HZDR-124

TEMPORAL CONTRAST-DEPENDENT MODELING OF LASER-DRIVEN SOLIDS

STUDYING FEMTOSECOND-NANOMETER INTERACTIONS AND PROBING

Marco Garten

Wissenschaftlich-Technische Berichte
HZDR-124 · 2022 · ISSN 2191-8708

WISSENSCHAFTLICH-
TECHNISCHE BERICHTE

HZDR
HELMHOLTZ ZENTRUM
DRESDEN ROSSENDORF

Wissenschaftlich-Technische Berichte
HZDR-124

Marco Garten

TEMPORAL CONTRAST-DEPENDENT MODELING OF LASER-DRIVEN SOLIDS

STUDYING FEMTOSECOND-NANOMETER INTERACTIONS AND PROBING

HZDR
HELMHOLTZ ZENTRUM
DRESDEN ROSSENDORF

Druckausgabe: ISSN 2191-8708

Elektronische Ausgabe: ISSN 2191-8716

https://www.hzdr.de/publications/Publ-34551

urn:nbn:de:bsz:d120-qucosa2-839958

Die vorliegende Arbeit wurde sowohl als Dissertation an der Fakultät Mathematik und Naturwissenschaften der Technischen Universität Dresden sowie als Wissenschaftlich-Technischer Bericht des Helmholtz-Zentrum Dresden – Rossendorf mit der Berichtsnummer **HZDR-124** veröffentlicht.

2023

Herausgegeben vom

Helmholtz-Zentrum Dresden - Rossendorf

Bautzner Landstraße 400

01328 Dresden

Germany

TECHNISCHE UNIVERSITÄT DRESDEN

HZDR
HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Fakultät Mathematik und Naturwissenschaften Fachrichtung Physik

Institut für Strahlenphysik

Temporal contrast-dependent modeling of laser-driven solids

studying femtosecond-nanometer interactions and probing

Marco Garten

Born on: 24th September 1990 in Wolfen

Dissertation

to achieve the academic degree

Doctor rerum naturalium (Dr. rer. nat.)

Referees

Prof. Dr. Ulrich Schramm
Prof. Dr. Thomas E. Cowan
Prof. Dr. Jörg Schreiber

Supervisor

Dr. Thomas Kluge

Submitted on: 20th January 2022
Defended on: 13th May 2022

DRESDEN concept
Exzellenz aus Wissenschaft und Kultur

Für meine Familie und besonders für meine lieben Großeltern.

TECHNISCHE UNIVERSITÄT DRESDEN

HZDR
HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Fakultät Mathematik und Naturwissenschaften Fachrichtung Physik

Institut für Strahlenphysik

Kurzfassung

Die Erlangung präziser Kontrolle über die einzigartigen Strahlparameter von laserbeschleunigten Ionen aus relativistischen Ultrakurzpuls-Laser-Festkörper-Wechselwirkungen ist ein wesentliches Ziel der letzten 20 Jahre. Während die räumlich-zeitliche Kopplung von Laserpuls- und Targetparametern transiente Phänomene auf Femtosekunden- und Nanometerskalen erzeugt, die für den Beschleunigungsprozess entscheidend sind, waren diese Skalen der experimentellen Beobachtung bisher weitgehend unzugänglich. Computersimulationen von lasergetriebenen Plasmen liefern dabei wertvolle Einblicke in die zugrunde liegende Physik. Dennoch mangelt es noch an Vorhersagemöglichkeiten aufgrund des massiven Rechenaufwands, um Parameterstudien in 3D mit hoher Auflösung für längere Simulationszeiten durchzuführen. In dieser Arbeit wird die optimale Beschleunigung von Protonen aus ultradünnen Folien nach der Wechselwirkung mit einem ultrakurzen Ultrahochintensitäts-Laserpuls unter Einbeziehung realistischer Kontrastbedingungen bis zu einer Pikosekunde vor dem Hauptpuls untersucht. Hierbei ermöglichen neu implementierte fortschrittliche Ionisierungsmethoden für den hoch skalierbaren, quelloffenen Partikel-in-Zelle-Code PIConGPU von nun an Studien dieser Art. Bei der Unterstützung zweier Experimentalkampagnen führten diese Methoden zu einem tieferen Verständnis der Laser-Wakefield-Beschleunigung bzw. des Schmelzens kolloidaler Kristalle, da nun experimentelle Beobachtungen mit simulierter Ionisations- und Plasmadynamik erklärt werden konnten. Im Anschluss werden explorative 3D3V-Simulationen verbesserter Laser-Ionen-Beschleunigung vorgestellt, die auf dem Schweizer Supercomputer Piz Daint durchgeführt wurden. Dabei veränderte die Einbeziehung realistischer Laserkontrastbedingungen die Intrapulsdynamik des Beschleunigungsprozesses signifikant. Im Gegensatz zu einem perfekten Gauß-Puls erlaubte eine bessere räumlich-zeitliche Überlappung der Protonen mit dem Ursprung der Elektronenwolke die volle Ausnutzung des Beschleunigungspotentials, was zu höheren maximalen Energien führte. Die Adaptation bekannter analytischer Modelle erlaubte es, die Ergebnisse qualitativ und in ausgewählten Fällen auch quantitativ zu bestätigen. Trotz der in den 1D-Modellen nicht abgebildeten komplexen 3D-Plasmadynamik zeigt die Vorhersage erstaunlich gut das obere Limit der erreichbaren Ionen-Energien im TNSA-Szenario. Strahlungssignaturen, die aus synthethischen Diagnostiken von Elektronen, Protonen und Bremsstrahlungsphotonen gewonnen wurden, zeigen, dass der Target-Zustand bei maximaler Laserintensität einkodiert ist, was einen Ausblick darauf gibt, wie Experimente Einblicke in dieses bisher unbeobachtbare Zeitfenster gewinnen können. Mit neuen Freie-Elektronen-Röntgenlasern sind Beobachtungen auf Femtosekunden-Nanometerskalen endlich zugänglich geworden. Damit liegt ein Benchmarking der physikalischen Modelle für Plasmasimulationen bei Festkörperdichte nun in Reichweite, aber Experimente sind immer noch selten, komplex, und schwer zu interpretieren. Zuletzt werden daher in dieser Arbeit die ersten Start-zu-End-Simulationen der Pump-Probe Wechselwirkungen von optischem sowie Röntgenlaser mit Festkörpern mittels des Photonenstreu-Codes ParaTAXIS vorgestellt. Darüber hinaus dienten die zugehörigen PIC-Simulationen als Grundlage für die Planung und Durchführung eines LCLS-Experiments zur erstmaligen Beobachtung einer durch nah-relativistische Kurzpuls-Laserpulse getriebenen Festkörper-Plasma-Dichte, dessen Auflösungsbereich gleichzeitig bis auf Femtosekunden und Nanometer vordrang.

TECHNISCHE UNIVERSITÄT DRESDEN

HZDR
HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Fakultät Mathematik und Naturwissenschaften Fachrichtung Physik

