of a marker to a cell i but evaluates to zero at neighboring cell positions $i - \Delta x$ and $i + \Delta x$. As particles move, the zero-order shape assignment appears as discontinuous particle flux on the Eulerian cells and is therefore not used in practice.² In three dimensions, the marker shape is a multiplicative union of B-splines $S^m(\mathbf{x}) = S^m(x)S^m(y)S^m(z)$. Particle shapes mediate between the Eulerian field quantities and Lagrangian markers for particles and higher orders lead to smaller numerical ringing effects from grid aliasing [132].

Each marker further carries a momentum vector (3V), contributing a delta-function for its location in momentum space. This is simply a practical choice, keeping the spatial extent of a marker constant over time. Otherwise, particle shapes would require flexible adjustment over arbitrary orders or a method for dynamic particle splitting. Trajectories of particle markers are followed over time, which is a Lagrangian frame of reference.

Depicted in Fig. 2.8, the distribution function f_s is sampled with particle markers. Each marker carries a finite spatial extent and a discrete momentum, similar to a rigid body

$$f_s(\mathbf{x}, \mathbf{p}, t) \approx \sum_{i=1}^N w_i \cdot \underbrace{S^m(\mathbf{x} - \mathbf{x}_i(t))}_{3D} \underbrace{\delta(\mathbf{p} - \mathbf{p}_i(t))}_{3V}. \quad (2.63)$$

²Instead, $m \geq 1$ shape orders can be implemented with both momentum-conserving and charge-conserving properties, as shown in the following section.

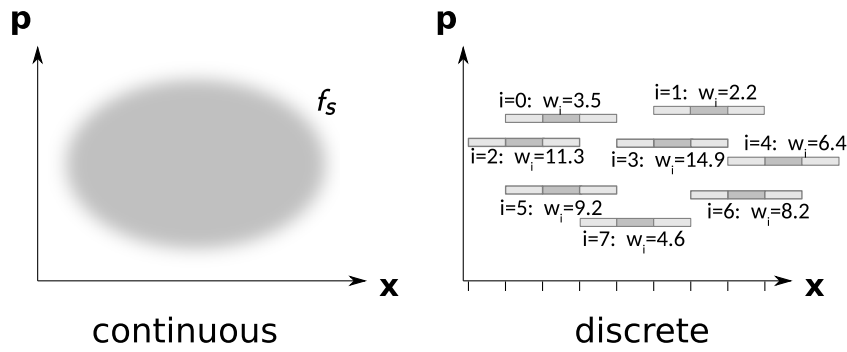


Figure 2.8.: Scheme of Lagrangian particles approximating the distribution function f_s after discretization into markers. In PIC, a phase space distribution such as the ellipse of a particle beam (left) is approximated with representative particle markers (right), each with finite spatial extent and discrete value in momentum.

2.3.3. Explicit Temporal Evolution

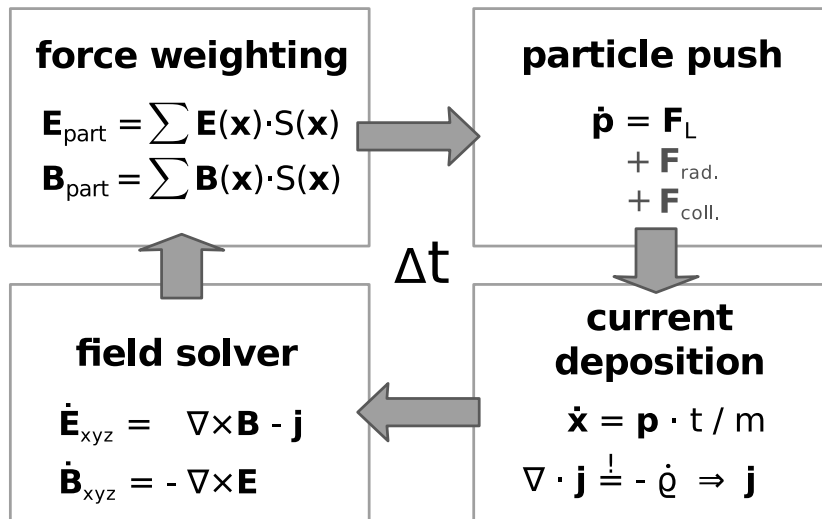


Figure 2.9.: Explicit particle-in-cell cycle for an electro-magnetic PIC code.

The numerical schemes relevant for this thesis are explicit methods which advance the system in n discrete time steps of size Δt [131, 132]. Figure 2.9 shows a typical explicit time integration step of an explicit PIC simulation. Starting in the bottom left of the image with the field solver, the rotation terms in Faraday's Law (2.55) and Ampere-Maxwell's Law (2.56) are locally solved to update the electro-magnetic field in time. A typical choice of discretization for \mathbf{B} and \mathbf{E} fields is a component-wise staggered cell, as seen in Fig. 2.7, with half-step temporal offset between \mathbf{B} and \mathbf{E} . With staggered cells, a second-order accurate finite-difference time-domain (FDTD) method (or Yee's method) [130] can be defined, updating

local cells for Eqs. (2.55) and (2.56) from next-neighbor differences

$$\mathbf{E}_{ijk}^n = \mathbf{E}_{ijk}^{n-1} + \nabla^{(-)} \times \mathbf{B}_{ijk}^{n-1/2} \cdot \frac{\Delta t}{\epsilon_0 \mu_0}, \quad (2.64)$$

$$\mathbf{B}_{ijk}^{n+1/2} = \mathbf{B}_{ijk}^{n-1/2} - \nabla^{(+)} \times \mathbf{E}_{ijk}^n \cdot \Delta t. \quad (2.65)$$

Here, $\nabla^{(-)} \times$ and $\nabla^{(+)} \times$ are discretized linear operators (stencils) over cells with indices ijk and size Δx , Δy , and Δz [130]

$$\begin{aligned} \nabla^{(-)} \times \mathbf{B} = & \left(\frac{B_{z,ijk} - B_{z,ij-1k}}{\Delta y} - \frac{B_{y,ijk} - B_{y,ijk-1}}{\Delta z} \right) \mathbf{e}_x \\ & + \left(\frac{B_{x,ijk} - B_{x,ijk-1}}{\Delta z} - \frac{B_{z,ijk} - B_{z,i-1jk}}{\Delta x} \right) \mathbf{e}_y \\ & + \left(\frac{B_{y,ijk} - B_{y,i-1jk}}{\Delta x} - \frac{B_{x,ijk} - B_{x,ij-1k}}{\Delta y} \right) \mathbf{e}_z, \end{aligned} \quad (2.66)$$

$$\begin{aligned} \nabla^{(+)} \times \mathbf{E} = & \left(\frac{E_{z,ij+1k} - E_{z,ijk}}{\Delta y} - \frac{E_{y,ijk+1} - E_{y,ijk}}{\Delta z} \right) \mathbf{e}_x \\ & + \left(\frac{E_{x,ijk+1} - E_{x,ijk}}{\Delta z} - \frac{E_{z,i+1jk} - E_{z,ijk}}{\Delta x} \right) \mathbf{e}_y \\ & + \left(\frac{E_{y,i+1jk} - E_{y,ijk}}{\Delta x} - \frac{E_{x,ij+1k} - E_{x,ijk}}{\Delta y} \right) \mathbf{e}_z. \end{aligned} \quad (2.67)$$

Explicitly solving these partial differential equations (PDEs) leads to stable solutions if the Courant-Friedrichs-Lewy (CFL) condition [133] is fulfilled, which reads for the Yee scheme

$$c_0 \Delta t \leq \Delta x \quad (1D),$$

$$c_0 \Delta t \leq \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} \right)^{-0.5} \quad (2D),$$

$$c_0 \Delta t \leq \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right)^{-0.5} \quad (3D). \quad (2.68)$$

Explicitly forward updating \mathbf{E}_{ijk}^n and $\mathbf{B}_{ijk}^{n+1/2}$ over n time steps is called leapfrog integration, since both fields are updated via each others time-centered values [132].

Next in Figure 2.9 follows the force weighting in which the discrete fields in the Eulerian frame are interpolated at particle positions. Most common is a trilinear interpolation via the particle assignment function to all cells with non-zero density contribution of the particle. Alternative interpolation methods exist, yet differ mainly in the limit of infinitesimal small cells [134]. With interpolated fields, particle momenta are updated according to the Lorentz force ("particle push" in Fig. 2.9). Due to temporal interpolation of the fields with half-step offset, the magnetic rotation term needs to be applied at half-push in order to conserve momentum, first defined by Boris [135]

$$\mathbf{p}^- = \mathbf{p}^{n-1} + \frac{q\mathbf{E}}{2} \cdot \Delta t \quad \text{1st half acceleration,} \quad (2.69)$$

$$\mathbf{a} = \frac{q}{2\gamma(\mathbf{p}^-)m} \mathbf{B} \cdot \Delta t, \quad (2.70)$$

$$\mathbf{b} = \frac{2\mathbf{a}}{1 + |\mathbf{a}|} \quad \text{energy conserving} \quad (2.71)$$

$$\mathbf{p}' = \mathbf{p}^- + \mathbf{p}^- \times \mathbf{a} \quad \text{velocity space rotation,} \quad (2.72)$$

$$\mathbf{p}^+ = \mathbf{p}^- + \mathbf{p}' \times \mathbf{b}, \quad (2.73)$$

$$\mathbf{p}^n = \mathbf{p}^+ + \frac{q\mathbf{E}}{2} \cdot \Delta t \quad \text{2nd half acceleration.} \quad (2.74)$$

Experience from applications showed that momentum-conservation alone, also known as structure- or phase-space volume-preservation, is not sufficient for systems with highly relativistic particle bunches and the compensation of the relativistic E-field with the self-generated B-field of beams is desired as well, e.g. to model appropriate beam emittance [136]. Recently, Higuera and Cary derived a slightly modified particle push scheme which fulfills both requirements [137]. Within the particle push step, further momentum-modifying effects such as radiation reaction, QED-photon emission, or collision operators can be added (section 2.3.6) [54, 138–143].

What follows is the update of the particle position and so-called current deposition (last box in Fig. 2.9). The goal of this step is preserve the discretized form of Gauss's law (2.53) which is equivalent to self-consistently fulfilling the continuity equation

$$\frac{\Delta\rho_{s,ijk}}{\Delta t} = -\nabla^{\text{Yee}} \cdot \mathbf{j}_{ijk}. \quad (2.75)$$

This step depends on the chosen numerical discretization of the field solver for the choice of ∇^{Yee} and the particle shape for the determination of $\Delta\rho_s$ (section 2.3.2). With the limit of the temporal step Δt from the FDTD field solver, see Eq. (2.68), a charged particle is limited to a change in position of less than one cell per iteration. This allows to avoid solving a global Poisson equation for the scalar potential of the new electric field \mathbf{E} , which only changes locally within Δt . Instead, a self-consistent current \mathbf{j} can be derived for the change of particle position and according Eulerian density $\Delta\rho_s$. With that current, one updates the local \mathbf{E} -field via Ampere-Maxwell's Law (2.56) and at next half-step the local \mathbf{B} -field. Current deposition is usually the most time-consuming step of the PIC cycle, as the local particle assignment function needs to be integrated over all potentially crossed cell boundaries and accumulated from all contributing particles [144–148].³

The overall scheme in Fig. 2.9 is repeated until enough time steps Δt are simulated to cover the time span of the plasma process of interest. Initial conditions and applied boundary conditions must fulfill Gauss's Law (2.53) and Gauss's Law for magnetism (2.54) as well since they are only implicitly included in the individual solver steps of the PIC cycle [145].

³The efficient trajectory splitting scheme in Ref. [146] was also implemented in PIConGPU and it was found that the terms for contributions at the corners of a cell are missing, which violates charge conservation. One can easily add those terms as in Ref. [144].

2.3.4. Reduced Dimensionality




Figure 2.10.: Typical geometrical symmetries and their physical implications in PIC simulations.

PIC simulations can be performed in reduced geometry if the studied system is sensibly described with less dimensions, which saves computational costs. Examples of typical symmetries shown in Fig. 2.10. A widely used reduction is so called 2D3V modeling. In the most common 2D3V geometry, a translational symmetry in the third spatial dimension is assumed. Field gradients and particle position updates in the third partial dimensions are neglected, which is equivalent to periodic boundary conditions for that dimension. Hence, since one assumes a symmetry of infinitely many "planes" in the periodic dimension, the modeled charge carriers are not of point-geometry⁴ anymore but infinitely long wires (translational symmetric in the neglected dimension). Such "wire"-particles might even gain momentum in the third dimension which will self-consistently generate a B-field in the two other dimensions, just like a current in a wire.

Self-consistent fields \mathbf{E} and according scalar potentials Φ , e.g. from charge separation of an electron and ion, are therefore implicitly modified in translational symmetry to

$$\mathbf{E}(r)_{\text{point}} \propto r^{-2} \quad \Phi(r)_{\text{point}} \propto r^{-1} \quad 3\text{D3V}, \quad (2.76)$$

$$\mathbf{E}(r)_{\text{wire}} \propto r^{-1} \quad \Phi(r)_{\text{wire}} \propto \ln(r) \quad 2\text{D3V}, \quad (2.77)$$

$$\mathbf{E}(r)_{\text{plane}} \propto \text{const} \quad \Phi(r)_{\text{plane}} \propto r \quad 1\text{D3V}. \quad (2.78)$$

⁴With finite shape order from section 2.3.2, general particle markers represent density cloud shapes of ellipsoidal geometry in 3D3V.

Figure 2.10 visualizes typical symmetries implemented in PIC simulations. The upper row shows rotational symmetries and the lower the already introduced translational symmetries. Rotational symmetries [149] are unsuitable for modeling laser-matter interaction with polarized laser pulses, simply because the \mathbf{E} -field polarization of a propagating wave cannot be expressed adequately. Nevertheless, further geometric developments exist, e.g. decompositions in Fourier-Bessel modes, which allow quasi-cylindrical treatment [150, 151]. Unfortunately, geometric decompositions around a symmetry axis still suffer from significant particle noise originating from particle trajectories close to the rotational axis for applications with overdense plasma.

Therefore, the most common reduced geometry for laser-ion acceleration, e.g. for interaction with flat foils, remains the translational 2D3V symmetry. When simulating in 2D3V, one must be aware that, e.g. the laser focus is an infinite line along the z axis and that the electro-static fields between charge carriers are overestimated, thus overestimating the accelerating fields in TNSA and resulting ion energy (see section 2.2.2). Therefore, proper 3D3V modeling is nearly unavoidable for predictive studies. The development of PIConGPU within the framework of this thesis (chapter 3) enables such studies, especially for overdense plasmas with naturally high resolution requirements.

2.3.5. Resolution and Computational Estimates

When modeling laser-plasma interactions with the PIC algorithm, smallest-scale processes need to be sampled in both space and time. The smallest wavelength in underdense plasmas is determined by the laser wavelength, while overdense plasmas are dominated by plasma oscillations from collisionless processes (see section 2.1.2). Moreover, numerical methods introduce discretization artifacts in derived physical properties modeled on a PIC grid. For example, a plane wave of the form $\mathbf{E}(x, t) = \mathbf{E}_0 \cdot \exp(ik_x x - i\omega t)$ propagates along the x axis with the wavevector $\mathbf{k} = (k_x, 0, 0)^T$ and the wavenumber $k_x = 2\pi/\lambda_0$, λ_0 as the wavelength, t as the time, $\omega = 2\pi f$ as the angular frequency, and f as the frequency. The physical dispersion relation for electro-magnetic waves in vacuum is given with speed of light c_0 by $\omega = c_0 \cdot |\mathbf{k}|$. Yet, when calculating the numerical dispersion relation for the Yee field solver [130] from Eqs. (2.64) and (2.65) follows the modified relation [152]

$$\frac{\sin(\omega\Delta t/2)^2}{c_0^2\Delta t^2} = \frac{\sin(k_x\Delta x/2)^2}{\Delta x^2} + \underbrace{\frac{\sin(k_y\Delta y/2)^2}{\Delta y^2}}_{=0 \text{ if } k_y=0} + \underbrace{\frac{\sin(k_z\Delta z/2)^2}{\Delta z^2}}_{=0 \text{ if } k_z=0}. \quad (2.79)$$


Consequently, the group velocity $v_g(k_x) = \partial\omega/\partial k_x$ and phase velocity $v_\phi(k_x) = \omega/k_x$ in simulations are artificially slower than the vacuum speed of light. They read for given time step Δt and cell size Δx ,

$$v_\phi(\Delta x, \Delta t, k_x) = \frac{2}{k_x\Delta t} \arcsin \left[\frac{c_0\Delta t}{\Delta x} \sin(k_x\Delta x/2) \right], \quad (2.80)$$

$$v_g(\Delta x, \Delta t, k_x) = \frac{c_0 \cdot \cos(k_x\Delta x/2)}{\sqrt{1 - \frac{c_0^2\Delta t^2}{\Delta x^2} \sin(k_x\Delta x/2)^2}}. \quad (2.81)$$

Figure 2.11 visualizes the numerical dispersion of electro-magnetic waves in 1D to 3D PIC with the Yee solver. Time steps are limited by the CFL condition in Eqs. (2.68) which introduces strong dispersion effects which in turn can be shifted by higher resolution.

One particular problem arises from the limited phase velocity called numerical Cherenkov radiation [153]. Just as in a dispersive medium, charged particles start to emit electro-magnetic radiation if they propagate faster than the (artificial) numerical phase velocity of electro-magnetic waves in the simulation. The mitigation of this effect is a topic of ongoing research and especially spectral and pseudo-spectral field solver methods yield promising properties [134, 143, 151, 153–155]. However, modifications in the field solver require a self-consistent coupling to further parts of the PIC cycle, especially to current deposition (see section 2.3.3). Therefore, recent numerical research focuses on methods to express this step with local methods (for efficient execution) and algorithms that do not violate charge conservation [156].



(a) Numerical phase velocity from Eq. (2.80).

(b) Numerical group velocity from Eq. (2.81).

Figure 2.11.: Numerical dispersion in the Yee scheme [130] for electro-magnetic waves propagating along an axis. Plotted for "best" choices of time step Δt that satisfies the inequality of the CFL condition in Eqs. (2.68) to 99.9% for given cell size Δx . Assumed are squared cells in 2D and cubic cells in 3D.

Literature usually advises to discretize the laser central wavelength λ_0 with at least 25 – 35 cells in order to avoid aliasing and strong numerical dispersion over many time steps [131, 132]. Likewise, the plasma frequency should ideally be resolved with $\omega_{p,e} \cdot \Delta t \leq 0.1$, averting the accumulation of large errors from the linear second-order accurate explicit methods in section 2.3.2 [157]. Resolving the Debye length λ_D in Eq. (2.21) is often challenging, e.g. for scenarios with sub- to mildly-relativistic laser (pre-)pulses which generate only few-keV average kinetic electron energies. Yet, when studying surface absorption effects and plasma instabilities in detail, proper Debye length resolution is necessary to avoid particle energy noise from numerical heating [50, 132, 158, 159]. For explicit solvers, an optimal correlation between spatial and temporal resolution is set close to the CFL condition in Eqs. (2.68).

Choosing the number of particle markers per cell to sample the distribution function f_s (see section 2.3.2) depends on the dimensionality of the simulation, the average kinetic energy and density of the modeled processes, among others. The underlying physical prob-

lem can be seen as a sampling approach for the ergodicity of a system, properly resolving trajectories of initially close particles (in phase space) as they evolve over time. This often requires to test the stability of a solution with repeated simulations. Lagrangian particle markers are sample trajectories in a simulation and enough such markers are chosen if the same observable is consistently obtained with further improvement of resolution. Lower-dimensional simulations require more particles per cell as less spatially close-by cells contribute to the momentum sampling of f_s . Generally, more particle markers per cell will reduce numerical aliasing (numerical heating and noise in f_s) and represent important observables such as the cutoff energy of an ion energy distribution with higher precision. Empirically, one to tens of markers per cell for 3D3V simulations, tens to hundreds of markers in 2D3V and hundreds to thousands of markers in 1D3V are typical choices for overdense plasma simulations [157, 160].




Figure 2.12.: Exemplified 3D3V simulation box providing sufficient space for ambipolar expansion and acceleration of ions over the interaction time with an ultrashort ultrahigh intensity laser such as DRACO at Helmholtz-Zentrum Dresden – Rossendorf (HZDR).

Summarizing the accuracy criteria above, a sample system is presented in Fig. 2.12. A solid target of moderate density, e.g. a plastic foil, shall be irradiated with a $\tau_0 = 30$ fs, $\lambda_0 = 800$ nm ultrahigh intensity laser pulse ($a_0 \gg 1$) with narrow focal spot of several micrometers. The laser pulse will ionize the target under idealized conditions to a free electron density of $n_e = 400 n_c$. Resolving plasma oscillations from Eq. (2.16) with $\Delta t \cdot \omega_{p,e} \leq 0.1$ imposes a time step $\Delta t = 2.1$ as and via Eq. (2.68) a 3D spatial cell resolution of $\Delta x = \Delta y = \Delta z = 1.1$ nm. The simulation box shall be large enough for the accelerating field and the laser focal spot, e.g. $15 \mu\text{m} \times 15 \mu\text{m} \times 30 \mu\text{m}$ which requires $13\,600 \times 13\,600 \times 27\,000$ cells. Initially, particles might be placed in 5 % of the simulation with 20 particles per cell (distributed over electron and ion species), resulting in $5 \cdot 10^{12}$ particle markers in total. If one estimates the required memory resources with 3×3 scalar values per cell for \mathbf{E} , \mathbf{B} , \mathbf{j} and at least 7 scalar attributes per particle marker for \mathbf{x} , \mathbf{p} , w (section 2.3.2), such a simulation will operate on $8 \cdot 10^{13}$ scalar values. With proper normalization, one is able to store each value in single precision (4 bytes each) which requires 290 TiByte of constantly available, fast memory. Furthermore, the acceleration process may be covered over a time span of at least 0.25 ps, equivalent to 120 000 iterations.

The required computational resources even for these moderate densities are immense and increase with the modeled free electron density (see section 2.1.2 for further examples).

One could try to build a single computer with adequate memory and compute such a system serially, which would take a long time to generate a result and does not yet account for processing output for scientific understanding of the dynamic evolution of the plasma. Or one might obtain a faster result if one instead assigns 32 GiByte chunks of the system to 9 300 computers working on the problem in parallel. Chapter 3 of this thesis addresses exactly this challenge. By advancing simulations via parallel (super)computing one can provide fast results for well-resolved laser-plasma scenarios.

2.3.6. Extended Fundamental Processes

For plasmas with short mean free path between collisions compared to the plasma wavelength (see section 2.1.2), particle-particle collisions are not inherently modeled when resolving the plasma frequency in a PIC simulation. In such scenarios, an additional operator must be implemented in codes which applies an adequate modification of momentum on the moving electrons. Such operators must account for the weighted "marker" nature of PIC particles and several binary collision Monte-Carlo approaches exist [54, 140–143, 161].

Similarly, atomic ionization dynamics can be added to the PIC cycle (section 2.3.3). Effective models for probabilistic tunneling ionization [162, 163], barrier-suppression ionization [164, 165] and collisional ionization [52, 54] are commonly used to estimate free electron generation [166]. Treatment of recombination channels is yet to be implemented in current PIC codes and would require additional steps for the conservation of Gauss's law (2.53), since Lagrangian particles suitable for recombination are not necessarily located one upon the other in PIC (see section 2.3.2). Nevertheless, recombinations contribute to the charge state distribution in and radiation from high density plasmas, especially under non-equilibrium conditions [167–171].

Numerical ionization models were added to PIConGPU (see chapter 3) during the time of this thesis in a Master's project [166]. For demonstration purposes, three state-of-the-art ionization schemes commonly used in PIC codes are shown in Fig. 2.13. There, the influence of ionization dynamics on free electron density is presented with a cylindrical target of $r_0 = 0.6 \mu\text{m}$ radius (similar to chapters 4 and 5) and nominal free electron density $n_e = 100 n_c$. Figure 2.13a shows a pre-ionized plasma with immediate plasma response on the surface to any laser intensity, Figure 2.13b displays the free electron density created by a threshold-based ionization model for barrier-suppression effects [164], and Figure 2.13c features a probabilistic tunneling ionization model [162, 163].

On first glance, the difference between pre-ionization and barrier-suppression ionization is not severe for mildly relativistic laser intensity and given target size. With regards to the target rear, realistic ionization and delayed rear expansion might shift the temporal onset of rear ion acceleration and strength of the accelerating TNSA field. The probabilistic tunneling ionization model seems to allow deeper penetration of the laser pulse at the target front and less laser fields can reach into the sides of the target. Nevertheless, both tunneling ionization and barrier suppression occur at the same time in reality, likely suppressing the front-side "ripple" effects seen in the electric field and electron density of Fig. 2.13c for high laser contrast. Numerical schemes for concurrent ionization processes might therefore be used in combination [172].

All extensions of the PIC cycle need additional computational resources on top of the requirements from section 2.3.5. It is therefore desirable to base these additions on a fast implementation of the basic PIC algorithm.

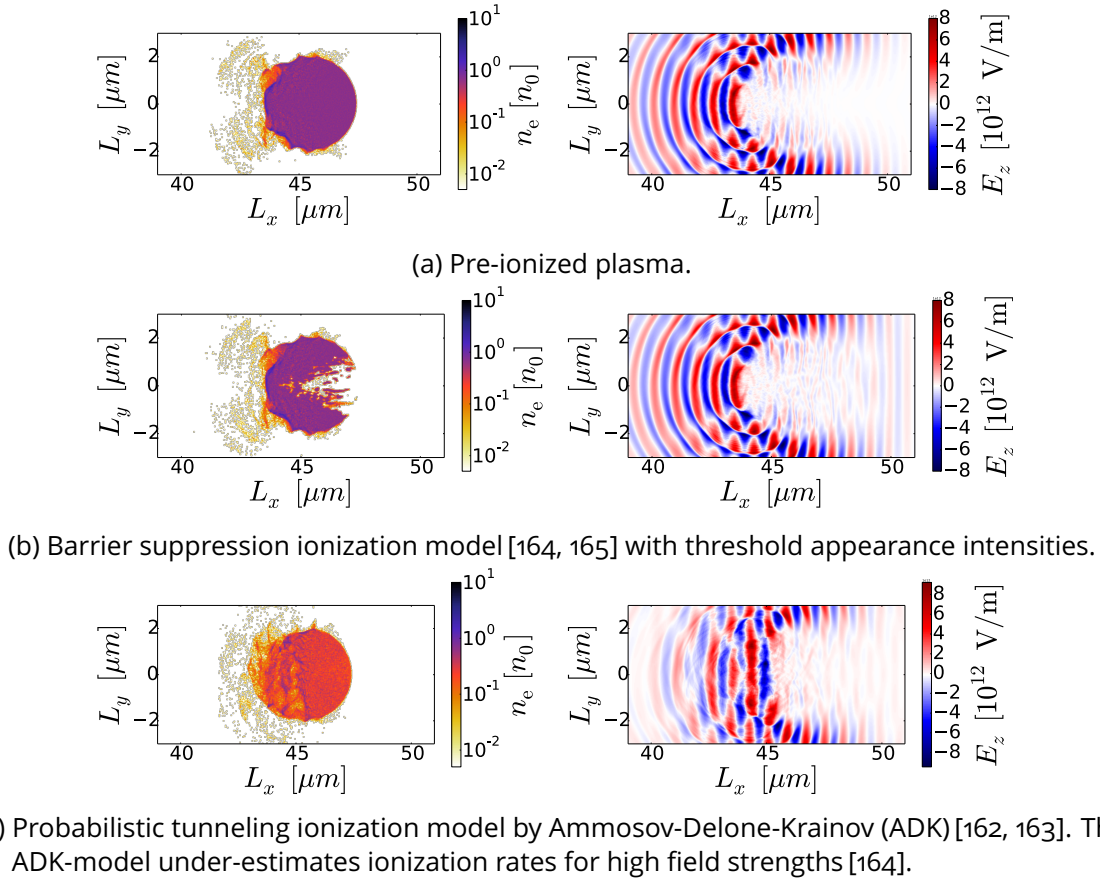


Figure 2.13.: Comparison of electron density slices with newly implemented ionization models in 2D3V for a cylindrical target with radius $r_0 = 0.6 \mu\text{m}$, nominal free electron density $n_e = 100 n_c$ and laser pulse amplitude $a_0 = 2$ and pulse length $\tau_0 = 12 \text{ fs}$. Snapshots shown for $t = 6.8 \text{ fs}$ after the laser pulse's peak intensity reaches the target.

2.3.7. Alternative Numerical Methods

For completeness, the previously described discretization of the distribution function on a regular mesh of fields and in discrete particle markers with fixed spatial shape is a robust and, due to local updates per iteration, scalable⁵ approach for the modeling of kinetic effects such as crossing trajectories in phase space, something that a fluid field description would be unable to assess. A non-exhaustive list of viable kinetic alternatives commonly used in laser-plasma physics includes adaptive refinement of the mesh [173], mesh-free algorithms calculating multi-pole moments between domains of the particle distribution [174–177], and Vlasov codes advancing phase space elements [178, 179]. Furthermore, the frame of reference can be changed from the lab frame to a boosted frame [64, 180] and spectral as well as

⁵Scalable means a problem can be distributed well over parallel computers, see section 3.3.3 for details.

stochastic PDE integrators can be used [151, 181]. Also, a recent geometric generalization of the PIC method has been proposed in Ref. [182]. Additionally, some plasma problems are suitable for treatment in quasi-static approximation [183]. Differences to electro-magnetic PIC in terms of numerical stability, preservation of conservation laws, computational costs scaling with the number of particles and resolution for modeled density and density gradients, and numerical noise are explained in the referenced literature and in comprehensive works on computational many-particle physics [184].

2.4. Basics of Scientific Computing

This section introduces selected terms often used in chapter 3 and basic concepts for efficient scientific computing. The focus is set on basic aspects relevant for application development and does not aim to be conclusive. Recent comprehensive introductions are provided in Refs. [185–187] and the glossary of this thesis contains further descriptions for important terms introduced in later sections.

High-performance computing (HPC) is often referred to as a sub-domain of computer science solving research questions with supercomputers. A supercomputer is a cluster or collection of interconnected computers (nodes) solving a problem collaboratively that exceeds the computational resources of an individual node. Furthermore, a supercomputer provides a fast filesystem for storing non-volatile data for later analysis.

Hardware

Each compute node of a cluster provides a so-called compute platform, which is the software and hardware environment in which a computer program is executed, e.g. an operating system such as Linux, executed on a processor such as a CPU (central processing unit). The latter is named hardware architecture and each hardware vendor, prominent examples in HPC are Intel,⁶ AMD,⁷ Nvidia,⁸ NEC,⁹ and ARM,¹⁰ defines a specific set of capabilities and rules to execute on it. Computing devices with a certain architecture usually provide several processing units called cores. One differentiates between multicore processors with currently up to ≈ 72 cores [188] and manycore processors with thousands of cores. A single operation that is executable by a core of a specific architecture is an instruction. While multicore cores can execute instructions independently, manycore cores are designed to operate as larger groups, performing the same operation simultaneously on multiple data, called single-instruction multiple-data (SIMD). SIMD operations execute an instruction for an array of values instead of a single scalar date, which is also known as vectorization. The latter also exists on multicore devices, yet with significantly shorter data array lengths. Between operations, volatile data is stored in random-access memory (RAM) which is fast storage connected to a computing device. Further intermediate levels of memory in a device

⁶Intel Corporation, Santa Clara (CA), USA

⁷Advanced Micro Devices, Inc., Sunnyvale (CA) & Austin (TX), USA

⁸Nvidia Corporation, Santa Clara (CA), USA

⁹NEC Corporation, Tokyo, Japan

¹⁰Arm Holdings plc, Cambridge, United Kingdom (focuses on hardware design)

are termed caches. Compared to RAM, a cache provides higher bandwidth and shorter time to access (latency) but significantly smaller storage size. The memory closest to a single compute core is a processor register providing a small amount of very fast storage for either data or instructions.

Floating point operations per second (Flop/s) count the multiplication-equivalent floating-point operations per time on a device as a metric for processing power. A typical CPU is a multicore device and currently provides a processing power of about 1 TFlop/s and connects between tens and hundreds of GBytes of volatile RAM (bandwidth ≈ 30 GByte/s) [189]. Manycore devices such as a graphics processing unit (GPU) provide approximately 8 TFlop/s but only up to 32 GBytes of RAM with high bandwidth (≈ 300 GByte/s) [189].

Software

Algorithms are designed for computing devices via a programming language. A programming language follows a formal specification for source code which is written by the programmer. Tools named compiler translate the source code into instructions for a specific architecture, which is then termed assembly code. During this step, also referred to as *at compile-time*, the compiler can optimize the given algorithm to a certain extent, e.g. reduce redundant operations, reorder independent functionality, etc.. The goal of such optimizations are a shorter run-time of the generated program due to improved utilization of architectural capabilities and often require manual annotations by the programmer in the source code. In order to efficiently run on a distinct platform, programmers often rely on a wider programming model, which consists of the source language, auxiliary functionality provided in an application programming interface (API) and workflows on how to express algorithms efficiently, e.g. in terms of data access patterns. Examples for the latter are CUDA [190], OpenMP [191], and OpenACC [186].

Algorithms

In order to solve a physical problem such as laser-plasma interaction with computers, one needs a governing set of equations for the underlying fundamental processes and a discrete numerical scheme to solve those, as presented in section 2.3. Translating a numerical method into an algorithm for a given hardware depends significantly on data access characteristics. Computing hardware is designed with the assumption that a loaded byte from a memory pool, for instance shared RAM of a device or any kind of intermediate cache, will be used for a certain number of instructions before the result is written back. The theoretical performance a given algorithm with fixed operations per loaded byte ratio can achieve must therefore be theorized, e.g. in a roofline model [192]. In order to utilize the compute potential of today's hardware, a double-precision¹¹ algorithm must perform at least 7 floating point operations on each loaded number, otherwise its performance will be *memory transfer bound* and under-utilizes the device's compute potential (in Flop/s) [193]. The opposite case, an algorithm that is limited solely by processor speed, is referred to as *compute bound*.

¹¹A double is a floating point number represented with 8 bytes.

3. Exascale-Era Simulations with PIConGPU

Simulations modeling the highly non-linear processes in laser-plasmas from first principles accompany laser-driven particle accelerator research since its inception [57, 64, 68, 194–196]. The development of the world’s fastest particle-in-cell code PIConGPU¹ and its advancement to an integrated scientific instrument for laser-plasma science is a central result of this thesis. Based on previous work published in the author’s diploma thesis [152], the long-term development of PIConGPU was continuously led and driven by the author of this thesis. Furthermore, workflows have been established that foster modularity in an open ecosystem. The latter is necessary to continuously add further fundamental processes (multi-physics) and establish a strong connection to computer science research in order to take advantage of the most recent developments in supercomputing and influence its evolution from requirements for predictive plasma simulations.

This chapter describes modern methods, theoretical design and implementations of the project. Furthermore, student projects conceived and advised during this thesis are presented as integrated parts in PIConGPU, which contain significant involvement by the author.

3.1. Scope and Applications

PIConGPU is designed as a general, 3D3V particle-in-cell (PIC) code (see section 2.3) [131, 132], applicable to general plasma and laser-plasma physics. While also providing 2D3V support, specific focus is set on 3D3V modeling for high-resolution, full-geometry studies of plasma accelerators. The driving applications for PIConGPU are simulations for laser-ion acceleration [30, 44, 46, 91, 95, 198–200], intrinsically requiring high density resolutions for short-scale plasma frequencies $\omega_{p,e}$, and novel concepts for laser-electron acceleration and radiation sources such as presented in Ref. [35, 37, 201, 202]. Additionally, PIConGPU aims

¹in terms of measured floating point operations per second (Flop/s) on supercomputers [197]

towards applicability in laboratory astrophysics and high-energy density physics [36, 152, 197, 203–207]. PIConGPU is also a widely studied application in computer science [189, 208–218].

PIConGPU's development is driven by the vision to strongly enhance the predictive capabilities for laser-plasma interaction. In order to improve the quality of simulations and to predict experimental setups reliably, an environment of short turn-around (see section 2.3.5), robust, tested, reproducible numerical studies must be established. Such an environment can best be achieved following an open science paradigm, based on open source and sustainable development, open data, interactive data analysis, transport of knowledge and open reviews. Yet the most distinctive feature of PIConGPU is speed, most visible in the required time to present a solution (time-to-solution).

Simulations for today's physics under extreme conditions with ultrahigh intensity, ultrashort laser pulses and XFELs are rendered complicated by the fact that little diagnostic capability exists for the actual target conditions in the laser focus (see section 2.2.5). Explicit PIC simulations are able to model the laser-target interaction around the peak intensity of the laser pulse on a time scale of few picoseconds of the overdense plasma evolution. In reality, 100s of picoseconds of expansion physics in the rising edge of a laser pulse and even pre-pulses can modify the assumed initial conditions of a simulation significantly (see chapter 4). Consequently, a single simulation cannot provide a comprehensive result, dozens or rather hundreds of simulations with variations of initial conditions, laser intensity, laser contrast, even deployed numerical schemes are necessary to capture the generally non-linear response of a plasma.

Before deploying PIConGPU in daily research, a 3D3V PIC simulation on a conventional CPU cluster could take several weeks from start to end and, as of today still requires computation times in the order of a week. With the rise of graphics processing unit (GPU) computing, devices originally designed for graphics processing with massive parallelism are deployed for general purpose computing (GPGPU) [219]. GPUs are also referred to as accelerator hardware, as they are hosted in a conventional CPU system as an additional computing device. With GPUs, simulation turn-around times can be reduced to half a day, with a 3D3V example for a novel target shown in chapter 4. In some cases, e.g. 2D3V simulations as presented for high-throughput studies in chapter 5, recent GPU hardware allows to simulate about half a dozen of simulations in parallel with each about 1.5 hours simulation time on a medium sized cluster. In fact, simulations with moderate resolution are already fast enough for workflows of interactive computing such as exploratory studies for example with interactive 3D visualization with appealing frame rates (see section 3.5).

Furthermore, with improved speed arises the possibility to study systems that were prohibitively expensive in the past. Higher resolutions on longer time scales are one example (see section 2.3.5), yet equally important are novel algorithms, allowing to add physics effects not covered in the basic PIC cycle, introduced in section 2.3. This chapter shows the fundamental building blocks of PIConGPU, its community approach and extensions, challenges that arise from high-throughput computing and how they are addressed today and in ongoing projects.

3.2. Community Code and Automation

Within the framework of the author’s diploma thesis [152], PIConGPU was open sourced under GPLv3+ license, one of few other open PIC codes in the world at the time [220, 221].² Uniquely suited for the newly released Titan system at Oak Ridge National Laboratory (ORNL), which was also entering a new territory as the world’s fastest and dominantly GPU-powered supercomputer, PIConGPU was able to perform first plasma simulations using the full size of the system and was nominated as a Gordon Bell prize finalist [197]. Over the years of this thesis, the scientific community around PIConGPU has grown into a world-spanning network presented in Fig. 3.1.

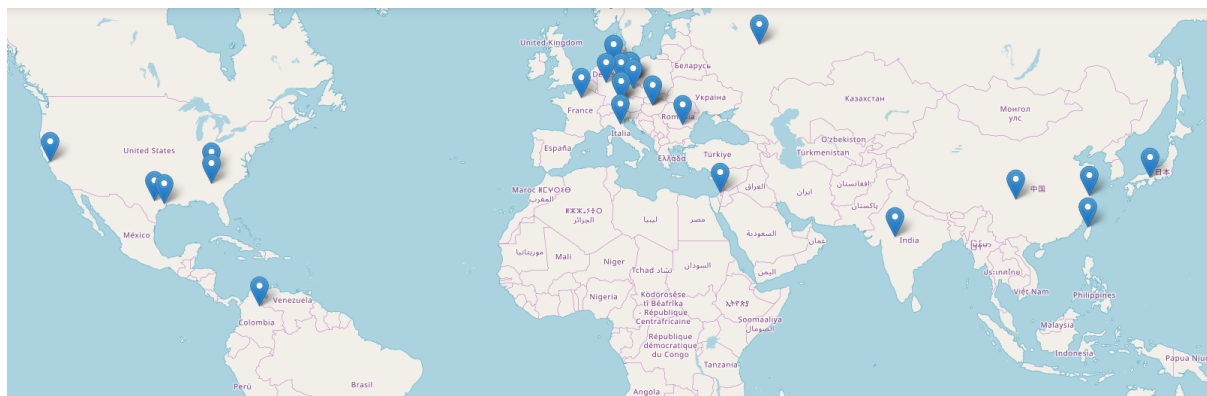


Figure 3.1.: Map of PIConGPU contributors, collaborators and users worldwide.
Map tiles ©OpenStreetMap contributors under CC BY-SA license.

The following principles and workflows are established with regards to sustainable, open development. Every release of PIConGPU is audit-proof archived with a digital object identifier (DOI) for citation purposes, crediting all contributors to the source code. Changes of the source code undergo an open review process, in which consent by the maintainers of the project should be achieved and all developers as well as external observers can consider to weigh in (GitHub pull request reviews).³ Decisions and meetings concerning strategic extensions and changes of the code are documented in the same manner (GitHub issues). Once a year, a major release is drafted from the added features of a development branch which receives on average four compatible bugfix releases until the next feature release. All changes are semi-automatically documented in a human readable changelog generated from the review process’ title, allowing to reconstruct even years later if a newly found bug might have affected previous simulations.

Code updates are automatically tested for common issues before they are accepted (continuous integration). An extensive manual is auto-generated from within comments inside the source code, allowing to keep documentation and code in sync while archiving preceding manuals [172]. Installing and running its modular software, also known as software stack, PIConGPU relies on automation and well-known tools. The build system is based on CMake [222], searching for dependent projects and preparing system-specific

²open source releases: EPOCH (2008), WARP (2013)

³see <https://github.com/ComputationalRadiationPhysics/picongpu/>

instructions for compilation. A recipe for the modern high-performance computing (HPC) package manager Spack [223] automatically installs all software, which otherwise can take up to a working day of a scientist when executed manually. Furthermore, ready-to-use software images are provided for Docker [224], including PIConGPU as one of the few highlight projects in the official Nvidia GPU Cloud catalog [225]. Since 2018, two official workshops teaching PIConGPU to new users were conducted with international attendees, addressing both experimentalists and theoreticians.

Since PIConGPU's initial open source release in 2013, other projects like ALaDyn [226], Chimera [227], FBPIC [151], pica [228], PICCANTE [229], PICTOR [230], PlasmaPy [231], PSC [232],⁴ SMILEI [233], Starfish [234], VPIC [235],⁵ and WarpX [173] followed, not least resulting from the favorable development of public funding guidelines that add open science requirements.

3.3. Zero-Overhead Abstractions

3.3.1. Algorithmic Agility

PIConGPU is built with the goal of algorithmic agility, allowing to combine all implemented functionality at the same time and on all supported platforms. This follows from the underlying concept, that a specific science case might simultaneously require modern compute resources for speed (GPUs), a specific variant of a field solver (e.g. to mitigate numerical effects such as numerical Cherenkov radiation [153]), ionization physics [52, 163, 166], radiation reaction [138], and effects of quantum electrodynamics (QED) [139, 236], etc.. A code that cannot deploy its full functionality in an integrated run is of little use for research campaigns. In order to avoid re-writing each algorithm for different use cases and combinations one needs to reuse functionality via abstraction techniques.

Fundamental numerical schemes of the PIC cycle define properties of all dependent algorithms. The field solver's order and staggering scheme influences particle interpolation and current deposition, as well as communication patterns. The order of a particle shape modifies the former algorithms as well. In PIConGPU, such algorithmic dependencies are abstracted at compile time, allowing to write code once and in a generic way determined by the underlying dependencies from discrete mathematical solvers. This saves lines of code to maintain, since redundancies are avoided, yet is harder to design as underlying symmetries, mathematical and computational dependencies need to be explored and described first.