Institut für Strahlenphysik

Abstract

Establishing precise control over the unique beam parameters of laser-accelerated ions from relativistic ultra-short pulse laser-solid interactions has been a major goal for the past 20 years. While the spatio-temporal coupling of laser-pulse and target parameters create transient phenomena at femtosecond-nanometer scales that are decisive for the acceleration performance, these scales have also largely been inaccessible to experimental observation. Computer simulations of laser-driven plasmas provide valuable insight into the physics at play. Nevertheless, predictive capabilities are still lacking due to the massive computational cost to perform these in 3D at high resolution for extended simulation times. This thesis investigates the optimal acceleration of protons from ultra-thin foils following the interaction with an ultra-short ultra-high intensity laser pulse, including realistic contrast conditions up to a picosecond before the main pulse. Advanced ionization methods implemented into the highly scalable, open-source particle-in-cell code PIConGPU enabled this study. Supporting two experimental campaigns, the new methods led to a deeper understanding of the physics of Laser-Wakefield acceleration and Colloidal Crystal melting, respectively, for they now allowed to explain experimental observations with simulated ionization- and plasma dynamics. Subsequently, explorative 3D3V simulations of enhanced laser-ion acceleration were performed on the Swiss supercomputer Piz Daint. There, the inclusion of realistic laser contrast conditions altered the intra-pulse dynamics of the acceleration process significantly. Contrary to a perfect Gaussian pulse, a better spatio-temporal overlap of the protons with the electron sheath origin allowed for full exploitation of the accelerating potential, leading to higher maximum energies. Adapting well-known analytic models allowed to match the results qualitatively and, in chosen cases, quantitatively. However, despite complex 3D plasma dynamics not being reflected within the 1D models, the upper limit of ion acceleration performance within the TNSA scenario can be predicted remarkably well. Radiation signatures obtained from synthetic diagnostics of electrons, protons, and bremsstrahlung photons show that the target state at maximum laser intensity is encoded, previewing how experiments may gain insight into this previously unobservable time frame. Furthermore, as X-ray Free Electron Laser facilities have only recently begun to allow observations at femtosecond-nanometer scales, benchmarking the physics models for solid-density plasma simulations is now in reach. Finally, this thesis presents the first start-to-end simulations of optical-pump, X-ray-probe laser-solid interactions with the photon scattering code ParaTAXIS. The associated PIC simulations guided the planning and execution of an LCLS experiment, demonstrating the first observation of solid-density plasma distribution driven by near-relativistic short-pulse laser pulses at femtosecond-nanometer resolution.

Publications by the Author

[G1] U Schramm et al. "First results with the novel petawatt laser acceleration facility in Dresden". In: *Journal of Physics: Conference Series*. Vol. 874. 1. 2017. DOI: 10.1088/1742-6596/874/1/012028 (cit. on pp. 3, 4, 25–27, 77, 78).

[G2] I Prencipe, J Metzkes-Ng, A Pazzaglia, C Bernert, D Dellasega, L Fedeli, A Formenti, M Garten, T Kluge, S Kraft, A Laso Garcia, A Maffini, L Obst-Huebl, M Rehwald, M Sobiella, K Zeil, U Schramm, TE Cowan, and M Passoni. "Efficient laser-driven proton and Bremsstrahlung generation from cluster-assembled foam targets". In: *New Journal of Physics* 11.20 (2021), pp. 5035–5040. ISSN: 1367-2630. DOI: 10.1088/1367-2630/ac1fcd (cit. on pp. 3, 38).

[G3] 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: *Nature Communications* 8.1 (2017), p. 487. ISSN: 2041-1723. DOI: 10.1038/s41467-017-00592-7 (cit. on pp. 4, 7, 64, 66).

[G4] N Mukharamova, S Lazarev, JM Meijer, OY Gorobtsov, A Singer, M Chollet, M Bussmann, D Dzhigaev, Y Feng, M Garten, A Huebl, T Kluge, RP Kurta, V Lipp, R Santra, M Sikorski, S Song, G Williams, D Zhu, B Ziaja-Motyka, TE Cowan, AV Petukhov, and IA Vartanyants. "Femtosecond laser produced periodic plasma in a colloidal crystal probed by XFEL radiation". In: *Scientific Reports* 10.1 (2020), p. 10780. ISSN: 2045-2322. DOI: 10.1038/s41598-020-67214-z. arXiv: 1911.03752 (cit. on pp. 7, 64, 70–73).

[G5] M Garten. "Modellierung und Validierung von Feldionisation in parallelen Particle-in-Cell-Codes". Master. TU Dresden, 2015, p. 97. DOI: 10.5281/zenodo.202500 (cit. on pp. 15, 38, 63, 67).

[G6] M Garten. "Comparing field ionization models in simulations of laser-matter interaction". Bachelor. TU Dresden, 2013, p. 71. DOI: 10.5281/zenodo.192109 (cit. on pp. 47, 63, 67).

[G7] 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, and M Bussmann. *ComputationalRadiationPhysics/picongpu: Perfectly Matched Layer (PML) and Bug Fixes*. 2020. DOI: 10.5281/ZENODO.3875374 (cit. on pp. 47, 132).

[G8] T Kluge, M Rödel, J Metzkes-Ng, A Pelka, AL Garcia, I Prencipe, M Rehwald, M Nakatsutsumi, EE McBride, T Schönherr, M Garten, NJ Hartley, M Zacharias, J Grenzer, A Erbe, YM Georgiev, E Galtier, I Nam, HJ Lee, S Glenzer, M Bussmann, C Gutt, K Zeil, C Rödel, U Hübner, U Schramm, and TE Cowan. "Observation of Ultrafast Solid-Density Plasma Dynamics Using Femtosecond X-Ray Pulses from a Free-Electron Laser". In: *Physical Review X* 8.3 (2018), p. 031068. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.8.031068 (cit. on pp. 65, 119, 120, 124, 125, 130, 137, 139).

[G9] T Kluge, M Rödel, J Metzkes-Ng, A Pelka, AL Garcia, I Prencipe, M Rehwald, M Nakatsutsumi, EE McBride, T Schönherr, M Garten, NJ Hartley, M Zacharias, J Grenzer, A Erbe, YM Georgiev, E Galtier, I Nam, HJ Lee, S Glenzer, M Bussmann, C Gutt, K Zeil, C Rödel, U Hübner, U Schramm, and TE Cowan. "Observation of Ultrafast Solid-Density Plasma Dynamics Using Femtosecond X-Ray Pulses from a Free-Electron Laser". In: *Physical Review X* 8.3 (2018), p. 031068. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.8.031068. arXiv: 1801.08404 (cit. on pp. 65, 124, 125, 131).

[G10] C Fortmann-Grote, A Andreev, R Briggs, M Bussmann, M Garten, A Hübl, T Kluge, S Pascarelli, A Sharma, S Yakubov, and AP Mancuso. *EUCALL Deliverable D4.3 - Interoperability of simulation workflows*. Tech. rep. 654220. 2017, pp. 1–18. DOI: 10.5281/zenodo.998647 (cit. on pp. 65, 135–138).

[G11] C Fortmann-Grote, JP Burchert, M Bussmann, M Garten, A Huebl, T Kluge, and AP Mancuso. *EUCALL Deliverable D4.4 Simulated coherent scattering data from plasma and non–plasma samples*. Tech. rep. 2017. DOI: 10.5281/ZENODO.998644 (cit. on pp. 65, 137).

[G12] C Fortmann-Grote, A Andreev, R Briggs, M Bussmann, M Garten, A Hübl, T Kluge, S Pascarelli, A Sharma, S Yakubov, and AP Mancuso. *EUCALL Deliverable D4.5 - Testing, validation, and example workflow*. Tech. rep. 654220. 2018, pp. 1–18 (cit. on pp. 65, 137).

[G13] L Gaus et al. "Probing ultrafast laser plasma processes inside solids with resonant small-angle X-ray scattering". In: (2020), pp. 1–12. arXiv: 2012.07922 (cit. on pp. 80, 131).

[G14] C Bernert, S Assenbaum, FE Brack, TE Cowan, CB Curry, M Garten, L Gaus, M Gauthier, S Göde, I Goethel, SH Glenzer, T Kluge, S Kraft, F Kroll, M Kuntzsch, J Metzkes-Ng, M Loeser, L Obst-Huebl, HP Rehwald M. Schlenvoigt, C Schoenwaelder, U Schramm, M Siebold, F Treffert, T Ziegler, and K Zeil. "Off-harmonic optical probing of high-intensity laser-plasma expansion dynamics in solid-density hydrogen jets". In: *Scientific Reports* (submitted 2021) (cit. on p. 83).

[G15] 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: *Plasma Physics and Controlled Fusion* 62.12 (2020), p. 124003. ISSN: 0741-3335. DOI: 10.1088/1361-6587/abbe33. arXiv: 1903.06428 (cit. on p. 106).

[G16] T Kluge, C Rödel, M Rödel, A Pelka, EE McBride, LB Fletcher, M Harmand, A Krygier, A Higginbotham, M Bussmann, E Galtier, E Gamboa, AL Garcia, M Garten, SH Glenzer, E Granados, C Gutt, HJ Lee, B Nagler, W Schumaker, F Tavella, M Zacharias, U Schramm, and TE Cowan. "Nanometer-scale characterization of laser-driven compression, shocks, and phase transitions, by x-ray scattering using free electron lasers". In: *Physics of Plasmas* (2017). DOI: 10.1063/1.5008289 (cit. on pp. 119, 120).