PIConGPU's development embraces modern computer science methods in the programming language C++11 [237] that allow implementing orthogonal aspects without increasing execution time for each introduced indirection. The deployed zero-overhead abstraction technique is called template meta-programming [238, 239]. It is part of the syntax of standard C++, which is widely supported in the HPC landscape. Compiled code can be as fast as hand-written code as compilers resolve template meta-programming abstractions in their entirety [238]. Furthermore, such code can be optimized by the compiler for minimal usage of processor registers (especially limited on GPUs) and sub-functionality can use function


⁴open source release in 2016 (C variant)

⁵open source release in 2015

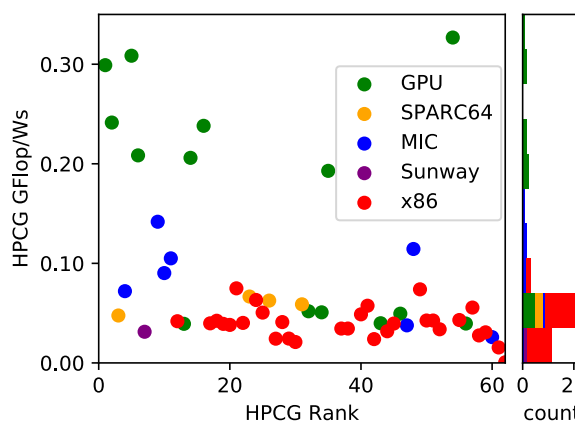
inlining extensively, replacing function calls with directly added instructions which avoids flushing of instruction pipelines [240]. Orthogonal aspects of a PIC routine for HPC include for example choice of a numerical integrator (section 2.3), targeted compute platform (section 2.4), data layout, data caching strategy, among others. The latter impose latency for data access and determine the efficiency of computing for a given processor architecture. For a simplified example about the policies of a PIC current deposition algorithm, see appendix A.1.1.

An additional benefit of template meta-programming is the concept of partial specialization. It allows to select specific combinations of the above mentioned orthogonal concepts of a solver of the PIC cycle and overwrite them with a manually tuned implementation. The concept of partial specialization also allows to write so-called type-traits [239], on which dependencies can be declared and exchanged at compile-time, e.g. if a particle species is eligible for momentum pointing filtering (e.g. projection onto a virtual detector with a pinhole aperture) or if it is a resting test particle without a momentum attribute.⁶ For an example on controlled compile-time code generation for simulation particles, see appendix A.1.2.

3.3.2. Supercomputing Systems



(a) HPC system performance in terms of the HPL benchmark over time. Green is the sum of all 500 systems in TOP500, yellow shows the fastest and blue the 500th system in the list. Lines are projections. Image from Ref. [241].



(b) Reported system performance per Watt, evaluated for performance numbers of the HPCG benchmark list (62 systems) in November 2018. Systems with accelerator hardware, especially GPUs, are more energy efficient [242].

Figure 3.2.: Exponential growth of theoretical supercomputing system performance as recorded in the TOP500 list [241, 243] and energy-efficiency of such systems with respect to the High Performance Conjugate Gradients (HPCG) benchmark [244].

A supercomputing system is commonly understood as a collection of computers (cluster), connected via a fast network connection and able to solve a large, non-trivial computing problem collaboratively. Traditionally, the largest supercomputers in high-performance

⁶Resting test particles are also called particle probes.

computing (HPC) are ranked by the TOP500 list [241], a public and voluntary measure of solving a large dense matrix problem in terms of floating point operations per second (Flop/s) with the High-Performance Linpack (HPL) benchmark [243]. Figure 3.2a presents various measures of this benchmark as recorded in the TOP500 list: the fastest supercomputer in the world as well as the sum of all systems and the 500th system. In all metrics, the computational capabilities of such systems grow exponentially with a tenfold performance increase approximately every four years. Continuing this trend, the first supercomputer that will exceed 1 ExaFlop/s of performance is expected in the upcoming year or two (2020/21).

In recent years, the High Performance Conjugate Gradients (HPCG) benchmark was established as an additional performance metric [244], which tests data access patterns and computations with higher sparsity than a dense matrix solve, e.g. such as stencils used in PIC codes. Its performance numbers for systems in terms of Flop/s are about two orders of magnitude below HPL results, as it measures support for algorithms with limited spatial and temporal data locality. Computing hardware with larger and faster memory throughput as well as low-latency network infrastructure profits in this measure, as well as real-world applications with complex data access patterns in many scientific domains.

In the past, exponential performance increase of computing hardware over time was achieved by miniaturization, which allowed overall faster (scalar) processors [245]. Nevertheless, as physical limits⁷ prevent this trend to continue to faster individual processors, parallelization over a large amount of computing units is responsible for a continued growth in supercomputer performance. This parallelization occurs on multiple levels, requiring to re-invent algorithms for fine-grained data decomposition and added access latencies for non-local contributions in each level of parallelization. For example, in order to utilize modern computing hardware efficiently, one must perform the same operation over multiple, closely stored data sets at once (vectorization). Several of such single-instruction multiple-data (SIMD) vector operations are performed on the same shared memory but with localized input and output that needs coordination and data packing. As motivated with the resolution requirements in section 2.3.5, parallelization of large problems spans over multiple computing devices and nodes without shared memory, which requires data communication.

As of today, GPU computing devices provide the most compact design and highest degree of parallelism. As shown in Fig. 3.2b, systems accelerated by GPUs are disproportionately more energy efficient than conventional architectures, even for the HPCG benchmark. Energy efficiency is relevant, since current supercomputers already consume 10-18 MW of electrical power [241], constituting a significant part of the running costs. Yet, in order to utilize their computing potential, algorithms and applications need to be redesigned in order to exploit both the latest parallel architecture as well as the multi-node performance of a parallel cluster. With exponential growth of overall performance and constantly changing architectures due to the limits set by overall power consumption, HPC target platforms are evolving continuously, which requires substantial commitment in order to keep pace.

⁷Specifically, the minimum possible transistor size and power consumption for data access.

3.3.3. Scalability

Computing a large physical system with proper resolution (section 2.3.5) requires more memory than a single node can offer. As introduced above, clusters provide parallel nodes that can store and compute a given problem collectively. Furthermore, within each of such a computer node, additional levels of parallelism exist such as multiple CPU and GPU devices and within each of the latter tens to thousands of threads which are executing instructions concurrently. Developing parallel algorithms that can efficiently utilize parallel compute resources and benefit (scale) proportionally with their size is key to translate computing power into fast simulations.

For each level of parallelism in today's supercomputers, the maximum speedup S an algorithm can achieve for a given problem size with theoretically infinite parallel computing power over serial execution is given by Amdahl's law [246]

$$S = \frac{t_{\text{serial}} + t_{\text{parallel}}}{t_{\text{serial}} + \frac{t_{\text{parallel}}}{N}}, \quad (3.1)$$

$$\lim_{N \rightarrow \infty} S = 1 + \frac{t_{\text{parallel}}}{t_{\text{serial}}}. \quad (3.2)$$

In this formula, t_{parallel} stands for the execution time of the part of the algorithm that benefits from parallelization over N contributing computing resources while a potential serial part t_{serial} does not benefit from parallelization and remains as minimum runtime at all times. For example, an algorithm that has a serial part of only 2% ($t_{\text{parallel}}/t_{\text{serial}} = 0.98/0.02$) cannot exceed a speedup more than a factor 50, independent of how much resources N are available for computing it. Assuming an algorithm can be expressed in small enough, independently computable sub-domains, an additional overhead exists as soon as communication needs to be performed between each computations, which one can add as well as a latency term to t_{serial} .

Figure 3.3 presents the distribution of a PIC simulation on a cluster in PIConGPU. This so-called domain-decomposition spans over multiple nodes, where in each node the domain is further decomposed over multiple compute devices (blue and red), and within the latter over grids of disjoint threads (orange) in which collaborating blocks of threads (light yellow) perform the actual computation. A device is defined as a group of compute units with shared memory. Between those, communication is needed over an intra- or interconnection to access each others data.⁸ Since updates per time step in an explicit, finite-difference time-domain (FDTD) PIC code are based on finite stencils (see section 2.3), only moving particles and field surface areas of local (per-device) domains need to be exchanged with all neighboring non-local devices once per iteration. This introduces communication latencies, which are hidden by overlapping computation and communication in time, see section 3.3.7.

For real-world applications, the speedup from Eq. (3.1) depends heavily on the chosen problem size to begin with. Generally, there are two reasons to parallelize an algorithm: solving a fixed problem faster or solving an $N \times$ larger problem, ideally at the same time as the initial problem. These are called strong and weak scaling respectively and both

⁸Examples for devices: a single GPU or a package of cores of a CPU.

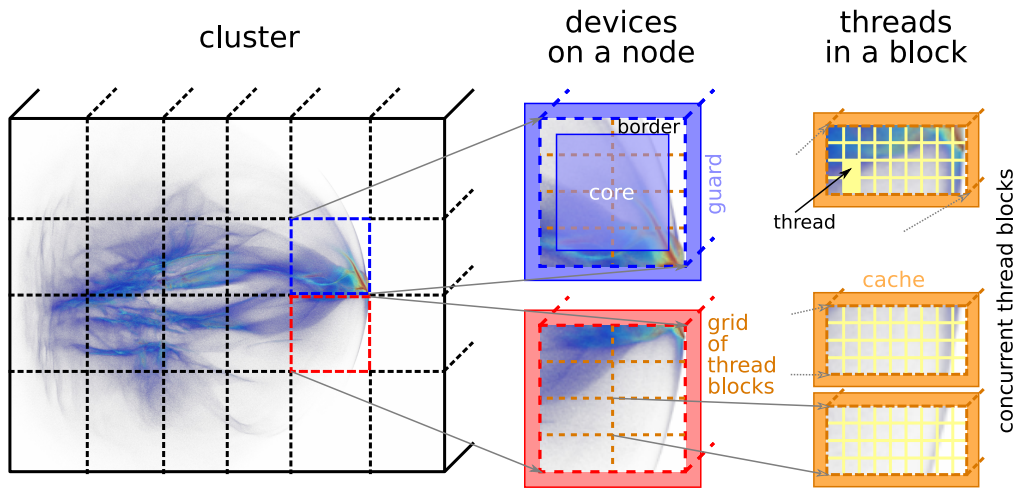



Figure 3.3.: Domain-decomposition in PIConGPU over the multiple levels of parallelism in a super-computing cluster. Compute updates for a PIC step (section 2.3) require next-neighbor information which are marked as guard volume for the device-level and cache for the thread-level, requiring communication. The picture on the left is a rendered 3D proton density distribution for a simulation presented in chapter 4.

can be expressed in Eq. (3.1), if one just changes the assumed reference times t . With massive in-device parallelism, as is provided by GPUs, one can reduce the time-to-solution (see section 3.1) compared to a CPU by approximately a factor 5-10 just by the increased density of computing units. Yet, this comparison is not entirely conclusive as it depends on the base size of the problem and starting with a too-small problem size of the number of contributing computing units in a device might even exhibit super-linear speedup if a device is underutilized initially, due to caching effects. Additionally, by computing on a large cluster with N devices, N times larger (e.g. better resolved) problems can be solved without waiting longer for a solution.

Considering all challenges above, Figure 3.4 displays the measured scaling of PIConGPU as measured on the Titan supercomputer (ORNL). Figure 3.4a shows that near-ideal weak-scaling is achieved to full machine size with 95 % efficiency for 18 252 GPUs versus a single device. Changes of efficiency between 16 and 32 nodes are visible, since communication patterns for small setups (less than $3 \times 3 \times 3$ nodes) are optimized to avoid unnecessary next-neighbor communication when possible. Further efficiency decrease is the result of load imbalances. Consequently, one can simulate laser-plasma problems with 18 252 times higher resolution or equally larger domain size than a single node could compute due to random-access memory (RAM) limitation, yet will only need 5.1 % more time. Strong scaling efficiency in Fig. 3.4b is naturally limited as domain decompositions are approached in which the communication overhead dominates over the local work assigned to a node. Nevertheless, one can reduce time-to-solution with sufficient computing devices by an order of magnitude if even faster turn-around time is needed, e.g. a $\approx 14\times$ shorter time-to-solution with $32\times$ more devices (44 % efficiency).

An application's time-to-solution for a given hardware depends further on the deployed algorithms and their memory transfer characteristics, as described in section 2.4. PIConGPU's



(a) Weak scaling: increased simulation size from 128^3 cells to $3456 \times 3328 \times 3328$ cells.

(b) Strong scaling: 32-fold increased number of devices until the onset of memory depletion.

Figure 3.4.: PIConGPU scaling on the formerly fastest supercomputer in the world, Titan (ORNL). The 3D3V setup simulated two particle species with each 16 particles per cell and triangular-shaped cloud (TSC) particle shape in a standard PIC cycle, modeling a Kelvin-Helmholtz instability [197, 218]. Efficiency is the relative overhead to ideal scaling.

performance in that regard is presented in section 3.3.5. Unfortunately, rapidly computing problems alone is not useful without the ability to gain scientific results from such computations, which are in turn derived from data. Parallel 3D PIC simulations, as show in section 2.3.5, generate large amounts of data due to the resolution of small-scale kinetic effects. Gathering, transferring and analyzing scientifically meaningful data at such high resolution and output frequency on the largest available supercomputers is a challenge on its own, presented in sections 3.4 and 3.5.1.

3.3.4. HPC Software Stack

Implementing the numerical models from section 2.3 in order to assemble a PIC simulation requires several layers of abstractions from the underlying compute platform. For example, physical quantities like three dimensional vector fields \mathbf{E} and \mathbf{B} can be naturally described with data container classes in a programming language. Such a container class hides the logic addressing three coordinates and updating information from neighboring domains in simple synchronization methods, presented as "guard" volume on a device in section 3.3.3.

Hence, PIConGPU is organized in modular software components displayed in Fig. 3.5. The level of abstraction increases from the bottom to the top. On the bottom level, the software stack depends on standardized, low-level programming models such as CUDA for Nvidia GPUs [190], OpenMP for CPUs [191], TBB for Intel multicore CPUs [248], HIP for AMD GPUs [249], or C++ threads [237]. The flexibility of choosing such a runtime programming model allows to explore upcoming compute hardware with simple configuration changes

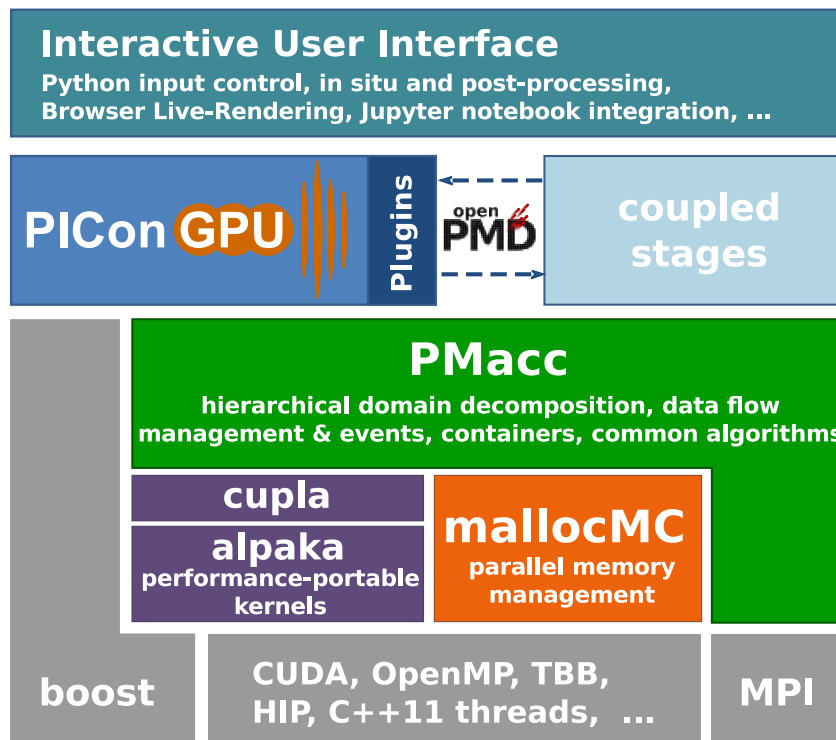


Figure 3.5.: Overview of the modular PIConGPU software stack. Colored components are projects developed in our group over the last years. Further software is in Refs. [190, 191, 247–250].

during compilation [189]. Some auxiliary functions for meta-programming are taken from the community library Boost [247]. Communication between multiple nodes on HPC systems is currently performed via the message-passing interface (MPI), which is widely supported in HPC.

Everything built on top of these fundamental building blocks are recent developments of our research group. GPU support is generalized to further computing architectures via the library Alpaka [189, 214, 215, 251], introduced in the following section 3.3.5. Memory allocations on GPUs are challenging, as particles such as electrons have a high mobility and move rapidly over the PIC domain (see Lagrangian frame in section 2.3). Parallel memory management within each device is therefore implemented in the mallocMC library (Memory Allocator for Many Core Architectures) [252]. mallocMC organizes a fast data pool for thousands of concurrently active threads in a device (so called memory heap) [253].

Shown in green in Fig. 3.5 is the library PMacc [209] (accelerated particles and meshes) which provides aforementioned memory containers for fields and particles. These containers recognize their distribution over multiple nodes and assist in communication of moving particles and updates of neighboring fields. For example, due to the spatial domain decomposition introduced in section 3.3.3, the outermost layer of cells of a device needs to be updated once per iteration in the order of the deployed field solver (section 2.3). Particles moving within a device and leaving from one device to another are organized in efficient local and exchange memory, described in more detail in section 3.3.6. In order to continue computations while such memory communications take place, PMacc provides latency-hiding functionality explained in section 3.3.7.

PIConGPU builds the actual application with domain-specific numerical solvers for laser-plasma physics, introduced in section 2.3. PIconGPU plugins, although displayed as a relatively small block in Fig. 3.5, are essential for productive simulations. One can imagine plugins as the virtual equivalent of an experimental detector observing the simulation and producing a stream of physical observables [203–205, 254]. Plugins can write raw data or generate filtered and derived output, couple to other codes or steer the simulation. Instead of copying all data as in typical data analysis workflows (post-processing), plugins operate in-memory with arbitrary PIC data of the simulation. In-memory operations have a minimal overhead in terms of processing time and data footprint, allowing to utilize device memory extensively for field and particle data.⁹ Additionally, plugins can also modify the simulation, e.g. in order to reduce the number of particles in a simulation through particle thinning or merging [236, 255]. Section 3.5.1 discusses PIconGPU plugins in detail.

A recent addition to PIconGPU is an interactive, scriptable, graphical user interface that allows to control parameter scans and analysis with a lightweight Python layer (section 3.5.2). The strategic vision for the last development is to integrate PIC simulations even deeper into joint experimental-theoretical campaigns with effortless data exchange and easy organization of parameter surveys. Also, it allows to automate scientific workflows with simulations, improving reproducibility and productivity. From setup over simulation run to analysis, archival and reuse of scientific results, this new frontend explores possibilities to provide a PIC simulation "as a service", accessible by laser-plasma scientists through a modern web interface, without requiring expert knowledge in PIC [256–258].

3.3.5. Abstract Accelerator Programming Model

Individual components of a scientific simulation can often be expressed as individual functions, that calculate part of a larger algorithm. In PIconGPU, due to its historical heritage from CUDA GPU programming, such functions are called kernels [190]. A kernel may update the electric field with a stencil function, push particles or prepare particle buffers for communication with a neighboring device.

Kernels traditionally operate on a single device with a shared memory pool (RAM) on which one to thousands of threads collaboratively compute an algorithm. With each new hardware architecture and emerging new programming model, kernels require porting to a new in-node parallelization strategy, which is an error prone, tedious process often involving hard to maintain code duplications. In the worst case, a programming model or the hardware it supports reaches its end of life and codes relying on it become useless. Consequently, applications should be written in a way such that algorithms are only implemented once but run efficiently on a wide range of significantly different architectures. This concept is called performance portability.

With the arrival of GPU hardware, computing paradigms for algorithms had to be reinvented for an optimum amount of parallelism within a single kernel, an evolution which was lead by our research group for the PIC algorithm [197, 208]. While common CPU hardware

⁹Device memory is very limited on manycore architectures and hence should mainly be used to store Eulerian fields and moving Lagrangian particles in order to solve large problem sizes (see section 2.3.5).

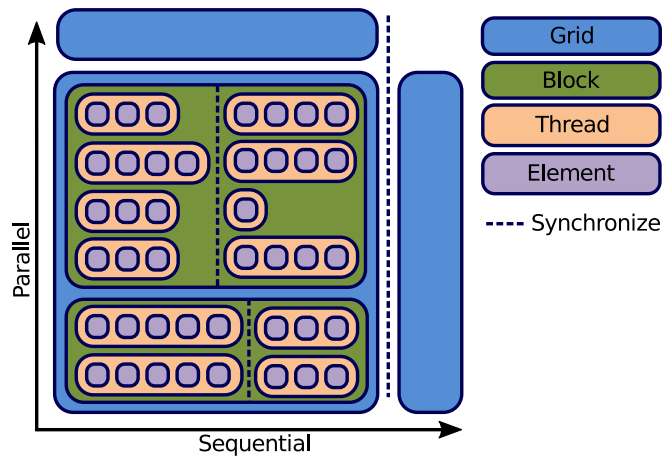


Figure 3.6.: Alpacka's abstract, hierarchical task parallelism model. Similarities to the CUDA programming model are intentional, as it suites well for an explicit description of multi-level device parallelism and eases porting. Image published in Refs. [215, 251].

provides collaborative work division between 1 to 288 threads¹⁰ with limited vectorization capabilities, GPUs work in collaborative multi-dimension blocks of tens of thousands of threads and mandate an vectorization-equivalent (warps) of 32 threads in Nvidia's CUDA model [190]. The idea behind the introduction of the Alpacka library was to release the strong dependence on CUDA with an abstract version of it as shown in Fig. 3.6, suitable to compile in standard C++ and to run equivalently well on legacy and potentially upcoming hardware.

With a high degree of parallelism from the design of PICongPU's algorithms for GPUs, most alternative hardware provides a lesser degree of in-node parallelism than GPUs (see section 2.4). Initial experiments with auto-serialization of CUDA code via compilers and just-in-time (JIT) runtimes such as PGI's CUDA-x86 product and GPUocelot [259] indicated that even serialization of already parallel algorithms to legacy hardware (e.g. CPUs) poses a non-trivial problem with respect to performance portability. Hence, Alpacka is designed as an explicit programming model with explicit control over serialization and parallelization.




Figure 3.7.: Task mapping with explicit control of serialization in Alpacka for various computing hardware. Image published in Refs. [215, 251].

As shown in Figure 3.7, the hierarchy "grid-block-thread-element" from Fig. 3.6 can be arbitrarily mapped for a specific computing hardware. The degree of parallelization in each of these levels is a compile-time configuration as well as the programming model that

¹⁰Intel Xeon Phi 7290 & 7295 with 72 cores and 4× hyperthreading.

implements it and can be fine-tuned to the individual kernel if needed. As demonstrated with the Intel Xeon Phi architecture [188] in the picture, some compute platforms or even kernels might be suitable to various approaches of mapped parallelism, which can easily be explored without rewriting kernels.

Alpaka deploys the same zero-overhead, single-source meta-programming strategies as PICongGPU and PMacc, introduced in section 3.3. This allows to aim for both platform and performance portability from the same source code, simply by changing options in provided C++ policies at compile-time. With that, a single-source abstraction that is written once is only configured for portability and runs efficiently on many platforms [189, 214, 215, 251]. Figure 3.8 provides evidence for this claim, with measured efficiency of PICongGPU on typical

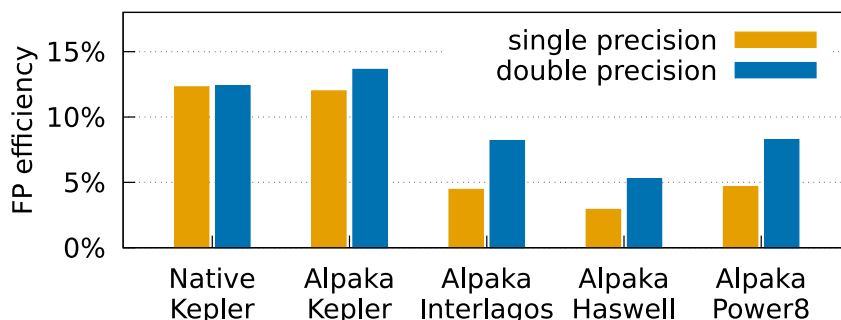


Figure 3.8.: Floating point (FP) efficiency with respect to theoretical achievable Flop/s on an architecture. Measurements show the central PICongGPU routines after three weeks of porting to Alpaka for single and double precision. The double precision efficiency on Nvidia Kepler hardware increased with the transition to Alpaka since an unoptimized atomic reduction was found. Image published in Ref. [214].

compute hardware before and after the introduction of Alpaka [214]. The overall efficiency is significant for the general purpose hardware it runs on, especially when considering that this graph normalizes to the theoretical (advertised) maximum performance of each hardware vendor. Such theoretical peak performance is determined by a dense matrix multiplication, which for a cluster defines its position in the TOP500 worldwide and mainly consists of fused multiply-add operations [241, 243].

However, the PIC method contains significantly more functionality than a dense matrix multiplication promoted by this metric, and it is memory transfer-bound for most kernels and hardware (see section 2.4). Recalling the numerical methods in section 2.3, stencil solvers for field updates are poor in number of operations per loaded value. Within the field update, simple local gradients are calculated, similar to steps of the conjugate gradient benchmark HPCG in section 3.3.3. Force weighting to particles is a gather operation, which can benefit from caching of the interpolated fields.¹¹ Particle pushing, i.e. the update of momenta, is exemplified in Eq. (2.74) and more compute intensive than the field solver per loaded byte. Higher-order push schemes mainly modify its intermediate calculation of $\gamma(\mathbf{p}^-)$, which involves calculating at least one, in some schemes even two, square roots [136, 137]. While increased computations per particle are in principle profitable to avoid memory

¹¹A special kind of read-only memory on GPUs is texture memory [190] with intrinsic interpolation functionality. It is currently not explicitly utilized in PICongGPU as only cloud-in-cell (CIC) particle shapes would benefit from it, due to hardware restrictions.

transfer bounds introduced in section 2.4, square root operations take significantly longer to calculate than simple operations such as additions and multiplications, which in turn reduce achievable floating point operations per second (Flop/s).

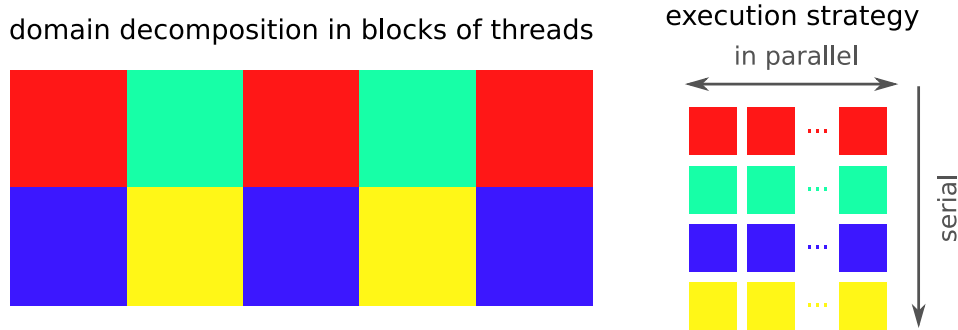


Figure 3.9.: Scheme of a multi-checkerboard domain decomposition. This example shows a four-color domain decomposition with each color representing a group of cells. During in-device parallelization of particle current scattering to the grid, only volumetric patches of the same color are active at the same time in parallel, avoiding memory collisions between non-collaborating thread blocks. Instead, block-wise reductions are performed and collectively written back. Depending on the order of the particle shape (see section 2.3) and the size of the group of cells in the decomposition (see supercell in section 3.3.6), four or up to nine colors are used by PIConGPU.

Current deposition is the most time consuming part, scattering local particle contributions back to the Eulerian mesh of the fields. For each particle, trajectories are evaluated by piecewise integration of the assignment function introduced in section 2.3 [145]. During the accumulation of current to a cell from many contributing particles, atomic operations (memory-safe operations between collaborative threads) are nearly unavoidable due to overlapping particle shapes. PIConGPU implements this step in local caches on supported hardware, e.g. in shared memory on Nvidia GPUs, which are then further added conjunction-free to the device main memory¹² via a multi-checkerboard algorithm. The latter, as shown in Fig. 3.9, is a multi-colored domain decomposition in which only same-colored, spatially non-overlapping domains are active for computation at the same time, in order to avoid memory conjunction. Depending on particle shape order, PIConGPU's implementation decomposes the concurrently active domain per device in 2×2 (2D3V with small shapes) or up to $3 \times 3 \times 3$ (3D3V with large shapes) partitions of same color. Each block of color consists of several cells forming a supercell, which is explained in detail in the next section.

Regarding the complexity of a PIC code, achieving more than 10% of the theoretical peak performance is a significant and unrivaled result, according to published literature [214]. Currently, additional improvements for even higher floating point efficiency focus on generalized low-level data types, assisting the compiler in further vectorization of generated code for CPUs [260, 261].

¹²CUDA term: global memory

3.3.6. Data Parallelism

Programming a PIC code as a physicist requires repeated access to typical quantities such as \mathbf{E} and \mathbf{B} fields and particles. In order to expressively design new algorithms for these, one tries to abstract implementation details via conveniently usable data structures that appear closer to their mathematical concept, e.g. a vector field or an individual (macro) particle.

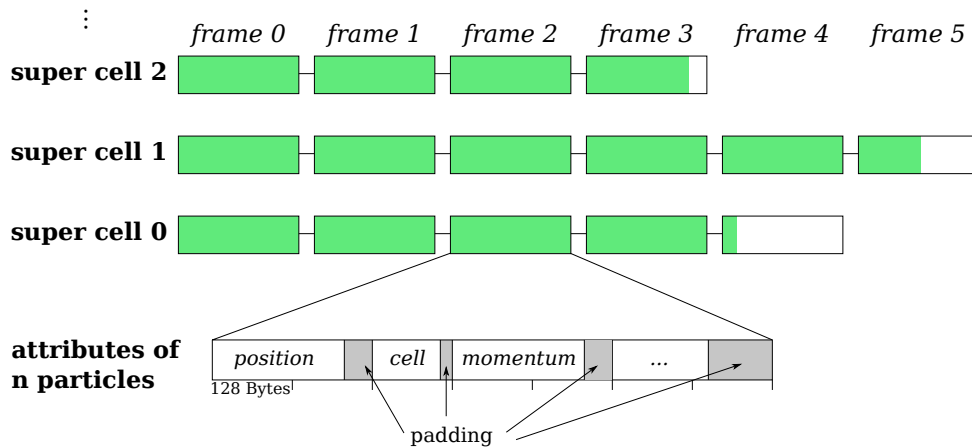


Figure 3.10.: PICConGPU's list of particle frames per supercell. Image published in Ref. [152].

Manycore-suitable data structures for particle and field data are provided by PMac containers with device-local views on the domain-decomposed data and next-neighbor ghost (guard or halo) surface buffers (section 3.3.3). Fields are represented as regular, multi-dimensional arrays which are sub-tiled in small blocks for execution, which we established as supercells [197]. A supercell has the additional correlation to an associated data structure of local particles as shown in Fig. 3.10, which we call particle frame lists. A particle frame is a fixed-size array of particles for a given spatial extent of a supercell. Particles move freely between cells and supercells and new frames are allocated and lock-free connected to a doubly linked list, as particles enter or leave a supercell. Within a supercell, particles are not sorted, which is equivalent to a coarse bucket sort by position. After the particle push and communication in each step, a particle frame list is once compacted from back to front, guaranteeing that all but the last frame of a list of frames of a supercell are entirely filled.

The size of particles in a frame as well as the size of a supercell are flexible configuration parameters in PICConGPU and can be chosen depending on the degree of parallelism or cache sizes of the specific hardware. On Nvidia GPUs, default 3D3V supercells are $8 \times 8 \times 4$ cells and a frame contains 256 particles, 2D3V supercells are usually 16×16 cells. Specific numbers might differ for higher-order solvers and single/double precision due to caches, as selected by the user. New frames are allocated via a parallel memory heap using an improved version of the so-called ScatterAlloc algorithm [253] on GPUs. This heap is implemented in the library mallocMC [252].

3.3.7. Task Parallelism

The PIC algorithm introduced in section 2.3 consists of individual functions per time step. Functions might advance a particle, communicate with neighboring devices, copy memory or derive a physical quantity such as average kinetic energy from the simulation. PIConGPU implements the PIC algorithm in fine granularity over 30-40 individual functions, depending on chosen extensions of the PIC cycle (see sections 2.3.3 and 2.3.6). Each function is implemented as a kernel (see section 3.3.5) via a C++ functor, a class with `operator()` overload [239]. Functors are generalized functions and fit well in the context of C++ meta programming (section 3.3.1), as they can be partially specialized and reused in a different context [239]. An instance of such a functor object can be passed and stored in data structures, which allows to create a queue of work for delayed execution.

In traditional programming, a function is called immediately and synchronously as described in the order of the source code by the programmer. This introduces the need to block the execution of the program until all requirements, such as communication updates from neighboring devices (section 3.3.3) are complete, leading to stalled computing resources. Due to domain decomposition inside computing devices (section 3.3.3), executing more than 200 individual kernels per iteration is common in PIConGPU. It is therefore advantageous to overlap independent computations on computing processors with communication tasks.¹³ Hence, instead of synchronous execution, work packages for computation, communication, memory copies, etc. are only added to a work queue for each device. One can delay the execution based on information that determines which dependencies must be run before a work package can be executed. This approach allows to detect independent algorithmic sections and start those as soon as computing resources become available instead of unnecessarily enforcing a strict order given by the programmer.

In its current state, PIConGPU developers can define regions of the source code as parallel transactions, which are sections explicitly marked as dependency-free or sequentially dependent. Per device, one thread is responsible for permanently checking the queue of work for newly fulfilled dependencies, e.g. when previous functors finished. In computer science, the described concept is also known as scheduling, and permanently polling implementation is referred to as busy-waiting. If a work package is observed to be eligible in terms of dependencies for execution, its functor is removed from the work queue and called.¹⁴ As demonstrated in section 3.3.3, PIConGPU scales well with increased communication partners when run on clusters due to the overlap of computations and communication with this scheduling system.

Manually describing code-sections of concurrently executable work packages is a verbose process which compromises on flexible adjustments of the code. Hence, currently developed work advances the described transaction concept to a graph-based description. The upper row in Fig. 3.11a shows the order of programming code given by the programmer. In the newly developed graph-based approach, each functor declares their requirement in terms

¹³Communication tasks can be send-receive operations via a network card, device-to-device copies of memory or filesystem I/O operations, among others.

¹⁴The controlling thread also starts kernels that can operate independently on various domain-decomposed sub-sections of a device in different CUDA-streams [190], in case of GPU hardware.

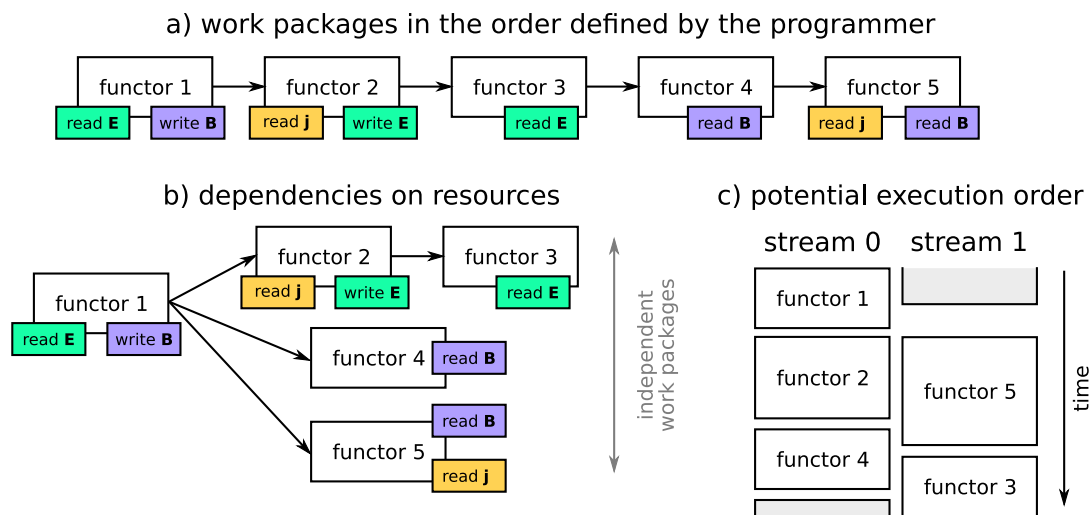


Figure 3.11.: Concept of a new on-device work scheduler following a declarative approach. Based on the required memory resources, work packages expressed as functors can be reordered into a dependency graph. A scheduling algorithm can read the dependency graph and prefer execution of some functors over others, e.g. to start long-running device computations or critical communication work first. The term "stream" [190] is just an example available on CUDA devices and demonstrates parallelly executable work.

of memory access for resources such as fields, particles, read or write access, sub-volume, etc. (colored boxes). From that, one can derive a dependency graph as in Figure 3.11b. For example, the execution order of two read operations on the magnetic field **B** (purple) can commute (functor 4 and 5) but a write operation and a read operation on **E** (green) must not be reordered (functor 2 and 3). Furthermore, an already requested resource requirement can be released at runtime, e.g. when processing a data stream of initially unknown size such as exchanged particles. We define this as "demotion" of a functor which allows to relax the graph of dependencies dynamically. In a last step, a scheduling algorithm decides the actual execution order as seen in Figure 3.11c, which allows to explore prioritization of specific work packages such as communication or known long-running computations like current deposition without changing any of the source code.

With the implementation of this new concept, we expect to achieve higher overlap of communication, I/O, in situ diagnostics and computation with significantly reduced manual description as well as an abstraction of on-device scheduling algorithms. The work for this modernization is ongoing as an undergraduate research project, for which a whitepaper was developed within the time of this thesis [262].

3.4. IO and Data Reduction

3.4.1. Theoretical Background

Results of this section are published in
*A. Huebl et al., ISC High Performance Workshops,
 Lecture Notes in Computer Science **10524** (2017)*




Figure 3.12.: Comparison of parallel I/O throughput for state-of-the-art I/O libraries on the Titan supercomputer (ORNL). Throughput is measured as given by Eq. (3.3) from an application's perspective, including data preparation and compression. Peak filesystem throughput was measured with 280 GByte/s (see Ref. [218] Fig. 2). Image published in Ref. [218]

During the first years of this thesis, the first large-scale 3D3V production simulations with PICongPU were performed. At the time, the world's largest open science supercomputer Titan at ORNL opened access to external, international researchers for scientific campaigns, one shown in chapter 4. The Titan supercomputer was furthermore the first GPU-powered machine occupying the leading position in the TOP500 list [241].

As demonstrated for the first time in Ref. [197], PICongPU is able to scale to the full size of such novel machines. When entering production phase with PICongPU, filesystem operations (I/O) for 3D3V simulations turned out to be an underestimated challenge, with its implications visualized in Fig. 3.12. Employing state-of-the parallel I/O techniques based on MPI-I/O, output with parallel HDF5 [263], a widely used hierarchical file format, does not scale to the requirements of such new class of supercomputing systems. The last data point of the green line in Fig. 3.12 is equivalent to 25 minutes of time to write a single checkpoint, time in which no computations for further advancement of the simulation can be executed because the required memory cannot be overwritten. Limitations of traditional I/O methods posed a severe challenge for the acquisition of simulation data in chapter 4.

Hence, in collaboration with a data science team at ORNL, the ADIOS I/O library was proposed as an alternative I/O solution, avoiding slow-down in global re-ordering and synchronization, contrary to aforementioned approaches [264]. Implementing this library, and addressing early-adopter issues such as support of heavily varying particle sizes from differing nodes,^{15,16} I/O throughput was improved by one order of magnitude as shown in the blue line of Fig. 3.12. Instead of measuring raw filesystem throughput, which peaked at 280 GByte/s [218], I/O operations are optimized and measured as real-world "apparent" I/O throughput from an application perspective defined as

$$T_{\text{eff}} = \frac{N \times S}{t_{\text{I/O}}}. \quad (3.3)$$

Here, N is the number of contributing parallel nodes with data, S the data size written by each node (on average) and $t_{\text{I/O}}$ the time the simulation has to wait in total because of the I/O operation, compared to just progressing with the computations.

For the studied system in chapter 4, overall checkpoint time, the time it takes to write the whole simulation state which might be used for restarts or data analysis, was reduced to about one minute. However, subsequent data dumps in order to study the dynamics of a system accumulated to about 1 PByte of raw data per 3D3V simulation. The size of such output turns out to be critical as well, since it equally needs to be post-processed in parallel and, critically, the overall shared storage for all users of the system amounted to merely 30 PByte.

Exploring data reduction techniques, an on-the-fly lossless compression algorithm was implemented in ADIOS using zlib [265]. PIC simulations with regular meshes are naturally well-compressible, as the resolution of the highest frequency ever appearing in the simulation leads to an over-resolution of all other areas of the simulation for most times, which was observed in a compressibility of PIC data to 25 % of its original size. Yet surprisingly, as shown by the orange line in Fig. 3.12, writing less data took more time than writing four times more data without compression. It is worth noting that this posed quite some confusion in the HPC community, as it was commonly thought that one can hardly spend too much time reducing data, since filesystems are orders of magnitude slower than any compression technique [266].

As proven by the measurement in Fig. 3.12, this expectation is not true and requires systematic analysis. In principle, one can consider every data analysis problem as a reduction problem with the following metrics: T_{R} the data throughput (e.g. in GByte/s) of a reduction operation and $0 < f_{\text{R}} \leq 1$ the ratio between output and input of the analysis.¹⁷ For the Titan system with $N = 18\,688$ compute nodes, the best case throughput of raw I/O is designed with $N \cdot T_{\text{R}} = 1$ TByte/s and $f_{\text{R}} = 1.0$ for disk data vs. simulation data size, or put in terms of I/O performance per parallel node, only 55 MByte/s.¹⁸

Consider the following typical PIC data analysis tasks as data reduction examples with the aforementioned metrics. Binning an energy histogram from 10^{11} particles [197] with

¹⁵Some nodes might contribute many particles to I/O, others zero.

¹⁶Even small memory leaks are problematic for repeated, large scale data processing tasks.

¹⁷Ideally, f_{R} is smaller than unity. This is not always the case for transformations of complex data streams.

¹⁸Writing constantly at full parallel speed would fill the whole, shared filesystem of 30 PByte within 9 hours.

each 4 floating point values for momentum \mathbf{p} and marker weight w to 1024 bins has a throughput close to the in-node memory copy speed $T_R \lesssim T_{\text{memcpy}}$ and a data reduction factor of $f_R = 2.6 \cdot 10^{-9}$. Even rendering a two hour long, ultra-HD video stream¹⁹ of a simulation versus writing only 10 times the simulation data from 8000 GPUs (5 GByte RAM) performs a data reduction of $f_R \approx 2.3 \cdot 10^{-4}$. And as GPUs were primarily designed for 3D graphics tasks, the rendering throughput still outperforms the 55 MByte/s per node of the filesystem.

As the last example might appear surreal at first, it will be revisited in the following sections. Following the reduction of raw data with compression algorithms, Figure 3.13 shows on-node benchmarks on the Titan system for real-world PIC data. Higher throughput and higher compression (toward the bottom right corner of the figure) is considered "better", as shown in the following. In all further symbols, data throughput \mathcal{T} is normalized to the in-node memory copy throughput T_{memcpy} . The dashed line, which is derived in a performance model, separates algorithms that would slow down overall I/O against algorithms that improve the time spent in I/O with compression.

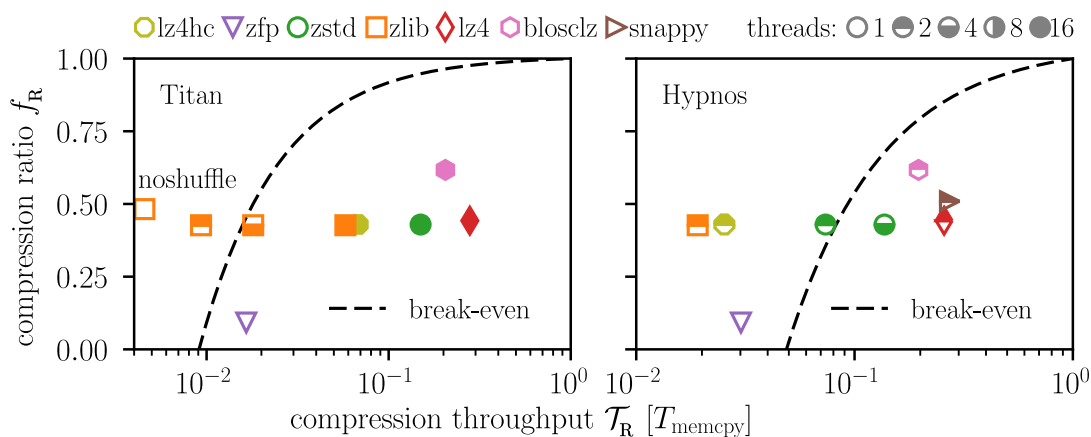


Figure 3.13.: Ex situ analysis on a single node of a cluster benchmarking reduction (compression) algorithms for the data reduction metrics \mathcal{T}_R and f_R [265, 267–269]. zfp is a lossy compression algorithm [270]. Hypnos is a 72 GPU cluster at Helmholtz-Zentrum Dresden – Rossendorf (HZDR) with similar hardware as Titan, yet less CPU cores per GPU and faster average file system per node ($\mathcal{T}_{\text{out}}^{\text{Hypnos}} = 22 \times \mathcal{T}_{\text{out}}^{\text{Titan}}$). Image published in Ref. [218].

In a data reduction workflow, the time to reduce data before starting the actual transport of the then-reduced data adds to the overall I/O time. As part of this thesis it was derived in Ref. [218] that any data reduction algorithm must fulfill at least the following inequality

$$\frac{\mathcal{T}_R \times (1 - f_R)}{1 - \mathcal{T}_R} > \mathcal{T}_{\text{out}} \quad (3.4)$$

in order to finish faster than directly writing the raw input data with data throughput \mathcal{T}_{out} . The latter might for example be given by the per-node filesystem throughput or any other data transport. The threshold in Eq. (3.4) for break-even is drawn as dashed line in all plots.

Furthermore, realistic data reduction algorithms cannot start immediately on given input

¹⁹This is currently the highest available quality setting on Netflix (25 MBit/s for UHD-videos).

data but require a preparatory phase, e.g. for device-to-host data transfers, data concatenation or any other form of data packing. The time t_{prep} for such data preparation is application-specific and can be again normalized to a data throughput by its local data size S and the in-node memory copy throughput T_{memcpy} .²⁰

The relative time saved (or lost) over direct output to \mathcal{T}_{out} due to a data reduction step is given by the unitless I/O performance ratio Γ

$$\Gamma = \frac{C_{\text{prep}} + f_R/\mathcal{T}_{\text{out}} + \mathcal{T}_R^{-1}}{C_{\text{prep}} + \mathcal{T}_{\text{out}}^{-1} + 1}, \quad (3.5)$$

$$C_{\text{prep}} = \frac{t_{\text{prep}}}{S} \times \mathcal{T}_{\text{memcpy}}. \quad (3.6)$$

For values $\Gamma < 1$, time spent in I/O is reduced proportionally to Γ while for $\Gamma > 1$ an additional overhead is introduced, which will further slow down the application waiting for I/O to finish.

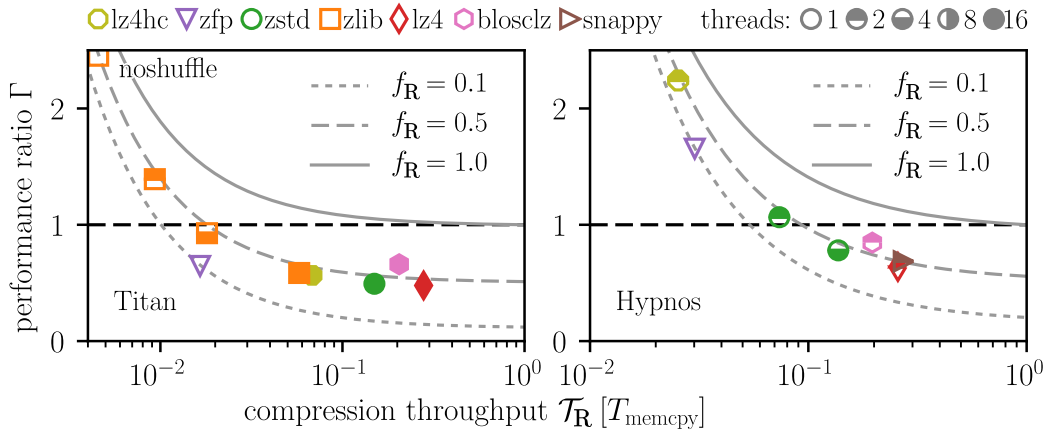


Figure 3.14.: Prediction of Eq. (3.5) for I/O performance gains or loss depending on ex situ measured performance of deployed reduction (compression) algorithms. In grey are iso-reduction lines (fixed f_R), e.g. for a given compression algorithm that improves throughput with threading. Image published in Ref. [218].

The ratio of host-side CPU performance compared to the accelerator-side of most clusters is in the order of 10 % or less. While it is usually not worth considering for load-balancing the main computations to this part, fully-GPU accelerated codes such as PICongPU have a few CPU cores left in the system (currently only one glscpu core is busy scheduling the GPU, see section 3.3.7). The prediction of Eq. (3.5) is therefore tested with a novel threaded implementation for compression, using the available host-CPU cores per GPU of the system. Figure 3.14 demonstrates that compression algorithms can be threaded on Titan, such that even zlib with relatively slow throughput (Fig. 3.13) reaches the break-even threshold for reduced I/O time.

On similar GPU clusters such as HZDR's Hypnos, yet with significantly less nodes, the per-node data throughput \mathcal{T}_{out} to its filesystem is usually much higher.²¹ Following the

²⁰In PICongPU, this time is typically one second.

²¹ $\mathcal{T}_{\text{out}}^{\text{Hypnos}} = 22 \times \mathcal{T}_{\text{out}}^{\text{Titan}} = 1.2 \text{ GiByte/s}$

presented model, this leads to higher requirements on the data reduction algorithm in terms of foremost throughput performance \mathcal{T}_R , until one reaches break-even over raw I/O. Zlib is insufficient for the Hypnos cluster, which only has four spare CPU cores per GPU for threaded compression. Faster algorithms such as zstd [268], blosclz [269], snappy [267] or lz4 [271] with pre-conditioners²² are required in order to provide enough data reduction throughput \mathcal{T}_R to be suitable for I/O performance gains.

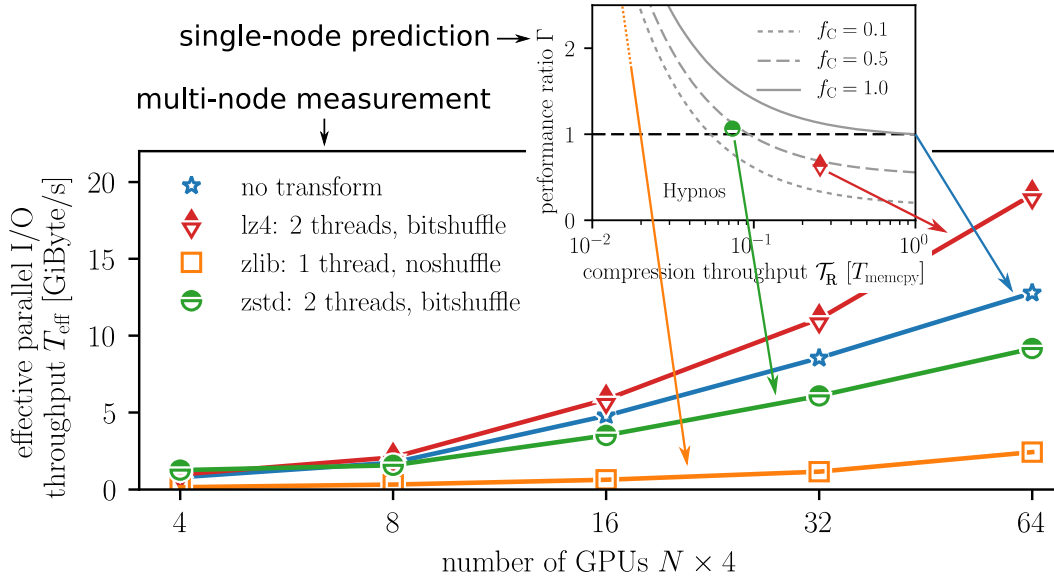


Figure 3.15.: Comparing single-node, ex situ measurements passed into the performance model from Eq. (3.5) (see inset) with parallel application I/O throughput T_{eff} on the Hypnos cluster at HZDR. Comparing to the blue line without data reduction, two-threaded zstd [268] and zlib [265] performance above the break-even threshold (dashed line) decreases I/O performance, while lz4 [271] increases the applications I/O throughput T_{eff} .

In order to summarize, one can predict the I/O performance of a data reduction algorithm at scale from ex situ measurements. For a reduction throughput \mathcal{T}_R and compression ratio f_R measured on a single node and knowledge of the overall filesystem performance \mathcal{T}_{out} , one can choose suitable algorithms for a given application before even running on multiple nodes of a system, e.g. shown for Hypnos in Fig. 3.15. Furthermore, this performance model was implemented and tested in PIconGPU and ADIOS, introducing threaded compression before I/O. With this implementation, the leftover 10 % of performance on the system’s CPUs can be traded for shorter time-to-solution of a real application at scale.

For the recently introduced successor of the Titan cluster, Summit, data reduction modeling becomes even more important. Compared to Titan, Summit has 10 times the compute performance (i.e. applications generate 10 times more data) yet the parallel filesystem throughput \mathcal{T}_{out} only grew by a factor 2.5. Therefore, any time span an application spent in I/O on Titan is now four times longer when running on the full scale of the new machine. Consequently, further data reduction techniques that allow bridging this continued challenge are presented in section 3.5, some were already mentioned briefly with the initial examples from the current section.

²²Compression pre-conditioners are simple, high-throughput operations that change the Byte-order [269].

3.4.2. openPMD

The generation of data in natural sciences is the fundamental basis for empirical evidence from observations in experiments, simulations, as well as the evaluation of theoretical models. Consequently, when generating scientific data one must not only think about how well one can measure an observable under given uncertainties, but also how the perception of the data can be preserved for fellow researchers. When handling extremely large data,²³ the focus should not be on writing (generating, or measuring) such data but rather on reading, in order to extract the features that provide answers for the scientific questions at hand.

openPMD is a self-describing meta-standard (sometimes called schema) for scientific data, addressing those challenges [272]. Conceived during this thesis, openPMD adds a description about the origin and purpose of data, its scales in terms of units and dimensions, and provides common information for proper author crediting. Contrary to many existing approaches that are either domain-specific [273] or visualization-focused [274, 275], openPMD defines a base standard that is agnostic of the scientific domain or purpose of the data, which enables versatile data exchange. Moreover, commonly needed supplementary information for a specific domain is in turn standardized in optional extensions alongside the base standard, e.g. for PIC simulations, elementary particles and isotopes, or particle accelerator simulations. openPMD can then be implemented in so-called "self-describing" data formats and frameworks that add portability in terms of data types [263, 264, 276, 277].

With the above properties and its independence of specific frameworks, openPMD is portable, exchangeable and scalable, yet at the same time not limited to actual files but designed with potential streaming (or sometimes called staging) solutions in mind [264, 278, 279]. Suitable for workflows with data pipelines that cross boundaries of previously disconnected communities, openPMD provides a basis for open-data and open-access workflows, an example is discussed in section 3.4.3. In its current version [280], supported data structures are intentionally limited to the very common cases of regular n-dimensional meshes and tabular data (data frames or "particles"). As openPMD is a versioned schema, constraints are relaxed in newer releases while existing older data can be automatically updated. Since such updates only concern meta-data changes, even PBytes of data can be updated, resulting in a future-proof standard.

Conversion between data sources and data sinks (i.e. writer and reader) are defined via a unit system that allows to define both a human and a machine-readable description without the need of specific naming schemes. Physical entities that describe a quantity like a vector or scalar field, or a particle attribute are called records in openPMD. Instead of enforcing a unit system such as SI and storing the string representation of a record's unit, e.g. "Volt/meter" for an electric field, openPMD stores its dimensionality. The dimensionality is currently expressed as a seven-dimensional array with powers of seven base quantities in the international system of quantities (ISQ) for dimensional analysis, while orientational dimensions might be added in the future [281]. In order to avoid enforcement of data transformations, e.g. between CGS and SI, and to enable scaling to an absolute unit system, defining a multiplicative conversion factor `unitSI` to a reference system (SI) is

²³see section 3.4.1 and chapter 4

recommended.

Just like the PIConGPU community (section 3.2), the openPMD community is organized around open proposals and open reviews to extend the standard. Independent adoptions are carried out by individual scientists, acting as multipliers for a much larger group of users benefiting from compatibility. Implementations as of today include PIC codes such as PIConGPU, WarpX [173], FBPIC [151], Osiris [282], UPIC-Emma [283], XFEL science start to end simulations (see SIMEX in section 3.4.3), astrophysics frameworks such as PlasmaPy [231] and yt-project [284] as well as upcoming implementations in accelerator simulations such as Bmad [285]. Versions of the standard as well as the software are archived with audit-proof DOIs for citation in data repositories and reproducible publications, e.g. by stating "data in supplementary materials of this publication follow the openPMD standard in version 1.1.0 [280]". In order to ease adoption, open source tools for verification of created files in the most common file formats, as well as updaters, converters, popular visualizers and reference library implementations are shared.²⁴

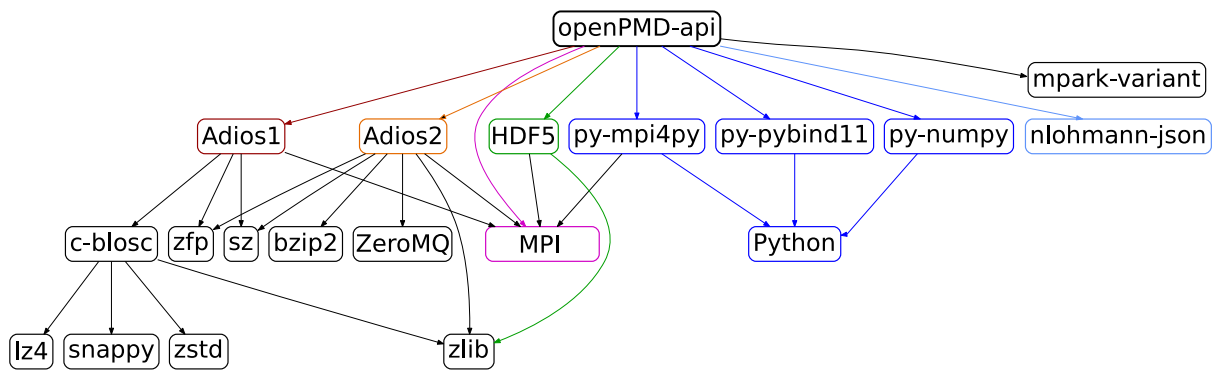


Figure 3.16.: Simplified I/O software stack, exemplified with openPMD-api. The graph is generated from its Spack [223] installation package for easy deployment on HPC systems.

As of PIConGPU 0.4.3 [286], data plugins account for as much as 50 % of the core code base, including the code for raw I/O to HDF5 and ADIOS1 which amount to 15 % (more details follow in section 3.5.1). In HPC, the I/O software stack is work-intensive and dependency-heavy, due to the complexities explained above. For illustration, the dependencies on external community libraries are shown in Fig. 3.16 for openPMD-api, our new library that eases adoption of scalable data workflows for the community. Designed again in a modular manner, its introduction allows fast adoption of new research libraries from computer science without the need of each application adopting and duplicating its I/O routines every few years. Furthermore, its bindings to scripting languages will enable productive development of parallel I/O stages for data intensive workflows with the benefits of openPMD annotated data [218, 264, 278, 279]. Such staging workflows provide a similar environment to the scientist as well-established post-processing routines. Yet, the main difference lies in the omission of raw-data generation which is often limited by the size and bandwidth of parallel filesystems (see sections 3.4.1 and 3.5).

²⁴See software repositories at github.com/openPMD.

3.4.3. SIMEX Platform

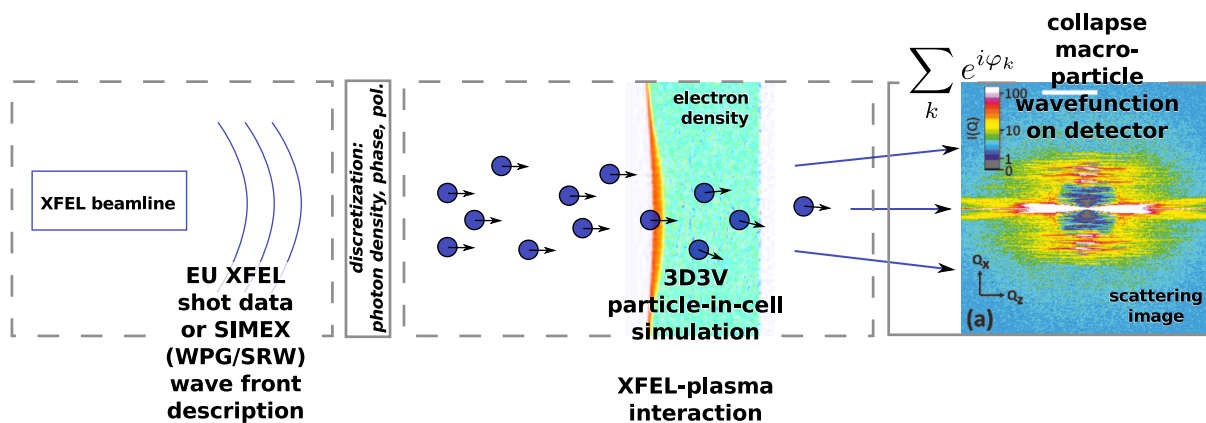


Figure 3.17.: Schematic depiction of multiple calculators constituting a laser-plasma interaction setup with XFEL probing in `simex_platform` [206, 207]. From left to right: photon beam generation and transport, translation into a kinetic picture suitable for PIC simulations, XFEL-plasma scattering, and in situ diagnostics as a detector. Image published in Ref. [207].

Automated and well-described scientific data exchange is the basis for high-throughput workflows that model upcoming scientific instruments. Pump-probe experiments studying high energy density (HED) physics, a regime of matter and radiation with energy densities exceeding approximately 100 GJ/m^3 , with unique ultrashort and extremely bright X-ray laser pulses start operating soon at the European XFEL [287, 288]. Its HED endstation allows probing of the complex interaction dynamics of ultra-intense laser pulses with solid targets at the femtosecond time scale [104, 107].

Within the European Horizon 2020 program EUCALL [289], coupled start-to-end simulations were developed. The latter approach modeling of experiments in an integrated manner, from the generation and transport of light sources over the interaction up to the detector response. Hence, a more realistic simulation description is possible than with an isolated simulation, as input and output do not need to be approximated with idealized conditions. The work package Simulations and Experiments (SIMEX) developed the start-to-end framework `simex_platform` which models an XFEL experiment, see components in Fig. 3.17. Its individual components include photon source and photon propagation, photon-matter interaction and resulting scattering, emission or absorption of X-rays to virtual detectors, and integrated data analysis [206, 207].

The PARAllel Tracer for Arbitrary X-ray Interaction and Scattering (ParaTAXIS) was developed based on the PIConGPU software stack, contributing to the design and implementation of the XFEL plasma interaction part in Fig. 3.17 [290]. ParaTAXIS discretizes the wavefront of an X-ray pulse into PIC photons for scattering on a plasma profile. The various software for XFEL beamline description, coupling to PIC and further stages of the interaction as well as the analysis, are described with openPMD (see section 3.4.2). With the new modeling capabilities, pump-probe experiments with simulated non-ideal XFEL pulses or even measured data from the European XFEL XPD database [291] are foreseeable. For a full integration, the ParaTAXIS software will need to be incorporated in PIConGPU in future work and non-elastic

scattering, absorption and emission must be implemented. The openPMD code-coupling strategy and PIconGPU's in situ data analysis approach (see sections 3.4 and 3.5) enable combining of software packages that differ highly in computational requirements, such as single-node wavefront propagation codes and 3D3V PIC simulations with up to 10^{11} particles [197]. This improves predictive capabilities of simulations beyond traditional boundaries of individual domains of laser-plasma and photon science [206, 207].

3.5. Interactive Simulations

3.5.1. In Situ Coupling

Having gained considerable experience running simulations on the world's largest supercomputers, seasoned workflows in simulation science reveal the need for a fundamental rethinking for high-resolution and high-throughput computing.²⁵ Traditional simulations would pre-process data or approximate experimental conditions with idealized, analytical conditions, compute the problem of interest on a HPC system, write out the data and then post-process microscopic data in order to assess its influence on macroscopic, experimentally accessible quantities. With ever-growing computational capabilities in the TOP500 [241] and slower-growing data-handling capabilities (see section 3.4), one cannot expect to continue this modus operandi with multi-PBytes or ExaBytes of data, already limited by storage capacity and throughput alone.

Instead, PIconGPU approaches this challenge by developing synthetic diagnostics for laser-plasma science, also known as setups of digital twins. Synthetic diagnostics approach a simulation like a virtual experiment with integrated diagnostics [131]. Hence, they bridge the gap between macroscopic observables measured in experiments and the fundamental plasma scales on the order of femtoseconds, only accessible in simulations. In situ implementation of synthetic diagnostics is a scalable solution, defining virtual detectors and observables as well as to-be-correlated quantities before even starting a simulation and calculating their detector output as the simulation is running. PIconGPU implements synthetic diagnostics as plugins, see Table 3.1. Plugins are coupled in-memory and in-code to the main application, allowing in situ calculations without data copies. Additionally, plugins can be configured to run several times with varying dynamic range, changed sub-selection of, e.g. filtered particles, and can be placed without the technical constraints present in real-world experiments.

For example, one might want to model two synthetic particle calorimeters, located at the same spatial location with temporal resolution, yet differentiating particles by origin inside a target in order to correlate those with an acceleration mechanism. Moreover, snapshots of the longitudinal phase space (1D1V) over time from a 3D3V simulation, filtered by pointing to the acceptance of a real-world pinhole aperture of a detector, can reveal the plasma dynamics leading to an observed energy spectra. Furthermore, one might derive field quantities from particles such as charge density ρ , current density \mathbf{j} , energy density, etc.

²⁵This also includes workflows performing hundreds or thousands of simulations for improved understanding of non-linear processes under variations of input and numerical methods.

and property-filtered variants of those, which are easier to describe in effective theoretical models than Lagrangian particles (see section 2.3). The research presented in chapter 4 still relied on traditional post-processing, which results for 3D3V simulations in challenging, PByte-scale data tasks. Later works presented in chapter 5 deploy such in situ derived observables heavily for systematic scans of a novel target for laser-ion acceleration, which produces significantly less virtual detector data. Plugins are data reduction methods, as discussed in section 3.4, and therefore provide a scalable approach for data analysis which is applicable for both reduced PIC geometries as well as full 3D3V modeling (see section 2.3.4).

Plugin	Description
ADIOS [218]	ADIOS files in openPMD schema
energy histogram	energy histograms
charge conservation	$\max(\nabla \cdot \mathbf{E}_{i,j,k} - \rho_{i,j,k}/\epsilon_0)$ (tests)
checkpoint	checkpoint-restarts for primary simulation data
count particles	total number of macro particles (debug)
count per supercell	macro particles per supercell (debug)
energy fields	electro-magnetic field energy
energy particles	kinetic particle energy
HDF5	HDF5 files in openPMD schema
ISAAC [216]	interactive 3D live visualization
intensity	$ \mathbf{E} $ and $\max(\mathbf{E})$ along propagation axis
particle calorimeter [236]	spatially resolved particle calorimeter at $\mathbf{r} \rightarrow \infty$
particle merger [236]	macro particle merging
phase space [152]	reduce particles to 2D phase space
PNG	slices as preview images
positions particles	single-particle tracking (tests)
radiation [203, 204, 292]	angularly-resolved far field emission spectra
resource log	memory utilization per device (debug)
slice field printer	slices of \mathbf{E} , \mathbf{B} , or \mathbf{j}
sum currents	$\sum_{i,j,k} \mathbf{j}_{i,j,k}$ (tests)

Table 3.1.: Developed and contributed plugins in PIConGPU for in situ data analysis. ADIOS and HDF5 are able to write raw data (\mathbf{E} , \mathbf{B} and particle species), as well as online-derived fields from particles and probes of fields from test particle locations.

Another prominent example of synthetic diagnostics is a plugin developed by Richard Pausch calculating angularly-resolved far field emission spectra of electrons [203, 204, 292]. Conventional workflows would store particle trajectories, e.g. of electrons injected in a laser wakefield accelerator [195, 293], and perform an analysis post-simulation to calculate emitted radiation from Liénard-Wiechert potentials. Storing many particle trajectories over a long time would be an extremely data intensive task. Remembering section 2.4, it is therefore favorable to convert this memory transfer bound algorithm into a compute bound problem. Indeed, it is possible and implemented in PIConGPU to calculate a discrete Fourier transform (DFT) by solving a double-integral for selected particles over time and change in momentum for the Liénard-Wiechert potentials. Projecting the resulting emitted radiation on a synthetic, angle- and frequency-resolved detector allows to investigate collective radiative effects with many particles in a plasma [152, 205].

Following the here described workflow, it might take significantly longer to plan a simulation than running an individual scenario, time which is in turn saved in post-processing. Nevertheless, if important metrics are not anticipated before the simulation, e.g. if one performs exploratory studies as described in the following section 3.5.2, one might need to repeat the calculation. An important question arises regarding reproducibility of such simulations, since raw data that can be archived is not created in the sense of traditional post-processing workflows. Coming back to the initial design goals of PIConGPU, we believe that sharing all necessary tools to reproduce a scientific study as open source, maintained, fully documented repositories can achieve such reproducibility through transparency. In order to re-run the simulation on other machines, this naturally requires to publish simulation input and analysis steps in a scientific publication in an open science manner (for example, see Refs. [11, 46, 214, 218, 254]).

3.5.2. Interactive Control and Visualization

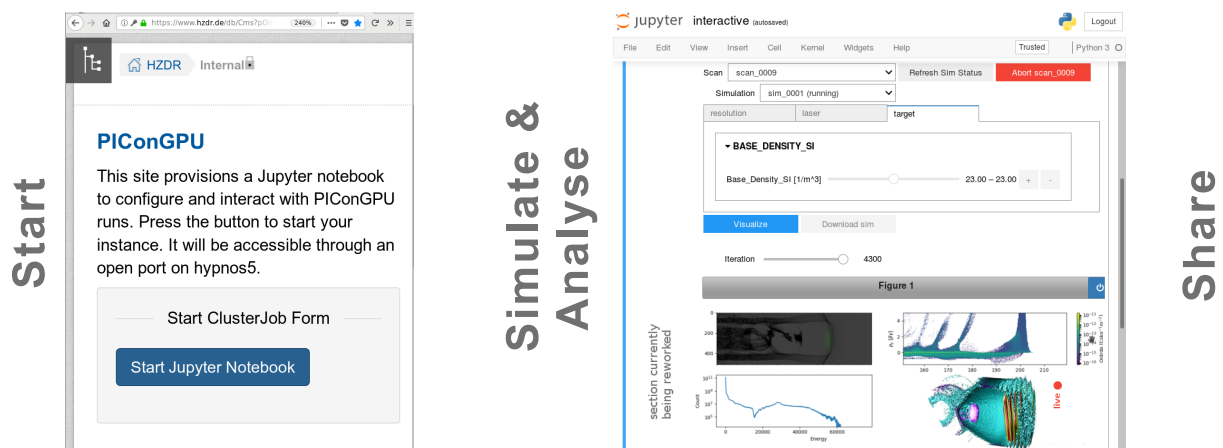


Figure 3.18.: Graphical user interface of PIConGPU as a scientific service. Based on open web technologies and lightweight abstraction of the PIConGPU input and in situ diagnostics, this service can be accessed by all collaborators.

Building on extensive in situ diagnostics introduced in the previous section, a collaborative software service for PIC simulations based on PIConGPU was developed [258]. Conceptually, this service focuses on a workspace studying one science case at a time. As its user group, integrated teams of theoreticians and experimentalists working together on a common research question are considered, e.g. HZDR teams working with the DRACO laser system. In such large, multi-disciplinary teams, work is often performed in an asynchronous manner regarding beam times, breaks of individual participants attending conferences, etc. which may disrupt collaborations over multi-month projects. One can overcome such hurdles with a common, easily accessible workspace, allowing access to simulations and their results for all collaborators at all times with automated data management.

Designing PIConGPU as an open science service, a light-weight user-frontend was added that exposes a subset of PIConGPU functionality tailored to the current research question. Using widely adopted technologies such as access via web browsers and Jupyter note-

books [256], this abstraction in Python is transparent in order to avoid lock-in to a specific solution and allows re-use of all components in a highly adaptable manner for power-users and developers without duplicating work. HPC resources as well as PIconGPU input specifics are virtualized and teams focus on the scientific input parameters (e.g. variations of laser and target settings) and synthetic diagnostics, configuring a digital twin of an experimental setup as seen in Fig. 3.18. All simulations are designed for ensemble studies, submitting combinations of parameters collectively for computation with automated storage and organization of synthetic results for correlation analysis. Users can quit and re-connect to this service at any time, add further simulations to existing ensembles²⁶ or follow the live progress of in situ analysis as a simulation is performed.

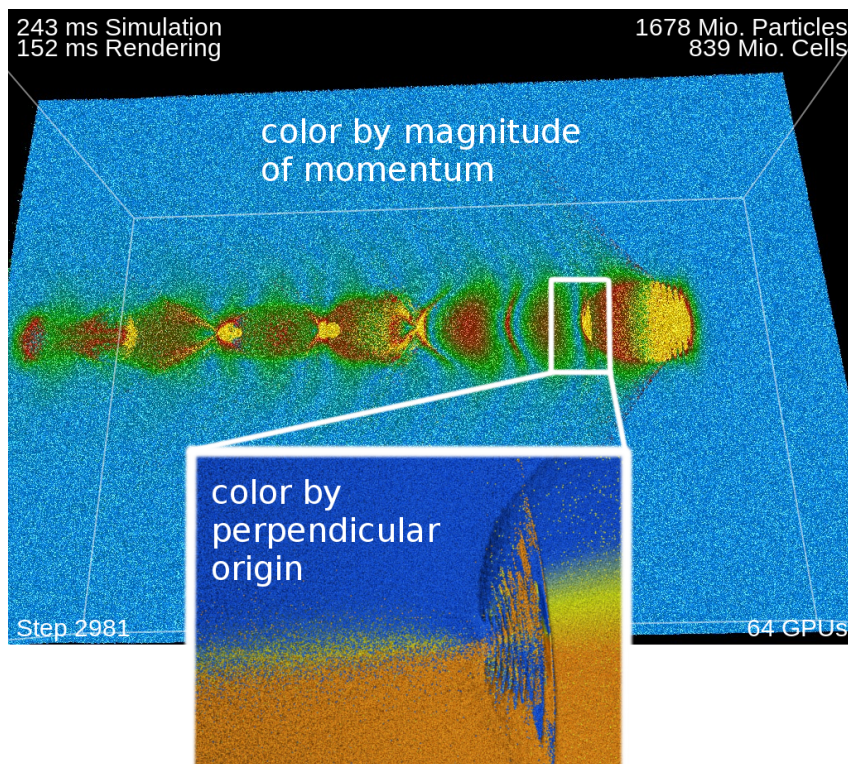


Figure 3.19.: In situ, in-memory rendering with the PIconGPU plugin ISAAC [216] of a laser-electron acceleration setup [195, 293]. Visualized are randomly sampled electron markers as solid spheres. The inset is taken at a later time in the simulation, interactively switching the representation. Colorbars are implemented prototypical in the user interface (only the rendered screen is shown here). Image credit: A. Matthes and F. Meyer.

Providing a domain scientist without prior HPC knowledge with such an interface, controlling potentially very powerful computational resources, is only possible since one consequently avoids storing raw data. Necessary data reduction techniques are described with analytical theory in section 3.4 and are implemented in section 3.5.1. Interactivity is feasible as all data is processed on-the-fly and only lightweight modifications for data presentation are performed in the user interface. Contemporary research envisions such approaches for data handling in upcoming Exascale machines,²⁷ see section 3.3.2, even for the extreme

²⁶For reproducibility, one cannot remove simulations.

²⁷Supercomputers with the ability to reach 1 ExaFlop/s in the HPL linear algebra benchmark [241].

case of storing millions of movie-representations of a single simulation [294]. However, such an approach would still reduce the data from simulations dramatically and allow exploratory user workflows, without recomputation of the whole system for follow-up questions. Nevertheless, such ideas will probably need support from automated prediction systems, e.g. based on machine learning, in order to record likely relevant observables and perspectives for a user [295].

In that sense, a tightly coupled, performance portable, in-memory operating, in situ 3D visualization called ISAAC (In Situ Animation of Accelerated Computations) was developed [216]. ISAAC renders fields and particles of interest from various perspectives over the cause of a simulation and allows to interact with a simulation for exploratory workflows or teaching. ISAAC's user interface is also implemented in standardized web technologies (websockets, HTML5, Javascript) [296]. Users can define a pipeline of parametrized transformations on exposed PIC data sources, allowing to derive non-trivial observables without recompilation. ISAAC supports both field and particle based visualizations, with an example presented in Fig. 3.19 for a laser wakefield acceleration (LWFA) [195, 293]. As shown therein, the goal is to shift the traditional qualitative impression of 3D visualizations to quantitative scientific observables. For example, by coloring particles injected into an LWFA by longitudinal origin in a gas jet, one can try to connect laser self-focusing and femtosecond injection dynamics with slices of the generated electron beam. By simply looking at the dominant color of such beam, one will be able to distinguish injected electrons from early self-focusing, continued injection over the propagation in the target and downramp injection. Likewise, by coloring according to perpendicular origin, one can investigate dominantly injected electron populations in a plasma channel [297].

On a final note, these modular additions to PIConGPU for interactive control and visualization also assist with a clear separation of concerns and responsibilities for users, application developer team and HPC centers. While users pay attention to the science case and a responsible use of resources, the core PIConGPU developers can concentrate on the numerical methods and synthetic diagnostics which are transparently exposed to the user interface. HPC centers can deploy PIConGPU as a consistent service, re-using existing scheduling, access, data analysis, monitoring and resource accounting infrastructure which is transparently abstracted in scripts and configuration files.

3.5.3. Cling CUDA

With the presented code base relying on single-source, zero-overhead C++ abstractions (see section 3.3.5), interactive data analysis tasks like the functor chains of ISAAC mentioned in the last sections are limited to pre-compiled parametrizations (see section 2.4). Yet for complex analysis tasks, replacing compiled code sections of an already running program as sketched in Fig. 3.20 would assist interactive simulation workflows and speed up the development of new algorithms. Compiling complex changes is time consuming and requires checkpointing and restarting the dynamic data of a parallel simulation. Alternative approaches could reload dynamic libraries used by a program or rely on scripting languages, but are either technically too unpractical for the average user, platform dependent, or do not offer the same degree of highly optimized kernels (see section 3.3.5) for optimal

performance.

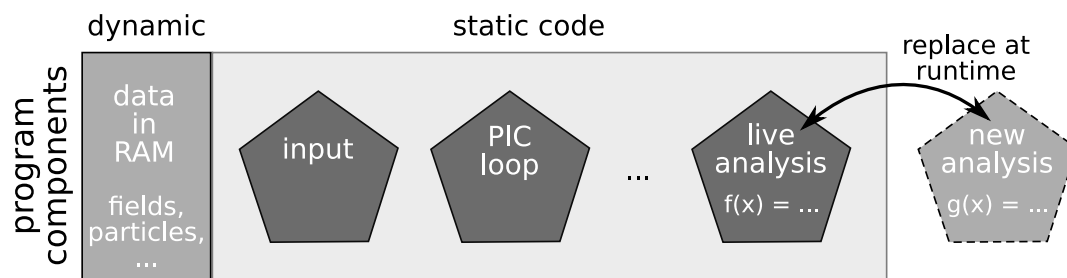


Figure 3.20.: A typical simulation consists of compiled code which is started once and holds dynamic data. Replacing individual code fragments at runtime enables fast changes to program logic and avoids time consuming program restarts.

Luckily, one of the largest data analysis frameworks in modern physics equally relies on modern C++ for interactive, high data-throughput analysis: the ROOT framework developed by CERN [298, 299]. At its core, the ROOT framework is based on a unique interpreter called Cling, which just-in-time (JIT) compiles C++ code while the program is already running [300]. As such, users can enter new code in an interactive interface, which is immediately compiled, appended to the already running program, called for execution and then returns the result. This workflow principle is also known as read-eval-print-loop (REPL), as new expressions are repeatedly read, evaluated and output is printed. Hence, Cling eliminates the need to switch constantly between C++ source code, tools for building, and the execution environment (see sections 2.4 and 3.2). The resulting binary code for common scenarios can be as efficient as re-compiling the whole application, adding or updating one entire kernel at a time, but does not interrupt the running program and its potentially large managed state.

The interactive exchange of highly-optimized kernels is the logical next step for the interactive workflows with high data rates, described in the last sections. In the Diploma thesis of Simeon Ehrig, an extension for CUDA [190] within the Cling framework was conceived, providing an interactive programming environment for the already presented manycore software from the last sections [301]. The central challenge for CUDA C++ code arises from its slight extension of the C++ standard within the CUDA runtime application programming interface (API), which is the API addressed within the source code of Alpaka (see sections 3.3.4 and 3.3.5). Hidden from the developer, the compiler splits CUDA device code from the single-source code and compiles it separately for the manycore device, which is then only called from the controlling host. Generated GPU instructions are named parallel thread execution (PTX) code.²⁸

Cling is based on the Clang/LLVM compiler project and has recently gained support for the CUDA language extensions in its frontend via the GPUCC project [302], which is nowadays an integrated part of Clang. As such, the Cling JIT pipeline was extended for a second compile path for CUDA code, shown in Fig. 3.21, which is then dynamically linked into the controlling interpreter instance. While the first prototype shown here compiles the whole source code (translation unit) for each change, a reworked approach with incremental PTX compilation is

²⁸PTX is actually an intermediate representation which is further translated by the GPU driver to architecture-specific Shader ASSEMBLER (SASS) code [190].

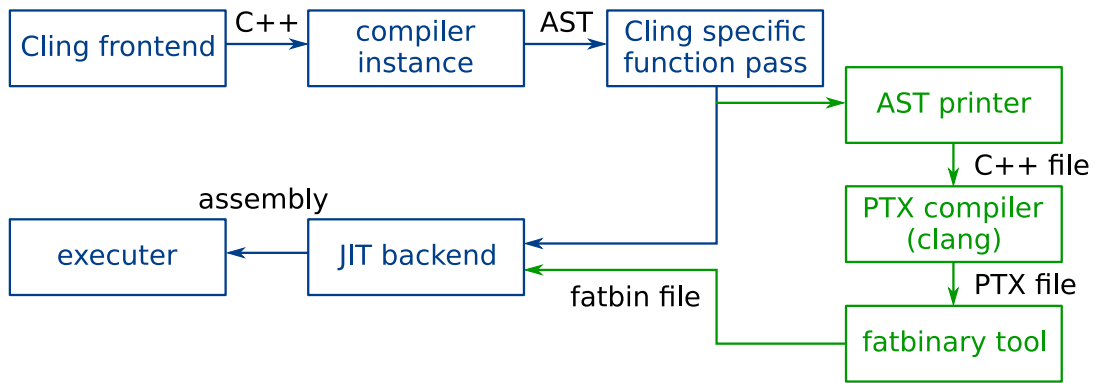


Figure 3.21.: Initial design of the Cling-CUDA compiler pipeline: device-specific functionality is extracted from the compiler’s abstract syntax tree, compiled separately to the intermediate representation understood by Nvidia GPUs, PTX, and then dynamically linked back into the interactive executer program [301].

currently being tested, allowing significantly faster compilation on large code bases over interactive, incremental changes.²⁹

In a final step, the newly extended Cling interpreter was run within interactive Jupyter notebooks, see Figure 3.22, providing seamless access to HPC resources as demonstrated in section 3.5 for the interactive PIconGPU interface [256, 303]. With the new ability to load and deploy in-house software libraries interactively, a wide range of potential applications is foreseen. In daily work with undergraduate students and their projects, our fundamental software libraries can be explored and learned with a focus on development and data analysis. Especially in short-time projects such as Bachelor and Master projects, usually interacting only with a small functional subset of a group’s research tools, time-consuming entry hurdles can be avoided with a seamless access to computing hardware and software from a unified interface. Time spent waiting for complex meta-programming libraries to

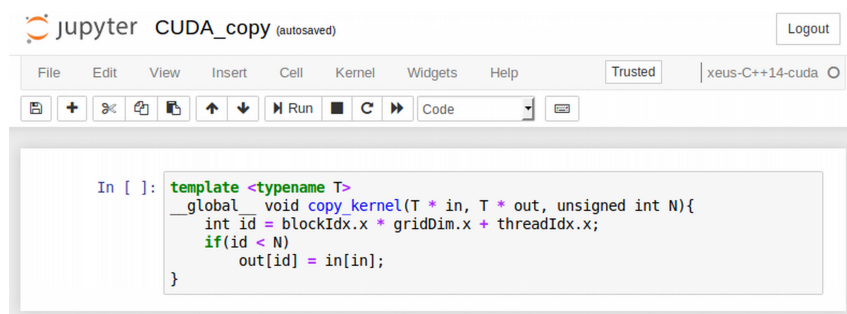


Figure 3.22.: Cling-CUDA integration in Jupyter notebook [256, 301, 303]. A simple browser interface can expose the computational power of manycore development to users, solely requiring C++/CUDA knowledge, and lets researchers focus on their science case.

compile can be reduced with provided incremental compilation. For libraries such as Alpaka (section 3.3.5), which can generate code for accelerator and legacy hardware with the change of a single policy, Cling’s CUDA support will assist further adoption in the community. Presented in-memory coupled data analysis routines can now be rethought for complex

²⁹See github.com/root-project/cling/pull/284 for details.

and live-adjustable data analysis, executable on host and accelerator devices, migrating and even fine-tuning performance on demand. Furthermore, the ROOT project at CERN is now finally able to explore GPU-accelerated algorithms for data analysis, without waving the interactivity of their workflows.

As this project builds on the latest trends in data science and accelerated computing, the implications for teaching, new user groups accessing HPC hardware, performance optimization, high-throughput computing, parallel simulations and experimental data analysis have to be explored in the upcoming years [304]. Regarding underlying technology, it is expected that open standards such as SYCL [305] and HIP [249] become available in the LLVM/Clang ecosystem soon, which might be easily adoptable in the spirit of this project and more sustainable than the proprietary Nvidia CUDA ecosystem [190]. Our software libraries such as Alpaka and everything we build on it are designed for such alternative accelerator backends, it is henceforth up to the community to benefit from its potential in interactive scientific workflows.

3.6. Conclusions and Perspectives

PIC simulations are essential scientific instruments in laser-plasma physics and require state-of-the-art supercomputers to solve large-scale 3D3V simulations and to perform parameter surveys (see sections 2.3.5 and 3.3.2). Naturally, the closer the simulation capabilities of the plasma physics community are to leading computer science capabilities, in terms of world-class computing hardware, the more productive computational surveys become. Scientific software needs long-term development to reach production quality (see section 3.2) and must be flexible enough in order to run on multiple generations of machines without constant reimplementations (see section 3.3.5). HPC evolution is dominantly led by manycore accelerators, which will likely provide the platform for the first Exascale machines in the following years.

PIConGPU provides a solution to the aforementioned challenges with an open, modular and computing architecture independent software ecosystem which runs on most currently available supercomputers. In combination with significant hardware utilization from a single-source code base and near-perfect parallel scalability, it provides a unique PIC research platform capable of running efficiently on the world's fastest supercomputers [11, 197, 218]. Interdisciplinary research on the interface of computer science and theoretical physics is important for its success, corroborated by well-received publications and reusable components developed for PIconGPU [189, 197, 214, 215].