[G17] A Grund, A Huebl, and M Garten. *PARAllel — Tracer for Arbitrary X-ray Interaction and Scattering*. 2016. URL: https://github.com/ComputationalRadiationPhysics/parataxis (cit. on p. 133).

[G18] C Fortmann-Grote, AA Andreev, R Briggs, M Bussmann, A Buzmakov, M Garten, A Grund, A Hübl, 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 PLATFORM*. [Online; accessed 10-March-2020]. 2017 (cit. on pp. 133, 136, 146).

[G19] C Fortmann-Grote, AA Andreev, R Briggs, M Bussmann, A Buzmakov, M Garten, A Grund, A Hübl, 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: October (2016). DOI: 10.17199/NOBUGS2016.2121. arXiv: 1610.05980 (cit. on p. 133).

[G20] 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 Zastrau, B Ziaja-Motyka, and AP Mancuso. "Simulations of ultrafast x–ray laser experiments". In: *Proceedings of SPIE - The International Society for Optical Engineering*. Ed. by T Tschentscher and L Patthey. Vol. 10237. 2017, 102370S. ISBN: 9781510609754. DOI: 10.1117/12.2270552 (cit. on pp. 133, 136).

[G21] M Garten, A Huebl, A Grund, C Fortmann-Grote, T Kluge, and M Bussmann. *ParaTAXIS X-ray Scattering Input & Output*. 2017. DOI: 10.5281/ZENODO.885033 (cit. on p. 137).

[G22] 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: *Nature Communications* 9.1 (2018), p. 5292. ISSN: 2041-1723. DOI: 10.1038/s41467-018-07756-z (cit. on p. 145).

Contents

Glossary

ADIOS	Adaptable I/O System
ADK	tunnel ionization model developed by Ammosov, Delone, and Krainov
AMR	Adaptive Mesh Refinement
AOPDF	Acousto-Optic Programmable Dispersive Filter
ASCII	American Standard Code for Information Interchange
ASE	Amplified Spontaneous Emission
BOA	Break-Out Afterburner
BSI	barrier suppression ionization model
CAAR	Center for Accelerated Application Readiness
CFL	Courant–Friedrichs–Lewy condition
CP	circularly polarized
CPA	Chirped Pulse Amplification
CPU	central processing unit
CSA	Collisionless Shock Acceleration
CSCS	Swiss National Supercomputing Center (*Centro Svizzero di Calcolo Scientifico*)
CSPAD	Cornell–SLAC Pixel Array Detector
CUDA	Compute Unified Device Architecture
DCE	Directed Coulomb Explosion
DESY	Deutsches Elektronen-Synchrotron
DOI	Digital Object Identifier
DRACO	Dresden laser acceleration source
EoS	Equation of State
EUCALL	European Cluster of Advanced Laser Light Sources
FDTD	finite-difference-time-domain method
FS	File System
FTL	Fourier transform limited
FWHM	full width at half maximum

GPU	graphics processing unit
GUI	Graphical User Interface
GVD	group velocity dispersion
HB	Hole-Boring RPA
HDF5	Hierarchical Data Format 5
HED	High Energy Density
HIBEF	Helmholtz International Beamline for Extreme Fields
HPC	high-performance computing
HZDR	Helmholtz-Zentrum Dresden - Rossendorf
I/O	[Data] Input / Output
IR	infra-red
ISAAC	In Situ Animation of Accelerated Computations
KFR	Keldysh-Faisal-Reiss
KHI	Kelvin-Helmholtz Instability
LCLS	Linac Coherent Light Source
LCT	liquid crystal target
LP	linearly polarized
LS	Light-Sail RPA
LTE	Local Thermal Equilibrium
LWFA	laser wakefield acceleration
MD	Molecular Dynamics
MEC	Matter in Extreme Conditions
MHD	Magneto-Hydrodynamics
MPI	Message Parsing Interface
MPI	Multi-Photon Ionization
MVA	Magnetic Vortex Acceleration
OAP	Off-axis Parabola
OLCF	Oak Ridge Leadership Computing Facility
OPCPA	Optical Parametric Chirped Pulse Amplification
openPMD	Open Particle Mesh Data
ORNL	Oak Ridge National Laboratory
OSS	Object Storage Server
OST	Object Storage Target
ParaTAXIS	PARAllel Tracer for Arbitrary X-ray Interaction and Scattering
PDE	Partial Differential Equation
PIC	particle-in-cell
PICLS	2D Particle-in-Cell Code by Sentoku et al.
PIConGPU	Particle-in-Cell on Graphics Processing Units
PM	plasma mirror
PRACE	Partnership for Advanced Computing in Europe

QED	Quantum Electrodynamics
RAM	random-access memory
RCF	Radiochromic Film
RCXD	Resonant Coherent X-ray Diffraction
RF	radio-frequency
RGB	Red Green Blue color space
RIT	Relativistically Induced Transparency
RNG	Random Number Generator
RODARE	Rossendorf Data Repository
RPA	Radiation Pressure Acceleration
SAXS	Small Angle X-ray Scattering
SFA	Strong-Field Approximation
SI	International System of Units
SIMD	Single Instruction Multiple Data
SIMEX	Simulation of Experiments at advanced light sources
SLAC	Stanford Linear Accelerator Center
SLURM	Simple Linux Utility for Resource Management
SPIDER	Spectral Phase Interferometry for Direct Electric-field Reconstruction
SRSI-ETE	Self-Referenced Spectral Interferometry with Extended Time Excursion
STII	self-truncated ionization-injection
TCP	Transmission Control Protocol
TF	Thomas-Fermi
Ti:Sa	Titanium-Sapphire crystal - gain medium in contemporary ultra-high intensity short-pulse laser systems
TNSA	Target-Normal Sheath Acceleration
TOD	third order dispersion
TPS	Thomson-Parabola Spectrometer
UHI	ultra high intensity
URI	Uniform Resource Identifier
WDM	Warm Dense Matter
WHELMI	Weizmann-Helmholtz Laboratory for Laser Matter Interaction
XFEL	X-ray Free-Electron Laser
XML	Extensible Markup Language
ZIH	Zentrum für Informationsdienste und Hochleistungsrechnen der Technischen Universität Dresden

1. Introduction

1.1. Motivation

When in the early 2000s, the first dedicated experiments demonstrated the acceleration of ions to several 1–10 MeV of kinetic energy per nucleon[1–3] by focusing ultra-high intensity (UHI) laser pulses onto thin foil targets, many research groups around the world set out to replicate these results and a new and fast-evolving research field was opened. About 20 years later, the first record for the maximum observed energy of 58 MeV by Snavely et al. [1] has only been improved by less than twofold[4, 5]. While this is true for the class of kilojoule long-pulse (several 100 to 1000 fs) lasers, the progress for the now more readily available joule-class short-pulse (10s of fs) laser systems[G1] has been more significant. Since they were continuously developed to reach petawatt (PW) laser powers[6, 7], they now reliably produce ion beams with energies of several 10 MeV[8–10, G2]. However, none of the available systems has so far been able to push into the regime beyond 150 MeV, as they would be required for, e.g., medical applications in radiation oncology[11–13].

Figure 1.1 shows a schematic representation of a typical laser-ion acceleration experiment. A UHI laser pulse is produced by means of chirped pulse amplification[16], is recompressed, shaped and transported in the optical beamline that eventually focuses it onto the target where it ionizes matter into the plasma state. The interaction region typically measures only a few cubic micrometers and the relevant physics of plasma heating, particle– and energy transport, charge separation and ion acceleration take place within a short time window measuring only several femto- to picoseconds[13]. The interaction with the target is highly sensitive to its density and thickness, pre-formed plasma conditions as well as the laser pulse duration, temporal shape and maximum intensity. Over the years, many different acceleration mechanisms have been identified in elaborate simulations, promising that higher energies are in principle possible[1–3, 17–20]. However, unambiguous separation of mechanisms in laboratory experiments is usually not possible since diagnostic methods do not provide the simultaneously temporal and spatial resolution that is needed to observe the characteristic signatures of the dynamics at play. On the other hand, numerical simulations usually employ simplifications and idealized parameter combinations, particularly with regard to the driving laser pulses, that cannot be matched by current experiments.