Upcoming developments will have to maintain such a performance-portable algorithmic approach and progress further on data-intensive workflows arising at scale (see section 3.4). Essential data reduction challenges have been demonstrated for specific scenarios and with analytical descriptions within this thesis [36, 204, 218]. Future works must extend presented workflows with an even higher degree of automation and flexible programmability, as demonstrated for coupled simulation pipelines. With regards to device memory available in manycore devices per Flop/s and domain decomposition, dynamic load balancing should be explored over currently implemented static balancing [232, 233].

Interactive simulations foreseen in the context of data reduction will likely receive further attention with increasing computational capabilities. Although approaches have been demonstrated successfully within PIConGPU over the course of this work [36, 204, 216], HPC systems and computing workflows today are not yet ready for efficient interactive use cases. For example, schedulers are not designed for interactive (e.g. briefly pausing) simulation workloads that could potentially share resources between user controlled decisions and applications may need extensions for stateless APIs [306], allowing easier addressing and correlation of results while running on distributed systems. Unsupervised machine learning methods [307] may provide guidance in parametric scans through hundreds of simulations and could help suggesting regimes in ensembles suitable for in-detail analysis by a domain scientist. With the presented approaches providing PIC simulations as a service, potential community applications to large-scale computing facilities could become part of experimental planning phases and synergies may arise from integrated surveys.

4. Volumetric PW-Level Ion Acceleration on the Micrometer-Scale

*Results of this chapter are published in
P. Hilz, T.M. Ostermayr, A. Huebl et al.,
Nature Communications **9**, 423 (2018)*

Currently achievable ion energies from laser-plasma accelerators depend fundamentally on the efficiency of laser-electron coupling. As shown in section 2.2.2, the maximum kinetic ion energy K_{\max} is a function of both laser-accelerated electron density n_e and average kinetic energy $\langle T_e \rangle$. With flat target geometries, laser-accelerated electrons can spread into cold sections of the target and become unavailable for ion acceleration. Mass-limited targets on the other hand limit the available volume for electrons to escape and can lead to increased laser pulse absorption into high-energy electrons, as introduced in section 2.2.4. In particular, the average energy density of laser-accelerated electrons can be increased by either limiting the flux of accelerated electrons that leave the area of the laser focal spot, efficiently increasing n_e of energetic electrons, or re-accelerating recirculating electrons, thereby increasing $\langle T_e \rangle$. Moreover, near-wavelength target sizes can lead to the onset of transparency as the laser pulse penetrates an overcritical plasma up to the plasma skin depth even though its electro-magnetic fields are exponentially damped (see section 2.2.4)

An experimental campaign with mass-limited targets at Helmholtz-Zentrum Dresden – Rossendorf (HZDR) was able to demonstrate an energy increase of up to 87 % for maximum proton energies over reference foil targets [97]. Lithographically manufactured targets with $100 \times 100 \mu\text{m}^2$ surface area, mounted on a $5 \mu\text{m}$ stalk were irradiated at the DRACO laser system with $8 \cdot 10^{20} \text{ W/cm}^2$. However, mount structures holding the target in place act as conductors for energetic electron currents, leading to edge-fields and a disturbance of the accelerating sheath.

A target system for mass-limited isolated targets was developed by the group of Prof. Jörg

Schreiber at Ludwig-Maximilians-Universität München (LMU) [308]. Based on the well-known concept of a Paul-trap [309], it allows high precision-positioning of spherical, micrometer-scale targets in an electric quadrupole field [308]. Spherical, micrometer-scale geometry of trapped targets allows a laser pulse interaction with not only the surface but also the sides of the target and keeps accelerated electrons in the interaction region [80, 82, 83]. Experimentally, individually trapped spheres further provide the advantage of a controlled number of spatially confined atoms interacting with the laser pulse, contrary to previous experiments with gas clusters [5, 310, 311]. Compared to liquid droplet targets systems with a target diameter tens of micrometers, Paul-trap targets can be an order of magnitude smaller [312–314]. The Paul-trap positions new targets semi-automatically, with currently achievable repetition rate up to 1 shot every two minutes [308, 315]. Furthermore, the experimental design provides wide opening angles for target diagnostics [316]. Target materials need to be slightly charged, currently performed via an ion gun [308].

In order to explore the new target system for laser-ion acceleration, the LMU group performed campaigns at various ultrahigh-power laser systems [11, 123, 308]. The Petawatt campaign central to the work presented here was performed at a 150 Joule glass laser system, PHELIX at Gesellschaft für Schwerionenforschung (GSI) Darmstadt, with 500 fs pulse length, and 1054 nm central wavelength. The new target system and experimental choice of diagnostics required elaborate evaluation and analysis after the on-site measurements were conducted. Therefore, a collaboration with HZDR regarding simulations was started in parallel.

PIConGPU simulations for this experimental setup were performed in the framework of this thesis and substantially contributed to the interpretation of experimental signatures [11]. As simulation studies for highly non-linear systems tend to have a significant number of free parameters (e.g. in initial conditions) which might unconsciously lead to a "desired" (biased) result if experimental results are already known. It was therefore decided to perform the simulation study in parallel to the ongoing experimental campaigns and their evaluation, without knowledge of detailed results such as observed particle numbers, ion energy spectra or their cutoff energies.

In this "blind-study", trends and regimes were systematically investigated with parameter scans in 2D3V before performing 3D3V particle-in-cell (PIC) simulations of high computational costs. Although absolute ion energy values cannot be assessed with 2D3V simulations (see section 2.3.4), a broad overview of the system can be acquired and general trends such as dominant ion angular emission distribution and many repeated runs with varied initial conditions are feasible. As inputs, the expected experimental conditions on target and its material were analyzed over a large series of preparatory simulation runs (around 100 repeated and varied 2D3V simulations followed by the 3D3V results presented here). Since experimental quantities such as intensity or target density profile were not measurable in the laser focus or shortly before the main laser pulse arrives, uncertain parameters in assumed laser contrast and its influence on the target before main pulse interaction were varied.

As such, the collaboration surrounding the Paultrap target stimulated the implementation of a variety of new PIConGPU methods, presented in chapter 3. Static load balancing was added for the domain decomposition (section 3.3.3) to distribute particles evenly over

computing devices. Composite materials with multiple ion species were implemented and required new memory management for particles (section 3.3.6). Three dimensional high-resolution modeling generated large data output, in the order of 1 PByte per simulation, which required new I/O methods, theoretical performance modeling (section 3.4), and novel routines for parallel post-processing. Experiences and workload with the latter provided further incentives for advanced in situ analysis in later developments (section 3.5).

4.1. Assessment of Realistic Interaction Conditions





Figure 4.1.: Overview of the experimental setup at the PHELIX laser (GSI). Image credit: LMU.

Figure 4.1 shows an overview of the corresponding experimental target area at the PHELIX laser system. The central wavelength of the laser pulse is $\lambda_0 = 1054 \text{ nm}$ and the pulse length is $\tau_l^{\text{FWHM}} = 500 \text{ fs}$, providing 150J of pulse energy. Spherical targets were aligned slightly out of focus ($w_{l,0}^{\text{FWHM}} = 3.7 \pm 0.3 \mu\text{m}$) by $y_0 \approx 1.5 y_R$ Rayleigh lengths, minimizing pointing jitter by the laser for stable overlap with the target. At best focus, the laser peak intensity amounted to $2 \cdot 10^{21} \text{ W cm}^{-2}$, equivalent to $a_0 = 40$. In the campaign, the reference target was a polymethyl methacrylate (PMMA) ($\text{C}_5\text{H}_8\text{O}_2$) sphere with 1 μm diameter and $n_e = 400 n_c$ while some targets were hollow (with a wall diameter of 100 nm).

Realistic modeling of high-intensity laser-matter interaction poses several challenges for PIC simulations. As introduced in section 2.3.5, the dominating resolution constraint is given by the plasma frequency as the smallest wavelength of the system and determines



(a) Auto-correlator measurement of PHELIX' laser intensity contrast ($\tau_0^{\text{FWHM},I} = 500$ fs). The red line marks the typical ionization appearance intensity ($I \approx 7 \cdot 10^{13}$ W/cm²) [38]. (b) Regime covered by a typical PIC simulation. The green line shows the temporal approximation of the peak intensity via a Gaussian pulse.

Figure 4.2.: Experimental laser intensity contrast with peak-intensity of $I_0 = 7 \cdot 10^{20}$ W/cm² in comparison to PIC simulation capabilities. The light blue marked time span in both figures denotes the time scale accessible with PIC simulations (see section 2.3.5).

computational costs. Hence, PIC simulations can only cover a limited temporal window of the overall interaction, at maximum a few picoseconds around the peak intensity of ultrashort laser pulses. Figure 4.2a displays the measured intensity contrast curve of the PHELIX laser. Before the rising edge of the laser pulse starts at $t = -100$ ps, the laser intensity on the amplified spontaneous emission (ASE) level (plateau) and pre-pulses (e.g. at -210 ps) are low enough to avoid preliminary target ionization [38, 317]. Nevertheless, during the last 100 ps before the main pulse the absorption of laser pulse energy initiates hydrodynamic pre-expansion and ionization of the target [96]. The temporal window accessible with a PIC simulation is highlighted (light blue) in Figs. 4.2a and 4.2b, in which over 99.8 % of the laser pulse energy interacts with the target.

Direct measurements of the target condition before the main pulse interaction require sophisticated diagnostics resolving the picosecond temporal and sub-micrometer spatial scale, which were not available at the time. Therefore, the target conditions after pre-expansion had to be estimated. The wide angle spectrometer shown in Fig. 4.1 was positioned in laser propagation direction with a screen behind it, which was imaged to a camera recording the transmitted light of the laser pulse. For a target of overcritical or transparent material, one can estimate the diffraction of light via interaction-free wavefront propagation similar to inline-holography [318]. It was found that the recorded fringe patterns of the integrated transmitted laser light matched best to a target size corresponding to an expansion to near-critical density, assuming a Gaussian density profile (see methods of Ref. [11] for details). Such a substantial decrease in target core density, from initially $400 n_c$ to near-critical, due to hydrodynamic pre-expansion, may initially seem counterintuitive. However, it is important to note that the isotropic three-dimensional expansion of a mass limited target lowers the target core density substantially faster than in the case of a flat foil. The density of which decreases linearly with expanding thickness, assuming a flat density profile persists. In

a cylinder target, isotropic expansion in two dimensions lowers the density quadratically with thickness and in a spherical expansion, density shifts even cubically under isotropic expansion. An initially solid sphere with $n_{e,0} = 400 n_c$ and $r_0 = 0.6 \mu\text{m}$ corresponds to the following particle numbers under expansion to radius r


$$N_e^{\text{solid}} = n_{e,0} \cdot \frac{4}{3} \pi r^3 \quad \text{solid sphere,} \quad (4.1)$$

$$N_e^{\text{Gaussian}} = n_{e,0} \cdot \sqrt{8\pi^3} \sigma^3 \quad \text{Gaussian profile.} \quad (4.2)$$

Assuming conservation of particle numbers

$$N_e^{\text{solid}}(r_0) = N_e^{\text{solid}}(r) = N_e^{\text{Gaussian}}(\sigma), \quad (4.3)$$

an expansion to a core density of $0.8 n_c$ corresponds to a solid sphere of $r = 4.76 \mu\text{m}$ or a sphere with Gaussian profile of $\sigma = 3.06 \mu\text{m}$, which is taken for PIC simulations. Both profiles are still covered by the large defocused laser spot in the experiment.



(a) Proton beam divergence from simulations in plane of polarization.

(b) Comparison of predicted proton energy from simulations vs. experiment.

Figure 4.3.: Beam divergence and final energy spectra of the simulated proton beam for 2D3V and 3D3V simulations with initial density of $0.8 n_c$. Details of the 3D simulation are discussed in the following section.

Furthermore, the analysis of particle detectors indicated that a significant high forward charge of protons per solid angle had been observed, while a transversely deployed detector was unable to detect ion signals above its threshold of a few MeV. One can link this observation to 2D3V PIC simulations, comparing the particle count in forward direction to lateral emission as seen in the angular emission in Fig. 4.3a (grey). Hence, the series of 2D3V PIC simulations was extended for scans at 0.016 to $8 n_c$ which also indicated that forward-acceleration of ions slightly increases towards near-critical density ($n_e = 0.8 n_c$ shown in Fig. 4.3a), but is still influenced by large contributions from isotropic emission. In 3D PIC geometry of the next section (red line), forward directionality is more pronounced, as in experiments. Proton energy spectra in 2D3V modeling, as shown in Fig. 4.3b, overestimate the maximum proton energy significantly for currently achievable ion energies in

experiments [7–10]. As explained in section 2.3.4, the reduced dimensionality introduces the geometric assumption of periodicity in third spatial dimension which overestimates electro-static fields between charge carriers.

In order to summarize, the performed 2D3V simulation scans over target density confirmed a slight increase in laser-forward directed ion emission for a substantially expanded target with $n_e \simeq n_c$, consistent with the in-line holography results obtained from the experiment. While 2D3V simulations for the given target geometry cannot provide predictive results for comparison with the experiment, they reliably confirmed the relevant target density range to be investigated in the following with 3D3V, of which only few well selected runs were performed due to limited computation resources. The results of this study are presented in the following section.

4.2. Petabyte-Scale 3D3V Modeling

Due to the spherical geometry of the target, 3D3V simulations are unavoidable for reliable modeling of the plasma dynamics. Laser pulse conditions were estimated as an ideal, diffraction-limited Gaussian beam with off-focus intensity on target according to a defocused spot of $w_{l,0}^{\text{FWHM}} = 13.5 \mu\text{m}$ and $a_0 = 17$. The target plasma density was assumed to be pre-expanded to a Gaussian profile of $\sigma = 3.06 \mu\text{m}$ and $n_e = 0.8 n_c$ core density, as given in Eq. (4.2). The target radius before expansion was $r_0 = 0.6 \mu\text{m}$ for comparability with further performed simulations and its material polymethyl methacrylate (PMMA) ($\text{C}_5\text{H}_8\text{O}_2$) was initialized with pre-ionization. As for all mass-limited targets with extent below the focal spot size, this might seem like a generally inefficient starting condition for laser-target coupling. Nevertheless, one can recall from section 2.2.4 that the energy deposited in electrons escaping the interaction region in flat targets does not contribute to the acceleration of ions either [5]. Thus, geometrically enhanced electron energy density can potentially compensate for the limited laser-target cross-section.

The simulation box size was chosen such that the acceleration was largely complete, i.e. ion energies converged, before the particles left the simulation box. Thus, a volume of $104 \times 139 \times 104 \mu\text{m}$ over $3456 \times 4624 \times 3456$ cells was chosen. Corresponding to the target density, the spatial resolution was set to $\Delta x = \Delta y = \Delta z = 30 \text{ nm}$ and the time step was $\Delta t = 57.5 \text{ as}$, resolving the plasma frequency at peak density with $\omega_{p,e} \cdot \Delta t < 0.09$ and the laser wavelength with 35 cells. The complete simulation input with further numerical parameters was published in Ref. [319]. Corresponding memory requirements are satisfied by static load balancing over 8 000 Nvidia K20x graphics processing units (GPUs) and required the world’s largest supercomputer Titan at Oak Ridge National Laboratory (ORNL). The runtime until particle energy converges (42 000 steps) was 11.5 hours, using 0.3 million GPUh ($\simeq 10$ million CPUh). Output of the simulation for later post-processing amounted to approximately 1 PByte of raw data, which was compressed on-the-fly with zlib [265] and would have been impossible without the presented developments in section 3.4 and Ref. [218].

4.2.1. Temporal Evolution

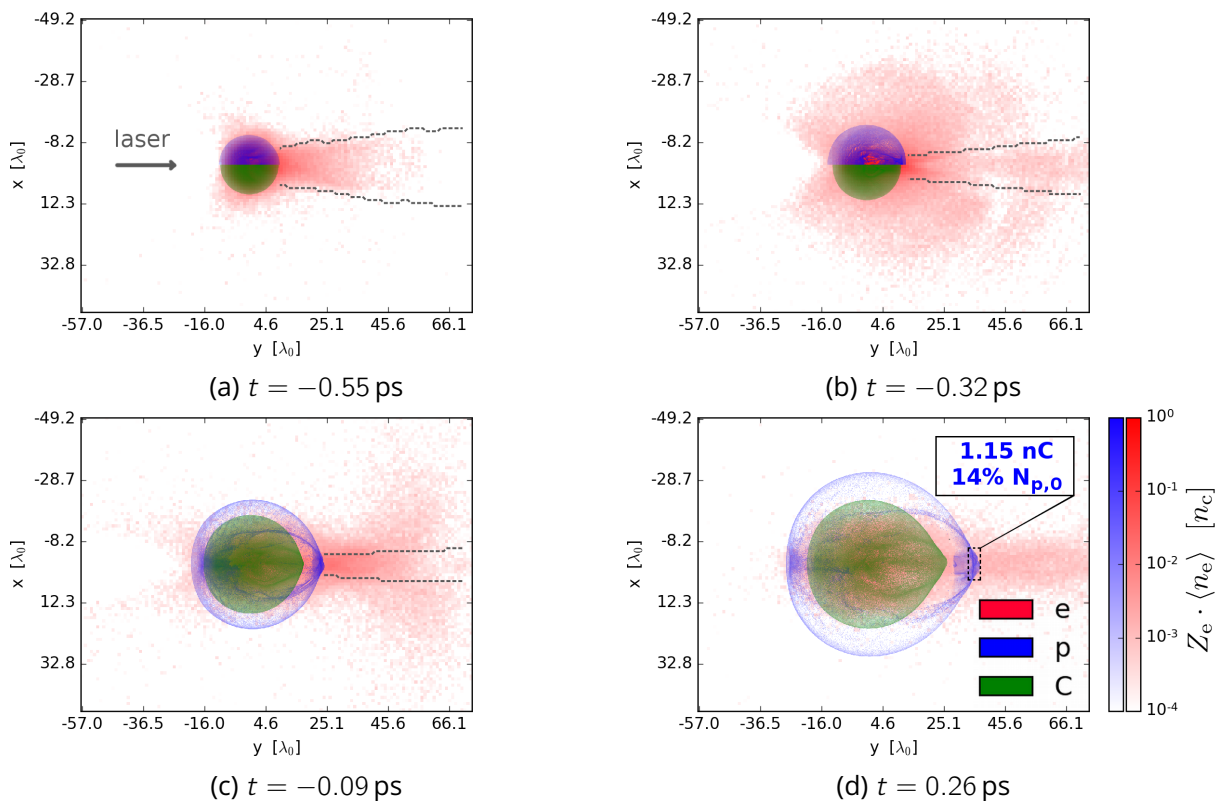


Figure 4.4.: Temporal evolution of ion and electron density for a slice in the plane of laser polarization. Dotted lines in (a)-(c) mark the peak intensity of diffracted laser light (see also Fig. 4.6a). Time t is relative to the instant of peak intensity on target at $t = 0$. O^{8+} density not shown here for simplicity, due to its close similarity to the C^{6+} distribution. The upper figures show proton and carbon ion density side-by-side for better visibility. In Figure (d), the spatial selection measures 14 % of all initial target protons, accelerated in forward direction.

Figure 4.4 presents ion and electron density slices of the 3D simulation. Following the temporal evolution in Fig. 4.4 for a first overview, some electrons (red) are expelled from the transparent target as early as 550 fs before the laser pulse's peak intensity reaches the target. Due to target transparency, the laser pulse interacts with electrons volumetrically over all depths of the target. This process extends continuously over the whole laser pulse rising edge, with a visible preference in forward direction. Ions (green and blue) follow the directed electron density while an isotropic background expansion from Coulomb repulsion forces persists, yet as displayed in Figure 4.3a, with approximately one order less particles per solid angle compared to forward acceleration. This is consistent with experimental results, as a significant increase in ion energy and particle numbers is observed along the laser-forward direction. Due to their favorable charge to mass ratio, protons gain higher velocity in the accelerating fields than carbon ions (oxygen ions at identical positions are not shown for simplicity), separating their expansion fronts as early as 320 fs before peak intensity on target. Additionally, the heavier ions' expansion front seen in Figs. 4.4c and 4.4d assists the acceleration of protons from deeper layers of the target, known as multi-

species effect [320–324], which contributes to the generation of a quasi-monoenergetic peak in the proton energy spectra (Fig. 4.3b). A detailed investigation of this effect for a hydrogen-deuterium target is presented in chapter 5.

Adding to the presented density slices, Figure 4.5 displays a 3D volume rendering of the proton density from simulation data at $t = -0.17$ ps prior to laser peak intensity on target. There, an accelerated short, high-charge proton bunch is as well visible on the right.




Figure 4.5.: Volume rendering of the three-dimensional proton density for time $t = -0.17$ ps. The generated, compact proton beam is seen on the right.

4.2.2. Ponderomotive Focusing of Electrons

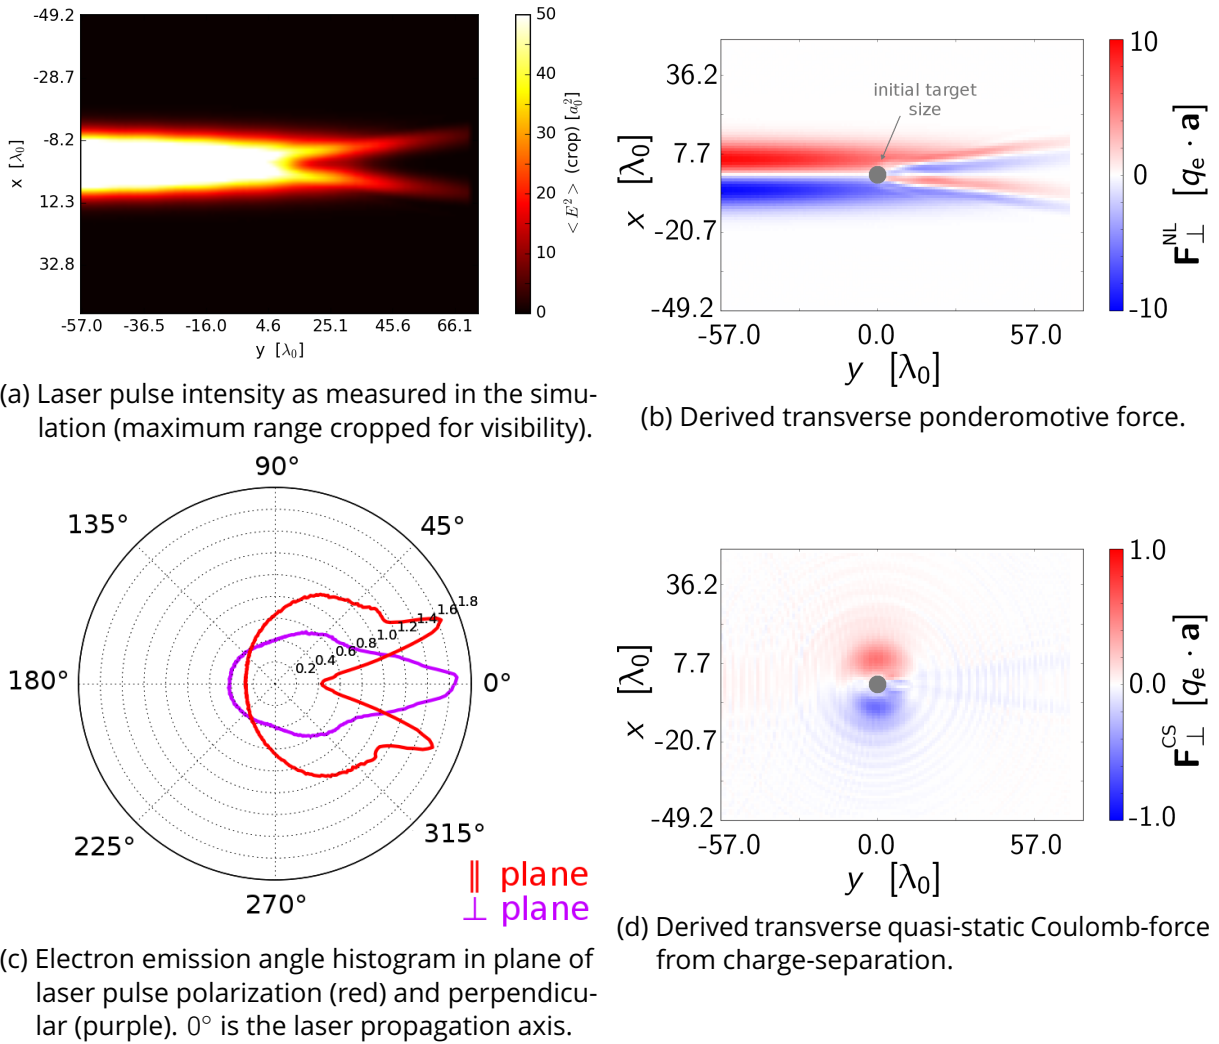


Figure 4.6.: Onset of the ponderomotive focusing effect via diffracted laser light on electrons expelled from the transparent target. Snapshots taken for time $t = -0.32$ ps. (a), (b) and (d) are shown for a slice in the plane spanned by laser pulse polarization x and propagation direction y .

Figure 4.6 presents the state of the target 320 fs before the instant of peak laser intensity on target. On the upper left, Figure 4.6a displays the laser intensity which clearly shows the refractive response of the near-critical plasma target, diffracting the laser field. Figure 4.6c displays the electron emission spectrum with respect to the two lateral axes. On the right, Figures 4.6b and 4.6d are derived force fields, showing the effective non-linear (NL) force electro-static fields from charge separation (CS), both in perpendicular direction.

A significant difference to 2D3V dynamics (see section 2.3.4) manifests in the correct balance of electro-static forces from charge separation (Coulomb-explosion) versus the strong focusing dynamics of the laser pulse. As seen in Fig. 4.6a, the transparent target still contains enough electrons in order to visibly diffract the laser pulse on its central axis. One can now estimate the combined force on an electron \mathbf{F}^Σ from ponderomotive force F^{NL}

of the laser pulse and the Coulomb-force \mathbf{F}^{CS} exerted by the electric field \mathbf{E}^{CS} from ions with $\mathbf{F}^{\Sigma} = -q_e \cdot \mathbf{E}^{\text{CS}} + \mathbf{F}^{\text{NL}}$ (see page 307 in Ref. [39]) by averaging over the electric field in our simulation. Here, one can approximate the temporal average over an oscillation period T_0 of the laser with a spatial average $\langle \dots \rangle \equiv \langle \dots \rangle_{T_0} \propto \langle \dots \rangle_{\lambda_0}$ of the electric field \mathbf{E}^{PIC} in the simulation in order to save frequent data output

$$\mathbf{F}^{\text{CS}} \approx q_e \cdot \langle \mathbf{E}^{\text{PIC}} \rangle, \quad (4.4)$$

$$\mathbf{F}^{\text{NL}} = -\frac{q_e^2}{2m_e\omega_0^2} \nabla \langle \mathbf{E}^2 \rangle, \quad (4.5)$$

$$\nabla \langle \mathbf{E}^2 \rangle \approx \nabla \langle (\mathbf{E}^{\text{PIC}} - \mathbf{E}^{\text{CS}})^2 \rangle. \quad (4.6)$$

In the formulas above, the fast oscillating contribution of the laser electric field is cycle averaged in order to derive electro-static fields from charge separation. The ponderomotive force is the gradient of the intensity minus the contribution from the Coulomb-field, estimated for a non-relativistic test-electron.¹ Comparing the absolute contributions from electro-static, isotropic Coulomb fields in Fig. 4.6d and the ponderomotive focusing in Fig. 4.6b, the non-linear focusing force dominates by an order of magnitude. This explains the electron density evolution in Figs. 4.4a-c, where a large portion of electrons (red) is confined within the opening cone of the diffracted laser pulse (dotted lines). The figures until now present the plane parallel to the laser polarization in x . Laser polarization naturally influences the symmetry of the system, therefore, slightly varying electron emission signatures are observed in the planes parallel and perpendicular to the laser polarization (Fig. 4.6c). Regardless, the dominant electron emission is directed forward in both observation axes and electrons are focused even stronger perpendicular to the polarization.

Clearly, a large amount of electrons is expelled from the target volume during laser interaction, as observed in Figs. 4.4. Whether sufficient electrons are swiftly expelled to induce a Coulomb explosion of the remaining ion population can be assessed analytically. For a perfectly charged sphere of ions, the total necessary energy U_{sphere} to strip off all electrons can be calculated from the energy density w of the electric field as [325]

$$w = \frac{1}{2} \epsilon_0 \cdot |\vec{E}|^2, \quad (4.7)$$

$$U_{\text{sphere}} = \frac{4\pi}{15\epsilon_0} q_e^2 n_0^2 r_0^5, \quad (4.8)$$

which for the target size considered here amounts to a potential energy of $U_{\text{sphere}} = 30\text{J}$. This is the necessary laser pulse energy to evacuate all electrons from the target volume. As mentioned before, the laser spot is larger than the target, such that only a fraction of the laser energy T_{tot} is applied. The energy overlap T_0 of the transverse laser beam profile with

¹The exact ponderomotive force depends also on an electron's $\gamma(t)$, see section 2.1.1, but this average serves as an estimate.

1 σ_{sphere} diameter of the target is

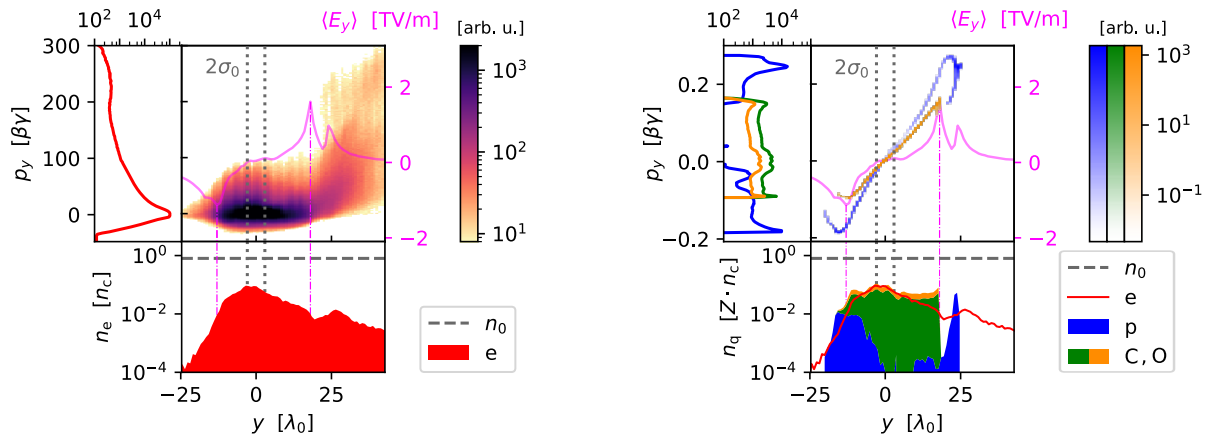
$$T_0/T_{\text{tot}} = \int_{-1\sigma_{\text{sphere}}}^{1\sigma_{\text{sphere}}} \frac{1}{\sigma_0^{\text{Int}}\sqrt{2\pi}} e^{-\frac{1}{2}(s/\sigma_0^{\text{Int}})^2} ds \quad (4.9)$$

$$= \text{erf}\left(\frac{1\sigma_{\text{sphere}}}{\sigma_0^{\text{Int}}\sqrt{2}}\right). \quad (4.10)$$

With the Gaussian beam width of $\sigma_0^{\text{Int}}\sqrt{2} = 13.5 \mu\text{m}$, the total energy interacting with the target until peak intensity does not exceed 13J. As already displayed in the temporal evolution in section 4.2.1, ions are accelerated significantly as the rising edge of the laser pulse interacts with the target, so only the energy contribution up to $t = 0$ is estimated.

Reference [98] provides an alternative calculation for the required laser field amplitude expelling electrons quickly from a homogenous cluster, delivering a similar conclusion like the one discussed in Eqs. (4.7)-(4.10), namely that the here studied target is too large for a Coulomb-explosion dominated plasma response which would translate eventually in an isotropic proton emission [99]. Hence, the laser field amplitude is not strong enough to drive a rapid Coulomb-explosion, but the laser pulse transverse profile is steep enough for dynamic focusing of expelled electrons (Fig. 4.6b). In summary, electrons are constantly removed in forward direction over the temporal rising edge of the laser pulse [96].

4.2.3. Directed Ion Acceleration



(a) Phase space for electrons with gradual removal of target electrons over the laser pulse upramp.

(b) Phase space for ions with visible multi-species effect. Further time steps are presented in appendix A.2.

Figure 4.7.: Longitudinal phase space snapshots for time $t = -0.09$ ps. Each figure is structured in (central) phase space with longitudinal electric field, (left) momentum histogram, and (bottom) stacked charge density. $2\sigma_0$ marks the initial target size. All data is taken along the laser propagation axis, perpendicularly averaged for one λ_0 .

Figure 4.7 presents snapshots of the longitudinal electron and ion phase space at a time shortly before the peak intensity of the laser pulse arrives at the target. A momentum histogram is shown left of each phase space plot and below a density lineout is presented.

The original target position is annotated as dashed grey lines and all data is perpendicularly averaged over one laser wavelength λ_0 around the laser propagation axis. As seen in Fig. 4.7a, electrons are volumetrically accelerated by the laser pulse from all depths of the transparent target. Electrons of highest momentum along the quasi-exponential energy spectrum can escape the ion potential.

For the given time close to the laser peak intensity on target, ions in Figure 4.7b already expand up to $25 \mu\text{m}$. Observing the spatial charge balance in the bottom inset of Fig. 4.7b, core regions of the target stay quasi-neutral and undergo self-similar expansion. Due to the high average kinetic energy of electrons, the averaged equivalent of a Debye length at the ion fronts extends to multiple micrometers, visible as charge imbalance in the density lineout. Hence, the accelerating electric field (purple) peaks at each ion species' front and influences several micrometers spatially.

A clear difference is visible between the proton density distribution as compared to the other ion species. In particular, a proton density peak is observed, preceding the comparably flat distribution of carbon and oxygen ions. The ratio between heavy ions with high charge state in fully ionized PMMA ($\text{C}_5\text{H}_8\text{O}_2$) is large compared to protons. 46 out of 54 electrons originate from the carbon-oxygen ion species, which results in a strong accelerating field at their front. Protons acquire higher velocities than heavier ions due to their favorable charge to mass ratio. Thus, protons from deeper target layers can pass the peaked accelerating field at the carbon-oxygen front and experience an acceleration which is stronger than for protons accelerated only at the proton front. Hence, a significant part of mid-momentum protons shifts to higher momentum during this process which is visible in the phase-space as a proton excess or "roll-up" (blue) above the carbon-oxygen front in Fig. 4.7b. The proton density is spatially enhanced as deeper-layer protons catch up with protons originating from outer layers of the target. This so-called multi-species effect and details on its influence on the resulting ion spectra is investigated in chapter 5.




Figure 4.8.: Evolution of the longitudinal electric field over time. $2\sigma_0$ marks the initial target size and arrows point to the position of the ion expansion fronts. Data is taken along the laser propagation axis, averaged over the laser wavelength λ_0 .

Simultaneously, the laser pulse's intensity still increases over the acceleration time of ions. As shown for the longitudinal electric field in Fig. 4.8 over time, the accelerating fields rise until the expansion front reaches approximately $15 - 20 \mu\text{m}$ from its original target center.

This increase in laser intensity and long-standing fields on the oxygen-carbon ion front enables even deeply originating protons to be accelerated via the described mechanism and exceed the energy of protons originating from outer target layers. Ultimately, the mass-limited target depletes its proton reservoir, seen for $0 < y < 10 \mu\text{m}$ in Fig. 4.7b and effectively all protons in the target rear can pass the accelerating carbon-oxygen front. This signature is clearly represented in the proton momentum distribution in blue on the left of Fig. 4.7b. The protons in the plateau below a momentum of $p_y < 0.2$ (as in $\beta\gamma$) originate from higher layers of the target that did not undergo the full additional acceleration in the carbon-oxygen front due to early demixing of surface target layers. These protons will appear in the final proton energy spectra as second peak at lower energy.

4.2.4. Discussion

The presented study reveals a regime of ion acceleration with surprisingly directed beam properties at realistic temporal laser contrast conditions. All experimental proton energy spectra are compared to the presented 3D3V simulation in Fig. 4.9. Before the main laser pulse interacts with the target, the pulse leading edge leads to substantial target pre-expansion, lowering the target density to approximately the critical density, as was estimated experimentally with fringe patterns and verified in 2D3V simulation scans. Nevertheless, due to limited control of the expansion dynamics, initial target conditions were not precisely known and variations in density profile between shots are likely. Important spectral characteristics of the resulting proton beam are surprisingly well reproduced by our simulation for shot #11, including a quasi-monoenergetic peak, a beam charge which reaches up to 14 % of all target protons, a double-peak structure and a reasonable agreement in cut-off energy for the given assumptions. Unfortunately, due to high computational costs at the time, only a single exploratory 3D3V simulation could be performed. Systematic variations of initial target parameters will likely have a great influence on the beam energy and might be able to explain the observed fluctuations up to a factor two in peak energy as seen in the other measurements in Fig. 4.9. Consequently, future PIC simulation campaigns might focus on 3D parameter scans varying initial target size, density, density profile, segregation of protons and heavier ions, expanded ionization distribution, and initial target temperature. Furthermore, special effort should be dedicated to the precise modeling of the temporal laser contrast at peak intensity which was demonstrated to be of high importance to the generated proton beam.

During the main interaction of the laser pulse with the transparent target, electrons are volumetrically accelerated by the laser pulse. Electrons form an exponential spectrum and high-energy electrons are ahead of the ion front, forming a several micrometer long Debye-like sheath. For the ultrahigh intensity and relatively long pulse of 500 fs in this scenario, the number of electrons expelled from the target rises over the laser pulse upramp leading to a constant flux through the ion surface. Furthermore, these electrons are strongly directed on-axis due to the ponderomotive potential of diffracted laser light, but do not form a critical layer as would be an important conditions for radiation pressure acceleration (RPA) [27, 113].

Ions experience a mixture of acceleration mechanisms which can be summarized as follows. Core-target carbon and oxygen ions mainly expand self-similarly to the target

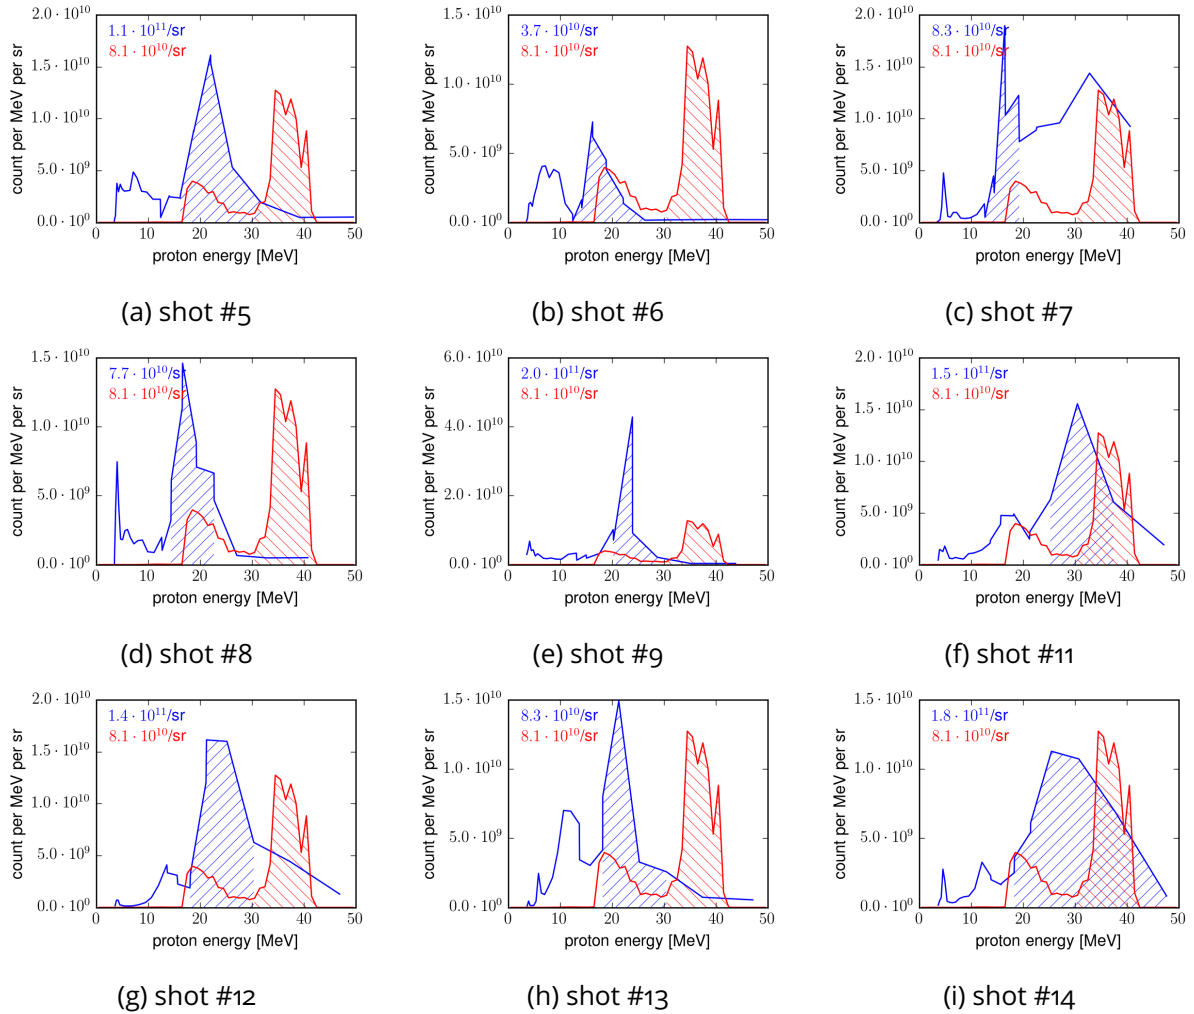

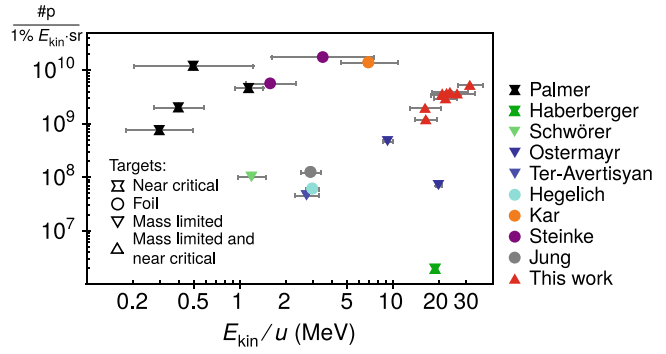


Figure 4.9.: Overview of all available experimental proton spectra (blue) versus the same 3D3V simulation (red). Numbers in the top left corner are integrated particle numbers from the shaded selection of the spectrum for an order of magnitude estimate.

normal. Outer layer ions experience a Coulomb explosion-assisted push and are further accelerated in the constant forward directed flux of electrons escaping the target with increased laser intensity. The latter is similar to target normal sheath acceleration (TNSA) acceleration with a prompt electron population, yet due to target transparency the electrons are driven constantly by the laser pulse. Furthermore, protons from deeper layers of the target undergo multi-species effects, shifting mid-energy protons to higher energies. In phase space, this effect is visible as a proton "roll-up" signature and causes a quasi-monoenergetic ion peak with high particle numbers per solid angle. Forward-accelerated proton numbers are summarized in Fig. 4.10. As shown in Fig. 4.10a, protons per solid angle significantly exceed isotropic emission levels (blue line), highlighting the difference to a perfect Coulomb explosion or TNSA expansion with thermalized electrons. With regards to experimentally measured shots, the presented 3D3V simulation (red line) reproduces the high directionality of the proton beam well within the margin of error by the experimental measurement. Moreover, Figure 4.10b compares the experimental measurements to other



(a) Proton number per solid angle for experimental shots \diamond in comparison to the 3D3V simulation (red line) and isotropic emission (blue line). The green band marks the standard deviation of experimental data points.



(b) Comparison of the differential ion yield for experiments observing a narrow ion energy spread (FWHM in horizontal bar): \times CO₂-laser with gas target [116, 117], \circ foil target [3, 109–111], ∇ spherical target (e.g. droplet) [4, 96, 123].

Figure 4.10.: Yield of forward accelerated protons. Comparison between experiment and simulation of the PHELIX campaign and comparison to experimental publications from other groups. Figures under CC-BY 4.0 [11].

experimental campaigns that reported narrow energy spreads. With the presented scenario, a regime of highest energies and simultaneously very high particle numbers was found.




Figure 4.11.: Proton energy evolution over time. Presented are forward-accelerated protons with $\pm 2^\circ$ acceptance to the laser propagation axis.

Multi-species effects contribute significantly to the narrow shape of the proton beam in this campaign but overlap with further effects of mass-limitation and rising intensity of the laser pulse which complicate the isolated investigation of the beneficial influence of multiple ion species. Therefore, chapter 5 of this thesis presents a system suitable for a systematic analysis in high-repetition rate experiments. Follow-up studies with the Paul-trap target would benefit from controlled pre-expansion, e.g. by deploying plasma-mirror setups [85–91] and variable pre-pulses. Since the important aspects of the acceleration mechanism are complete before even the main laser pulse intensity reaches the target (summary in Fig. 4.11), laser systems with shorter pulse lengths such as the upcoming PW-class PENELOPE laser with $\tau_0 = 150$ fs might be suited to explore this regime in more detail as well.

4.3. Summary and Perspectives

In this chapter, a mass-limited, micrometer-sized spherical target was investigated for PW-class laser-ion acceleration. In a collaborative effort, experiments were performed by a team from Ludwig-Maximilians-Universität München (LMU) and simulations by the author of this thesis. Important results were obtained under irradiation with $\tau_l^{\text{FWHM}} = 500$ fs, PW-class laser pulses as provided by the PHELIX laser system at Gesellschaft für Schwerionenforschung (GSI). With finite temporal laser intensity contrast, the 100 ps long rising edge of the laser pulse induced a strong pre-expansion of the target which was estimated to lower the plasma density to near-critical conditions. Although non-optimal laser contrast conditions are often limiting the quality of laser-ion acceleration [95, 326], here a regime was established generating an energetic proton bunch with narrow energy spread and high charge [11]. Driven by a 500 fs laser pulse, simulations predict a beam length of only $4 \mu\text{m}$ (≈ 54 fs) with charge of 1.15 nC, equivalent to 14 % of the initial target hydrogen ions. Experimentally, proton energies between 20 – 40 MeV were observed over nine measurements with higher differential ion yield than in previously reported experiments observing a narrow ion energy spread [3, 4, 96, 109–111, 116, 117, 123]. The author’s simulations were carried out at the formerly fastest supercomputer in the world and for the first time used manycore computing hardware for laser-ion acceleration at a scale of 8 000 GPUs. Simulations provided a substantial contribution to the interpretation of the research campaign, confirming the experimental results and attributing the favorable proton acceleration conditions to volumetric laser-electron interaction, acceleration during the rising edge of the laser pulse and multi-species effects. These 3D3V PIConGPU results further agree remarkably well with spectral shape, particle emission distribution and maximum energy deduced from the experiment under optimal conditions, while 2D3V simulations were unable to describe the acceleration process sufficiently, due to inherent geometric simplifications.

Due to computational costs and exploratory nature of the campaign, 3D3V simulations presented herein could only be conducted with the largest available supercomputers at the time. After these first 3D3V results, systematic parameter scans for variations in, e.g. laser intensity and pointing, expansion size and shape, ionization dynamics, etc. should follow. Knowing the highly directed ion beam generated in the process, simulation box sizes can be slightly reduced in the future, but full resolution and geometry will still be required for predictive capabilities. With new machines that provide one order of magnitude more computing power than available at the beginning of this thesis, one can investigate the scaling and critical parameters of this favorable acceleration regime in greater detail.

Simulations suggest that the crucial time of the acceleration process is finished as the main intensity of the 500 fs PHELIX laser pulse reaches the target, opening the possibility for studies with ultrashort-pulse laser systems such as DRACO ($\tau_0 = 30$ fs) or the upcoming PW-class 150 fs system PENELOPE. In view of such potential ultrashort-pulse experiments, simulations with this target for optimized contrast conditions were conducted. Here, the laser pulse is modeled as perfect Gaussian pulse with high temporal contrast, assuming no target pre-expansion occurred. Figure 4.12 displays a first target size scan with regards to




Figure 4.12.: 3D3V simulations results displaying the accelerated proton spectra with a high-contrast laser pulse of amplitude $a_0 = 17$ and pulse length $\tau_0 = 30$ fs.

maximum proton energy. The investigated range of targets with radius between $0.2\text{--}1.2\ \mu\text{m}$ seems to have little influence on the maximum energy under high laser contrast and yield reasonable maximum energies for conditions similar to the DRACO 150 TW system. For the smallest target in consideration, a transition to a Coulomb-explosion dominated regime similar to the observations in Ref. [123] is anticipated, as already visible in the reduction of low-energy Carbon ions, which leads to potentially increased maximum proton energies with less directionality.

Next steps for this target system might focus on systematic studies which control the expansion dynamics for further exploration of the favorable, near-critical acceleration regime. For such studies, readily available laser systems (e.g. DRACO) can introduce pre-pulses with varying delay to the main pulse, controllably pre-heating the target in order to find a potentially optimal configuration in terms of maximum proton energy.

5. Multi-Species Effects in PW-Class Laser-Ion Acceleration

*Results of this chapter are under review
A. Huebl et al., submitted (2019)*

Near-critical and mass-limited targets, as the one presented in the previous chapter, are attractive potential ion sources. Nevertheless, in order to improve identification and control of acceleration mechanisms in the ultrahigh intensity regime, systematic understanding of the interplay of simultaneously occurring microscopic mechanisms in laser-plasmas needs to be improved. In most targets, a mix of hardly controllable multi-physics processes such as ionization and absorption dynamics [66, 327], electron transport through the overdense target bulk [50, 328], or multi-species effects [127, 324, 329, 330], render specialized studies on an isolated physical effect challenging.

In this chapter, a novel cryogenic target based on cryogenic hydrogen is explored [43, 44, 46, 47]. With a low density of $30 n_c$, it provides a well-suited platform for fast turn-around simulations and a number of benefits for experimental studies. Simulations in close collaboration to experiments at Helmholtz-Zentrum Dresden – Rossendorf (HZDR) are performed continuously, preparing this target type as a suitable replacement for mechanically aligned metal or plastic foil targets. Furthermore, starting with a single ion-species material allows to explore whether multi-species effects can be harnessed with specific target mixtures and systematically understood in a dedicated system. As presented in the previous chapter, multi-species effects have a significant influence on the spectral shape of laser-generated ion beams and might lead to better interpretation of the underlying acceleration mechanism. Therefore, this chapter studies multi-species effects in a detailed manner. A new theoretical model is derived that predicts the influence of target composition on the spectrum of the generated ion beam and draws a connection to microscopic plasma properties driving it.

5.1. Unique Target Properties

Ongoing research towards applicable laser-driven ion beams investigates novel target technologies suitable for repetition rates in the order of 0.1 to 1 Hz. Contrary to flat foil targets or the mass-limited target presented in the last chapter, continuous targets without the need of realignment between shots are of high interest in order to achieve this goal. Improved repetition rates also improve systematic understanding due to better shot statistics with higher confidence levels, allowing to investigate fluctuations in laser and target conditions that may have critical influence on the acceleration performance [331].

Hence, our research group carried out experiments with a novel cryogenic target system [43, 44, 46, 47]. One fundamental issue under high shot rates is the generation of debris, which gradually degrades costly optics in the target chamber. Cryogenic targets overcome this limitation since the target material, e.g. hydrogen, is naturally gaseous at room temperature and does not accumulate in the form of contamination layers, which makes them intrinsically debris-free [41, 42, 45, 48]. The investigated jet target is further renewable, in the sense that no realignment between shots is necessary as new target material streams continuously through the focal spot volume, avoiding the need for mechanical repositioning between shots. Adding to these favorable experimental properties, it was shown in the last chapter that near-critical targets are of high interest for improved absorption of the laser pulse. The cryogenic hydrogen target operates with a mildly overcritical density of $30 n_c$ and can be generated with few-micron thickness in order to potentially benefit from mass-limitation effects [43, 44, 46, 47].

With regards to simulations, cryogenic jets need little resolution ($\Delta x \approx 3 - 5$ nm) in order to cover the microscopic plasma dynamics properly and are in principle geometrically suited for 2D3V treatment with respect to modeling only a plane orthogonal to the jet axis. Nevertheless, the neglected third dimension is assumed as periodic (see section 2.3.4) which is equivalent to an infinite line focus along the jet axis. For relativistic intensities, plasma collisional rates are negligible (see section 2.1.2). One can therefore achieve very fast simulation turn-around time even on medium-sized supercomputers, allowing to investigate up to four simulation setups per hour for systematic studies on one of HZDR's 64 graphics processing unit (GPU) systems [254].¹ Finally, with the particle-in-cell (PIC) techniques introduced in chapter 3, simulation and analysis time becomes comparable to the duration of experimental campaigns, fostering productive scientific exchange in real-time instead of sequential, multi-month efforts. It is worth noting, that equivalent 3D3V simulations do not take significantly longer ($\approx 2 - 3$ hours) with PIConGPU on a current TOP10 machine [241], on account of parallelization and scalable data reduction. Nevertheless for daily 3D simulations, there are still too few high-performance computing (HPC) resources in Europe for the required on-demand multi-million (parallel) compute-hours.² As of today, PIConGPU's computational software stack is capable of handling such workflows in 3D, as was demonstrated in chapters 3 and 4, and due to exponential capability growth (section 3.3.2) the corresponding scientific computing instruments will become available in the next years.

¹Nvidia P100

²A typical computing grant is currently about 50 million compute-hours per year, which one would spend within a week when working tightly integrated with experiments.

5.2. Target Geometry

Results of this section are published in
*L. Obst, ..., A. Huebl et al.,
 Scientific Reports 7, 10248 (2017)*

Before performing the central study on multi-species effects, the pure hydrogen jet's suitability for proton acceleration and predictive capabilities of the corresponding PIC simulations had to be established. As such, it was investigated if a cryogenic hydrogen jet can be a suitable replacement of flat foils in terms of generated ion beam energy and particle numbers [44]. Two specific target shapes were considered, a cylindrical geometry with diameter $d = 5 \mu\text{m}$, which is experimentally relatively easy to produce, and a technically more challenging foil-like geometry with a width $w = 20 \mu\text{m}$ and $t = 2 \mu\text{m}$ thickness.

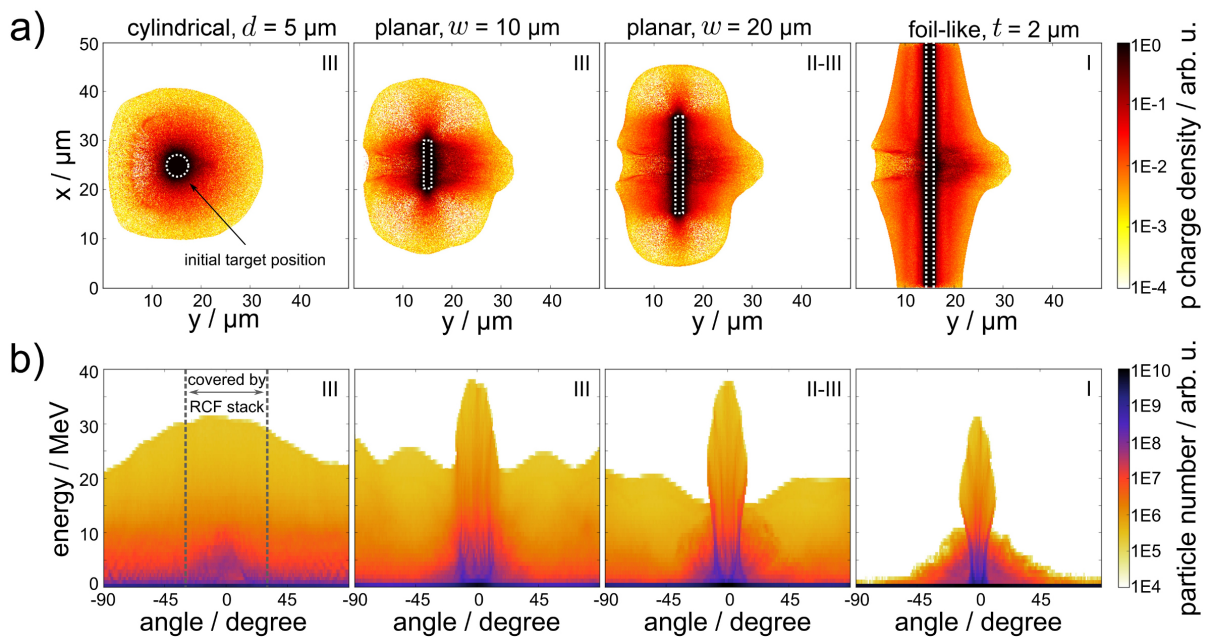


Figure 5.1.: 2D3V PIConGPU simulations for varied cryogenic hydrogen target geometries for a snapshot 260 fs after the laser pulse's peak intensity reached the target. Roman numerals denote mass-limitation regimes from Tab. 2.2, section 2.2.4. a) Proton charge density varied from cylindrical geometry to planar and flat geometry. b) Angular proton emission in forward direction for the same geometries. Figure under CC-BY 4.0 [44].

Experiments were performed with the 150 TW frontend of the DRACO laser at HZDR, focused to a spot size of $w_l^{\text{FWHM}} = 3 \mu\text{m}$. On-target laser pulse intensity was approximately $a_0 = 16$, equivalent to $\approx 6 \cdot 10^{20} \text{ W cm}^{-2}$. A first overview from simulations is given in Fig. 5.1, depicting proton density evolution and emission characteristics (for all parameters of the simulation setup see Ref. [44]). The cylindrical target geometry and planar targets with limited transverse width emit into larger angles consistent with target normal signatures from a bent surface, since the expanding electron sheath can surround the whole target. Changing from cylindrical to planar geometry, a significant forward directed emission characteristic arises with $10 - 20^\circ$ opening half-angle. With increased planar target width, the

Debye sheath is less influenced by electrons reflected at the target edges and recirculating through the entire target volume, but more by the size of the focal spot and the electron divergence angle. Although laterally limited planar geometries generate slightly less directed proton beams than the flat geometry, the maximum ion energy is increased in 2D3V. The underlying effect is caused by mass-limitation of the target, avoiding a perpendicular diffusion of laser-accelerated electrons far from the focal region, which in turn increases the energy density available for acceleration of ions in the sheath (see section 2.2.4).




Figure 5.2.: Experimental proton emission as measured with RCF stacks. a) and b) The proton angular emission distribution is shown perpendicular to the jet axis for comparison to simulation data in Fig. 5.1b. c) A titanium foil was irradiated by the laser pulse under an oblique incidence angle of 45° and the ion emission center is shifted by -15° , likely due to contributions from prompt electrons to the electron sheath. Figure under CC-BY 4.0 [44].

Corresponding experimental results in Fig. 5.2 show the emission characteristics measured from RCF stacks. For Figure 5.2a and b, lineouts in the horizontal plane correspond to cylindrical and planar geometry in Figure 5.1b. Qualitatively agreeing with simulation results, the cylindrical geometry leads to a broad emission and planar geometries exhibit higher directionality of the proton beam. A Titanium (Ti) foil (Fig. 5.2c) is shown for experimental comparison of emitted particle numbers and was irradiated under 45° (all other under 0°). An important result for experiments are the comparable particle numbers exceeding 10^9 protons $\text{MeV}^{-1} \text{sr}^{-1}$ from the planar hydrogen jet geometry under optimized target normal sheath acceleration (TNSA) conditions in comparison with the $2 \mu\text{m}$ titanium foil, which proves the desired applicability of the cryogenic target as replacement for solid density targets. Moreover, proton cutoff energies are comparable between all geometries and conform with optimal acceleration conditions achievable at DRACO 150 TW [79].

Figure 5.3 displays the simulated hydrogen spectra along lineouts in 0° direction. Here, particle numbers are compared with respect to change in target geometry. Under the given periodic target geometry along the jet axis, performing 2D3V simulations is acceptable for qualitative comparisons but quantitative predictions require proper three-dimensional treatment [176].³ Hence, the predicted maximum ion energy is inflated by a factor $1.7 - 2$ when comparing 2D3V simulated and experimental hydrogen cutoff energies (see section 2.3.4) due to an implicit line-focus in the periodic direction. Yet, trends of particle numbers in emission spectra under changed lateral target geometry from cylindrical to planar jets show

³Unfortunately, the available memory on manycore accelerators in our local cluster was not sufficient for fully dimensional treatment without sacrificing adequate resolution.

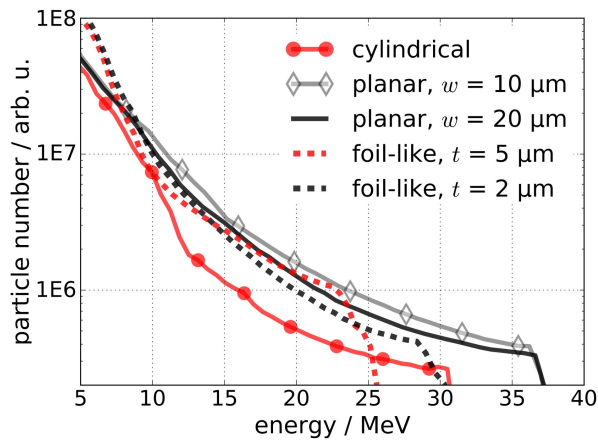


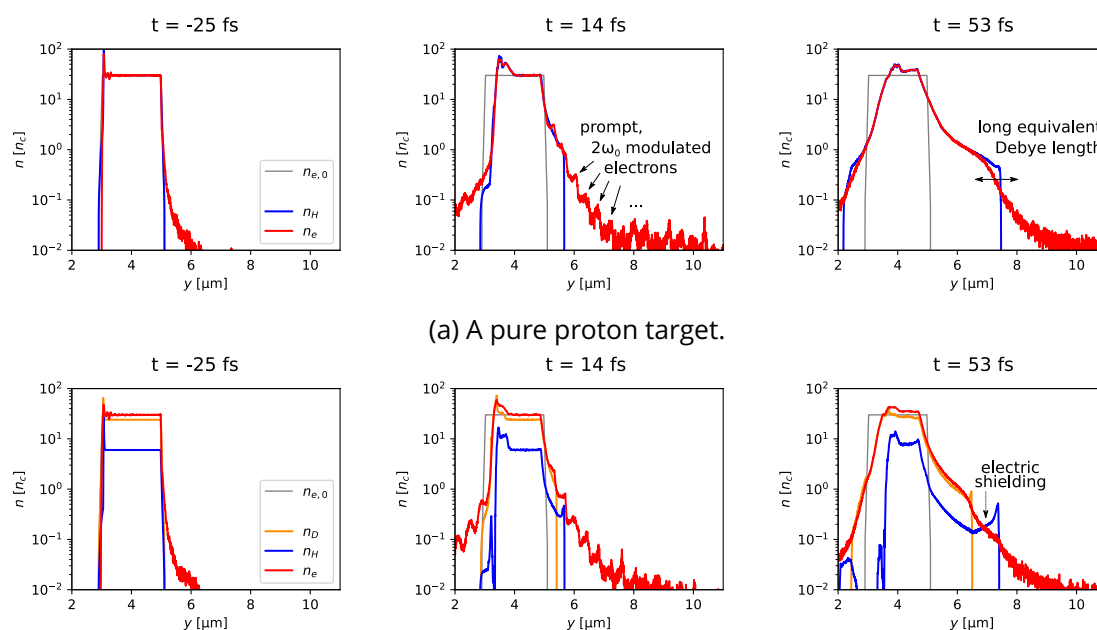
Figure 5.3.: 2D3V PIConGPU simulation results comparing proton spectra with changed geometry for the cryogenic hydrogen target. Particle numbers are scaled to the same reference. Figure under CC-BY 4.0 [44].

good qualitative agreement between simulation and experiment. Both simulated and experimentally measured emission functions show increased directionality and increased forward particle numbers for narrow planar geometry, corroborating the predictive capabilities of the presented simulations for the given target. As such, credible systematic numerical investigation of processes and effects unfolding during the acceleration, the assessment of which is highly challenging in experiments, becomes possible.

5.3. Systematic Characterization of Multi-Species Effects

As introduced in section 2.2.5, energy spectra are of central importance for the identification of acceleration mechanisms and are easily accessible in experiments compared to microscopic properties of the plasma evolution. Identification of acceleration regimes is often justified by correlating spectral shape and maximum ion energy with assumed conditions in the laser focus [5, 6, 108–111]. Yet, the shape of energy spectra is influenced by many effects (see section 2.2.5) even for well-known mechanisms such as TNSA, which is studied here.

One important aspect is the presence of multiple ionic charge states and species in virtually all target materials but hydrogen. Although also relevant for thin surface contamination layers [12, 83, 108, 127, 332], the following focuses on homogeneously mixed target materials. So-called multi-species effects are the result of charge shielding between varying charge-to-mass ion components as they co-propagate during or after acceleration [320–324]. Charged "lighter" particles, in terms of q/m , gain more velocity than heavier species in the same accelerating field, which leads to a charge separation of the expanding ion fronts. For example in a cryogenic methane (CH_4) target [47], the spatial TNSA expansion front is led by protons, followed by C^{6+} , then C^{5+} , C^{4+} , etc. [320]. Those expanding ion populations overlap spatially, as seen for two species in density lineouts derived from simulations and displayed in Fig. 5.4.



(a) A pure proton target. (b) A target of 20 % protons and 80 % "heavier" ions with Deuterium-like charge-to-mass ratio.

Figure 5.4.: Ion and electron density evolution for a single-species versus a multi-species target with flat geometry. Target and laser parameters: $n_e = n_H = 30 n_c$, with thickness of $2 \mu\text{m}$ and an exponential plasma pre-scale length of 20 nm . The peak laser amplitude in focus is $a_0 = 16$ and the pulse length is $\tau_0 = 30 \text{ fs}$.

The accelerating fields at each ion species expansion front of TNSA expands over the scale length of the local Debye sheath. Especially under pre-thermal conditions (electrons are red in Fig. 5.4; see section 2.2.2), the Debye-like screening lengths is large compared to the

expansion difference of the two ion species. The farther expanded lighter species shields part of the accelerating field and, conversely, this means that the heavier species' front loses momentum to spatially close, mid-energy lighter ions which are shifted to higher energies. The onset of this effect is readily visible in Fig. 5.4, including the important observation that the maximum position of the proton expansion front and thus maximum energy remains unaffected by the process. Hence, exploring multi-species effects in TNSA with two-species targets is not expected to compromise the overall acceleration performance of the lighter ion species as long as the laser-accelerated electron density is comparable.

5.3.1. Analytical Model

In the following, an analytical model for the multi-species effect is derived by describing the shielding of a lighter ion species upon the expansion front of a heavier species during their expansion. As nomenclature, the lighter species is denoted as "hydrogen/protons" (H) and the heavier species as "deuterium" (D), since the Deuterium charge-to-mass ratio is prototypical for the charge-to-mass ratio of many ions but hydrogen. Nevertheless, no assumptions are made about the target composition ratio or the charge-to-mass proportion between the two ion species.

Following the typical TNSA picture from section 2.2, each ion front expands approximately with a velocity proportional to their charge to mass ratio $v \propto Z/M$. As already visible in Fig. 5.4, electrons accelerated in such targets are initially not thermalized and the screening length at the ion fronts is large. One can therefore derive the energy lost by the heavier species in terms of maximum cutoff energy resulting from the lighter species' charge screening. The heavier species front will reach a lower maximum energy than a pure (heavy-species) target and the position of its cutoff energy coincides with a dip in the spectrum of the lighter species. The dip results from a shift of light ion charge to higher energies as they are further accelerated by the fields on the heavier ion front.

The rear-side expansion of the light ions is well-approximated by an exponential density profile [69, 71]

$$n_H(y, t) = n_{0,H} \cdot \exp\left(-\frac{y}{c_{s,H}t} - 1\right). \quad (5.1)$$

y is the propagation from the target rear, t is time, $c_{s,i} \approx (Z_i k_B T_e m_i^{-1})^{1/2}$ is the ion acoustic speed of sound for ion species i , and $n_{0,H}$ is the hydrogen content in the target. For the overall target composition, the free electron density is kept constant for comparability and only the ion mixing ratio with relative deuterium content d is changed

$$n_{0,e} = Z_D n_{0,D} / d \quad (5.2)$$

$$= Z_H n_{0,H} / (1 - d) \quad (5.3)$$

$$= Z_H n_{0,H} + Z_D n_{0,D}. \quad (5.4)$$

Z_i is the ion charge state.

Solving Gauss's law (with the dielectric constant ϵ_0)

$$\frac{\partial E}{\partial y} = \frac{Z_H q_e n_H}{\epsilon_0}, \quad (5.5)$$

one can assess the electro-static shielding experienced by the deuterium front. q_e is the (positive) elementary charge. Therefore, the superimposed decelerating electric field due to the presence of hydrogen at the position of the deuterium cutoff (ctf) front is given by

$$\Delta E_D^{\text{ctf}} = \frac{Z_H q_e}{\epsilon_0} \int_{y_D^{\text{ctf}}}^{y_H^{\text{ctf}}} n_H \cdot dy. \quad (5.6)$$

For the position of an expanding ion front in TNSA over time, one can apply a central finding of the Mora-model [71], predicting the spatial ion front at

$$y_i^{\text{ctf}}(t) \cong c_{s,i} t \cdot [2 \ln(\omega_{p,i} t) + \ln 2 - 3]. \quad (5.7)$$

The approximate solution for times $\omega_{p,i} t \gg 1$ is adequate, since ion front separation and hence hydrogen shielding at small times is negligible [71]. The prompt ("hot") electron plasma frequency in the target rear is given by

$$\omega_{p,i} = \sqrt{\frac{n_e^{\text{prompt}} Z_i^2 q_e^2}{m_i \epsilon_0}} = Z_i \sqrt{\frac{m_e}{m_i}} \omega_{p,e}, \quad (5.8)$$

scaled to the ionic frequency scale in terms of mass and charge.⁴ In all following equations, densities are therefore identified with respect to the prompt electron density $n_e \equiv n_e^{\text{prompt}} \leq n_{0,e}$ (see section 2.2.2). The according normalized position for the hydrogen density in Eq. (5.1) reads

$$\frac{y_H^{\text{ctf}}(t)}{c_{s,H} t} = \left[2 \ln \left(Z_H \sqrt{\frac{m_e}{m_H}} \omega_{p,e} t \right) + \ln 2 - 3 \right], \quad (5.9)$$

$$\begin{aligned} \frac{y_D^{\text{ctf}}(t)}{c_{s,H} t} &= \left[2 \ln \left(Z_D \sqrt{\frac{m_e}{m_D}} \omega_{p,e} t \right) + \ln 2 - 3 \right] \cdot \sqrt{\frac{m_H}{m_D}} \\ &\quad - \frac{1}{c_{s,H} t} \int_0^t \Delta v_D^{\text{ctf}}(d, \tau) \cdot d\tau. \end{aligned} \quad (5.10)$$

As the heavier species' front is slowed down due to the decelerating field of the lighter species, further protons can overtake the deuterium front compared to an undisturbed expansion given by Eq. (5.7). The second term in Eq. (5.10) accounts for this effect. One can solve Eq. (5.6) numerically for the resulting integro-differential equation of $y_D^{\text{ctf}}(t)$ or continue analytically by dropping the second term in Eq. (5.10). The latter approach permits to derive an analytical upper estimate for the reachable deuterium cutoff energy. The

⁴This scale is already needed for proper interpretation of Ref. [71] (see section 2.2.2). The dependence on electron density reflects the fact that the number of self-similar expanding ions is determined by the amount of laser-accelerated electrons driving the process. The self-similar expansion dynamics itself occurs on the ion time scale.

change of electric field from Eq. (5.6) at the deuterium front is

$$\Delta E_D^{\text{ctf}} = \frac{q_e}{\epsilon \cdot \epsilon_0} n_e (1-d) \cdot c_{s,H} t \cdot \left[-\frac{\epsilon^3}{2 \frac{m_e}{m_H} \cdot (Z_H \omega_{p,\epsilon} t)^2} + \left(\frac{\epsilon^3}{2 \frac{m_e}{m_D} \cdot (Z_D \omega_{p,\epsilon} t)^2} \right)^{\sqrt{\frac{m_H}{m_D}}} \right]. \quad (5.11)$$

Symbol ϵ denotes Euler's number. The equation can be simplified by introducing the normalized time $\tau = \omega_{p,D} t$ and the square root of the relative mass ratio $m_r^2 = m_H m_D^{-1}$.

$$\Delta E_D^{\text{ctf}} = \frac{m_D}{\epsilon \cdot q_e Z_D^2} \cdot (1-d) \cdot c_{s,H} \cdot \omega_{p,D} \tau \cdot \left[-\frac{\epsilon^3 Z_D^2 m_r^2}{2 Z_H^2 \tau^2} + \left(\frac{\epsilon^3}{2 \tau^2} \right)^{m_r} \right]. \quad (5.12)$$

The equation of motion for the deuterium front follows from Lorentz force $dp_y(t)/dt = Z_D q_e \cdot \Delta E_D^{\text{ctf}}(t)$

$$\Delta \dot{v}_D^{\text{ctf}}(t) = \Delta E_D^{\text{ctf}}(\tau) \cdot \frac{Z_D q_e}{m_D}, \quad (5.13)$$

$$\int_{\omega_{p,D}^{-1}}^t dt \rightarrow \int_1^{t \cdot \omega_{p,D}} d\tau, \quad (5.14)$$

$$\frac{d\tau}{dt} = \omega_{p,D} \rightarrow dt = \frac{d\tau}{\omega_{p,D}}, \quad (5.15)$$

$$\Delta \dot{v}_D^{\text{ctf}}(\tau) = \frac{(1-d)}{\epsilon \cdot Z_D} \cdot c_{s,H} \cdot \tau \cdot \left[-\frac{\epsilon^3 Z_D^2 m_r^2}{2 Z_H^2 \tau^2} + \left(\frac{\epsilon^3}{2 \tau^2} \right)^{m_r} \right]. \quad (5.16)$$

The latter is integrated to

$$\Delta v_D^{\text{ctf}}(\tau) = \frac{(1-d)}{\epsilon \cdot Z_D} \cdot c_{s,H} \cdot \left[\left(\frac{\epsilon^3}{2} \right)^{m_r} \cdot \frac{\tau^{2-2m_r}}{2-2m_r} - \frac{Z_D^2 \epsilon^3 m_r^2}{Z_H^2 2} \cdot \ln(\tau) \right]_1^{\tau_D} \quad (5.17)$$

$$= \frac{(1-d)}{\epsilon \cdot Z_D} \cdot c_{s,H} \cdot \left[\left(\frac{\epsilon^3}{2} \right)^{m_r} \cdot \frac{\tau_D^{2-2m_r} - 1}{2-2m_r} - \frac{Z_D^2 \epsilon^3 m_r^2}{Z_H^2 2} \cdot \ln(\tau_D) \right]. \quad (5.18)$$

Finally, the cutoff energy K_D^{ctf} of the heavier species is calculated with

$$K_D^{\text{ctf}}(d) = m_D/2 \left(v_D^{d=1} - \Delta v_D^{\text{ctf}} \right)^2. \quad (5.19)$$

In the resulting equation (5.19), one calibrates against the cutoff energy of a pure ($d = 1$) target for the heavier species. This reference energy can be determined either by measuring, simulating, or utilizing an established theory for single component targets [51, 71, 77].

As already introduced above, the time $\tau_D = \omega_{p,D} t$ is the normalized time as in existing models and an effective acceleration time $\tau_D \approx 1.3 \cdot \tau_{\text{laser}}$ should be applied [71, 72]. The model derived above complements existing theories that predict the ion cutoff energy from given average prompt electron energy $\langle T_e \rangle$ and corresponding laser-accelerated electron density n_e . Both microscopic distribution parameters are unknown in an experiment and models today mainly predict values for $\langle T_e \rangle$ [51, 57, 62, 333]. Trying to guess a proper effective electron energy density with, e.g. cutoff from Eq. (2.42) alone, is therefore an under-

determined problem. With the new model and appropriate material choice for a clearly visible proton modulation from multi-species effects, an effective average electron energy and density has to fulfill Eq. (2.42) and Eq. (5.19). With two unknowns and two equations, the equation system is fully determined

$$K_D^{\text{ctf}}(d) = K_D^{\text{ctf}}(\langle T_e \rangle, n_e), \quad (5.20)$$

$$K_H^{\text{ctf}} = K_H^{\text{ctf}}(\langle T_e \rangle, n_e). \quad (5.21)$$

For homogeneous target materials without contamination layer, the material mixing ratio d is generally known in experiments.

5.3.2. Systematic Target Preparation

Up until now, an isolated study of the multi-species effect was experimentally challenging. Conventionally used flat foil targets are both influenced by surface layers and ionization dynamics of the target material, leading to an unknown target charge state distribution. With the unique target properties of cryogenic hydrogen jets, explained in section 5.1, one can potentially prepare near-ideal starting conditions for a systematic research campaign. In particular, a multi-species target can be attained by mixing hydrogen with deuterium gas in the cryostat. The production of hydrogen-deuterium mixed jets should be hydrodynamically feasible due to similar atomic physics of both constituents. Assuming high-enough intensity for sufficient ionization, such a target will only consist of two possible ion species, protons and deuterium ions, suppressing the influence of complex ionization dynamics. Contrary to composite materials, such as the campaign shown in chapter 4, the target constituents can be systematically changed by gas composition instead of being limited to the constituents of solid materials, e.g. polymethyl methacrylate (PMMA), which, moreover, cannot easily be exchanged without influencing the reference free electron density. Predictive capabilities of the simulation for this kind of target have been demonstrated in a campaign with pure hydrogen jet targets, see section 5.2.

5.3.3. Simulation Setup

In view of the experimental exploration of multi-species effects with hydrogen-deuterium jets, simulations were performed with PIconGPU (see chapter 3) in 2D3V [286]. Assumed conditions match the ion-acceleration end-station at HZDR with 150 TW and 1 PW stages of the DRACO laser. As in previous campaigns, the focal spot size is set to $w_{0,l}^{\text{FWHM}} = 3.0 \mu\text{m}$ and a Gaussian pulse profile of $\tau_{0,l}^{\text{FWHM}} = 30 \text{fs}$ and a central laser wavelength of $\lambda_{800} \text{nm}$ is assumed [46]. The selected settings are representative for ultrahigh intensity Ti:Sa laser systems used in contemporary, ultrashort pulse experiments [30].

Figure 5.5 provides an overview of the simulation domain. A foil-like jet geometry is applied with a box density profile including a slight exponential ramp, accounting for slight

pre-expansion and improving numerical stability

$$n_i(y) = n_0 \cdot \begin{cases} e^{-(y-2)/0.02 \mu\text{m}} & \text{if } y > 2 \mu\text{m}, \\ 1 & \text{if } 0 \mu\text{m} < y < 2 \mu\text{m}, \\ e^{y/0.02 \mu\text{m}} & \text{if } y < 0 \mu\text{m}, \end{cases} \quad (5.22)$$

$$n_0 = 30 n_c.$$

Focused on-target, the laser pulse amplitude is varied in steps of $a_0 = 8, 16, 23, 30, 42$. The focal position is set to the target center at $4 \mu\text{m}$ in the simulation box. Particles originating in the first "half" of the target, before $y_0 = 4 \mu\text{m}$, are tracked as target "front" and all other particles as target "rear" for later discussion.




Figure 5.5.: Simulation setup and applied naming conventions. In situ applied particle filters are modeled according to a synthetic RCF stack or pinhole, e.g. in a Thomson parabola. The latter is chosen for energy spectra in the following figures. The laser incident angle is zero degrees (from left along y) and laser pulse polarization is parallel to the simulated plane.

Specific algorithms used in the PIC simulations comprise the Yee Maxwell-solver, an optimized Esirkepov current deposition (ZigZag path splitting), randomized in-cell starting positions without temperature, 3rd order (piecewise cubic) particle assignment shape, Boris particle pusher, and weighted trilinear force interpolation (see section 2.3). Collision operators are neglected (see Table 2.1 in section 2.1.2 for justification). The spatial resolution in the simulation is $\Delta x = \Delta y = 3.33 \text{ nm}$ with $\Delta t = 7.85 \text{ as}$ on a grid of 7488×14720 cells (the axis of laser pulse propagation is y). Equivalently, the central laser wavelength is resolved with 240 cells and the non-relativistic plasma frequency $\omega_p \cdot \Delta t = 0.10$. The macro-particle distribution is initialized as pre-ionized with 20 particles per ion species and cell and one electron per ion. Each simulation computes 0.60 ps interaction time on 16 Nvidia P100 GPUs within 1:20-1:45 hours. Energy spectra were compared on final, converged results depending on laser pulse intensity and before particles leave the simulation box. Complete simulation input and data analysis routines are published in Ref. [334].

5.3.4. Influence of Deuterium Concentration




Figure 5.6.: Deuterium and hydrogen energy spectra for laser amplitude $a_0 = 16$ over varied target composition (deuterium ratio d). The inset compares cutoff energy over d . The analytical model from Eq. (2.42) is plotted as grey line for hydrogen and the new analytical model for deuterium from Eq. (5.19) as red line. Dashed lines are models from Refs. [329, 335] for deuterium cutoff (black) and hydrogen peak after modulation (grey). As input parameters for the analytical models the electron average kinetic energy $\langle T_e \rangle = 2.17$ MeV and $n_e^{\text{rear}} = 0.31 \cdot n_{0,e}$ are measured from the simulations.

Figure 5.6 presents the variation of proton and deuterium spectra under change of target composition for $a_0 = 16$ for forward ("pinhole") acceptance. Operation in the TNSA regime can be confirmed by an overall exponential spectral shape. A deep modulation of the proton energy spectra at the position of the deuterium cutoff energy per nucleon is caused by multi-species effects, the spectral position of which is predicted by the model in Eq. (5.19). A spectral hydrogen number excess after each dip in Fig. 5.6 implies that mid-energy protons are shifted to higher energies due to acceleration at the deuterium front. Integrating the overall proton signal reveals that parts of the spectrum are indeed solely shifted to higher energies and transverse displacement of protons is negligible. As the initial target density is kept constant with mixing ratio d , proton counts for a given energy do not exceed the pure hydrogen case. The overall proton charge is reduced linearly with mixing ratio.⁵

In the inset of Figure 5.6, the cutoff energies per nucleon are compared for the presented 12 simulations with $a_0 = 16$, from pure hydrogen ($d = 0$) to pure deuterium ($d = 1$). For high deuterium ratio, a narrow energy spread is generated in the hydrogen spectra.⁶ Following the prediction by the Mora-cutoff energy model in Eq. (2.42) for hydrogen (grey line), the maximum proton energy should be unaffected by a change in mixture. Observed variations are discussed later in detail. The new model from Eq. (5.19) for the deuterium cutoff is

⁵Confirmed by integrating all ions above one MeV in the spectrum.

⁶In comparison with the continuous, exponential TNSA spectrum, authors often refer to such narrow energy distributions as "quasi-monoenergetic" spectra [3, 4, 110].

plotted as red line, which agrees well with simulation results.

Moreover, results are compared to existing models [329, 335], plotted as dashed lines in Fig. 5.6. Translated to the nomenclature in this thesis, these multi-fluid models are given by

$$K_D^{\text{ctf}} = Z_D \langle T_e \rangle \ln^2 \left[4Z_D d \sqrt{2A_D Z_H / A_H Z_D} / (Z_H (d-1) \epsilon) \right] / 2, \quad (5.23)$$

$$K_H^{\text{peak}} = Z_H \langle T_e \rangle \ln \left[4Z_D d \sqrt{2A_D Z_H / A_H Z_D} / (Z_H (d-1) \epsilon) \right]. \quad (5.24)$$

Here, heavy species cutoff energy K_D^{ctf} and the spectral position K_H^{peak} after the modulation in which the lighter species shows an exaggeration (peak) of particle numbers are treated separately [335]. Existing models suffer non-physical divergences (see $d \rightarrow 0$ and $d \rightarrow 1$) and do not agree well with simulations, since they assume large ratios between both ion species in terms of charge and density ratio.

Additional spectra for laser strength parameters $a_0 = 8, 23, 30, 40$ are provided in Fig. 5.7

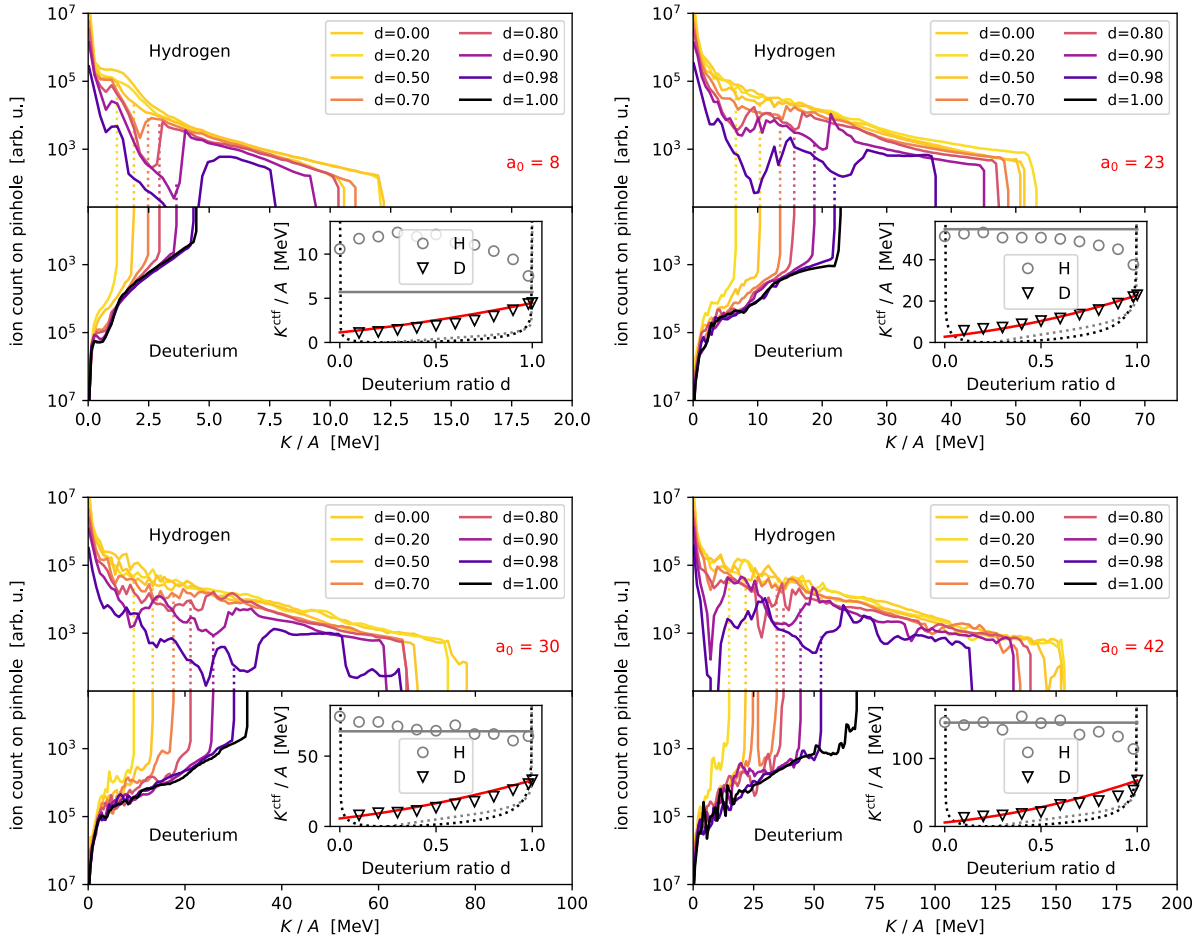




Figure 5.7.: Deuterium and hydrogen energy spectra for laser amplitude $a_0 = (8, 23, 30, 40)$ over varied target composition (deuterium ratio d). As in Fig. 5.6, the cutoff energy as a function of d is summarized in each inset with theoretical models from Eq. (5.19) and Refs. [71, 329, 335]. As input parameters for the analytical models the electron average kinetic energy $\langle T_e \rangle = (0.35, 3.52, 4.52, 8.64)$ MeV and non-perfect absorption into rear electrons $n_e^{\text{rear}} = (0.48, 0.44, 0.40, 0.57) \cdot n_{0,e}$ are measured from the simulations.

and discussed in section 5.3.6. Even for this wide range of laser pulse intensities, simulations are in excellent agreement with theory in Eqs. (2.42) and (5.19) while older models show systematic discrepancies.