Especially with the increasing onset of instabilities for acceleration mechanisms operating at laser intensities $\gtrapprox 10^{21}\,\mathrm{W\,cm^{-2}}$[21–23] it seems that a barrier has been reached that cannot be overcome by just a raw increase in focused laser energy. At the same time, using available methods for temporal intensity contrast cleaning with ever more laser energy would cause

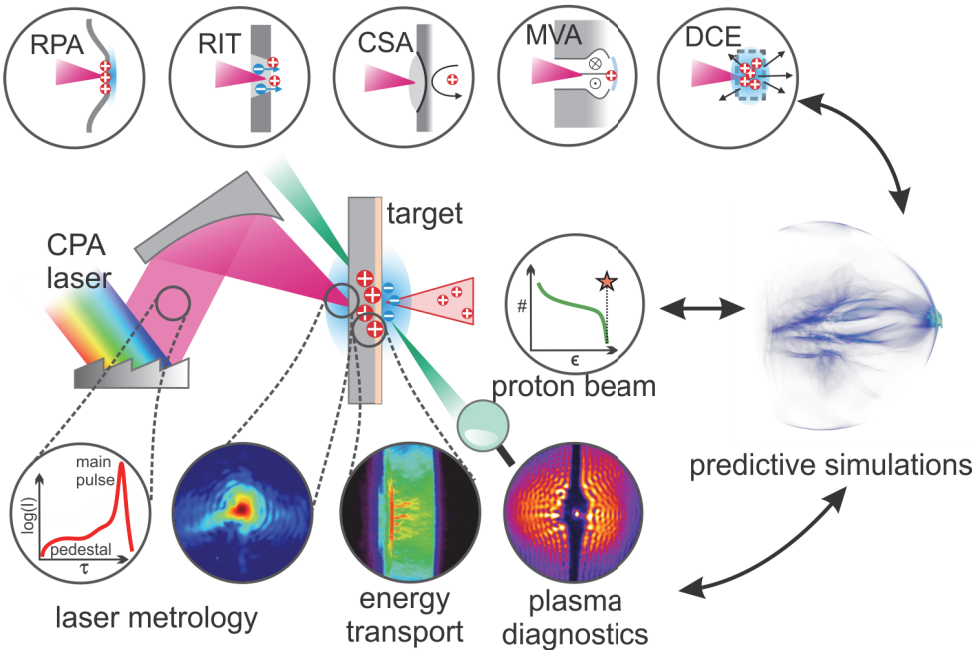

Figure 1.1.: Schematic overview over the stages, quantities and processes of interest in a UHI laser-ion acceleration experiment in the target-normal sheath acceleration (TNSA, *center image*) scenario. Depending on laser-target parameters and their spatio-temporal coupling a variety of other ion acceleration mechanisms (*in circles above*, see Sec. 2.2.1 for their disambiguation) can be dominant. Adapted from [14] and [15].

the target to ionize increasingly early before the arrival of the main pulse maximum, potentially resulting in a degraded target state that is not fit for efficient ion acceleration, anymore. Rather, the level of understanding about the intricate dynamics of laser-solid interactions needs to be improved and exploited like in the lower density counterpart of laser-electron acceleration from gas targets[24–27]. There, a promising degree of control[G3, 28] has been established that begins to put the long-standing vision of tabletop laser-electron accelerators as drivers[29] for next-generation high-brilliance light sources in reach[30].

To overcome current limitations in laser-ion acceleration, it is of vital importance to establish understanding and control over every aspect of the interaction shown in Fig. 1.1. One of the most crucial factors is the temporal intensity evolution of the laser pulse that drives the interaction. Figure 1.2 shows a so-called *laser contrast curve* of a UHI laser pulse, here the DRACO PW[G1] laser system at HZDR which creates ultra-short ($\sim$ 30 fs) pulses with peak intensities of $I_0 = 5.41 \cdot 10^{21}\,\mathrm{W\,cm^{-2}}$. It has a characteristic shape with features that do not just include the ultra-short (a few 10 fs) main pulse but also incoherent background light ($\sim$ ns), pre-pulses (at 10s of ps), intensity pedestals ($\sim$ ps) and the final ramp up to the peak intensity. Especially the influence of this last picosecond intensity ramp preceding the pulse maximum is one of the least explored parts of the interaction. That is due to the fact that the laser pulse is usually already intense enough to shape plasma densities that are opaque for optical probe light while the laser-driven plasma itself is also strongly emitting radiation in a broad spectral range that can easily outshine regular probing techniques like shadowgraphy[31, 32] or interferometry[33–36]. Only few works exist that begin to shed light on this interaction interval and what its effect is on the pre-plasma formation[37–39] and eventually on the final proton energies[40–42]. Additionally, the laser intensity is increasing by orders of magnitude within several hundred or even only tens of femtoseconds and the regions of interest measure only fractions of a micrometer. Resolving dynamics and substructures on such small scales is very challenging for many of the available (active and passive) plasma diagnostics since they.

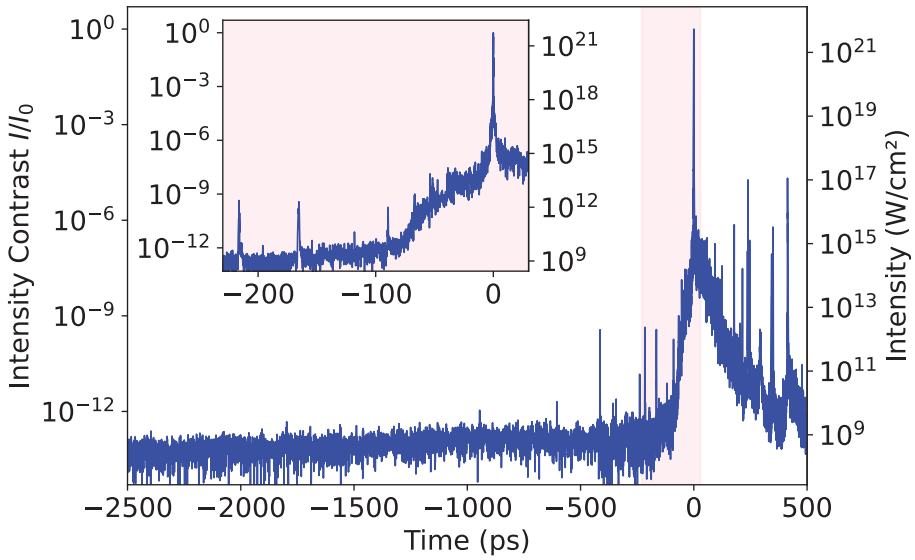

Figure 1.2.: Example temporal intensity contrast curve of the DRACO Ti:Sa laser system at HZDR. Measurements were taken on March 12th, 2019 with a scanning third-order autocorrelation (TOAC) device (*SequoiaHD*) by Ziegler et al. [10].

On the other hand, the numerical description of processes driving the acceleration of ions from solid-density plasmas is just as troublesome within the region of the last picosecond intensity ramp. Particle-in-cell (PIC) simulations, while widely used to describe the non-equilibrium plasma dynamics during the laser main pulse interaction, are computationally extremely demanding which usually limits the simulated time to only a few hundred femtoseconds surrounding the maximum intensity of the laser pulse. Additionally, they cannot exclusively be trusted to model the earlier parts of the interaction at lower intensities ($< 10^{17}\,\mathrm{W\,cm^{-2}}$) where the plasma just begins to form since they are unsuitable to describe all the occurring physics. Here, collisional and radiative effects dominate but especially the long-term stability of algorithms, general time-to-solution as well as modeling the multitudes of atomic excitations and other transitions remain challenges that usually limit the region of validity for PIC codes. While the lower intensity ranges are better represented by plasma fluid simulations, they rely heavily on equilibrium models and are only valid up to intensities of around 10^{15} to $10^{16}\,\mathrm{W\,cm^{-2}}$. This leaves a gap in both numerical modeling as well as experimental data.

Especially for highest ion energies following optimal acceleration, well-known theories[43–45] predict that ultra-thin solid-density foils at the threshold to transparency are required. Such targets are particularly sensitive to the laser contrast conditions and it is hypothesized that the last picosecond before the main pulse plays a decisive role for the acceleration performance. Recent experiments at the DRACO laser system at HZDR regarding optimum ion energies from controlled temporal pulse shaping[1] further motivated the project of this thesis.