5.3.5. Phase Space Dynamics



(a) Longitudinal phase space for protons in blue, plotted above deuterium ions in orange. Vertical grey lines denote the initial target surfaces.



(b) On-axis field lineouts for densities, longitudinal electric field (purple) and relativistic critical density $\bar{\gamma}n_c$ (black). The initial target position is plotted in grey. In all areas in which the local electron density (red) is larger than the relativistic critical density, the target is overdense. The longitudinal electric field strength is plotted on the right ordinate and is strongly peaked at the ion fronts.

Figure 5.8.: Phase space and field lineouts for $a_0 = 16$ and 80 % deuterium content. Given time is relative to the laser pulse's peak intensity on target.

Figure 5.8 displays the ion phase space, particle densities and electric fields for $a_0 = 16$. The phase space evolution displayed in Fig. 5.8a shows that the lighter hydrogen species expands faster than the heavier deuterium ions from the target rear, consistent with Eq. (5.7). The proton spectrum is modulated for mid-range energies by the deuterium front, shifting affected protons to higher energies. As is visible in the field lineout for the same times in Fig. 5.8b, the peaked electric field (purple) causes a corresponding spatial modulation in the proton density distribution.

The density line-outs in Fig. 5.8b further show that the electron density (red line) remains higher than the relativistic critical density (black line), hence the target is opaque, during the entire interaction. Self-similar front and rear-side expansion removes most of the target material over time. The highest gain in ion momentum is reached for rear-side originating particles in TNSA fields. Although ions from the front half of the target reach the rear, they do not contribute to high energy portions of the spectrum for $a_0 = 16$. This behavior

changes with increasing laser intensity, discussed in the next section.

One can also draw a connection to the phase space signature of the mass-limited, spherical target in chapter 4, section 4.2.3. Consistency with the therein described proton "roll-up" signature is observed, as protons from deeper layers of the target are accelerated due to multi-species acceleration to higher energies. It is worth noting that the effect is stronger in chapter 4 as the laser pulse intensity is still rising over the acceleration time. As compared to multi-species effects observed from PMMA spheres, the cryogenic target setup allows careful tuning and isolated investigation of the spectral modulation.

5.3.6. Scaling with Laser Pulse Amplitude




Figure 5.9.: **(a)** Variation of hydrogen (sphere) and deuterium (triangle) cutoff energy with target deuterium ratio d and laser strength a_0 . Each H-D pair of given d and a_0 combination corresponds to one simulation (60 in total). A simple TNSA scaling estimate is plotted with a dashed black line for constant $n_e^{rear} = 0.35$ and $\langle T_e \rangle$ from the literature [51, 71]. **(b)** Linear zoom into hydrogen cutoff energies from (a) in terms of relative energy increase/decrease compared to a pure hydrogen target ($d = 0$) under variation of deuterium content d . With changed target composition, enhancements up to 18% and reduction by 28% in hydrogen cutoff energy are observed.

The laser field amplitude on target as well as the target mixing ratio d are varied, recording each simulation as sphere-triangle (H-D) pair in Fig. 5.9a. As before, the increase of deuterium cutoff energy with higher d holds true, from TW- to PW-scale laser intensities. Indicated with red markers, front-side originating ions start to determine the highest energies for $a_0 \gtrsim 30$. As the hole-boring velocity for the target front increases with the laser strength a_0 [114]

$$\frac{a_0}{1 + a_0} \propto v_{HB} \left[\frac{\mu\text{m}}{30 \text{ fs}} \right] = \begin{cases} 0.2 & \text{if } a_0 = 8, \\ 0.6 & \text{if } a_0 = 23, \\ 1.0 & \text{if } a_0 = 42, \end{cases} \quad (5.25)$$

front originating ions are initially faster with higher a_0 and reach the target rear fields in shorter time. In order to undergo further acceleration, they must reach the target rear before the TNSA fields have substantially decayed. The latter process occurs on a time scale of the laser pulse length $\tau_0 = 30$ fs. With a short traverse time in Eq. (5.25), front-side acceleration ($K_{\text{HB}} = 0.26 - 6.0$ MeV) which is followed by a subsequent energy boost to high energies in rear-side TNSA fields is feasible.

Even for the highest laser intensity regime, modeled with $a_0 = 42$, the target stays opaque, see Figure 5.10. This is consistent with the relativistic skin depth $\delta = 98$ nm, which is smaller than the $2 \mu\text{m}$ initial target thickness. Nevertheless, in the 2D plane of simulation, transmitted light is detected in localized filaments breaking through the target for $a_0 \gtrsim 20$.




Figure 5.10.: On-axis field lineouts for $a_0 = 42$ and 50 % deuterium content. The initial target position is plotted in grey. The target is locally overdense if the local electron density (red) is larger than the relativistic critical density (black). The longitudinal electric field is plotted on the right ordinate. Given time is relative to the laser pulse's peak intensity on target.

Coming back to hydrogen cutoff energies displayed in detail in Fig. 5.9b, a variation with deuterium ratio d is noticeable. As the heavier ion species only modulates mid-energy light ions, this is unexpected for TNSA acceleration. When comparing to existing literature [127], an overall decrease of proton energy to high ratios of d seems consistent. With the extensive study performed in this thesis, mixing ratios between $d = 0.2 - 0.5$ reveal target conditions in which proton energy is even increased by a maximum of 18 %. This secondary effect might originate in a change of laser absorption at the front surface layer. An overall decrease for high d might be the result of a more rigid surface with higher deuterium content. An intermediate increase of absorption is likely attributed to a changed surface roughness, as front-side micro-structuring and plasma filaments are observable in most simulations, which is a realistic 2D effect [104].

One can confirm the aforementioned assumption on absorption variations with measured electron properties. Therefore, Figure 5.11 shows the laser-accelerated electron average kinetic energy $\langle T_e \rangle$ (black) and corresponding electron density n_e (blue). The average kinetic energy is furthermore compared to the theoretical scaling in Eq. (2.35), which agrees well and only underestimates the $a_0 = 42$ case as the theoretical model assumes an opaque target surface. For the $a_0 = 8$ case, the method of measurement (see appendix A.3) might underestimate the influence of thermalized bulk electrons.

As known from TNSA theory in Eq. (2.45), the proton cutoff energy depends approximately on the laser-accelerated electron energy density $\langle w_e \rangle = \langle T_e \rangle \cdot n_e$ in the electron sheath. The energy density of laser-accelerated electrons (red points) follows the trend of hydrogen

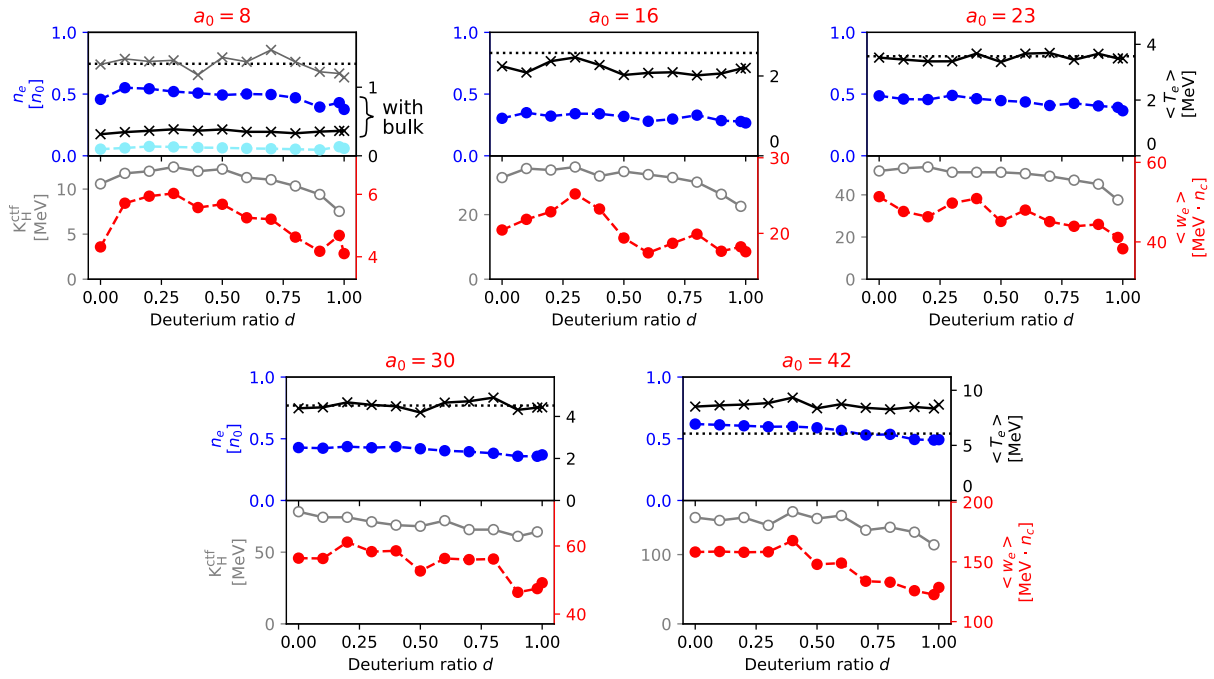


Figure 5.11.: Measured average laser-accelerated electron kinetic energy $\langle T_e \rangle$, corresponding electron density n_e , and energy density $\langle w_e \rangle$. For comparison with the average electron energy density trend, the maximum proton energy K_H^{ctf} is displayed. Results are presented as groups of laser pulse strength a_0 and within those as function of target mixing ratio d . For $a_0 = 8$, both laser-accelerated as well as target bulk electrons are presented, indicating a contribution of thermalized electrons to the acceleration process [334]. For comparison, the black dashed line shows the theoretical scaling from Eq. (2.35).

cutoff energy K_H^{ctf} (grey points) to reasonable extent, supporting that a change in absorption dynamics for variations in mixing ratio d is likely causing the variations in proton cutoff energies.

5.3.7. Ion Emission Characteristics

Assessing whether multi-species effects are observable in a potential experimental realization, a sensitivity analysis is performed with simulations. As already demonstrated in section 5.3.4, high deuterium content creates a modulation in the hydrogen spectrum in the mid to high energy part. Particle numbers vary by one to three orders of magnitude per MeV and steradian for an on-axis pinhole aperture with $\pm 2^\circ$ opening angle, which is well observable with the dynamic range of a Thomson parabola or an RCF stack.

Figure 5.12 shows the radial emission distribution in the plane of laser polarization. When comparing ion pointing for varying laser field strengths, both hydrogen and deuterium emission distributions are more directed for low a_0 than higher a_0 . In this specific setup, front-side protons add a diffuse emission signature to the overall distribution and emit into wider angles. The emission characteristics of hole-boring, pre-accelerated ions from the target front surface (e.g. wider or narrower than the TNSA rear and structuring) depends on the given temporal laser contrast and target pre-expansion and requires a detailed param-

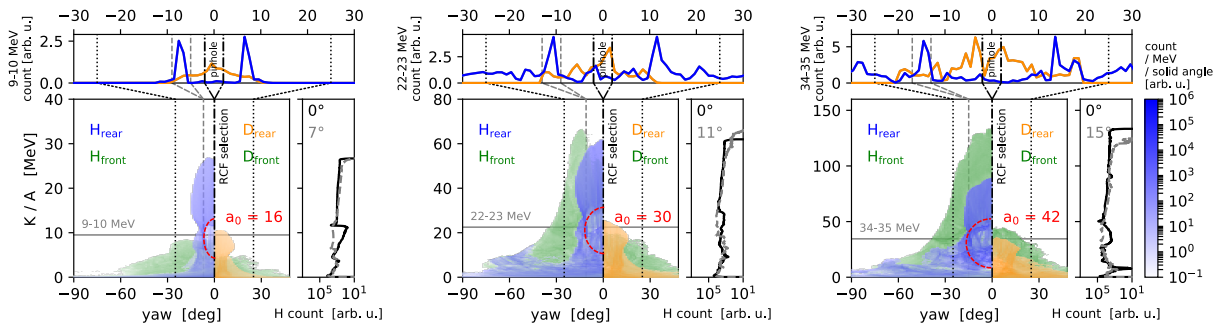


Figure 5.12.: **Lower plots:** Angular emission distribution of ions plotted for laser intensities $a_0 = 16, 30, 42$ (left-right) and target deuterium ratio $d = 90\%$. Left half is hydrogen (blue) and right half deuterium ions (orange) with green denoting front-originating ions in both cases. Rear-side ions are plotted on top of front-side ions. For large enough angles to the laser pulse propagation axis and energy, one can infer that only front-side ions contribute. **Upper plots:** A sample lineout, zoomed into smaller emission angles, shows the emission characteristics for a selected energy ($\Delta K/A = 1$ MeV), similar to an RCF layer (summed over front and rear contributions). **Side plots on right:** Energy lineout for pointing sensitivity with a pinhole aperture taken for acceptance on axis (black line: $(0 \pm 2)^\circ$) and off-axis (grey line: $(7, 11, 15 \pm 2)^\circ$), similar to a typical Thomson parabola.

eter scan for a given experiment. Note that the absolute values for energy and emission angle are expected to be smaller in 3D3V simulations and experiments, an experimental identification of pure front-side contributions might be possible when clear structural changes in the emission characteristic are visible. With increasing deuterium content in the target, the visible "hole" (marked red in Fig. 5.12) in the mid-energy range protons shifts to higher energies as more deuterium ions can displace the decreasing number of protons more efficiently. Experiments are expected to detect a spectral "gap" at low energies and quasi-monoenergetic features when observing strictly along the target normal with narrow angle acceptance.

When measuring ion spectra experimentally with RCF stacks, protons deposit their energy on a series of stacked radio-chromatic film and absorber layers. In order to separate various ion species, one might need to divert the ion beam in an electric field to differentiate q/m populations. After subtracting contributions on film layers traversed by higher energy protons and correcting for the according activation, one obtains radially resolved proton counts in an energy band. Such a band is exemplified as lineout in the upper plots of Figure 5.12, which will show the aforementioned lateral "hole" the in proton signal.

With small acceptance aperture, such as the $\pm 2^\circ$ in a typical Thomson parabola, beam jitter might have to be taken into account. Under limited cryogenic target orientation stability, such measurements might suffer from shot-to-shot fluctuation of target normal alignment with respect to the fixed diagnostic axis [44]. The lineouts on the right side of each plot in Fig. 5.12 exemplify how a positioning jitter influences the measured spectra. For example with $a_0 = 16$ and 7° off-axis beam-pointing, the modulation signal might not be visible. Experimental campaigns can mitigate this issue by wider angular pinhole acceptance, which may reduce the peak-to-valley ratio of the spectral modulation due to averaging, or by collecting enough statistics with a sufficient number of laser shots.

5.4. Summary and Outlook

In this chapter, an analytical model was derived predicting the spectral position of multi-species features in laser-ion acceleration in the TNSA regime [254]. Due to charge shielding of lighter ions such as protons, heavier species shift mid-energy protons to higher energies which results in a spectral dip of proton numbers at the energy of the heavier species cutoff energy per nucleon. With existing single-species theories for target electron average energy and density, the predicted spectral position of multi-species modulation complements models for proton cutoff energies. As such, one can connect macroscopic, experimentally accessible energy spectra with underlying microscopic, hardly accessible properties of the laser-driven electron population, in particular average kinetic energy and density. This allows a higher quality of predictions, as assumptions for both cutoff and spectral dip have to match self-consistently.

Furthermore, a potential experimental realization with a readily available target [44, 46] and laser system at HZDR is prepared via a systematic simulation scan in the applicable parameter range. Given that cryogenic deuterium-hydrogen jets are technologically feasible, the presented simulations predict the straight-forward detection of multi-species signatures with existing ion diagnostics. For flexibility, the simulations cover a broad set of setups, allowing to chose a subset that can be realized experimentally. For detailed predictions of realistic ion energies, additional 3D3V modeling and detailed studies on the response to temporal laser contrast and target pre-plasma are advisable.

With regards to the described spectral signatures, a review of previously conducted experiments which were evaluated concerning novel acceleration mechanisms from homogeneously mixed targets might be insightful. Following studies could explore the sensitivity of the predicted laser-accelerated electron properties from the proton spectral modulation position and the cutoff energy. Furthermore, the theory and studies in this chapter also serve as a basis for future adoptions, e.g. for thin targets. As such, modifications to the presented integration ranges can be applied.

6. Conclusions and Perspectives

In this work, predictive capabilities for studying laser-driven plasmas were established by means of computational modeling and theory. In order to achieve control of ultrafast processes, the world's fastest electro-magnetic particle-in-cell code¹ PIconGPU was developed into a scientific instrument suitable for productive research in laser-plasma physics. Generalized and future proof computational methods and data analysis workflows have been realized, addressing current and upcoming challenges in high-performance computing and reproducible simulations. Implemented methods are abstracted to be flexible and extensible for current and upcoming manycore hardware and programming models [189, 214, 215].

Furthermore, an international community around open science workflows, open data and open source was established that is able to collectively approach future developments. Within this community, student projects were mentored to transfer knowledge and improve existing capabilities. The author devised openPMD [272, 280], an open data standard that is now adopted widely in the laser-plasma accelerator community and starts being adopted in photon-science [206, 207], accelerator physics [285] and astrophysics [231, 284]. For accompanying data-driven, high-throughput computing in simulations and experiments, data pipelines and algorithms that are suitable for heterogeneous computing needed to be re-designed. A theoretical scaling model defining requirements for data reduction algorithms in current and upcoming data pipelines was derived, which allows to predict algorithm performance and applicability before running large scientific campaigns at scale [218]. With the latter, a special emphasis is set on interactive workflows and synthetic diagnostics, mirroring experimental observables in virtual setups.

As the author led the scientific development of PIconGPU, its first production-scale 3D3V simulations were performed for laser-ion acceleration [11], utilizing world-class supercomputers powered by graphics processing units (GPUs) [152, 197, 218], which are promising candidates for Exascale high-performance computing (HPC) in upcoming years. With such simulations, urgent questions in contemporary laser-plasma experiments were answered [11, 35, 44, 46]. One campaign focused on a novel, micron-scale, fully isolated target whose

¹in terms of sustained Flop/s [197]

unexpected pre-expansion dynamics under realistic temporal laser contrast led to the discovery of a regime of highly directed laser-ion acceleration with PW-class laser systems [11]. 3D3V modeling with PByte-scale data processing was able to confirm the experimental results and created the basis for a systematic understanding of underlying microscopic processes. Non-optimal temporal laser contrast portions before the main pulse initiate target expansion to a near-critical target density, which favors a volumetric laser-target interaction, generating highly charged, directed, quasi-monoenergetic proton bunches [11]. In future research, ultrashort-pulse campaigns with controlled pre-expansion might allow systematic understanding of the initial plasma conditions and first 3D simulation results with high contrast predict promising maximum proton energies in the range of 30 MeV for 150 TW DRACO laser parameters.

Near-critical and mass-limited target geometries were further studied with a novel, debris free, cryogenic jet target at Helmholtz-Zentrum Dresden – Rossendorf (HZDR) [44, 46]. With systematic simulations from TW- to PW-class laser intensities, good agreement between experiments and simulations under geometric variations were achieved. It is further explored how this target system can be adopted for a dedicated study of multi-species effects. The latter occur in virtually all but pure hydrogen targets used today due to different target constituents and ion charge states. With a deuterium-hydrogen gas fed jet, the ratio of the target constituents is highly tunable, while the influence of hardly controllable parameters, such as ionization dynamics and contamination layers, is suppressed. A novel theoretical model for the location of spectral modulations with such multi-component materials was derived, that are expected to be detectable experimentally in such a setup. The proposed model agrees well with our simulation results and as such is suitable to draw conclusions about microscopic electron ensemble properties based on macroscopic experimental observables [254].

Future research might focus on an even closer integration of simulation and experimental campaigns. One aspect includes simulation readiness for planned campaigns at upcoming light sources such as the European XFEL, adding code extensions for advanced physical models for X-ray interaction. Experiments under high-repetition rates will improve systematic understanding and rising experimental data rates can reuse data analysis methods already utilized for large-scale simulations in order to extract maximum information. With respect to exponentially growing computational capabilities, near-term supercomputers will permit to run ensemble studies with accelerated algorithms of this thesis in full resolution and geometry. The physical and computational studies presented in this thesis demonstrate and establish fast turn-around, highly scalable simulation workflows [11, 36, 218, 254]. Such studies should be treated like virtual beam times in experiments, planned in parallel and run at the same time in order to stimulate each other in a feedback loop. Modeling in simulations what is measured in experiments means exploring the influence of uncertainties especially introduced by temporal laser contrast, pulse intensity and target conditions with reliable simulation pipelines. As demonstrated in first start-to-end simulation coupling [206, 207], this goal requires virtual beam lines starting from measured laboratory conditions to virtual detectors, including limitations such as pointing jitter and temporal signal integration.

As demonstrated in this work, closely intertwined experiments and predictive simulations are the key to a quantitative microscopic description of laser-proton acceleration. For such

integrated laser-plasma campaigns, appropriate software was introduced in this thesis that allows to engage with upcoming Exascale supercomputers, enabling many parallel simulations at full dimensionality on the timescale of an experimental run. Such multidisciplinary efforts are central drivers in transferring laser-proton acceleration to a reliable technology.

A. Appendix

A.1. Template Meta-Programming Examples

A.1.1. Generalized Particle Shapes and Current Deposition

The following example shows generic particle-in-cell (PIC) shape implementation, without code duplication, and its generic use for current deposition. Optimizations such as perfect inlining and vectorization, due to known sizes at compile time, are detectable by or can be hinted at the compiler with pragmas.

Listing A.1: C++11 code example implementing a numerical particle shape.

```
namespace picongpu {
namespace particles {
namespace shapes {
    struct PCS {
        static constexpr int support = 4;

        struct ChargeAssignment {
            inline float_X operator()( float_X const x );
        };

        struct ChargeAssignmentOnSupport /* ... */ ;
    };
} // namespace shapes
} // namespace particles
```

Listing A.2: C++11 code example with compile-time code generation for current deposition from a policy of the numerical particle shape order.

```
namespace picongpu {
namespace currentSolver {
    template< typename T_ParticleShape >
    struct Esirkepov< T_ParticleShape, DIM3 > {
        using Assignment = typename T_ParticleShape::ChargeAssignment;
        static constexpr int supp = Assignment::support;
    };
}
```

```

static constexpr int lowerMargin = - supp / 2 + (supp + 1) % 2;
// ...
{
    // optional: unroll or vectorize the loop of known size
    for( int i = lowerMargin ; i < upperMargin; ++i )
        // ...
        // evaluate Assignment::operator()
    }
    // ...
};
} // namespace picongpu

```

A.1.2. Particle Species and Compile-Time Code Generation Switch

The following example demonstrates a compile-time switch (trait) that determines if code shall be generated or not. Such traits allow generic implementations of physical algorithms in declarative style, announcing requirements for allowed combinations with template policies. When generating the code, e.g. via a compile-time loop over a type list, the trait can be queried for valid combinations and code generation for invalid species-solver pairs is silently skipped (instead of aborting in a compile error).

Listing A.3: C++ code example with compile-time trait to determine if a particle species is eligible for filtering by pointing of its momentum vector.

```

namespace picongpu {
namespace particles {
namespace traits {
    // in principle, all particle species and solvers
    // are compatible
    template< typename T_Species, typename T_Solver >
    struct SpeciesEligibleForSolver
    {
        using type = boost::mpl::bool_< true >;
    };

    // additional constraints for calculating a
    // "ParticlesForwardPinhole" filter:
    // particle species must have momentum attribute
    template< typename T_Species >
    struct SpeciesEligibleForSolver <
        T_Species,
        filter::ParticlesForwardPinhole
    > {
        using type = typename pmacc::traits::HasIdentifiers <
            typename T_Species::FrameType,
            MakeSeq_t< momentum >
        >::type;
    };
} // namespace traits

// ... further code ...

```

```

// query for species electrons, probes and momentum filter
using ReadyForDetector = typename traits :: SpeciesEligibleForSolver <
    Electrons, filter :: ParticlesForwardPinhole >::type; // true
using ReadyForDetector = typename traits :: SpeciesEligibleForSolver <
    Probes, filter :: ParticlesForwardPinhole >::type; // false

} // namespace particles
} // namespace picongpu

```

A.2. Phase Space Evolution

Complementary to section 4.2.3, a series of longitudinal phase space images is presented in Fig. A.1 which leads to a quasi-monoenergetic, highly charged proton bunch (see chapter 4).

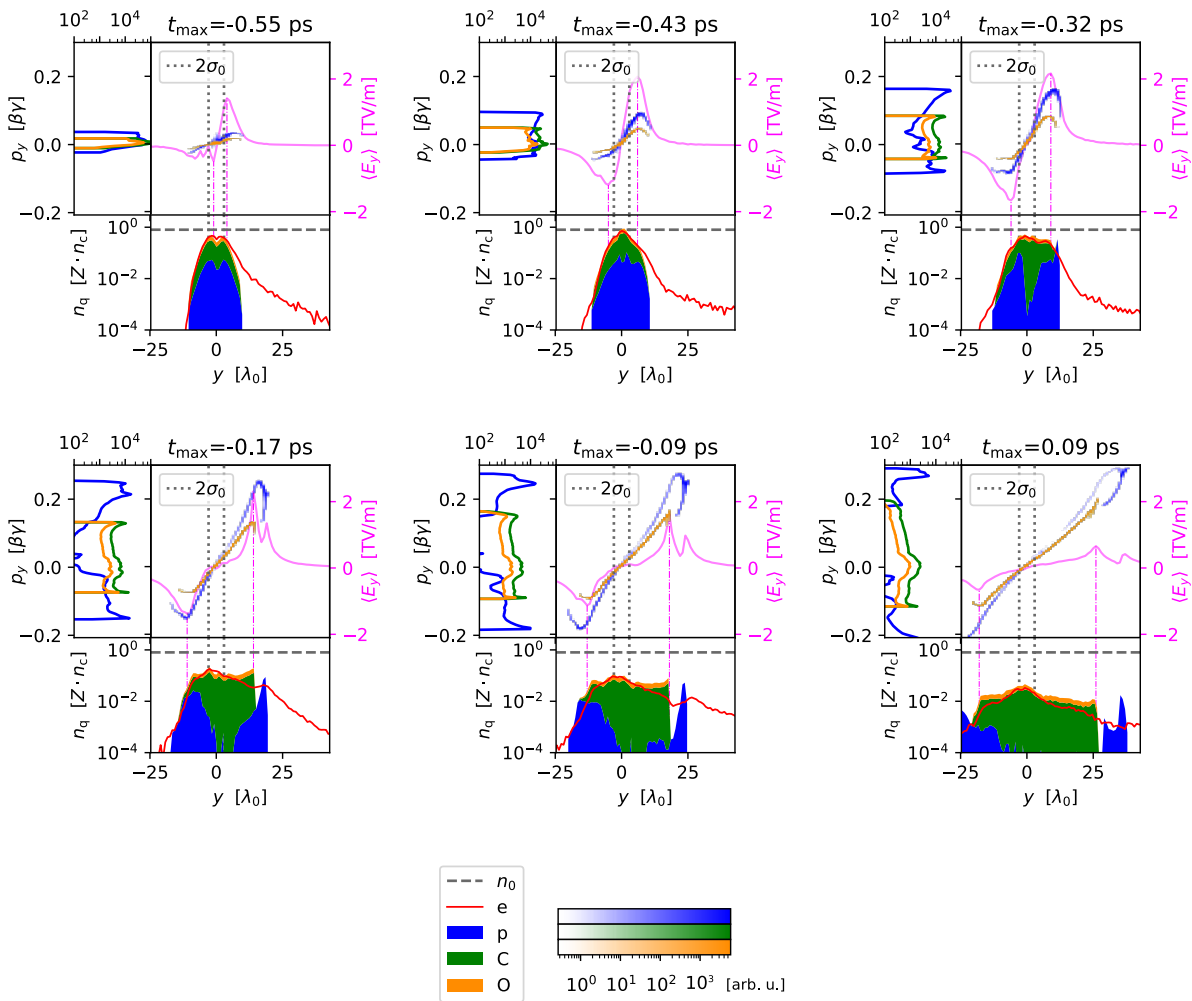


Figure A.1.: Temporal evolution of the longitudinal ion phase space as in Fig. 4.7b. Each figure is structured in (central) phase space with longitudinal electric field, (left) momentum histogram, and (bottom) stacked charge density. $2\sigma_0$ marks the initial target size. All data is taken along the laser propagation axis y , perpendicularly averaged for one λ_0 .

A.3. Measuring Prompt Electrons in a Simulation

As introduced in section 2.2.2, target normal sheath acceleration (TNSA) theories for relativistic laser strength a_0 often require an estimate of the prompt, non-thermal electron distribution in the target [51, 73]. In order to calculate the prompt average kinetic energy T_e and density n_e^{prompt} from electrons in a simulation, various methods are used by authors, yet rarely documented. Some apply one of the theoretical models for T_e [51] and assume the absorption in the *whole* target estimates the corresponding density $\eta \approx n_e^{\text{prompt}}/n_{0,e}$.

The following method turns out to be robust and is therefore documented for future readers. For thick-enough targets, one can focus analysis on the spatial rear half of the target to electrons that drive the expansion. Furthermore, depending on the divergence of the electron beam, one further selects a subset of electrons which are close to the laser propagation axis. In the analysis of chapter 5, a wide selection around this axis (up to the focal spot size) for high $a_0 \gtrsim 20$ and narrower selection for low $a_0 \lesssim 20$ were suitable for flat foils, but will likely depend on specific target geometry. Measuring shortly after the laser pulse's peak intensity reaches the target ensures the highest occurring electron energies are included. In all further analysis one may consider the longitudinal momentum component, as it is driving the target expansion. Cropping for relativistic electrons with $\beta_e \gamma_e \gtrsim 1$ excludes contributions from the "cold" electron bulk and return currents. Assuming that the resulting spatial average of electrons resembles in good approximation still a Boltzmann distribution [66, 77]

$$\frac{dN}{dE} = f(E) = \frac{N_e}{k_B T_e} \cdot \exp\left(-\frac{E}{k_B T_e}\right), \quad (\text{A.1})$$

one can calculate the average kinetic energy and density for the cropped distribution, e.g. from an energy histogram with m bins containing N (real) electrons. With discrete bins with width ΔE_i and count N_i follows

$$\frac{N_i}{\Delta E_i} = f_i(E_i) = \frac{N_e}{k_B T_e} \cdot \exp\left(-\frac{E_i}{k_B T_e}\right), \quad (\text{A.2})$$

$$\Delta E_i = E_{i+1} - E_i, \quad (\text{A.3})$$

$$N_e = \sum_{i=0}^m N_i. \quad (\text{A.4})$$

Alternatively, a measurement in a simulation may start from a (spatially filtered) phase space histogram as in Fig. A.2. In such a case, one normalizes the counts per bin via the




Figure A.2.: Electron phase space for a near-critical flat hydrogen target with $n_e = 30 n_c$, thickness $d_0 = 2 \mu\text{m}$ and a laser pulse with amplitude $a_0 = 23$ and length $\tau_0^{\text{FWHM},I} = 30 \text{ fs}$ (same conditions as in chapter 5). The snapshot is taken at time $t = 6 \text{ fs}$, relative to the laser peak intensity on target.

relativistic energy-momentum relation to energy bins

$$E_i = m_e c^2 \cdot \left(\sqrt{1 + \left(\frac{p_{y,i}}{m_e c} \right)^2} - 1 \right), \quad (\text{A.5})$$

$$\Delta p_{y,i} = p_{y,i+1} - p_{y,i}, \quad (\text{A.6})$$

$$\frac{N_i}{\Delta E_i} = \frac{N_i^{\text{mom.}}}{\Delta p_{y,i}} \cdot \frac{\Delta p_{y,i}}{\Delta E_i}. \quad (\text{A.7})$$

Note that bins for negative momentum should be excluded. Then calculate the first moment (mean) of the cropped probability distribution $f_i(E_i)$ for energy $\Delta V \cdot w_e^{\text{crop}}$ and number of electrons N_e^{crop}

$$\begin{aligned} N_e^{\text{crop}} \cdot k_B T_e^{\text{crop}} &= \sum_{i=j}^m E_i \cdot f_i(E_i) \cdot \Delta E_i \\ &= \Delta V \cdot w_e^{\text{crop}}, \end{aligned} \quad (\text{A.8})$$

$$N_e^{\text{crop}} = \sum_{i=j}^m N_i = \sum_{i=j}^m f_i(E_i) \cdot \Delta E_i. \quad (\text{A.9})$$

Here, j is the lower bin with energy E_j which fulfills the requirement $\beta_i \gamma_i > 1$ for cropping and ΔV is the spatial selection (on-axis, rear-target average) described above. The expressions

$$\sum_{i=j}^m \dots \cdot \Delta E_i \hat{=} \int_{E_j}^{\infty} \dots \cdot dE \quad (\text{A.10})$$

denote discrete integrals and should be calculated with, e.g. trapezoidal or Simpson's rule.

The parameters of the underlying Boltzmann distribution can be decomposed as follows.

$$N_e \cdot k_B T_e = \Delta V \cdot w_e = \sum_{i=0}^{j-1} E_i \cdot f_i(E_i) \cdot \Delta E_i + \sum_{i=j}^m E_i \cdot f_i(E_i) \cdot \Delta E_i \quad (\text{A.11})$$

$$= N_e \cdot k_B T_e - N_e \cdot (E_j + k_B T_e) \cdot \exp\left(-\frac{E_j}{k_B T_e}\right) + \frac{\Delta V \cdot w_e^{\text{crop}}}{}, \quad (\text{A.12})$$

$$N_e = \sum_{i=0}^{j-1} f_i(E_i) \cdot \Delta E_i + \sum_{i=j}^m f_i(E_i) \cdot \Delta E_i \quad (\text{A.13})$$

$$= N_e \cdot \left[1 - \exp\left(-\frac{E_j}{k_B T_e}\right)\right] + \frac{N_e^{\text{crop}}}{}. \quad (\text{A.14})$$

Finally, the missing low-energy contribution due to our spectral cutoff in E_j can be accounted for from the measured properties of the cropped distribution in Eqs. (A.8) and (A.9)

$$T_e = \frac{\Delta V \cdot w_e^{\text{crop}}}{N_e^{\text{crop}}} - E_j, \quad (\text{A.15})$$

$$N_e = N_e^{\text{crop}} \cdot \exp\left(\frac{E_j}{T_e}\right), \quad (\text{A.16})$$

$$n_e = \frac{N_e}{\Delta V}. \quad (\text{A.17})$$

In highly relativistic scenarios, electrons measured in such manner provide results that fit well with theory for both electron average kinetic energy and ion cutoff in flat foils [51, 71]. Nevertheless, for moderate $a_0 \lesssim 10$, one observes that only accounting for prompt electrons underestimates the overall electron density (and energy density) with respect to observed ion cutoff energy. In such cases, bulk electrons contribute significantly to the expansion dynamics. Average kinetic energy and density should then be estimated for the latter as well (with upper-energy range cropping) and both distributions might be combined with proper statistical averaging. Example scripts for the relativistic case are provided in Ref. [334].

Glossary

abstract syntax tree directed graph representing the abstract structure of source code within a compiler, *see* compiler & source code

API application programming interface, the user-facing part a software that is designated for accessing its implemented functionality

architecture a set of capabilities and rules of a certain computing device, *see* platform, CPU & GPU

ASE amplified spontaneous emission

assembly low level code instructions with close relation to the executing hardware, usually generated by a compiler, *see* compiler

atomic operation parallel computing operation that locks a shared resource exclusively during access

benchmark a defined test case measuring correctness, precision and/or performance of an executed algorithm or workflow, *see* HPC & platform

cache storage to serve frequently accessed data in shorter time compared to alternatives available in a system

CERN European Organization for Nuclear Research, a European research organization

for particle physics located in Geneva

CFL Courant-Friedrichs-Lewy

CGS centimetre-gram-second system of units, a unit system

CIC cloud-in-cell

cluster arrangement of interconnected computers that work collaboratively to solve a common problem

compiler computer program that translates program code from one language to another, usually converts a high-level language to specific machine code, *see* source code & assembly

core a processing unit inside a processor, *see* architecture

CPA chirped pulse amplification, Nobel Price-winning technique for amplification of an ultrashort laser pulse via temporal stretching, amplification and recompression

CPU central processing unit, *see* multicore

data layout arrangement of compound and array variables for efficient processing

data packing copying to or designing of a specific data layout, usually to increase read and write performance or to achieve vectorization, *see* data layout & vectorization

DFT discrete Fourier transform

DOI digital object identifier

double a floating point type occupying 64 bits in memory, formally specified in IEEE 754, *see* single & RAM

DRACO Dresden laser acceleration source, an ultrashort pulse highpower laser at HZDR, *see* laser & HZDR

European XFEL research laboratory capable of providing X-ray laser pulses for experiments; located in Schenefeld, Germany, *see* FEL & laser

FDTD finite-difference time-domain

FEL free-electron laser, photon beam with laser-properties generated from relativistic electrons in an undulator structure

filesystem non-volatile storage for data, usually with parallel access on a cluster, *see* cluster

Flop/s a unit counting multiplication-equivalent floating-point operations per time on a computer hardware, usually refers to double-precision performance, *see* architecture & double

framework software product and environment providing a complete collection of APIs, compiler, tools and workflows to build a specific application, *see* API, source code, compiler & platform

function inlining compiler optimization technique replacing a function call with the instructions of the function body, *see* compiler, assembly & instruction

fused multiply-add a floating-point operation performing a multiplication and addition in one step, *see* instruction

GPGPU general-purpose GPU, graphics card that is also capable of performing general purpose computations; all GPUs in considered in this thesis are GPGPUs, *see* many-core & GPU

GPU graphics processing unit

GSI Gesellschaft für Schwerionenforschung, a German research laboratory of the

Helmholtz Association of German Research Centres located in Darmstadt

HED high energy density, term for matter and radiation physics with energy densities in excess of approximately 100 GJ/m^3

HIBEF the Helmholtz International Beamline for Extreme Fields at the European XFEL is an international collaboration led by HZDR; provides an endstation for experiments with optical, short-pulse high-power lasers and the European XFEL, *see* FEL & European XFEL

HPC sub-domain of computer science solving research questions with supercomputers, *see* cluster

HPCG High Performance Conjugate Gradients, a performance benchmark testing the data access patterns with communication intensive sparse matrix operations, *see* HPC, cluster & benchmark

HPL High-Performance Linpack, a performance benchmark solving compute intensive dense matrix operations, *see* HPC, cluster & benchmark

HZDR Helmholtz-Zentrum Dresden – Rossendorf, a German research laboratory of the Helmholtz Association of German Research Centres located in Dresden

I/O input/output, short-hand for communication in a computing system, refers to memory transfer in general and filesystem operations in particular, *see* filesystem & RAM

instruction command that can be executed by a processor, *see* assembly & architecture

instruction pipeline execution technique of modern processors overlapping instructions such as fetch, decode, execute and store of independent steps in order to increase throughput, *see* architecture & instruction

ISQ international system of quantities, definition of base quantities for natural sciences and technology

JIT just-in-time, computer program execution strategy in which source code is translated at runtime directly before execution, *see* & compiler

kernel vectorized function, executed on a manycore device, *see* vectorization, manycore & GPU

laser light amplification by stimulated emission of radiation, device that emits spatially and temporally coherent electro-magnetic radiation (light) through optical amplification and stimulated emission

LMU Ludwig-Maximilians-Universität München, a German university located in Munich

LWFA laser wakefield acceleration

manycore computing processors with hundreds to thousands of processor cores, *see* GPGPU

memory container class concept in C++ that safely organizes memory and provides functionality to manipulate therein owned data, *see* data layout

memory heap memory pool providing functionality for dynamic allocations and deallocations of data at program runtime, *see* RAM

MPI message-passing interface, library interface specification for communication between multiple nodes in HPC, *see* HPC, API & cluster

multicore computing processors with more than one processor core, *see* CPU

NGP nearest-grid-point

ORNL Oak Ridge National Laboratory, an American national laboratory located in Oak Ridge, Tennessee

P4S piecewise quartic shape

P5S piecewise quintic shape

PCS piecewise cubic shape

PDE partial differential equation

PENELOPE Petawatt Energy-Efficient Laser for Optical Plasma Experiments, an upcoming ultrashort pulse highpower laser at HZDR, *see* laser & HZDR

performance portability ability to run a software efficiently on substantially different computing architectures from a single source code without algorithmic duplications, *see* HPC & architecture

PHELIX Petawatt Hoch- Energie Laser für Schwerionenexperimente, a high-power laser at GSI, *see* laser & GSI

PIC particle-in-cell

platform the software and hardware environment in which a computer program is executed, *see* architecture

PMMA polymethyl methacrylate, a material also known as acrylic glass

PQS piecewise quadratic cloud shape

PQS piecewise quadratic shape

processor register a small amount of fast, core-local storage storing either variables or instructions; some registers also provide specific processor functionality, *see* architecture & instruction

programming model combination of a programming language, auxiliary functionality provided in a (runtime) API and workflows on how to express algorithms efficiently, *see* source code, API & compiler

PTX parallel thread execution language, an assembly-like language used as intermediate representation for Nvidia GPUs, *see* GPGPU & assembly

QED quantum electrodynamics

RAM random-access memory, volatile memory with fast connection to a computing device and support for non-contiguous access to data

RCF radiochromic film

REPL computing environment that accepts source code expressions interactively, evaluates them and returns the results with-

out leaving the running program, *see* source code, compiler & JIT

RPA radiation pressure acceleration

SI international system of units, the unit system used in this thesis

SIMD executing a single-instruction on multiple-data, *see* instruction & vectorization

single a floating point type occupying 32 bits in memory, formally specified in IEEE 754, *see* double & RAM

source code high-level instructions for a computer program (software) as written by

a programmer, *see* compiler

time-to-solution required time for a computer program to present a solution, *see* HPC

TNSA target normal sheath acceleration

translation unit single source code file that can be compiled at once, *see* compiler & source code

TSC triangular-shaped cloud

vectorization execution of an instruction for an array of values instead of a single scalar data, *see* instruction & SIMD

Bibliography

- [1] SP Hatchett, CG Brown, TE Cowan, et al. "Electron, photon, and ion beams from the relativistic interaction of Petawatt laser pulses with solid targets". In: *Phys. Plasmas* 7.5 (2000), pp. 2076–2082. DOI: 10.1063/1.874030.
- [2] RA Snavely, MH Key, SP Hatchett, et al. "Intense high-energy proton beams from petawatt-laser irradiation of solids". In: *Phys. Rev. Lett.* 85.14 (2000), pp. 2945–2948. DOI: 10.1103/PhysRevLett.85.2945.
- [3] BM Hegelich, BJ Albright, J Cobble, et al. "Laser acceleration of quasi-monoenergetic MeV ion beams". In: *Nature* 439.7075 (2006), pp. 441–444. DOI: 10.1038/nature04400.
- [4] H Schworer, S Pfoth, O Jäkel, et al. "Laser-plasma acceleration of quasi-monoenergetic protons from microstructured targets". In: *Nature* 439.7075 (2006), pp. 445–448. DOI: 10.1038/nature04492.
- [5] T Ditmire, J Tisch, E Springate, et al. "High-energy ions produced in explosions of superheated atomic clusters". In: *Nature* 386.6620 (1997), pp. 54–56. DOI: 10.1038/386054a0.
- [6] SC Wilks, AB Langdon, TE Cowan, et al. "Energetic proton generation in ultra-intense laser–solid interactions". In: *Phys. Plasmas* 8.2 (2001), pp. 542–549. DOI: 10.1063/1.1333697.
- [7] SA Gaillard, T Kluge, KA Flippo, et al. "Increased laser-accelerated proton energies via direct laser-light-pressure acceleration of electrons in microcone targets". In: *Phys. Plasmas* 18.5 (2011), p. 56710. DOI: 10.1063/1.3575624.
- [8] F Wagner, O Deppert, C Brabetz, et al. "Maximum Proton Energy above 85 MeV from the Relativistic Interaction of Laser Pulses with Micrometer Thick CH₂ Targets". In: *Phys. Rev. Lett.* 116.20 (2016), pp. 1–5. DOI: 10.1103/PhysRevLett.116.205002.
- [9] JB Kim, S Göde, and SH Glenzer. "Development of a cryogenic hydrogen microjet for high-intensity, high-repetition rate experiments". In: *Rev. Sci. Instrum.* 87.11 (2016), 11E328. DOI: 10.1063/1.4961089.