At the same time, novel pump-probe experiments using optical UHI drive lasers in combination with recently developed X-ray Free-Electron Lasers (XFELs)[46–48] can provide simultaneously high resolution in space and time, promising to unlock the ability for observations at the femtosecond-nanometer level. Recently constructed XFEL facilities produce ultra-short ($\sim 1-10\,\mathrm{fs}$), hard [2] X-ray pulses that contain up to 10^{12} photons at very narrow spectral band-

[1] Early results motivated this work and during the time of this thesis the thorough experimental studies that followed were eventually published by Ziegler et al. [10].

[2] photon energy $\mathcal{E}_{\mathrm{ph}} > 5-10\,\mathrm{keV}$

width, classifying them as ultra-*brilliant*. Like the European XFEL[48, 49], modern machines are 10^4 to 10^{12} times brighter than the best conventional X-ray sources and can deliver peak brilliances of up to 10^{33} (photons / s / mm^2 / mrad^2 / 0.1 % bandwidth). Thus, they are able to penetrate optically opaque solid-density material, outshine laser-driven plasma self-emission, resolve nanometer structures due to wavelengths of a few Ångströms and produce strong signals from direct imaging or scattering[50, 51]. While XFELs seem to be the perfect tool for performing studies on laser-driven solids with unprecedented resolution, the novelty of such experiments introduces several caveats. X-ray Free-Electron Lasers are still very new, and with only a few machines available, experiments are also rare. However, the highest spatial resolution is achieved with X-ray scattering from structures in the irradiated targets. As a result, detector images contain signals in Fourier space that would elude straightforward interpretation, mainly because the dynamics from relativistic laser-solid interactions are violent, transient and complex in nature. Multiple scattering, the loss of depth information, temporal integration of the signal as the plasma evolves and as the X-ray pulse traverses the target, all including the influence of the plasma state by the high-intensity X-ray probe pulse itself, further complicate the analysis. For these reasons, novel methods and workflows need to be developed that can numerically predict the expected plasma dynamics but also the results of its probing to guide the planning, execution, and analysis of these new experiments.

1.2. Content of this thesis

This thesis contributes to the deeper understanding of the interaction between ultra-high intensity (UHI), ultra-short laser pulses with solid density matter through numerical simulation and theory. In particular, this work sheds light on the acceleration of ions from ultra-thin (submicrometer) foil targets and how the acceleration process and performance are affected by the shape of the intensity ramp that a realistic UHI laser pulse is preceded by. Numerical simulations and analysis work were mainly performed using the particle-in-cell (PIC) codes PIConGPU[52] and PICLS[53, 54] on the supercomputers *Hypnos* and *Hemera* (HZDR[3]), *Taurus* (ZIH[4]), and *Piz Daint* (CSCS[5]).

The interaction of a solid-density target with a laser (main) pulse (i.e. during a few ~ 10 fs) is highly sensitive to target density and thickness, pre-plasma conditions, pulse duration, shape and maximum intensity. After an introduction into the theory of laser-plasma acceleration processes and their theoretical modeling (Secs. 2.1 – 2.2), existing experimental and numerical simulation work on the influence of laser pre-pulse and contrast features is briefly reviewed. The few existing works that are reviewed in this section investigate only very narrow slices through the multi-dimensional laser-target parameter space. Furthermore, only 1D or 2D simulations are employed, the results of which cannot simply be adapted to the 3D case for reasons later explained in Sec. 2.3 where the numerical modeling of plasmas is described. While fully 3D PIC simulations are extremely costly and simulations of lower dimensionality remain a valid approach to study qualitative trends, the existing literature falls short of a methodical approach to cover and also cleanly separate the various influences of the above parameters.

The highly scalable, open-source, fully-relativistic, particle-in-cell code PIConGPU is well suited to shed light in this unsatisfactory situation with its fast time-to-solution which allows for broad 3D parameter scans which are needed for in-depth studies of realistic systems and quantitative predictive capabilities. PIConGPU is introduced in Chapter 3 and the technical details

[3] Helmholtz-Zentrum Dresden – Rossendorf
[4] TU Dresden - Zentrum für Informationsdienste und Hochleistungsrechnen
[5] Centro Svizzero di Calcolo Scientifico

regarding the advanced ionization framework enabling this thesis work are described in Sec. 3.1. Fully three-dimensional PIC simulations of laser-driven solid-density plasmas require enormous computational resources to resolve physics at the nanometer level with a temporal resolution of only a few attoseconds for tens of billions of particles. As such, large-scale numerical campaigns require careful planning but also compromises with respect to available resources (see Sec. 3.2) even on the largest supercomputers available on Earth. Especially explorative simulations will produce massive amounts of raw data as workflows only become more efficient when it becomes clear what to focus virtual diagnostics on. Sections 3.3 and 3.4 describe the methods used to overcome the challenges that the production, management, transport and final storage of several petabytes of data presented. As a result, new workflows have been established between the high-performance computing centers that the data was created, analyzed and archived at.

Before the main results of this thesis are presented, a brief excursion in Section 4.1 shows how the improved ionization framework in PIConGPU is used to support two experimental campaigns with simulations that provide more detailed insight into the temporal evolution of electron densities following ionization processes from specific atomic levels. The first campaign studies Laser-Wakefield acceleration (LWFA) of electrons from gas-jet targets in a Self-Truncated Ionization Injection (STII) setup[55] for the optimization of both the accelerated bunch charge as well as its maximum energy[G3]. Especially the second campaign, however, which concerns laser-induced melting of colloidal polystyrene crystal targets[G4], is vital for the main study of this thesis work since it presented the test bed for the advancement of modeling capabilities towards laser-solid density interactions via the implementation of new collisional ionization methods. Section 4.2 encompasses the main result of this work. The large-scale simulation campaign regarding the influence of the last picosecond leading intensity ramp on laser-ion acceleration performance from ultra-thin foil targets is motivated. Both the underlying theoretical models as well as recently obtained experimental results from DRACO at HZDR justify the promising expectation to find the highest proton energies in this parameter region, and the particular role that the leading pulse edge plays with respect to these findings. It will be shown how optimum proton energies are achieved from harnessing the full accelerating potential for protons that benefit from being injected into the very origin of the rear-side charge separation (sheath) field of the foil target after the interaction with the leading pulse ramp. The achieved maximum energies are higher than in the case of a perfectly Gaussian shaped laser pulse and, as such, the results are in contrast to the common expectation that the most ideal laser pulses deliver the highest proton energies. In the wake of discussing the underlying detailed plasma dynamics, the well-known 1D analytic model by Schreiber et al. [44] is slightly adapted to explain the results qualitatively and compare quantitative predictions to the results.

In the third section of Chapter 4, novel diagnostics are discussed that finally promise to shed light on the nanoscopic ultrafast processes of laser-driven solid interactions. In the first half, radiation diagnostics of electrons, protons and bremsstrahlung photons from the large-scale simulation campaign described above are displayed. The results at hand show that the intricate interaction dynamics and to an extent also the state of the target at the instance of peak laser intensity are encoded in these radiation signatures, emphasizing the need for more detailed studies of these “passive” probes.

Finally, the second half presents a preparatory simulation campaign that paved the way for a class of novel pump-probe experiments using infrared-laser-irradiated surface-structured foil targets that are probed with X-ray Free-Electron Laser (XFEL)-generated pulses. Studies like this are an essential first step towards measuring the plasma conditions in a laser-ion acceleration scenario in their full complexity by generating a clear and predictable signal from a target

with a pre-inscribed grating structure. This way, the plasma surface expansion, a key process in the Target-Normal-Sheath Acceleration (TNSA) of ions, can be directly characterized. In the experiment that followed, the first nanometer-femtosecond resolution measurement of the spatial density distribution of an expanding near-relativistic laser-driven solid-density plasma were performed. The simulation campaign furthermore revealed non-linear dynamics leading to a hitherto unexpected transient, but periodic, plasma structure for which the experimental results gave first evidence. The simulation work accompanying this campaign contributed also to the European EUCALL project, in particular first start-to-end simulations for optical-/X-ray laser pump-/probe experiments with a tool-chain connected by a single simulation framework. As such, the particle-in-cell simulations provided input for Monte-Carlo photon scattering simulations using the code ParaTAXIS[56], completing the forward-simulation chain of the virtual experiment which was also demonstrated and reported on here by the author.

Chapter 5 summarizes the results that were achieved, concluding consequences in Ch. 6. In the outlook in Chapter 7, a brief view into ongoing work regarding temporal pulse shaping for optimized proton acceleration is given.

2. Theory of Laser-Ion Acceleration From Thin Foils

In this chapter, the relevant quantities that describe the interaction of high-power ultra-short laser pulses with solid density matter are introduced. Laser-ion acceleration is an indirect acceleration process that relies on the transfer of laser energy to free electrons. It is first reviewed how free electrons are produced via ionization in the laser field and a plasma is created. Then, the motion of a free electron inside a laser field of relativistic strength is described. With increasing density, the laser pulse cannot penetrate into the target anymore and collective electron motion inside a plasma gains importance for the transfer of energy. Among these, plasma waves effect the further heating of the target while collisional ionization causes the solid target to turn into plasma in regions that cannot be reached by the laser.

This general introduction is followed by the most common acceleration mechanisms for ions from a laser-driven plasma, with particular focus on thin (µm) and ultrathin (nm) foil targets. Since the main aim of this work is to study the influence the last picosecond leading intensity ramp has on laser ion acceleration, the known influences of typical features in the temporal contrast of a high-power laser system are reviewed. The mainly used tool for the author's studies are highly scalable particle-in-cell simulations with PIConGPU. Therefore, the kinetic plasma modeling with PIC is introduced and extensions to the basic PIC-cycle which are needed to describe the physics of solid-density plasmas are presented. Finally, the influence of the simulation dimensionality is portrayed which gives motivation for the need of fully 3D explorative simulations to reach quantitative predictive capabilities of the laser-ion acceleration process.

2.1. Laser-Created Plasmas

The first section briefly describes the single and the collective motion of charged particles in the laser field and how ultra-high intensity (UHI) lasers create plasmas. Relevant terms and processes for field ionization, plasma waves and electron heating are introduced which are the basic concepts necessary for all relativistic laser-solid interactions, especially the laser-ion acceleration processes that are described in the following section 2.2.

2.1.1. Charged Particles in a Laser Field

To accelerate any type of charged particles in a directed beam by focusing a UHI laser pulse onto a so-called *target*, a transfer of energy from the pulse to the particles is necessary. Particles of charge q in electromagnetic fields are subject to the Lorentz-force $\vec{F}_\mathrm{L}$.

$$\vec{F}_\mathrm{L} = q\left(E + \vec{v} \times \vec{B}\right) \tag{2.1}$$

Due to energy and momentum conservation in the interaction of a plane wave laser field with a freely moving charged particle in vacuum the net gain of energy is zero, however. The charged particle performs oscillations mainly due to the electric $\vec{E} = E_0 \exp\left(i(kx - \omega_\mathrm{L} t)\right) \vec{e}_y$ field component of the laser pulse. With increasing laser intensity I_L, the magnetic component of the Lorentz-force causes the particle to also oscillate in the propagation direction of the wave which, in the rest frame of the oscillation center, yields the famous *figure-8 motion*[57].

Single Particle Motion

Integrating the equation of motion for an electron in the electric field yields the oscillation velocity $v_\mathrm{osc} = eE_0/m_\mathrm{e}\omega_\mathrm{L}$. When v_osc approaches the speed of light, i.e. $v_\mathrm{osc}/c \approx 1$, the electron motion becomes relativistic within a single half-wave. Such a laser pulse is often just dubbed *relativistic* and the transition into the relativistic regime is described by the normalized laser amplitude a_0.

$$a_0 = \left(\frac{v_\mathrm{osc}}{c}\right)_\mathrm{non\text{–}rel.} = \frac{eE_0}{m_\mathrm{e}\omega_\mathrm{L} c} \tag{2.2}$$

$$= 0.85\sqrt{I_\mathrm{L}\left[10^{18}\,\mathrm{W\,cm^{-2}}\right] \cdot \lambda_\mathrm{L}^2\left[\mu\mathrm{m}^2\right]} \tag{2.3}$$

The conversion between laser intensity $I_\mathrm{L} = \frac{1}{2} c\varepsilon_0 E_0^2$ is given by Eq. 2.3. In UHI laser-matter interactions, the light waves also exert a significant pressure on the plasma that can easily exceed 1 GPa. This radiation pressure originates in a non-linear force that is caused by the $\vec{v}\times\vec{B}$-term of the Lorentz-force and from evaluating the $\vec{E}$-field at the current position of the particle during its trajectory. The first effect leads to a push in $\vec{k}$-direction, resulting in a constant drift of $\beta_\mathrm{drift} = a_0^2/(4+a_0^2)$ overlaid with a longitudinal quiver motion of frequency $2\omega_\mathrm{L}$, while the second directs particles towards regions of lower field strengths when the field is non-uniform. In a cycle average, the net force on a single charged particle is given by

$$\vec{f}_\mathrm{pond} = -\frac{1}{4}\frac{e^2}{m_\mathrm{s}\omega_\mathrm{L}^2}\vec{\nabla} E_0^2. \tag{2.4}$$

This force is called the *ponderomotive force*[58] and Eq. 2.5 denotes the expression for an ensemble of electrons in a plasma

$$\vec{F}_\mathrm{pond} = -\frac{\omega_p^2}{\omega_L^2}\vec{\nabla}\frac{\langle\varepsilon_0 E^2\rangle}{2}. \tag{2.5}$$

Collective Electron Motion

Equation 2.4 shows that the ponderomotive force is proportional to the gradient of intensity. As it depends inversely on the mass of the particle species m_s, it immediately follows that the ponderomotive force is much weaker for ions than for electrons with a factor of $m_\mathrm{e}/m_\mathrm{i}$.

Therefore, ions cannot be accelerated efficiently to MeV energies in the direct interaction with a super-intense laser pulse of currently up to 10^{22} W cm^{-2} and the transfer of energy via the plasma electrons is needed.

2.1.2. Plasma Waves

When an electromagnetic wave interacts with a plasma, a wave is excited whose frequency depends on the number density of electrons n_e and ions n_i. Due to the much larger mass, and therefore higher inertia of ions, the situation can be reduced to electrons moving in front of a positively charged ionic background. In the picture of a hydrodynamic electron fluid without thermal motions and in absence of a magnetic field, the *plasma frequency* ω_p can be derived[59]. As stated before, the low-frequency ionic oscillations are typically neglected and ω_p is just expressed with respect to the density of electrons n_e

$$\omega_{pe} = \sqrt{\frac{n_e e^2}{\gamma_e m_e \varepsilon_0}} \tag{2.6}$$

In the interaction with ultra-high intensity lasers, electrons quickly reach relativistic velocities and therefore Eq. 2.6 contains the relativistically corrected mass $\gamma_e m_e$ which leads to lowered plasma frequencies. Collectively, the oscillating electrons radiate and whether the laser light can penetrate into the plasma or not depends on the capability of the particles to follow the the electric field synchronously. This capability is density- and wavelength-dependent which makes plasma a dispersive medium and is usually expressed in the ratio of laser frequency $\omega_L = 2\pi c/\lambda_L$ and plasma frequency ω_p. For $\omega_L/\omega_p > 1$, the wave can fully penetrate and the medium is called *underdense*. In the reverse $\omega_L/\omega_p < 1$, or *overdense*, case a part of the light wave can penetrate into the plasma evanescently as its intensity is exponentially damped while another part is reflected. The threshold at which these two regimes meet is defined by the so-called critical (electron) density n_c.

$$n_c = \frac{\gamma_e m_e \varepsilon_0}{e^2} \omega_{pe}^2 \tag{2.7}$$

Free electric charges only affect the plasma locally since they are shielded by other charge carriers around them. Plasmas are characterized by the property of *quasi-neutrality* that is expressed by the so-called *Debye-length* that describes the radius around an electric charge at which its electric field is damped by a factor of 1/e.

$$\lambda_D = \sqrt{\frac{\varepsilon_0 k_B T_e}{e^2 n_e}} \tag{2.8}$$

Eq. 2.8 assumes a local thermal equilibrium and again neglects the influence of ion motion and their temperature.