- [10] A Higginson, RJ Gray, M King, et al. "Near-100 MeV protons via a laser-driven transparency-enhanced hybrid acceleration scheme". In: *Nat. Commun.* 9.1 (2018), p. 724. DOI: 10.1038/s41467-018-03063-9.
- [11] P Hinz, TM Ostermayr, A Huebl, et al. "Isolated proton bunch acceleration by a petawatt laser pulse". In: *Nat. Commun.* 9.1 (2018), p. 423. DOI: 10.1038/s41467-017-02663-1.
- [12] SV Bulanov and VS Khoroshkov. "Feasibility of Using Laser Ion Accelerators in Proton Therapy". In: *Plasma Phys. Reports* 28.5 (2002), pp. 453-456. DOI: 10.1134/1.1478534.
- [13] B Damato, A Kacperek, M Chopra, et al. "Proton beam radiotherapy of choroidal melanoma: The Liverpool-Clatterbridge experience". In: *Int. J. Radiat. Oncol. Biol. Phys.* 62.5 (2005), pp. 1405-1411. DOI: 10.1016/j.ijrobp.2005.01.016.
- [14] KWD Ledingham, W Galster, and R Sauerbrey. "Laser-driven proton oncology - A unique new cancer therapy?" In: *Br. J. Radiol.* 80.959 (2007), pp. 855-858. DOI: 10.1259/bjr/29504942.
- [15] SD Kraft, C Richter, K Zeil, et al. "Dose-dependent biological damage of tumour cells by laser-accelerated proton beams". In: *New J. Phys.* 12.8 (2010), p. 085003. DOI: 10.1088/1367-2630/12/8/085003.
- [16] K Zeil, M Baumann, E Beyreuther, et al. "Dose-controlled irradiation of cancer cells with laser-accelerated proton pulses". In: *Appl. Phys. B Lasers Opt.* 110.4 (2013), pp. 437-444. DOI: 10.1007/s00340-012-5275-3.
- [17] BA Remington, D Arnett, RP Drake, et al. "Modeling Astrophysical Phenomena in the Laboratory with Intense Lasers". In: *Science* 284.5419 (1999), pp. 1488-1493. DOI: 10.1126/science.284.5419.1488.
- [18] SV Bulanov, TZ Esirkepov, M Kando, et al. "On the Problems of Relativistic Laboratory Astrophysics and Fundamental Physics with Super Powerful Lasers". In: *Plasma Phys. Reports* 41.1 (2015), pp. 1-51. DOI: 10.1134/S1063780X15010018.
- [19] M Roth, A Blazevic, M Geissel, et al. "Energetic ions generated by laser pulses: A detailed study on target properties". In: *Phys. Rev. Spec. Top. - Accel. Beams* 5.6 (2002), pp. 31-38. DOI: 10.1103/PhysRevSTAB.5.061301.
- [20] M Borghesi, AJ Mackinnon, DH Campbell, et al. "Multi-MeV Proton Source Investigations in Ultraintense Laser-Foil Interactions". In: *Phys. Rev. Lett.* 92.5 (2004), p. 055003. DOI: 10.1103/PhysRevLett.92.055003.
- [21] M Roth, TE Cowan, MH Key, et al. "Fast ignition by intense laser-accelerated proton beams". In: *Phys. Rev. Lett.* 86.3 (2001), pp. 436-439. DOI: 10.1103/PhysRevLett.86.436.
- [22] M Koenig, A Benuzzi-Mounaix, A Ravasio, et al. "Progress in the study of warm dense matter". In: *Plasma Phys. Control. Fusion* 47.12 B (2005). DOI: 10.1088/0741-3335/47/12B/S31.
- [23] M Tabak, P Norreys, VT Tikhonchuk, et al. "Alternative ignition schemes in inertial confinement fusion". In: *Nucl. Fusion* 54.5 (2014), p. 054001. DOI: 10.1088/0029-5515/54/5/054001.

- [24] P McKenna, KW Ledingham, T McCanny, et al. "Demonstration of Fusion-Evaporation and Direct-Interaction Nuclear Reactions using High-Intensity Laser-Plasma-Accelerated Ion Beams". In: *Phys. Rev. Lett.* 91.7 (2003), pp. 25–28. DOI: 10.1103/PhysRevLett.91.075006.
- [25] P McKenna, KW Ledingham, S Shimizu, et al. "Broad energy spectrum of laser-accelerated protons for spallation-related physics". In: *Phys. Rev. Spec. Top. - Accel. Beams* 8.March (2005), pp. 1–4. DOI: 10.1103/physRevLett.94.084801.
- [26] D Habs, PG Thirolf, M Gross, et al. "Introducing the fission-fusion reaction process: Using a laser-accelerated Th beam to produce neutron-rich nuclei towards the N=126 waiting point of the r-process". In: *Appl. Phys. B Lasers Opt.* 103.2 (2011), pp. 471–484. DOI: 10.1007/s00340-010-4261-x.
- [27] A Macchi, M Borghesi, and M Passoni. "Ion acceleration by superintense laser-plasma interaction". In: *Rev. Mod. Phys.* 85.2 (2013), pp. 751–793. DOI: 10.1103/RevModPhys.85.751.
- [28] ES Sarachik and GT Schappert. "Classical theory of the scattering of intense laser radiation by free electrons". In: *Phys. Rev. D* 1.10 (1970), pp. 2738–2753. DOI: 10.1103/PhysRevD.1.2738.
- [29] S Gales. "Laser driven nuclear science and applications: The need of high efficiency, high power and high repetition rate Laser beams". In: *Eur. Phys. J. Spec. Top.* 224.13 (2015), pp. 2631–2637. DOI: 10.1140/epjst/e2015-02575-7.
- [30] U Schramm, M Bussmann, A Irman, et al. "First results with the novel petawatt laser acceleration facility in Dresden". In: *J. Phys. Conf. Ser.* 874.1 (2017), p. 012028. DOI: 10.1088/1742-6596/874/1/012028.
- [31] J Menapace, ES Fulkerson, C Haefner, et al. "All Diode-Pumped, High-repetition-rate Advanced Petawatt Laser System (HAPLS)". In: *Conf. Lasers Electro-Optics* 60.4 (2017). DOI: 10.1364/cleo_si.2017.sth11.2.
- [32] E Esarey, CB Schroeder, and WP Leemans. "Physics of laser-driven plasma-based electron accelerators". In: *Rev. Mod. Phys.* 81.3 (2009), pp. 1229–1285. DOI: 10.1103/RevModPhys.81.1229.
- [33] OS Lieu and EE Shin. "Nonlinear forward drift of electrons in a laser beam". In: *Appl. Phys. Lett.* 20.12 (1972), pp. 511–512. DOI: 10.1063/1.1654038.
- [34] P Gibbon. *Short Pulse Laser Interactions with Matter*. Imperial College Press, 2005. DOI: 10.1142/p116.
- [35] JP Couperus, R Pausch, A Köhler, et al. "Demonstration of a beam loaded nano-coulomb-class laser wakefield accelerator". In: *Nat. Commun.* 8.1 (2017), p. 487. DOI: 10.1038/s41467-017-00592-7.
- [36] R Pausch, A Debus, A Huebl, et al. "Quantitatively consistent computation of coherent and incoherent radiation in particle-in-cell codes—A general form factor formalism for macro-particles". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 909.4 (2018), pp. 419–422. DOI: 10.1016/j.nima.2018.02.020.

- [37] A Debus, R Pausch, A Huebl, et al. "Breaking the dephasing and depletion limits of laser-wakefield acceleration". In: *Submitt. to Phys. Rev. X* (2019).
- [38] M Brewczyk, K Rzażewski, and CW Clark. "Appearance intensities for multiply charged ions in a strong laser field". In: *Phys. Rev. A* 52.2 (1995), pp. 1468–1473. DOI: 10.1103/PhysRevA.52.1468.
- [39] FF Chen. *Introduction to Plasma Physics and Controlled Fusion Volume 1: Plasma Physics*. 2nd ed. New York: Plenum Press, 1984.
- [40] F Sylla, M Veltcheva, S Kahaly, et al. "Development and characterization of very dense submillimetric gas jets for laser-plasma interaction". In: *Rev. Sci. Instrum.* 83.3 (2012), p. 033507. DOI: 10.1063/1.3697859.
- [41] S Garcia, D Chatain, and JP Perin. "Continuous production of a thin ribbon of solid hydrogen". In: *Laser Part. Beams* 760.4 (2014), p. 569. DOI: 10.1017/S0263034614000524.
- [42] D Margarone, A Velyhan, J Dostal, et al. "Proton Acceleration Driven by a Nanosecond Laser from a Cryogenic Thin Solid-Hydrogen Ribbon". In: *Phys. Rev. X* 6.4 (2016), p. 041030. DOI: 10.1103/physrevx.6.041030.
- [43] S Göde, C Rödel, K Zeil, et al. "Relativistic Electron Streaming Instabilities Modulate Proton Beams Accelerated in Laser-Plasma Interactions". In: *Phys. Rev. Lett.* 118.19 (2017), p. 194801. DOI: 10.1103/physrevlett.118.194801.
- [44] L Obst, S Göde, M Rehwald, et al. "Efficient laser-driven proton acceleration from cylindrical and planar cryogenic hydrogen jets". In: *Sci. Rep.* 7.1 (2017). DOI: 10.1038/s41598-017-10589-3.
- [45] SD Kraft, L Obst, J Metzkes-Ng, et al. "First demonstration of multi-MeV proton acceleration from a cryogenic hydrogen ribbon target". In: *Plasma Phys. Control. Fusion* 60.4 (2018), p. 044010. DOI: 10.1088/1361-6587/aaae38.
- [46] L Obst-Huebl, T Ziegler, FE Brack, et al. "All-optical structuring of laser-driven proton beam profiles". In: *Nat. Commun.* 9.1 (2018), p. 5292. DOI: 10.1038/s41467-018-07756-z.
- [47] JB Kim, C Schoenwaelder, and SH Glenzer. "Development and characterization of liquid argon and methane microjets for high-rep-rate laser-plasma experiments". In: *Rev. Sci. Instrum.* 89.10 (2018). DOI: 10.1063/1.5038561.
- [48] GG Scott, DC Carroll, S Astbury, et al. "Dual Ion Species Plasma Expansion from Isotopically Layered Cryogenic Targets". In: *Phys. Rev. Lett.* 120.20 (2018), p. 204801. DOI: 10.1103/PhysRevLett.120.204801.
- [49] H Okuda and CK Birdsall. "Collisions in a plasma of finite-size particles". In: *Phys. Fluids* 13.8 (1970), pp. 2123–2134. DOI: 10.1063/1.1693210.
- [50] P Gibbon. "Resistively enhanced proton acceleration via high-intensity laser interactions with cold foil targets". In: *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* 72.2 (2005), pp. 1–8. DOI: 10.1103/PhysRevE.72.026411.

- [51] T Kluge, M Bussmann, T Cowan, et al. "Electron Temperature Scaling in Laser Interaction with Solids". In: *Phys. Rev. Lett.* 107.20 (2011). DOI: 10.1103/physrevlett.107.205003.
- [52] RM More. "Pressure Ionization, Resonances, and the Continuity of Bound and Free States". In: *Adv. At. Mol. Phys.* 21.C (1985), pp. 305–356. DOI: 10.1016/S0065-2199(08)60145-1.
- [53] NE Frankel, KC Hines, and RL Dewar. "Energy loss due to binary collisions in a relativistic plasma". In: *Phys. Rev. A* 20.5 (1979), pp. 2120–2129. DOI: 10.1103/PhysRevA.20.2120.
- [54] F Pérez, L Gremillet, A Decoster, et al. "Improved modeling of relativistic collisions and collisional ionization in particle-in-cell codes". In: *Phys. Plasma* 19 (2012), p. 083104. DOI: 10.1063/1.4742167.
- [55] JE Allen and JG Andrews. "A note on ion rarefaction waves". In: *J. Plasma Phys.* 4.1 (1970), pp. 187–194. DOI: 10.1017/S0022377800004906.
- [56] WL Kruer and K Estabrook. "JXB heating by very intense laser light". In: *Phys. Fluids* 28.1 (1985), pp. 430–432. DOI: 10.1063/1.865171.
- [57] SC Wilks, WL Kruer, M Tabak, et al. "Absorption of Ultraintense Laser Pulses". In: *Phys. Rev. Lett.* 69.4 (1992), pp. 1383–1386.
- [58] P Gibbon. "Efficient production of fast electrons from femtosecond laser interaction with solid targets". In: *Phys. Rev. Lett.* 73.5 (1994), pp. 664–667. DOI: 10.1103/PhysRevLett.73.664.
- [59] P Mulser, D Bauer, and H Ruhl. "Collisionless laser-energy conversion by anharmonic resonance". In: *Phys. Rev. Lett.* 101.22 (2008), p. 225002. DOI: 10.1103/PhysRevLett.101.225002.
- [60] DW Forslund, JM Kindel, and K Lee. "Theory of Hot-Electron Spectra at High Laser Intensity". In: *Phys. Rev. Lett.* 39.5 (1977), pp. 284–288. DOI: 10.1103/PhysRevLett.39.284.
- [61] T Kluge. "Enhanced Laser Ion Acceleration from Solids". PhD thesis. Technische Universität Dresden, 2012. URL: <https://www.hzdr.de/db/Cms?p0id=38401>.
- [62] FN Beg, AR Bell, AE Dangor, et al. "A study of picosecond laser–solid interactions up to 10^{19} W cm⁻²". In: *Phys. Plasmas* 4.2 (1997), pp. 447–457. DOI: 10.1063/1.872103.
- [63] F Brunel. "Not-so-resonant, resonant absorption". In: *Phys. Rev. Lett.* 59.1 (1987), pp. 52–55. DOI: 10.1103/PhysRevLett.59.52.
- [64] P Gibbon and AR Bell. "Collisionless absorption in sharp-edged plasmas". In: *Phys. Rev. Lett.* 68.10 (1992), pp. 1535–1538. DOI: 10.1103/PhysRevLett.68.1535.
- [65] L Chopineau, A Leblanc, G Blaclair, et al. "Identification of coupling mechanisms between ultraintense laser light and dense plasmas". In: *Phys. Rev. X* 9.1 (2019), p. 011050. DOI: 10.1103/PhysRevX.9.011050.

- [66] T Kluge, M Bussmann, U Schramm, et al. "Simple scaling equations for electron spectra, currents, and bulk heating in ultra-intense short-pulse laser-solid interaction". In: *Phys. Plasmas* 25.7 (2018). DOI: 10.1063/1.5037753.
- [67] B Bezzerides, DW Forslund, and EL Lindman. "Existence of rarefaction shocks in a laser-plasma corona". In: *Phys. Fluids* 21.12 (1978), pp. 2179–2185. DOI: 10.1063/1.862176.
- [68] J Denavit. "Collisionless plasma expansion into a vacuum". In: *Phys. Fluids* 22.7 (1979), pp. 1384–1392. DOI: 10.1063/1.862751.
- [69] JE Crow, PL Auer, and JE Allen. "The expansion of a plasma into a vacuum". In: *J. Plasma Phys.* 14.1 (1975), pp. 65–76. DOI: 10.1017/s0022377800025538.
- [70] D Strickland and G Mourou. "Compression of Amplified Chirped Optical Pulses". In: *Opt. Commun.* 55.6 (1985), pp. 447–449. DOI: 10.1016/0030-4018(85)90151-8.
- [71] P Mora. "Plasma Expansion into a Vacuum". In: *Phys. Rev. Lett.* 90.18 (2003), p. 4. DOI: 10.1103/PhysRevLett.90.185002.
- [72] J Fuchs, P Antici, E D'Humières, et al. "Laser-driven proton scaling laws and new paths towards energy increase". In: *Nat. Phys.* 2.1 (2005), pp. 48–54. DOI: 10.1038/nphys199.
- [73] K Zeil, J Metzkes, T Kluge, et al. "Direct observation of prompt pre-thermal laser ion sheath acceleration". In: *Nat. Commun.* 3 (2012), p. 874. DOI: 10.1038/ncomms1883.
- [74] J Yu, Z Jiang, JC Kieffer, et al. "Hard x-ray emission in high intensity femtosecond laser-target interaction". In: *Phys. Plasmas* 6.4 (1999), pp. 1318–1322. DOI: 10.1063/1.873372.
- [75] P Mora. "Collisionless expansion of a Gaussian plasma into a vacuum". In: *Phys. Plasmas* 12.11 (2005), pp. 1–8. DOI: 10.1063/1.2134768.
- [76] T Grismayer and P Mora. "Influence of a finite initial ion density gradient on plasma expansion into a vacuum". In: *Phys. Plasmas* 13.3 (2006), p. 032103. DOI: 10.1063/1.2178653.
- [77] J Schreiber, F Bell, F Grüner, et al. "Analytical Model for Ion Acceleration by High-Intensity Laser Pulses". In: *Phys. Rev. Lett.* 97.4 (2006). DOI: 10.1103/physrevlett.97.045005.
- [78] MH Key, MD Cable, TE Cowan, et al. "Hot electron production and heating by hot electrons in fast ignitor research". In: *Phys. Plasmas* 5.5 (1998), pp. 1966–1972. DOI: 10.1063/1.872867.
- [79] K Zeil, SD Kraft, S Bock, et al. "The scaling of proton energies in ultrashort pulse laser plasma acceleration". In: *New J. Phys.* 12.4 (2010), p. 045015. DOI: 10.1088/1367-2630/12/4/045015.
- [80] J Limpouch, J Psikal, AA Andreev, et al. "Enhanced laser ion acceleration from mass-limited targets". In: *Laser Part. Beams* 26.2 (2008), pp. 225–234. DOI: 10.1017/S0263034608000268.

- [81] E D’Humières, E Lefebvre, L Gremillet, et al. “Proton acceleration mechanisms in high-intensity laser interaction with thin foils”. In: *Phys. Plasmas* 12.6 (2005), pp. 1–13. DOI: 10.1063/1.1927097.
- [82] J Pšikal, J Limpouch, S Kawata, et al. “PIC simulations of femtosecond interactions with mass-limited targets”. In: *Czechoslov. J. Phys.* 56 (2006), pp. 515–521. DOI: 10.1007/s10582-006-0246-8.
- [83] T Kluge, W Enghardt, SD Kraft, et al. “Enhanced laser ion acceleration from mass-limited foils”. In: *Phys. Plasmas* 17.12 (2010), p. 123103. DOI: 10.1063/1.3519512.
- [84] M Murakami, YG Kang, K Nishihara, et al. “Ion energy spectrum of expanding laser-plasma with limited mass”. In: *Phys. Plasmas* 12.6 (2005), pp. 1–8. DOI: 10.1063/1.1928247.
- [85] C Ziener, PS Foster, EJ Divall, et al. “Specular reflectivity of plasma mirrors as a function of intensity, pulse duration, and angle of incidence”. In: *J. Appl. Phys.* 93.1 (2003), pp. 768–770. DOI: 10.1063/1.1525062.
- [86] G Doumy, F Quéré, O Gobert, et al. “Complete characterization of a plasma mirror for the production of high-contrast ultraintense laser pulses”. In: *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* 69.2 (2004), pp. 1–12. DOI: 10.1103/PhysRevE.69.026402.
- [87] B Dromey, S Kar, M Zepf, et al. “The plasma mirror—A subpicosecond optical switch for ultrahigh power lasers”. In: *Rev. Sci. Instrum.* 75.3 (2004), pp. 645–649. DOI: 10.1063/1.1646737.
- [88] S Steinke, A Henig, M Schnürer, et al. “Efficient ion acceleration by collective laser-driven electron dynamics with ultra-thin foil targets”. In: *Laser Part. Beams* 28.1 (2010), pp. 215–221. DOI: 10.1017/S0263034610000157.
- [89] C Rödel, M Heyer, M Behmke, et al. “High repetition rate plasma mirror for temporal contrast enhancement of terawatt femtosecond laser pulses by three orders of magnitude”. In: *Appl. Phys. B Lasers Opt.* 103.2 (2011), pp. 295–302. DOI: 10.1007/s00340-010-4329-7.
- [90] BH Shaw, S Steinke, J Van Tilborg, et al. “Reflectance characterization of tape-based plasma mirrors”. In: *Phys. Plasmas* 23.6 (2016). DOI: 10.1063/1.4954242.
- [91] L Obst, J Metzkes-Ng, S Bock, et al. “On-shot characterization of single plasma mirror temporal contrast improvement”. In: *Plasma Phys. Control. Fusion* 60.5 (2018), p. 54007. DOI: 10.1088/1361-6587/aab3bb.
- [92] A Henig, D Kiefer, K Markey, et al. “Enhanced laser-driven ion acceleration in the relativistic transparency regime”. In: *Phys. Rev. Lett.* 103.4 (2009), pp. 1–4. DOI: 10.1103/PhysRevLett.103.045002.
- [93] VA Vshivkov, NM Naumova, F Pegoraro, et al. “Nonlinear electrodynamics of the interaction of ultra-intense laser pulses with a thin foil”. In: *Phys. Plasmas* 5.7 (1998), pp. 2727–2741. DOI: 10.1063/1.872961.

- [94] J Fuchs, JC Adam, F Amiranoff, et al. "Transmission through highly overdense plasma slabs with a subpicosecond relativistic laser pulse". In: *Phys. Rev. Lett.* 80.11 (1998), pp. 2326–2329. DOI: 10.1103/PhysRevLett.80.2326.
- [95] PL Poole, L Obst, GE Cochran, et al. "Laser-driven ion acceleration via target normal sheath acceleration in the relativistic transparency regime". In: *New J. Phys.* 20.1 (2018), p. 013019. DOI: 10.1088/1367-2630/aa9d47.
- [96] S Ter-Avetisyan, B Ramakrishna, R Prasad, et al. "Generation of a quasi-monoenergetic proton beam from laser-irradiated sub-micron droplets". In: *Phys. Plasmas* 19.7 (2012), p. 073112. DOI: 10.1063/1.4731712.
- [97] K Zeil, J Metzkes, T Kluge, et al. "Robust energy enhancement of ultrashort pulse laser accelerated protons from reduced mass targets". In: *Plasma Phys. Control. Fusion* 56.8 (2014), p. 084004. DOI: 10.1088/0741-3335/56/8/084004.
- [98] K Nishihara, H Amitani, M Murakami, et al. "High energy ions generated by laser driven Coulomb explosion of cluster". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 464.1-3 (2001), pp. 98–102. DOI: 10.1016/S0168-9002(01)00014-6.
- [99] EY Echkina, SV Bulanov, T Esirkepov, et al. "Ion Acceleration during the Coulomb Explosion of the Multispecies Clusters". In: *31st EPS Conf. Plasma Phys. London 28G* (2004), P-2.032.
- [100] T Ziegler, M Rehwald, L Obst, et al. "Optical probing of high intensity laser interaction with micron-sized cryogenic hydrogen jets". In: *Plasma Phys. Control. Fusion* 60.7 (2018). DOI: 10.1088/1361-6587/aabf4f.
- [101] O Jäckel, J Polz, SM Pfoth, et al. "All-optical measurement of the hot electron sheath driving laser ion acceleration from thin foils". In: *New J. Phys.* 12.10 (2010), p. 103027. DOI: 10.1088/1367-2630/12/10/103027.
- [102] S Mondal, AD Lad, S Ahmed, et al. "Doppler spectrometry for ultrafast temporal mapping of density dynamics in laser-induced plasmas". In: *Phys. Rev. Lett.* 105.10 (2010). DOI: 10.1103/PhysRevLett.105.105002.
- [103] M Rehwald, TE Cowan, K Zeil, et al. "Reflective optical probing of laser-driven plasmas at the rear surface of solid targets". In: *Plasma Phys. Control. Fusion* 58.3 (2016), p. 034012. DOI: 10.1088/0741-3335/58/3/034012.
- [104] T Kluge, C Gutt, LG Huang, et al. "Using X-ray free-electron lasers for probing of complex interaction dynamics of ultra-intense lasers with solid matter". In: *Phys. Plasmas* 21.3 (2014). DOI: 10.1063/1.4869331.
- [105] T Kluge, M Bussmann, HK Chung, et al. "Nanoscale femtosecond imaging of transient hot solid density plasmas with elemental and charge state sensitivity using resonant coherent diffraction". In: *Phys. Plasmas* 23.3 (2016), p. 33103. DOI: 10.1063/1.4942786.
- [106] T Kluge, C Rödel, M Rödel, et al. "Nanometer-scale characterization of laser-driven compression, shocks, and phase transitions, by x-ray scattering using free electron lasers". In: *Phys. Plasmas* 24.10 (2017), p. 102709. DOI: 10.1063/1.5008289.

- [107] T Kluge, M Rödel, J Metzkes-Ng, et al. "Observation of Ultrafast Solid-Density Plasma Dynamics Using Femtosecond X-Ray Pulses from a Free-Electron Laser". In: *Phys. Rev. X* 8.3 (2018). DOI: 10.1103/PhysRevX.8.031068.
- [108] M Hegelich, S Karsch, G Pretzler, et al. "MeV ion jets from short-pulse-laser interaction with thin foils". In: *Phys. Rev. Lett.* 89.8 (2002), p. 085002. DOI: 10.1103/PhysRevLett.89.085002.
- [109] S Kar, KF Kakolee, B Qiao, et al. "Ion acceleration in multispecies targets driven by intense laser radiation pressure". In: *Phys. Rev. Lett.* 109.18 (2012). DOI: 10.1103/PhysRevLett.109.185006.
- [110] S Steinke, P Hinz, M Schnürer, et al. "Stable laser-ion acceleration in the light sail regime". In: *Phys. Rev. Spec. Top. - Accel. Beams* 16.1 (2013). DOI: 10.1103/physrevstab.16.011303.
- [111] D Jung, L Yin, BJ Albright, et al. "Efficient carbon ion beam generation from laser-driven volume acceleration". In: *New J. Phys.* 15.2 (2013), p. 023007. DOI: 10.1088/1367-2630/15/2/023007.
- [112] A Macchi, S Veghini, and F Pegoraro. "'Light Sail' Acceleration Reexamined". In: *Phys. Rev. Lett.* 103.8 (2009), p. 85003. DOI: 10.1103/PhysRevLett.103.085003.
- [113] A Henig, S Steinke, M Schnürer, et al. "Radiation-pressure acceleration of ion beams driven by circularly polarized laser pulses". In: *Phys. Rev. Lett.* 103.24 (2009), pp. 3–6. DOI: 10.1103/PhysRevLett.103.245003.
- [114] APL Robinson, P Gibbon, M Zepf, et al. "Relativistically correct hole-boring and ion acceleration by circularly polarized laser pulses". In: *Plasma Phys. Control. Fusion* 51.2 (2009), p. 024004. DOI: 10.1088/0741-3335/51/2/024004.
- [115] LO Silva, M Marti, JR Davies, et al. "Proton Shock Acceleration in Laser-Plasma Interactions". In: *Phys. Rev. Lett.* 92.1 (2004), pp. 1–4. DOI: 10.1103/physrevlett.92.015002.
- [116] CA Palmer, NP Dover, I Pogorelsky, et al. "Monoenergetic proton beams accelerated by a radiation pressure driven shock". In: *Phys. Rev. Lett.* 106.1 (2011), pp. 1–4. DOI: 10.1103/PhysRevLett.106.014801.
- [117] D Haberberger, S Tochitsky, F Fiuza, et al. "Collisionless shocks in laser-produced plasma generate monoenergetic high-energy proton beams". In: *Nat. Phys.* 8.1 (2012), pp. 95–99. DOI: 10.1038/nphys2130.
- [118] F Fiuza, A Stockem, E Boella, et al. "Laser-driven shock acceleration of monoenergetic ion beams". In: *Phys. Rev. Lett.* 109.21 (2012), pp. 1–5. DOI: 10.1103/PhysRevLett.109.215001.
- [119] A Stockem, F Fiuza, A Bret, et al. "Exploring the nature of collisionless shocks under laboratory conditions". In: *Sci. Rep.* 4 (2014), pp. 1–6. DOI: 10.1038/srep03934.
- [120] B Qiao, S Kar, M Geissler, et al. "Dominance of Radiation Pressure in Ion Acceleration with Linearly Polarized Pulses at Intensities of 10^{21} W cm $^{-2}$ ". In: *Phys. Rev. Lett.* 108.11 (2012), p. 115002. DOI: 10.1103/PhysRevLett.108.115002.

- [121] I Last, I Schek, and J Jortner. "Energetics and dynamics of Coulomb explosion of highly charged clusters". In: *J. Chem. Phys.* 107.17 (1997), pp. 6685–6692. DOI: 10.1063/1.474911.
- [122] L Di Lucchio, AA Andreev, and P Gibbon. "Ion acceleration by intense, few-cycle laser pulses with nanodroplets". In: *Phys. Plasmas* 22.5 (2015), pp. 1–10. DOI: 10.1063/1.4921667.
- [123] TM Ostermayr, D Haffa, P Hinz, et al. "Proton acceleration by irradiation of isolated spheres with an intense laser pulse". In: *Phys. Rev. E* 94.3 (2016), pp. 1–6. DOI: 10.1103/PhysRevE.94.033208.
- [124] E Boella, BP Paradisi, A D'Angola, et al. "Study on Coulomb explosions of ion mixtures". In: *J. Plasma Phys.* 82.1 (2016), pp. 1–12. DOI: 10.1017/S0022377816000179.
- [125] L Yin, BJ Albright, BM Hegelich, et al. "GeV laser ion acceleration from ultrathin targets: The laser break-out afterburner". In: *Laser Part. Beams* 24.5 (2006), pp. 91–298. DOI: 10.1063/1.2436857.
- [126] SS Bulanov, VY Bychenkov, V Chvykov, et al. "Generation of GeV protons from 1 PW laser interaction with near critical density targets". In: *Phys. Plasmas* 17.4 (2010), p. 043105. DOI: 10.1063/1.3372840.
- [127] AP Robinson, AR Bell, and RJ Kingham. "Effect of target composition on proton energy spectra in ultraintense laser-solid interactions". In: *Phys. Rev. Lett.* 96.3 (2006), p. 035005. DOI: 10.1103/PhysRevLett.96.035005.
- [128] SM Pfotenhauer, O Jäckel, A Sachtleben, et al. "Spectral shaping of laser generated proton beams". In: *New J. Phys.* 10 (2008). DOI: 10.1088/1367-2630/10/3/033034.
- [129] Y Sentoku, K Mima, P Kaw, et al. "Anomalous Resistivity Resulting from MeV-Electron Transport in Overdense Plasma". In: *Phys. Rev. Lett.* 90.15 (2003), p. 4. DOI: 10.1103/PhysRevLett.90.155001.
- [130] KS Yee. "Numerical Solution of Initial Boundary Value Problems Involving Maxwell's Equations in Isotropic Media". In: *IEEE Trans. Antennas Propag.* 14.3 (1966), pp. 302–307. DOI: 10.1109/TAP.1966.1138693.
- [131] RW Hockney and JW Eastwood. *Computer Simulation Using Particles*. McGraw-Hill, 1981.
- [132] CK Birdsall and AB Langdon. *Plasma Physics via Computer Simulation*. IOP Publishing, 1991. DOI: 10.1201/9781315275048.
- [133] A Taflove and ME Brodwin. "Numerical Solution of Steady-State Electromagnetic Scattering Problems Using the Time-Dependent Maxwell's Equations". In: *IEEE Trans. Microw. Theory Tech.* 23.8 (1975), pp. 623–630. DOI: 10.1109/TMTT.1975.1128640.
- [134] BB Godfrey and JL Vay. "Numerical stability of relativistic beam multidimensional PIC simulations employing the Esirkepov algorithm". In: *J. Comput. Phys.* 248 (2013), pp. 33–46. DOI: 10.1016/j.jcp.2013.04.006.
- [135] JP Boris. "Relativistic Plasma Simulation - Optimization of a Hybrid Code". In: *Numer. Simul. Plasmas*. Washington, D.C.: Naval Research Laboratory, 1971, pp. 3–67.

- [136] JL Vay. "Simulation of beams or plasmas crossing at relativistic velocity". In: *Phys. Plasmas* 15 (2008). DOI: 10.1063/1.2837054.
- [137] AV Higuera and JR Cary. "Structure-preserving second-order integration of relativistic charged particle trajectories in electromagnetic fields". In: *Phys. Plasmas* 24.5 (2017). DOI: 10.1063/1.4979989.
- [138] M Vranic, JL Martins, RA Fonseca, et al. "Classical radiation reaction in particle-in-cell simulations". In: *Comput. Phys. Commun.* 204 (2016), pp. 141–151. DOI: 10.1016/j.cpc.2016.04.002.
- [139] A Gonoskov, S Bastrakov, E Efimenko, et al. "Extended particle-in-cell schemes for physics in ultrastrong laser fields: Review and developments". In: *Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys.* 92.2 (2015), pp. 1–18. DOI: 10.1103/PhysRevE.92.023305.
- [140] K Nanbu. "Theory of Cumulative Small-Angle Collisions in Plasmas". In: *Phys. Rev. E* 55.4 (1997), pp. 4642–4652. DOI: 10.1103/PhysRevE.55.4642.
- [141] K Nanbu and S Yonemura. "Weighted Particles in Coulomb Collision Simulations Based on the Theory of a Cumulative Scattering Angle". In: *J. Comput. Phys.* 145.2 (1998), pp. 639–654. DOI: 10.1006/jcph.1998.6049.
- [142] Y Sentoku, K Mima, Y Kishimoto, et al. "Effects of Relativistic Binary Collisions on PIC Simulation of Laser Plasmas". In: *J. Phys. Soc. Japan* 67.12 (1998), pp. 4084–4088. DOI: 10.1143/JPSJ.67.4084.
- [143] Y Sentoku and AJ Kemp. "Numerical methods for particle simulations at extreme densities and temperatures: Weighted particles, relativistic collisions and reduced currents". In: *J. Comput. Phys.* 227.14 (2008), pp. 6846–6861. DOI: 10.1016/j.jcp.2008.03.043.
- [144] J Villasenor and O Buneman. "Rigorous charge conservation for local electromagnetic field solvers". In: *Comput. Phys. Commun.* 69.2-3 (1992), pp. 306–316. DOI: 10.1016/0010-4655(92)90169-Y.
- [145] TZ Esirkepov. "Exact charge conservation scheme for Particle-in-Cell simulation with an arbitrary form-factor". In: *Comput. Phys. Commun.* 135.2 (2001), pp. 144–153. DOI: 10.1016/S0010-4655(00)00228-9.
- [146] T Umeda, Y Omura, T Tominaga, et al. "A new charge conservation method in electromagnetic particle-in-cell simulations". In: *Comput. Phys. Commun.* 156.1 (2003), pp. 73–85. DOI: 10.1016/S0010-4655(03)00437-5.
- [147] T Umeda, Y Omura, and H Matsumoto. "Charge conservation methods for computing current densities in electromagnetic particle-in-cell simulations". In: *Proc. ISSS-7*. Matsue, Japan, 2005.
- [148] J Yu, X Jin, W Zhou, et al. "High-order interpolation algorithms for charge conservation in particle-in-cell simulations". In: *Commun. Comput. Phys.* 13.4 (2013), pp. 1134–1150. DOI: 10.4208/cicp.290811.050412a.
- [149] E Boella, G Coppa, A D'Angola, et al. "Use of the shell model for plasma physics simulation". In: *APS Meet. Abstr.* 2011, QRP1.092.

- [150] AF Lifschitz, X Davoine, E Lefebvre, et al. "Particle-in-Cell modelling of laser-plasma interaction using Fourier decomposition". In: *J. Comput. Phys.* 228.5 (2009), pp. 1803–1814. DOI: 10.1016/j.jcp.2008.11.017.
- [151] R Lehe, M Kirchen, IA Andriyash, et al. "A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm". In: *Comput. Phys. Commun.* 203 (2016), pp. 66–82. DOI: 10.1016/j.cpc.2016.02.007.
- [152] A Huebl. "Injection Control for Electrons in Laser-Driven Plasma Wakes on the Femtosecond Time Scale". Diploma thesis. Technical University Dresden, 2014. DOI: 10.5281/zenodo.15924.
- [153] R Lehe, A Lifschitz, C Thauray, et al. "Numerical growth of emittance in simulations of laser-wakefield acceleration". In: *Phys. Rev. Spec. Top. - Accel. Beams* 16.2 (2013), pp. 1–8. DOI: 10.1103/PhysRevSTAB.16.021301.
- [154] BB Godfrey and JL Vay. "Suppressing the numerical Cherenkov instability in FDTD PIC codes". In: *J. Comput. Phys.* 267 (2014), pp. 1–6. DOI: 10.1016/j.jcp.2014.02.022.
- [155] F Li, P Yu, X Xu, et al. "Controlling the numerical Cherenkov instability in PIC simulations using a customized finite difference Maxwell solver and a local FFT based current correction". In: *Comput. Phys. Commun.* 214 (2017), pp. 6–17. DOI: 10.1016/j.cpc.2017.01.001.
- [156] S Jalas, I Dornmair, R Lehe, et al. "Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods". In: *Phys. Plasmas* 24.3 (2017). DOI: 10.1063/1.4978569.
- [157] D Tskhakaya, K Matyash, R Schneider, et al. "The particle-in-cell method". In: *Contrib. to Plasma Phys.* 47.8-9 (2007), pp. 563–594. DOI: 10.1002/ctpp.200710072.
- [158] AJ Kemp, Y Sentoku, and M Tabak. "Hot-Electron Energy Coupling in Ultraintense Laser-Matter Interaction". In: *Phys. Rev. Lett.* 101.7 (2008), p. 075004. DOI: 10.1103/PhysRevLett.101.075004.
- [159] AG MacPhee, L Divol, AJ Kemp, et al. "Limitation on prepulse level for cone-guided fast-ignition inertial confinement Fusion". In: *Phys. Rev. Lett.* 104.5 (2010), pp. 1–4. DOI: 10.1103/PhysRevLett.104.055002.
- [160] RA Fonseca, SF Martins, LO Silva, et al. "One-to-one direct modeling of experiments and astrophysical scenarios: Pushing the envelope on kinetic plasma simulations". In: *Plasma Phys. Control. Fusion* 50.12 (2008). DOI: 10.1088/0741-3335/50/12/124034.
- [161] R Mishra, P Leblanc, Y Sentoku, et al. "Collisional particle-in-cell modeling for energy transport accompanied by atomic processes in dense plasmas". In: *Phys. Plasmas* 20.7 (2013). DOI: 10.1063/1.4812701.
- [162] MV Ammosov, NB Delone, and VP Krainov. "Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field". In: *Sov. JETP* 91.6 (1986), pp. 1191–1194. DOI: 10.1017/S0263574703005605.
- [163] N Delone and V Krainov. "Tunneling and barrier-suppression ionization of atoms and ions in a laser radiation field". In: *Physics-Usppekhi* 469 (1998). DOI: 10.1070/PU1998v041n05ABEH000393.

- [164] D Bauer and P Mulser. "Exact field ionization rates in the barrier-suppression regime from numerical time-dependent Schrödinger-equation calculations". In: *Phys. Rev. A - At. Mol. Opt. Phys.* 59.1 (1999), pp. 569–577. DOI: 10.1103/PhysRevA.59.569.
- [165] P Mulser and D Bauer. *High Power Laser-Matter Interaction*. Springer Tracts in Modern Physics. Springer Berlin Heidelberg, 2010.
- [166] M Garten. "Modellierung und Validierung von Feldionisation in parallelen Particle-in-Cell-Codes". Master's Thesis. Dresden: Technische Universität Dresden, 2015. DOI: 10.5281/zenodo.202500.
- [167] D Salzmänn. *Atomic Physics in Hot Plasmas*. International Series of Monographs on Physics. Oxford University Press, 1998.
- [168] HK Chung, MH Chen, WL Morgan, et al. "FLYCHK: Generalized population kinetics and spectral model for rapid spectroscopic analysis for all elements". In: *High Energy Density Phys.* 1.1 (2005), pp. 3–12. DOI: 10.1016/j.hedp.2005.07.001.
- [169] HK Chung, MH Chen, and RW Lee. "Extension of atomic configuration sets of the Non-LTE model in the application to the $K\alpha$ diagnostics of hot dense matter". In: *High Energy Density Phys.* 3.1-2 (2007), pp. 57–64. DOI: 10.1016/j.hedp.2007.02.001.
- [170] FL Wang, D Salzmänn, H Takabe, et al. "Calculation of photoionized plasmas with a detailed-configuration-accounting atomic model". In: *J. Phys. Soc. Japan* 78.6 (2009), pp. 1–4. DOI: 10.1143/JPSJ.78.064301.
- [171] Y Ralchenko. *Modern Methods in Collisional-Radiative Modeling of Plasmas*. Springer Series on Atomic, Optical, and Plasma Physics. Springer International Publishing, 2016.
- [172] A Huebl, R Widera, B Worpitz, et al. *Manual of PIconGPU 0.4.3*. 2019. URL: <https://readthedocs.org/projects/picongpu/downloads/pdf/0.4.3/>.
- [173] JL Vay, A Almgren, J Bell, et al. "Warp-X: A new exascale computing platform for beam-plasma simulations". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 909 (2018), pp. 476–479. DOI: 10.1016/j.nima.2018.01.035.
- [174] J Barnes and P Hut. "A hierarchical $O(N \log N)$ force-calculation algorithm". In: *Nature* 324.4 (1986), pp. 446–449. DOI: 10.1038/324446a0.
- [175] L Greengard and V Rokhlin. "A Fast Algorithm for Particle". In: *J. Comput. Phys* 73.2 (1987), pp. 325–348. DOI: 10.1006/jcph.1997.5706.
- [176] P Gibbon, FN Beg, EL Clark, et al. "Tree-code simulations of proton acceleration from laser-irradiated wire targets". In: *Phys. Plasmas* 11.8 (2004), pp. 4032–4040. DOI: 10.1063/1.1767096.
- [177] M Winkel, R Speck, H Hübner, et al. "A massively parallel, multi-disciplinary Barnes-Hut tree code for extreme-scale N-body simulations". In: *Comput. Phys. Commun.* 183.4 (2012), pp. 880–889. DOI: 10.1016/j.cpc.2011.12.013.

- [178] A Ghizzo, P Bertrand, MM Shoucri, et al. "A Vlasov code for the numerical simulation of stimulated raman scattering". In: *J. Comput. Phys.* 90.2 (1990), pp. 431–457. DOI: 10.1016/0021-9991(90)90174-Y.
- [179] F Filbet and E Sonnendrücker. "Comparison of Eulerian Vlasov solvers". In: *Comput. Phys. Commun.* 150.3 (2003), pp. 247–266. DOI: 10.1016/S0010-4655(02)00694-X.
- [180] JL Vay. "Noninvariance of space- and time-scale ranges under a Lorentz transformation and the implications for the study of relativistic interactions". In: *Phys. Rev. Lett.* 98.13 (2007), pp. 2–5. DOI: 10.1103/PhysRevLett.98.130405.
- [181] J Ameres. "Stochastic and Spectral Particle Methods for Plasma Physics". PhD thesis. Technische Universität München, 2018.
- [182] M Kraus, K Kormann, PJ Morrison, et al. "GEMPIC: Geometric ElectroMagnetic Particle-In-Cell Methods". In: *J. Plasma Phys.* 83.4 (2017). DOI: 10.1017/S002237781700040X.
- [183] C Huang, VK Decyk, M Zhou, et al. "QuickPIC: A highly efficient fully parallelized PIC code for plasma-based acceleration". In: *J. Phys. Conf. Ser.* 46.1 (2006), pp. 190–199. DOI: 10.1088/1742-6596/46/1/026.
- [184] H Fehske, R Schneider, and A Weiße. *Computational Many-Particle Physics*. Lecture Notes in Physics. Springer Berlin Heidelberg, 2007.
- [185] DB Kirk and WW Hwu. *Programming Massively Parallel Processors: A Hands-on Approach*. Elsevier Science, 2016.
- [186] S Chandrasekaran and G Juckeland. *OpenACC for Programmers: Concepts and Strategies*. Pearson Education, 2017.
- [187] P Czarnul. *Parallel Programming for Modern High Performance Computing Systems*. CRC Press, 2018.
- [188] J Jeffers and J Reinders. *Intel Xeon Phi Coprocessor High Performance Programming*. Elsevier Science, 2013.
- [189] A Matthes, R Widera, E Zenker, et al. "Tuning and optimization for a variety of many-core architectures without changing a single line of implementation code using the alpaka library". In: *Lect. Notes Comput. Sci.* Vol. 10524. 2017, pp. 496–514. DOI: 10.1007/978-3-319-67630-2_36.
- [190] NVIDIA Corporation. *CUDA Programming Guide - NVIDIA Developer Documentation*. URL: <https://docs.nvidia.com/cuda>.
- [191] OpenMP Architecture Review Board. *OpenMP Application Programming Interface, Version 5.0*. 2018.
- [192] S Williams, A Waterman, and D Patterson. "Roofline: An Insightful Visual Performance Model for Multicore Architectures". In: *Commun. ACM* April (2009), pp. 65–76. DOI: 10.1145/1498765.1498785.
- [193] K Rupp. *CPU, GPU and MIC Hardware Characteristics over Time*. 2013. URL: <https://web.archive.org/web/20190315143202/https://www.karlrupp.net/2013/06/cpu-gpu-and-mic-hardware-characteristics-over-time/>.