In the underdense case, a direct consequence of the ponderomotive force is *self-focusing*. Electrons are pushed down the spatial intensity gradient, vacating the central region around the beam axis and populating the surrounding volume. Accordingly, the plasma frequency in the center decreases while it increases on the outside. The index of refraction $n = \sqrt{1 - \omega_p/\omega_L}$ behaves inversely which acts as like a convex lense, channeling and focusing the laser beam. Self-focusing can also occur due to the decrease of ω_p with the growing relativistic mass of electrons $\gamma_e m_e$ (see Eq. 2.6) or due to local heating, expansion and further heating by the laser. The beneficial effect of self-focusing is deliberately exploited in laser electron acceleration.[1]

[1] A recent work by Levy et al. [60] used the increase of laser-intensity on target also for laser-ion acceleration by introducing tailored pre-plasma when the steel target was put into stream of a helium gas nozzle.

Laser Electron Acceleration

The acceleration of electrons from laser-driven plasmas is a field of research that has seen tremendous successes in recent years which lead to the development of well-controlled, table-top accelerators that readily produce ultrashort (fs) electron beams with kiloampere peak currents and tunable energies up to the GeV range[61]. The main concept is that transient accelerator structures are created inside an underdense plasma which allow for acceleration gradients exceeding the conventional accelerator cavity degradation threshold of 100 MV/m by three orders of magnitude and more. Typical targets are gas jets or gas cells with electron densities of 10^{17} cm^3 (kJ long-pulse lasers) and 10^{18} cm^3 (1 to 10 J short-pulse systems).

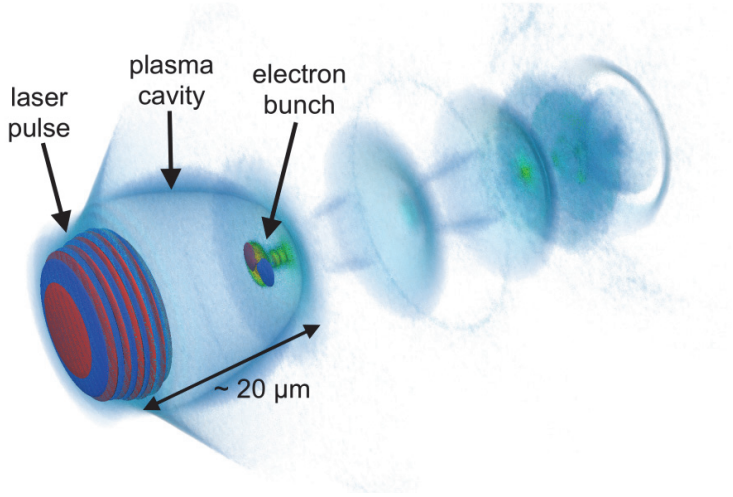

Figure 2.1.: Plasma cavities in a laser-wakefield accelerator (LWFA) simulated in 3D with PIConGPU and rendered with ISAAC. Transverse electric fields of the driving laser pulse and the accelerated electron bunch are shown as red and blue contour plots. Electron densities larger than the initial plasma density are displayed as blue and green clouds. Image originally in [62].

Laser-Wakefield Acceleration The most famous technique relies on the excitation of a plasma wave that forms in the wake of a UHI laser pulse which displaces the plasma electrons ponderomotively from the center of the beam propagation axis. A positively charged region remains due to the inert, mostly stationary ions. As the electrons relax back towards the beam axis they close the accelerating cavity that follows the laser pulse as it propagates through the plasma, just like the wake behind a motorboat. The electrons overshoot their initial position and multiple accelerating cavities can form. A part of the displaced electrons reaches the rear side of the plasma cavities with the right initial momentum conditions to be injected into the structure. These are subsequently accelerated quickly towards MeV energies upon which they co-propagate behind the laser with close to the speed of light. This mechanism, dubbed *Laser-Wakefield Acceleration* (LWFA) was proposed by Tajima and Dawson in 1979[24] and was demonstrated experimentally for the first time in 1994 by Nakajima et al. [63].

In 2004 the next breakthrough was achieved with the demonstration of the *blow-out regime* (or "bubble regime")[25–27] that had previously been proposed by Pukhov and Meyer-ter-Vehn[64]. With the expulsion of all electrons from the center region of the laser pulse, a stable plasma bubble can form which offers both extremely high longitudinally accelerating fields as well as focusing transverse fields that confine the accelerated electrons. For the first time, quasi-monoenergetic high-energy electron bunches could be produced which opened up a variety of novel applications, ranging from the generation of secondary radiation in laboratory-sized experiments towards possible drivers for new free electron laser facilities (FELs) that

produce X-ray beams of unprecedented brightness. The latter has very recently been demonstrated in a proof-of-principle experiment by Wang et al. [30].

A multitude of experimental methods has been developed in the past 20 years that deliver on-shot diagnostic data of both the wakefield structures as well as the electron beam quality[65]. A great benefit to these techniques is that the plasma cavities, with their average sizes on the order of $\sim$ 10 µm, are still accessible to optical probing for which the plasma is still transparent. Most laser-ion acceleration processes, however, take place on scales that are still too fast (fs) and too small (nm) such that the interaction region is inaccessible to most diagnostic techniques. As mentioned earlier, ions are accelerated indirectly in plasma structures that are created by energetic electrons and therefore the efficient transfer of laser energy into electrons is vital.

2.1.3. Laser Absorption and Heating

Fast electron generation and electron bulk heating are two crucial prerequisites for the acceleration of ions from solid density targets. Both effects require the efficient absorption of laser energy which first leads to the *acceleration* of electrons. The term *heating* refers the system moving towards a state of local thermal equilibrium (LTE) where all electrons or a subpopulation can be described by a *temperature* T_e. Especially in the interaction of ultrashort UHI laser pulses with matter, the plasma dynamics are highly non-linear and much faster than the typical electron-electron collision timescales ($\tau_{ee} = 1/\nu_{ee}$) that lead to thermalization. The interaction duration and laser intensity as well as the plasma density distribution and direction of the electric field with respect to gradients of the latter play a deciding role in which of the multitude of mechanisms is dominantly causing the absorption of laser energy into the target. [2]

The main focus of this thesis are the interactions of ultrashort ($\tau_L \sim 10$ fs), super-intense ($I_L > 10^{18}$ W/cm^2, i.e. $a_0 > 1$) laser pulses with ultra-thin ($d \sim 10-100$ nm) overcritical ($n_e > n_c$) targets. The two main absorption mechanisms for these laser pulses are *Brunel* (*vacuum*) heating[67] and $\vec{v}\times\vec{B}$-acceleration[68, 69]. The former relies on an electric field component $E_\perp$ that is perpendicular to the target surface which pulls an electron layer into the vacuum region where they are accelerated and pushed back into the over-critical region of the target. Under the condition that the excursion length $v_{osc,\perp}/\omega_L$ is larger than the front-side plasma density scale length L, these electrons transport energy to the regions that are screened from the laser field behind the skin layer ($L_s = c/\omega_p$) where they cannot be pulled back. That is in contrast to resonance absorption[70–72] where underdense plasma waves are excited in the region between the critical density surface and specular reflection depth (at which plasma permittivity becomes $\varepsilon_r = \sin^2\theta$) and are subsequently dampened, a high quiver velocity is needed for the process to be efficient. However, Brunel heating requires an obliquely incident pulse and is strongly suppressed for $\theta = 0°$. In a head-on irradiation scenario, it only gains importance once the target surface is significantly indented by radiation pressure such that $E_\perp$ increases.