- [194] AV Gurevich, L Pariiskaya, and L Pitaevskii. "Ion acceleration upon expansion of a rarefied plasma". In: *Sov. Phys. JETP* 63.2 (1973), p. 516. URL: <http://www.jetp.ac.ru/cgi-bin/e/index/e/36/2/p274?a=list>.
- [195] T Tajima and JM Dawson. "Laser Electron Accelerator". In: *Phys. Rev. Lett.* 43.4 (1979), pp. 267–270. DOI: 10.1103/PhysRevLett.43.267.
- [196] T Katsouleas, C Joshi, JM Dawson, et al. "Plasma accelerators". In: *AIP Conf. Proc.* Vol. 130. 1985, p. 63. DOI: 10.1063/1.35293.
- [197] M Bussmann, H Burau, TE Cowan, et al. "Radiative signatures of the relativistic Kelvin-Helmholtz instability". In: *SC'13 Proc. Int. Conf. High Perform. Comput. Networking, Storage Anal.* Vol. 60. 4. 2013, pp. 1–12. DOI: 10.1145/2503210.2504564.
- [198] A Sharma and A Andreev. "Effective laser driven proton acceleration from near critical density hydrogen plasma". In: *Laser Part. Beams* 34.02 (2016), pp. 219–229. DOI: 10.1017/s0263034616000045.
- [199] A Sharma, Z Tibai, J Hebling, et al. "Quasi-monoenergetic proton acceleration from cryogenic hydrogen microjet by ultrashort ultraintense laser pulses". In: *Phys. Plasmas* 25.3 (2018). DOI: 10.1063/1.5003353.
- [200] A Sharma. "High Energy electron and proton acceleration by circularly polarized laser pulse from near critical density hydrogen gas target". In: *Sci. Rep.* 8.1 (2018), pp. 1–10. DOI: 10.1038/s41598-018-20506-x.
- [201] A Sharma, Z Tibai, and J Hebling. "Intense tera-hertz laser driven proton acceleration in plasmas". In: *Phys. Plasmas* 23.6 (2016). DOI: 10.1063/1.4953803.
- [202] A Sharma, Z Tibai, J Hebling, et al. "Terahertz-driven wakefield electron acceleration". In: *J. Phys. B At. Mol. Opt. Phys.* 51.20 (2018), p. 204001. DOI: 10.1088/1361-6455/aadf50.
- [203] R Pausch, A Debus, R Widera, et al. "How to test and verify radiation diagnostics simulations within particle-in-cell frameworks". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 740 (2013), pp. 250–256. DOI: 10.1016/j.nima.2013.10.073.
- [204] R Pausch, H Burau, M Bussmann, et al. "Computing angularly-resolved far field emission spectra in particle-in-cell codes using GPUs". In: *Proceeding of IPAC2014*. Vol. MOPRI069. 2014, pp. 761–764. DOI: 10.18429/JACoW-IPAC2014-MOPRI069.
- [205] R Pausch, M Bussmann, A Huebl, et al. "Identifying the linear phase of the relativistic Kelvin-Helmholtz instability and measuring its growth rate via radiation". In: *Phys. Rev. E* 96.1 (2017), p. 013316. DOI: 10.1103/PhysRevE.96.013316.
- [206] C Fortmann-Grote, AA Andreev, R Briggs, et al. "SIMEX: Simulation of Experiments at Advanced Light Sources". In: *11th NOBUGS Conf.* Copenhagen, 2016, pp. 29–34. DOI: 10.17199/NOBUGS2016.21.
- [207] C Fortmann-Grote, AA Andreev, K Appel, et al. "Simulations of ultrafast x-ray laser experiments". In: *Proc. SPIE 10237, Adv. X-ray Free. Lasers Instrum. IV.* Vol. 10237. 2017, 102370S. DOI: 10.1117/12.2270552.

- [208] H Burau, R Widera, W Hönig, et al. "PIConGPU: A fully relativistic particle-in-cell code for a GPU cluster". In: *IEEE Trans. Plasma Sci.* 38.10 PART 2 (2010), pp. 2831–2839. DOI: 10.1109/TPS.2010.2064310.
- [209] W Hönig, F Schmitt, R Widera, et al. *Presentation and Paper at GTC2010: Developing Highly Scalable Particle-Mesh Codes for GPUs: A Generic Approach*. Santa Clara, 2010.
- [210] R Dietrich, T Ilsche, and G Juckeland. "Non-intrusive performance analysis of parallel hardware accelerated applications on hybrid architectures". In: *Proc. Int. Conf. Parallel Process. Work.* (2010), pp. 135–143. DOI: 10.1109/ICPPW.2010.30.
- [211] W Joubert, R Archibald, M Berrill, et al. "Accelerated application development: The ORNL Titan experience". In: *Comput. Electr. Eng.* 46 (2015), pp. 123–138. DOI: 10.1016/j.compeleceng.2015.04.008.
- [212] M Lieber, K Gössner, and WE Nagel. "The Potential of Diffusive Load Balancing at Large Scale". In: *Proc. 23rd Eur. MPI Users' Gr. Meet. EuroMPI 2016*. New York, NY, USA: ACM, 2016, pp. 154–157. DOI: 10.1145/2966884.2966887.
- [213] M Weber, R Brendel, T Hilbrich, et al. "Structural Clustering: A New Approach to Support Performance Analysis at Scale". In: *Proc. - 2016 IEEE 30th Int. Parallel Distrib. Process. Symp. IPDPS 2016* (2016), pp. 484–493. DOI: 10.1109/IPDPS.2016.27.
- [214] E Zenker, R Widera, A Huebl, et al. "Performance-portable many-core plasma simulations: Porting PIconGPU to OpenPower and beyond". In: *Lect. Notes Comput. Sci.* 9945.4 (2016), pp. 293–301. DOI: 10.1007/978-3-319-46079-6_21.
- [215] E Zenker, B Worpitz, R Widera, et al. "Alpaka - An abstraction library for parallel kernel acceleration". In: *Proc. - 2016 IEEE 30th Int. Parallel Distrib. Process. Symp. IPDPS 2016* 60.4 (2016), pp. 631–640. DOI: 10.1109/IPDPSW.2016.50.
- [216] A Matthes, A Huebl, R Widera, et al. "In situ, steerable, hardware-independent and data-structure agnostic visualization with ISAAC". In: *Supercomput. Front. Innov.* 3.4 (2016). DOI: 10.14529/jsfi160403.
- [217] A Goswami, Y Tian, K Schwan, et al. "Landrush: Rethinking In-Situ Analysis for GPGPU Workflows". In: *Proc. - 2016 16th IEEE/ACM Int. Symp. Clust. Cloud, Grid Comput. CCGrid 2016*. i. 2016, pp. 32–41. DOI: 10.1109/CCGrid.2016.58.
- [218] A Huebl, R Widera, F Schmitt, et al. "On the scalability of data reduction techniques in current and upcoming HPC systems from an application perspective". In: *Lect. Notes Comput. Sci.* 10524.4 (2017), pp. 15–29. DOI: 10.1007/978-3-319-67630-2_2.
- [219] M Harris. "Mapping computational concepts to GPUs". In: *ACM SIGGRAPH 2005 Courses*. Los Angeles, California: ACM, 2005, p. 50. DOI: 10.1145/1198555.1198768.
- [220] TD Arber, K Bennett, CS Brady, et al. "Contemporary particle-in-cell approach to laser-plasma modelling". In: *Plasma Phys. Control. Fusion* 57.11 (2015), p. 113001. DOI: 10.1088/0741-3335/57/11/113001.
- [221] A Friedman, DP Grote, and I Haber. "Three-dimensional particle simulation of heavy-ion fusion beams". In: *Phys. Fluids B* 4.7 (1992), pp. 2203–2210. DOI: 10.1063/1.860024.

- [222] K Martin and B Hoffman. *Mastering CMake Version 3.1*. Kitware Incorporated, 2015. URL: <https://cmake.org/>.
- [223] T Gamblin, M LeGendre, MR Collette, et al. "The Spack Package Manager: Bringing Order to HPC Software Chaos". In: *Proc. Int. Conf. High Perform. Comput. Networking, Storage Anal. - SC '15*. 2015. DOI: 10.1145/2807591.2807623.
- [224] C Boettiger. "An introduction to Docker for reproducible research". In: *ACM SIGOPS Oper. Syst. Rev. - Spec. Issue Repeatability Shar. Exp. Artifacts* 49.1 (2015), pp. 71–79. DOI: 10.1145/2723872.2723882.
- [225] NVIDIA Corporation. *NVIDIA GPU Cloud (NGC)*. 2018. URL: <https://docs.nvidia.com/ngc/>.
- [226] S Stefano, A Marocchino, D Terzani, et al. *ALaDyn: A high-accuracy PIC code for the Maxwell-Vlasov equations*. 2016. DOI: 10.1109/TPS.2008.927143.
- [227] IA Andriyash, R Lehe, and V Malka. "A spectral unaveraged algorithm for free electron laser simulations". In: *J. Comput. Phys.* 282 (2015), pp. 397–409. DOI: 10.1016/j.jcp.2014.11.026.
- [228] S Bastrakov and I Surmin. *pica: Kernel library for particle-in-cell plasma simulation*. 2017. URL: <https://github.com/pictools/pica>.
- [229] A Sgattoni, L Fedeli, S Sinigardi, et al. "Optimising PICCANTE - an Open Source Particle-in-Cell Code for Advanced Simulations on Tier-0 Systems". In: (2015), pp. 1–8. arXiv: 1503.02464.
- [230] R Kumar and P Crumley. *PICTOR: a particle-in-cell code for plasma simulations*. 2018. URL: <https://github.com/rahulbgu/PICTOR>.
- [231] PlasmaPy Community, NA Murphy, AJ Leonard, et al. *PlasmaPy: an open source community-developed Python package for plasma physics*. 2018. DOI: 10.5281/zenodo.1238132.
- [232] K Germaschewski, W Fox, S Abbott, et al. "The Plasma Simulation Code: A modern particle-in-cell code with patch-based load-balancing". In: *J. Comput. Phys.* 318 (2016), pp. 305–326. DOI: 10.1016/j.jcp.2016.05.013.
- [233] J Derouillat, F Pérez, M Flé, et al. "Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation". In: *Comput. Phys. Commun.* 222 (2017), pp. 351–373. DOI: 10.1016/j.cpc.2017.09.024.
- [234] Particle In Cell Consulting LLC. *Starfish: A two-dimensional PIC-DSMC plasma / rarefied gas simulation program*. 2016. URL: <https://github.com/particleinCELL/Starfish>.
- [235] KJ Bowers, BJ Albright, L Yin, et al. "Advances in petascale kinetic plasma simulation with VPIC and Roadrunner". In: *J. Phys. Conf. Ser.* 180. July 2017 (2009), p. 012055. DOI: 10.1088/1742-6596/180/1/012055.
- [236] H Burau. "Entwicklung und Überprüfung eines Photonenmodells für die Abstrahlung durch hochenergetische Elektronen". Diploma thesis. Dresden: TU Dresden, 2016. DOI: 10.5281/zenodo.192116.

- [237] B Stroustrup. *The C++ Programming Language*. 4th ed. Addison-Wesley, 2013.
- [238] TL Veldhuizen and ME Jernigan. "Will C++ Be Faster Than Fortran ?" In: *Lect. Notes Comput. Sci.* 1343 (1997), pp. 49–56. DOI: 10.1007/3-540-63827-X_43.
- [239] A Alexandrescu. *Modern C++ Design: Generic Programming and Design Patterns Applied*. C++ In-Depth Series. Pearson Education, 2001.
- [240] JP Shen and MH Lipasti. *Modern Processor Design: Fundamentals of Superscalar Processors*. Waveland Press, 2013.
- [241] HW Meuer, E Strohmaier, J Dongarra, et al. *TOP500 Supercomputer Sites*. 2018. URL: <https://www.top500.org/lists/2018/11/>.
- [242] A Huebl. *The Green HPCG List*. 2018. DOI: 10.14278/rodare.68. URL: https://plasma.ninja/blog/hpc/manycore/top500/computing/hardware/energy/efficiency/2018/11/18/HPCG_Green.html.
- [243] JJ Dongarra, P Luszczek, and A Petite. "The LINPACK benchmark: Past, present and future". In: *Concurr. Comput. Pract. Exp.* 15.9 (2003), pp. 803–820. DOI: 10.1002/cpe.728.
- [244] J Dongarra, MA Heroux, and P Luszczek. "High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems". In: *Int. J. High Perform. Comput. Appl.* 30.1 (2016), pp. 3–10. DOI: 10.1177/1094342015593158.
- [245] GE Moore. "Cramming more components onto integrated circuits". In: *Electronics* 38.8 (1965), pp. 114–117. DOI: 10.1109/N-SSC.2006.4785860.
- [246] GM Amdahl. "Validity of the single processor approach to achieving large scale computing capabilities". In: *AFIPS Conf. Proc.* 1967, pp. 483–485. DOI: 10.1145/1465482.1465560.
- [247] D Abrahams and A Gurtovoy. *C++ Template Metaprogramming: Concepts, Tools, and Techniques from Boost and Beyond (C++ in Depth Series)*. Addison-Wesley Professional, 2004.
- [248] J Reinders. *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O'Reilly Media, 2007.
- [249] Advanced Micro Devices (AMD). *HIP: C++ Heterogeneous-Compute Interface for Portability*. 2016. URL: <https://gpuopen.com/compute-product/hip-convert-cuda-to-portable-c-code>.
- [250] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard, Version 3.1*. 2015. URL: <https://www.mpi-forum.org/docs>.
- [251] B Worpitz. *Investigating performance portability of a highly scalable particle-in-cell simulation code on various multi-core architectures*. Master's Thesis. Dresden, 2015. DOI: 10.5281/zenodo.49768.
- [252] CHJ Eckert. "Enhancements of the massively parallel memory allocator ScatterAlloc and its adaption to the general interface mallocMC". *Großer Beleg*. 2014. DOI: 10.5281/zenodo.34461.

- [253] M Steinberger, M Kenzel, B Kainz, et al. "ScatterAlloc: Massively parallel dynamic memory allocation for the GPU". In: *2012 Innov. Parallel Comput. InPar 2012* (2012). DOI: 10.1109/InPar.2012.6339604.
- [254] A Huebl, M Rehwald, L Obst-Huebl, et al. "Spectral Control via Multi-Species Effects in PW-Class Laser-Ion Acceleration". In: *Submitted* (2019). arXiv: 1903.06428.
- [255] PT Luu, T Tückmantel, and A Pukhov. "Voronoi particle merging algorithm for PIC codes". In: *Comput. Phys. Commun.* 202 (2016), pp. 165–174. DOI: 10.1016/j.cpc.2016.01.009.
- [256] T Kluyver, B Ragan-Kelley, F Pérez, et al. "Jupyter Notebooks—a publishing format for reproducible computational workflows". In: ed. by F Loizides and B Schmidt. Vol. 16. 4. IOS Press. 2016, pp. 87–90. DOI: 10.3233/978-1-61499-649-1-87.
- [257] A Huebl, R Pausch, R Widera, et al. *Scalable Multi-Platform PIC Simulations as an Open Science Service*. Breckenridge, CO, 2018. DOI: 10.5281/zenodo.1345080.
- [258] S Rudat. "Laser Wakefield Acceleration Simulation as a Service". Master's Thesis. Technische Universität Dresden, 2019.
- [259] N Farooqui, A Kerr, G Damos, et al. "A framework for dynamically instrumenting GPU compute applications within GPU Ocelot". In: *Proc. Fourth Work. Gen. Purp. Process. Graph. Process. Units.* 2011, 9:1–9:9. DOI: 10.1145/1964179.1964192.
- [260] A Matthes. *LLAMA: Low Level Abstraction of Memory Access*. 2018. URL: <https://github.com/ComputationalRadiationPhysics/llama>.
- [261] M Kretz and V Lindenstruth. "Vc: A C++ library for explicit vectorization Matthias". In: *Softw. - Pract. Exp.* 42 (2011), pp. 1409–1430. DOI: 10.1002/spe.1149.
- [262] A Huebl, R Widera, A Matthes, et al. *Whitepaper - New Design PMacc Resource Register (unpublished)*. 2016.
- [263] The HDF Group. *Hierarchical data format version 5 (C-API: 1.8.14)*. 2000. URL: <https://www.hdfgroup.org>.
- [264] Q Liu, J Logan, Y Tian, et al. "Hello ADIOS: the challenges and lessons of developing leadership class I/O frameworks". In: *Concurr. Comput. Pr. Exp.* 26 (2014), pp. 1453–1473. DOI: 10.1002/cpe.
- [265] M Adler. *zlib - a massively spiffy yet delicately unobtrusive compression library*. URL: <https://zlib.net>.
- [266] "Papers and private communications." In: *1st Int. Work. Data Reduct. Big Sci. Data*. Ed. by S Klasky, G Liu, M Ainsworth, et al. Frankfurt, 2017. DOI: 10.1007/978-3-319-67630-2. URL: <https://web.njit.edu/~qliu/drbsd.html>.
- [267] Gunderson, S H., A Evlogimenos, et al. *Snappy 1.1.1 - a fastcompressor/decompressor*. 2011. URL: <https://github.com/google/snappy>.
- [268] Y Collet, P Skibinski, N Terrell, et al. *Zstandard (zstd) 1.1.4 - fast real-time compression algorithm*. 2014. URL: <https://github.com/facebook/zstd>.
- [269] F Alted. *blosc 1.11.4-dev*. 2017. URL: <https://github.com/Blosc/c-blosc>.

- [270] P Lindstrom. "Fixed-rate compressed floating-point arrays". In: *IEEE Trans. Vis. Comput. Graph.* 20.12 (2014), pp. 2674–2683. DOI: 10.1109/TVCG.2014.2346458.
- [271] Y Collet. *lz4 - extremely fast compression algorithm*. 2011. URL: <http://www.lz4.org/>.
- [272] A Huebl, R Lehe, JL Vay, et al. *openPMD 1.0.0: A meta data standard for particle and mesh based data*. en. 2015. DOI: 10.5281/zenodo.33624.
- [273] M Könnecke, FA Akeroyd, HJ Bernstein, et al. "The NeXus data format". In: *J. Appl. Crystallogr.* 48.1 (2015), pp. 301–305. DOI: 10.1107/S1600576714027575.
- [274] S Klasky, Q Liu, N Podhorszki, et al. *ADIOS Visualization Schema, Version 1.1: User's Manual*. 2013. URL: <https://users.nccs.gov/~pnorb/ADIOS-VisualizationSchema-1.1.pdf>.
- [275] S Shasharina, JR Cary, S Veitzer, et al. "VizSchema - visualization interface for scientific data". In: *IADIS Int. Conf. Comput. Graph. Vis. Comput. Vis. Image Process.* 2009, p. 49.
- [276] Internet Engineering Task Force. *XML Media Types, RFC 7303*. 2014. URL: <https://www.rfc-editor.org/rfc/rfc7303.txt>.
- [277] Opeengeospatial.org. *OGC standard netCDF Classic and 64-bit Offset*. 1989. URL: <https://www.opeengeospatial.org/standards/netcdf>.
- [278] H Abbasi, M Wolf, G Eisenhauer, et al. "DataStager: Scalable data staging services for petascale applications". In: *Cluster Comput.* 13.3 (2010), pp. 277–290. DOI: 10.1007/s10586-010-0135-6.
- [279] C Docan, M Parashar, and S Klasky. "DataSpaces: An interaction and coordination framework for coupled simulation workflows". In: *Cluster Comput.* 15.2 (2012), pp. 163–181. DOI: 10.1007/s10586-011-0162-y.
- [280] A Huebl, R Lehe, JL Vay, et al. *openPMD 1.1.0: Base paths for mesh- and particle-only files and updated attributes*. 2018. DOI: 10.5281/zenodo.1167843.
- [281] DB Siano. "Orientational analysis-a supplement to dimensional analysis-I". In: *J. Franklin Inst.* 320.6 (1985), pp. 267–283. DOI: 10.1016/0016-0032(85)90031-6.
- [282] RA Fonseca, LO Silva, FS Tsung, et al. "OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators". In: *Lect. Notes Comput. Sci.* 2331 (2002), pp. 342–351. DOI: 10.1007/3-540-47789-6_36.
- [283] P Yu, X Xu, VK Decyk, et al. "Modeling of laser wakefield acceleration in Lorentz boosted frame using EM-PIC code with spectral solver". In: *J. Comput. Phys.* 266 (2014), pp. 124–138. DOI: 10.1016/j.jcp.2014.02.016.
- [284] MJ Turk, BD Smith, JS Oishi, et al. "Yt: A multi-code analysis toolkit for astrophysical simulation data". In: *Astrophys. Journal, Suppl. Ser.* 192.1 (2011). DOI: 10.1088/0067-0049/192/1/9.
- [285] D Sagan. "Bmad: A relativistic charged particle simulation library". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 558.1 (2006), pp. 356–359. DOI: 10.1016/j.nima.2005.11.001.
- [286] A Huebl, R Widera, B Worpitz, et al. *PICongPU 0.4.3: System Updates and Bug Fixes*. 2019. DOI: 10.5281/zenodo.2565503.

- [287] *Helmholtz International Beamline for Extreme Fields at the European XFEL (HIBEF) Consortium*. URL: <https://www.hzdr.de/hgfbeamline>.
- [288] EA Schneidmiller and MV Yurkov. "Baseline Parameters of the European XFEL". In: *38th Int. Free Electron Laser Conf.* (2017), pp. 109–112. DOI: 10.18429/JACoW-FEL2017-MOP033.
- [289] The European Cluster of Advanced Laser Light Sources (EUCALL). *EUCALL - Achievements and Outcome*. 2018. URL: <https://www.eucall.eu/news/brochure/>.
- [290] A Grund. "A highly-parallel Monte-Carlo-Simulation of X-Ray- Scattering using a Particle-Mesh-Code on GPUs". Diploma thesis. Technische Universität Dresden, 2016. DOI: 10.5281/zenodo.1164654.
- [291] M Manetti, A Buzmakov, L Samoylova, et al. "FAST-XPD: XFEL photon pulses database for modeling XFEL experiments". In: *AIP Conf. Proc.* 2054 (2019), p. 030019. DOI: 10.1063/1.5084582. URL: <https://in.xfel.eu/xpd/>.
- [292] R Pausch. "Synthetic radiation diagnostics as a pathway for studying plasma dynamics - from advanced accelerators to astrophysical observations". PhD thesis. Technische Universität Dresden, 2019.
- [293] E Esarey and P Sprangle. "Overview of plasma-based accelerator concepts". In: *IEEE Trans. Plasma Sci.* 24.2 (1996), pp. 252–288. DOI: 10.1109/27.509991.
- [294] A Kageyama and T Yamada. "An approach to exascale visualization: Interactive viewing of in-situ visualization". In: *Comput. Phys. Commun.* 185.1 (2014), pp. 79–85. DOI: 10.1016/j.cpc.2013.08.017.
- [295] J Ahrens, S Jourdain, P O'Leary, et al. "An Image-based Approach to Extreme Scale in Situ Visualization and Analysis". In: *SC'14 Proc. Int. Conf. High Perform. Comput. Networking, Storage Anal.* Piscataway, NJ, USA: IEEE Press, 2014, pp. 424–434. DOI: 10.1109/SC.2014.40.
- [296] T Berners-Lee and E al. *World Wide Web Consortium (W3C)*. URL: <https://www.w3.org/standards/>.
- [297] CG Durfee and HM Milchberg. "Light pipe for high intensity laser pulses". In: *Phys. Rev. Lett.* 71.15 (1993), pp. 2409–2412. DOI: 10.1103/PhysRevLett.71.2409.
- [298] R Brun and F Rademakers. "ROOT - An Object Oriented Data Analysis Framework". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 389.1-2 (1997), pp. 81–86. DOI: 10.1016/S0168-9002(97)00048-X.
- [299] I Antcheva, M Ballintijn, B Bellenot, et al. "ROOT - A C++ Framework for PetaByte Data Storage, Statistical Analysis and Visualization". In: *Comput. Phys. Commun.* 180.12 (2009), pp. 2499–2512. DOI: 10.1016/j.cpc.2009.08.005.
- [300] V Vasilev, P Canal, A Naumann, et al. "Cling - The New Interactive Interpreter for ROOT 6". In: *J. Phys. Conf. Ser.* 396.5 (2012). DOI: 10.1088/1742-6596/396/5/052071.
- [301] S Ehrig. "Interaktive GPU Programmierung mit dem Cling-C++-Interpreter in Lehre und Wissenschaft". Diploma thesis. Technische Universität Dresden, 2018.

- [302] J Wu, A Belevich, E Bendersky, et al. "gpucc: An Open-Source GPGPU Compiler". In: *Proc. 2016 Int. Symp. Code Gener. Optim.* Barcelona, Spain, 2016, pp. 105–116. DOI: 10.1145/2854038.2854041.
- [303] S Corlay, L Gouarin, J Mabile, et al. *xeus-cling: Jupyter kernel for the C++ programming language*. 2017. URL: <https://github.com/QuantStack/xeus-cling>.
- [304] H Finkel, D Poliakoff, and DF Richards. "ClangJIT: Enhancing C++ with Just-in-Time Compilation". In: (2019). arXiv: 1904.08555.
- [305] Khronos Group. *SYCL: C++ Single-source Heterogeneous Programming for OpenCL*. 2015. URL: <https://www.khronos.org/sycl/>.
- [306] RT Fielding. "Architectural Styles and the Design of Network-based Software Architectures". PhD thesis. University of California, Irvine, 2000.
- [307] GE Hinton, TJ Sejnowski, and TA Poggio. *Unsupervised Learning: Foundations of Neural Computation*. A Bradford Book. MCGRAW HILL BOOK Company, 1999.
- [308] TM Ostermayr, J Gebhard, D Haffa, et al. "A transportable Paul-trap for levitation and accurate positioning of micron-scale particles in vacuum for laser-plasma experiments". In: *Rev. Sci. Instrum.* 89.1 (2018). DOI: 10.1063/1.4995955.
- [309] W Paul. "Electromagnetic Traps for Charged and Neutral Particles - Nobel Lecture". In: *Nobel Lect. Phys. 1981-1990*. Ed. by T Frängsmyr and G Ekspong. Singapore: World Scientific Publishing Co., 1989. URL: <https://www.nobelprize.org/prizes/physics/1989/paul/lecture/>.
- [310] T Ditmire, J Zweiback, VP Yanovsky, et al. "Nuclear fusion from explosions of femtosecond laser-heated deuterium clusters". In: *Nature* 398 (1999), pp. 489–492. DOI: 10.1038/19037.
- [311] Y Fukuda, AY Faenov, M Tampo, et al. "Energy Increase in Multi-MeV Ion Acceleration in the Interaction of a Short Pulse Laser with a Cluster-Gas Target". In: *Phys. Rev. Lett.* 103.16 (2009), pp. 1–4. DOI: 10.1103/physrevlett.103.165002.
- [312] S Karsch, S Düsterer, H Schwoerer, et al. "High-Intensity Laser Induced Ion Acceleration from Heavy-Water Droplets". In: *Phys. Rev. Lett.* 91.1 (2003), pp. 2–5. DOI: 10.1103/PhysRevLett.91.015001.
- [313] S Ter-Avetisyan, M Schnürer, PV Nickles, et al. "Quasimonoenergetic deuteron bursts produced by ultraintense laser pulses". In: *Phys. Rev. Lett.* 96.14 (2006), pp. 1–4. DOI: 10.1103/PhysRevLett.96.145006.
- [314] T Sokollik, M Schnürer, S Steinke, et al. "Directional Laser-Driven Ion Acceleration from Microspheres". In: *Phys. Rev. Lett.* 103.13 (2009), pp. 25–28. DOI: 10.1103/physrevlett.103.135003.
- [315] F Balling and J Gebhard. *private communication*. 2019.
- [316] T Sokollik, T Paasch-Colberg, K Gorling, et al. "Laser-driven ion acceleration using isolated mass-limited spheres". In: *New J. Phys.* 12 (2010). DOI: 10.1088/1367-2630/12/11/113013.

- [317] F Wagner, S Bedacht, A Ortner, et al. "Pre-plasma formation in experiments using petawatt lasers". In: *Opt. Express* 22.24 (2014), p. 29505. DOI: 10.1364/oe.22.029505.
- [318] D Gabor. "A new microscopic principle". In: *Nature* 4098 (1948), pp. 777–778. DOI: 10.1038/161777a0.
- [319] A Huebl, P Hilz, and T Kluge. *Supplementary Materials "Isolated proton bunch acceleration by a petawatt laser pulse": 3D3V Input Files and Plot Data Figure 3*. 2017. DOI: 10.5281/zenodo.1005729.
- [320] C Joshi, MC Richardson, and GD Enright. "Quantitative measurements of fast ions from CO₂laser-produced plasmas". In: *Appl. Phys. Lett.* 34.10 (1979), pp. 625–627. DOI: 10.1063/1.90640.
- [321] F Begay and DW Forslund. "Acceleration of multi-species ions in CO₂laser-produced plasmas: Experiments and theory". In: *Phys. Fluids* 25.9 (1982), pp. 1675–1685. DOI: 10.1063/1.863926.
- [322] M Allen, Y Sentoku, P Audebert, et al. "Proton spectra from ultraintense laser-plasma interaction with thin foils: Experiments, theory, and simulation". In: *Phys. Plasmas* 10.8 (2003), pp. 3283–3289. DOI: 10.1063/1.1592154.
- [323] AJ Kemp and H Ruhl. "Multispecies ion acceleration off laser-irradiated water droplets". In: *Phys. Plasmas* 12.3 (2005), pp. 1–10. DOI: 10.1063/1.1856933.
- [324] J Psikal, VT Tikhonchuk, J Limpouch, et al. "Ion acceleration by femtosecond laser pulses in small multispecies targets". In: *Phys. Plasmas* 15.5 (2008). DOI: 10.1063/1.2913264.
- [325] PA Tipler and G Mosca. *Physik für Wissenschaftler und Ingenieure*. 2nd ed. Springer, 2007.
- [326] M Kaluza, J Schreiber, MIK Santala, et al. "Influence of the laser prepulse on proton acceleration in thin-foil experiments." In: *Phys. Rev. Lett.* 93.4 (2004), p. 045003. DOI: 10.1103/PhysRevLett.93.045003.
- [327] P Gibbon and E Förster. "Short-pulse laser - plasma interactions". In: *Plasma Phys. Control. Fusion* 38.6 (1996), pp. 769–793. DOI: 10.1088/0741-3335/38/6/001.
- [328] Y Sentoku, E D'Humières, L Romagnani, et al. "Dynamic control over mega-ampere electron currents in metals using ionization-driven resistive magnetic fields". In: *Phys. Rev. Lett.* 107.13 (2011), pp. 1–5. DOI: 10.1103/PhysRevLett.107.135005.
- [329] VT Tikhonchuk, AA Andreev, SG Bochkarev, et al. "Ion acceleration in short-laser-pulse interaction with solid foils". In: *Plasma Phys. Control. Fusion*. Vol. 47. 12 B. 2005. DOI: 10.1088/0741-3335/47/12B/S69.
- [330] J Limpouch, J Psikal, VT Tikhonchuk, et al. "Laser acceleration of ions in mass-limited multi-species targets". In: *J. Phys. Conf. Ser.* 112.Part 4 (2008), p. 42033. DOI: 10.1088/1742-6596/112/4/042033.
- [331] I Prencipe, J Fuchs, S Pascarelli, et al. "Targets for high repetition rate laser facilities: Needs, challenges and perspectives". In: *High Power Laser Sci. Eng.* 5.March (2017). DOI: 10.1017/hpl.2017.18.

- [332] TZ Esirkepov, K Nishihara, VS Khoroshkov, et al. "Proposed Double-Layer Target for the Generation of High-Quality Laser-Accelerated Ion Beams". In: *Phys. Rev. Lett.* 89.17 (2002). DOI: 10.1103/physrevlett.89.175003.
- [333] Y Sentoku, TE Cowan, A Kemp, et al. "High energy proton acceleration in interaction of short laser pulse with dense plasma target". In: *Phys. Plasmas* 10.5 (2003), pp. 2009–2015. DOI: 10.1063/1.1556298.
- [334] A Huebl, M Rehwald, L Obst-Huebl, et al. *Supplementary Data: Spectral Control via Multi-Species Effects in PW-Class Laser-Ion Acceleration*. 2019. DOI: 10.14278/rodare.116.
- [335] AV Brantov, VT Tikhonchuk, O Klimo, et al. "Quasi-mono-energetic ion acceleration from a homogeneous composite target by an intense laser pulse". In: *Phys. Plasmas* 13.12 (2006). DOI: 10.1063/1.2404928.

Acknowledgments

Ich möchte mich herzlich bei allen Menschen bedanken, die mich in den Jahren der Doktorarbeit, aber auch auf dem langen Weg dahin unterstützt haben. Vielen Dank an Ulrich Schramm und Thomas Cowan für das hervorragende Arbeitsumfeld in unserem Institut und die Förderung unserer Aktivitäten über viele Jahre. Euer Feedback und weitergegebene Erfahrung in Diskussionen hat mich stets inspiriert und motiviert.

Unter Michael Bussmanns Leitung durfte ich 10 Jahre wissenschaftlich arbeiten, Forschen (lernen) und Physik mit Supercomputern verbinden. Michael, vielen Dank für all deine Unterstützung, unseren vielen Austausch und deine Inspiration, die ungewöhnliche Freiheit und Förderung zum Arbeiten, Verknüpfen, Vorausschauen, Fortschicken, Ent-, Verwerfen und Aufbauen. Rückblickend bin ich sichtlich bewegt was wir unter deiner Leitung alles in so kurzer Zeit erreicht haben.

Thomas Kluge danke ich für seine Geduld und sein großzügig geteiltes Wissen zu Laser-Ionenphysik, Feedback und Hilfe bei meinen Ergebnissen und Schreibereien, und eine richtig angenehme, freundschaftliche Zusammenarbeit. Ich habe viel gelernt bei dir, danke! Meinem langjährigen Kollegen René Widera danke ich für die gemeinsamen großartigen Projekte und sein immenses Engagement, was mir immer eine große Hilfe war. Deine freundschaftliche und fachlich großartige Kompetenz, sowie die vielen Ergebnisse die aus Diskussionen in gemeinsamen Projekten entstanden, sind eine tolle Erfahrung.

Richard danke ich für die vielen Diskussionen und tollen gemeinsamen Arbeiten, sowie den Kampfgeist aus dem KHI-"Gewobbel" etwas Ordentliches zu machen. Vielen Dank auch an Klaus, Marco, Alexander Debus, João, Ilja und Jan für die vielen Diskussionen die wir zu Physik und Simulationsprojekten geführt haben. Thank you as well to Sergei for interesting discussions. Jan, ebenso vielen Dank für konstruktives Feedback zu großen Teilen dieser Arbeit. Danke an meine vormaligen Kollegen Heiko und Felix, für die schöne und produktive Zeit als PIconGPU noch neben unserem Studium entwickelt wurde. Vielen Dank auch an alle Freunde vom Klettern, Fußball, Quizzen, GoT-Schauen und meine viel zu remote Nerdbeergruppe für die schöne Zeit ohne Arbeit. Danke an Jurjen, Rengin, Raphael, Josefine, Flo K., Viki, Axel J., Flo B., Fabian R., Malte, Max B., Nick, Alex K., Melanie, Martin R., Max L., Andrea, Christoph E., René W., René G., Karl, Daniel J., Clemens, Simon, Markus, Manfred, Lars, Julia, Cliò, Alex L., Goofy, Sophia, Jupp und Marius.

Vielen Dank an Alexander M., Erik und Carlchristian für die langjährige gemeinsame Arbeit und starken Projekte die ihr geschaffen habt. Die Zusammenarbeit mit euch war super, danke für die mallocMC Implementierung an Carlchristian. Thank you Ksenia, for the code coupling with openPMD and great work on particle reduction methods. Bei den vielen StudentInnen die im Laufe dieser Dissertation Projekte und Arbeiten erstellt haben und mit denen ich die Freude hatte zu arbeiten bedanke ich mich herzlich, es hat mir viel Spaß gemacht und ich hoffe ihr empfanget die Zusammenarbeit ebenso angenehm. Benjamin, deine Arbeit an Alpaka war und ist herausragend und es ist großartig was du immernoch daran schaffst. Danke auch an Matthias und noch einmal an Erik und Alexander M. für aktive Weiterentwicklungen. An Letzteren ebenso danke für die schöne Zusammenarbeit und Umsetzung von ISAAC. Danke noch einmal an Marco, für seine tolle Implementation

und Entwicklung von Ionisationsphysik in PIconGPU. Vielen Dank an Simeon, den ich nach dreijähriger Suche gefunden habe um Cling mit CUDA Fähigkeiten auszustatten. Dein Engagement und deine Ergebnisse sind großartig. Vielen Dank an Axel Naumann und Vassil Vassilev, die ihn am CERN fachlich betreuten und an Robert Dietrich am Lehrstuhl von Prof. Nagel an der Technischen Universität. Sophie, Sebastian und Jeffrey danke ich für die tolle Zusammenarbeit am Jupyter Frontend, aber auch ganz allgemein an Guido Juckeland und sein Team. Vielen Dank an Michael S. für die Umsetzung des neuen Eventsystems, ich bin überzeugt das wird richtig spannend für Anwendungen. Thanks to Aleksandar B. who worked with us on directional splitting as a summer student. Danke an Alexander G. für die Einarbeitung in unseren Softwarestack, seine Beiträge und die Umsetzung von ParaTAXIS. Fabian and Franz danke ich für die tolle Zusammenarbeit an openPMD-api, welches noch viele datenintensive Projekte ermöglichen wird.

Der Austausch mit unserem experimentellen Ionenteam hat mir große Freude bereitet und mich stets inspiriert. Danke an die spannenden Diskussionen und Realitätsabgleiche mit Karl, Martin, natürlich Lotti, Tim, Josefine, Flo und Flo, Stephan und HP. Speziell auch noch einmal Danke an Karl für dein hilfreiches inhaltliches Feedback bei gemeinsamen Projekten und viele spannende Gespräche. Nicht zuletzt auch Danke für leckere Mate. Vielen Dank auch für die Diskussionen und Gespräche mit dem Elektronenteam bestehend aus Arie, Jurjen, Thomas, Alexander Köhler, Omid und Susanne. Ebenso herzlichen Dank an Peter Hilz und Jörg Schreiber für die erfolgreiche und langfristige Zusammenarbeit an einem sehr spannenden Thema mit euren PHELIX Experimenten. Auch wenn es eine echte Herausforderung war ist das Ergebnis toll und ich danke auch Felix und Johannes für den spannenden Austausch zu weiteren Kampagnen.

Danke an Carsten Fortmann-Grote für die angenehme und produktive Zusammenarbeit in EUCALL. Thank you very much to Hyun-Kyung Chung for the great support with teaching me non-LTE methods. I hope we can continue the work on this topic, as it is complex and exciting. Thank you to our long-term collaborators Scott Klasky, Norbert Podhorszki, David Pugmire and Michael A. Matheson for continuous support and stimulating exchange on truly challenging data workflows. I am much obligated to Rémi Lehe for our longstanding pleasant correspondence and continued partnership on many projects. I am also grateful for the repeated warm welcome he and Jean-Luc Vay provided me in Berkeley whenever we found the time. Thank you to all people driving openPMD, especially Michael Bussmann, Rémi Lehe, Jean-Luc Vay, Stephan Kuschel as well as David C. Sagan, Christopher Mayes, Frédéric Pérez, and Fabian Koller. Thank you as well to all further implementers and adapting users. I would like to thank everyone whom I met at conferences and/or online that discussed scientifically with me, guided and inspired me, such as: Todd Gamblin (LLNL), Sylvain Corlay, Johan Mabilie, Wolf Vollprecht (all with QuantStack), Ray Donnelly (Anaconda), Laurent Villard, Stephan Brunner, Noé Ohana (all with EPFL), Eric Sonnendrücker, Michael Kraus, Katharina Kormann, Jakob Ameres, Nils Moschüring (all with IPP), Sunita Chandrasekaran (UDEL), and much more.

I would like to personally thank all already named and unnamed contributors to our open-source projects, early adopters, power users, former colleagues and recently jointed collaborators. Your input and improvements are invaluable, thank you for your contributions! Furthermore, a big thank you to all the existing open source projects on which I rely

for my daily work: the amazing ADIOS team at OLCF, the Python and SciPy ecosystem, the communities of CMake, Spack, Cling, LLVM/Clang, GCC, pybind11, conda-forge, Boost, MPICH, OpenMPI, and everyone who contributes to the evolution of modern C++ and scientific computing environments.

Meinen lieben Eltern, meinem Bruder und Großeltern danke ich für die langjährige Unterstützung, die wunderbare Zeit vor und in Dresden, die allzeit lieben Worte, das Fördern von Ansporn und Mut für alle meine Vorhaben. Vielen Dank für alles! Insbesondere freue ich mich über die herzliche Erweiterung unserer Familie und das aufgeschlossene Aufnehmen bei Jayne und Michael, Kimmy und Becky. Meiner liebsten Lotti danke ich für ihr gutmütiges Korrekturlesen, ihre bewunderswerte Organisiertheit und das wir Ionenfragestellungen selbst noch zum Abendbrot bereden können. Lotti, vielen Dank für deine Unterstützung, Weitsicht, Aufgeschlossenheit, deinen Humor und das Erden auf die wichtigen Dinge im Leben.

I would like to thank the Helmholtz-Zentrum Dresden – Rossendorf and Technical University Dresden for the excellent facilities and staff providing the environment for my studies in Dresden. Thank you to all user support staff in supercomputing centers across Dresden, Jülich, Munich, Lugano, Berkeley, Osaka, Bologna, Santa Clara, and Oak Ridge, enabling the computational campaigns in this thesis. The GPU Center of Excellence Dresden provided a framework for many fruitful projects and travel funding, sponsored by the Nvidia Corporation. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This research has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 654220. This project received funding within the MEPHISTO project (BMBF-Förderkennzeichen 01IH16006C).

Erklärungen

Hiermit versichere ich, dass ich die vorliegende Arbeit ohne unzulässige Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe; die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sind als solche kenntlich gemacht. Die Arbeit wurde bisher weder im Inland noch im Ausland in gleicher oder ähnlicher Form einer anderen Prüfungsbehörde vorgelegt.

Die Dissertation wurde unter der wissenschaftlichen Betreuung von Prof. Dr. Ulrich Schramm, Dr. Michael Busmann und Dr. Thomas Kluge am Helmholtz-Zentrum Dresden – Rossendorf angefertigt.

Die Promotionsordnung der Fakultät Mathematik und Naturwissenschaften an der Technischen Universität Dresden vom 23. Februar 2011 mit letzten Änderungen vom 18. Juni 2014 und 23. Mai 2018 erkenne ich an.

Dresden, 25. April 2019

Dipl.-Phys. Axel Huebl