Otherwise, $\vec{v}\times\vec{B}$-acceleration is the most dominant absorption mechanism in the ultra-high intensity short-pulse regime. The ponderomotive drift constantly pushes electrons into the plasma, creating a charge separation zone which can lead to a steepening of the density gradient. The electric potential at the front side effectively hinders electrons from passing into the overdense plasma region. But the quickly oscillating part of the $\vec{v}\times\vec{B}$-force creates directed, high energy electron bunches at a frequency of $2\omega_L$ that have their origin close to the $n_e = n_c$ layer. These, often called *hot* or *prompt* electron bunches, exhibit a broad, $1/\gamma_e$ energy distri-

[2] A concise overview of the most important heating mechanisms can be found in the thesis of Klimo [66], particularly Tab. 2.1.

bution that is peaked at the end (see [73]). When multiple bunches with different maximum energy, resulting from the temporal intensity envelope of the laser pulse, are superimposed the collective exhibits a Boltzmann-like energy spectrum for which a temperature value is often defined[3]. Kluge et al. [74] report

$$T_e^{\text{prompt}} = \left(\frac{\pi}{2\mathcal{K}(-a_0^2)} - 1 \right) m_e c^2, \tag{2.9}$$

where $\mathcal{K}(-a_0^2)$ is the complete elliptical integral of the first kind. In cases of $a_0 \ll 1$, Eq. 2.9 converges against the usual ponderomotive electron scaling $T_e = m_e c^2 \left(\sqrt{1 + \frac{a_0^2}{2}} - 1 \right)$ [75]. While the time-averaged spectrum of the $2\omega_L$–bunches is exponential, the tail of a relativistic Maxwellian is not and using the term "temperature" remains incorrect, even though it is still done quite often in the literature.

Even though ultrashort laser pulses convey almost all of their energy within the few 10 fs of their main pulse duration τ_L, the temporal pulse contrast on the nanosecond-ASE and picosecond pedestal phases can significantly determine the later interaction. Even prior to the creation of a plasma with intensities above the ionization threshold (see 2.1.4), perpetually high laser intensities, e.g. an unsuppressed ASE nanosecond-background at 10^{12} W/cm^2, can cause the melting and evaporation of the target, possibly creating an underdense pre-plasma which changes the absorption. Typically, metals and dielectrics exhibit a different behavior due to their electronic structure and resulting heat conduction properties. The radiant fluence (time-integrated irradiance) is the central quantity that decides if a laser pulse is ablating a target surface. To ablate the target to a certain depth, short pulses require higher intensities than longer pulses and already for femtosecond pulses of 10^{13} to 10^{14} W/cm^2, the differences between dielectrics and metals vanish[76]. For such sub-relativistic pulses, Gamaly et al. [76] derived an analytic expression for the electron temperature T_e in the skin layer of a pre-expanded target.

$$T_e = \frac{4}{3} \frac{\eta I_0 \tau_L}{L_s n_e} \exp\left(-2 \frac{x}{L_s} \right) \tag{2.10}$$

$$\eta \approx 4\pi \frac{L_s}{\lambda_L} \tag{2.11}$$

The absorption coefficient η is in general material-, intensity- and density-dependent. It is very important for analytic models that predict final energies of laser-accelerated ion beams (see 2.2.1) but, unfortunately, very challenging to measure in a laboratory experiment. Equation 2.11 gives a general expression for η in the sub-relativistic short-pulse regime. Here, I_0 is the (pre-) pulse intensity, τ_L its duration, x is the coordinate along the plasma gradient and L_s the skin depth.

2.1.4. Ionization

A prerequisite to laser-ion acceleration is that the laser pulse transfers energy to free electrons which in turn transfer their energy to the ions in the target. Free electrons in solid density targets are first produced from ionization processes that are caused by the direct influence of the electric field of the laser pulse and later mainly by collision processes inside the plasma.

[3] Note however, for 30 femtosecond UHI pulses the laser intensity changes much faster than the typical electron thermalization time. During every laser period there are two prompt electron bunches but with a quickly changing laser envelope each bunch is imprinted with a different spectral shape but a clear preference direction such that the concept of an overall temperature is not justified for such non-equilibrium dynamics.

In this section the most dominant field ionization processes are introduced and the main ionization regimes for the laser-target conditions in this thesis are identified. During the course of this thesis, the author continued the development of the ionization implementation in the particle-in-cell code PIConGPU[15, 52] (see 3.1) that was started in the author's Master's thesis [G5] prior to this work.

Field Ionization

The ionization of atoms under the influence of a strong external electromagnetic field depends mainly on the laser frequency ω_L, the ionization potential $\mathcal{E}_{ip}$ of the specific atom charge state and the magnitude E of the electric field. To best determine the regime of ionization and with that the simplified picture for the process and the model to describe it, it is customary to express all relevant quantities in so-called *atomic units* (for an overview, see Tab. 7.1 in [77]). This unit system describes typical mass-, length-, time-, and derived quantity scales inside an atom and its use simplifies the equations for atomic processes significantly. The base units are the electron mass $M_{AU} = 9.109 \cdot 10^{-31}$ kg, the first Bohr radius $L_{AU} = 5.292 \cdot 10^{-11}$ m, and the classical electron orbit period in hydrogen reduced by 2π, $T_{AU} = 2.419 \cdot 10^{-17}$ s, respectively. Some important derived units are the ones for energy $\mathcal{E}_{AU} = 27.21$ eV, electric field $E_{AU} = 5.14 \cdot 10^{11}$ V/m, and intensity $I_{AU} = 3.51 \cdot 10^{16}$ W/cm^2.

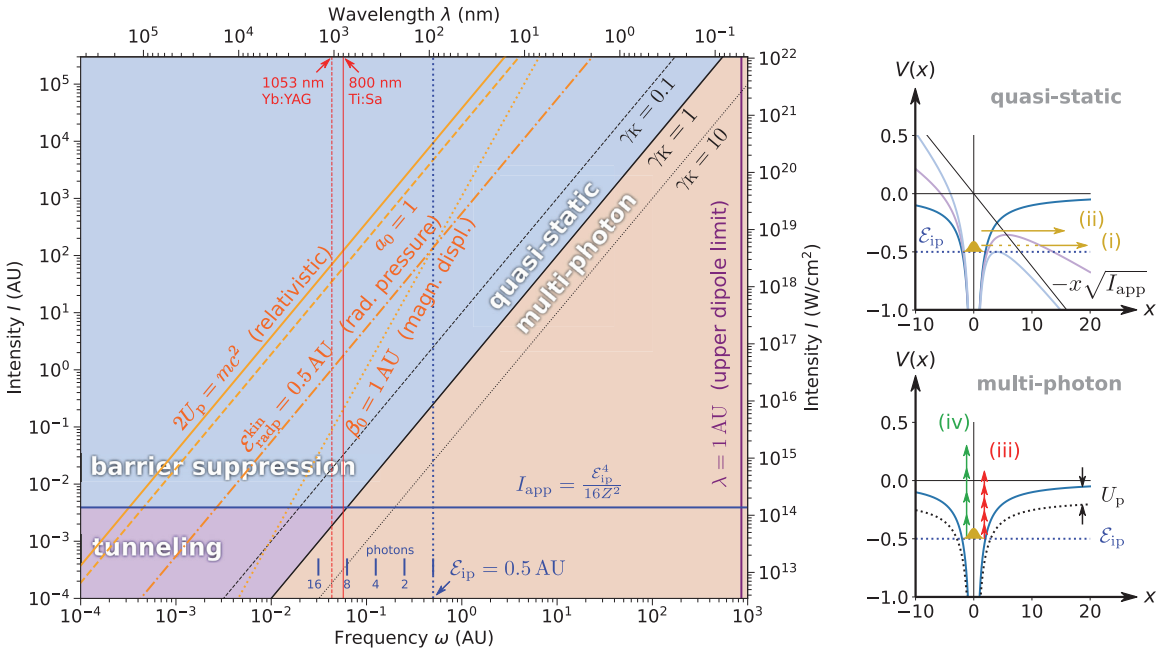

Figure 2.2.: Schematic visualization of field ionization regimes (left) and –mechanisms (right). Shaded areas mark the regimes known from Keldysh theory[78], where the Keldysh parameter $\gamma_K = \omega\sqrt{2\mathcal{E}_{ip}}/E = \sqrt{\mathcal{E}_{ip}/2U_p}$ separates the *quasi-static* ($\gamma_K < 1$) and the *multi-photon* ($\gamma_K > 1$) pictures. More recent considerations of Reiss [79][80] regarding the unsuitability of the Keldysh parameter especially in the low-frequency, high-intensity regime ($\gamma_K \to 0$) were marked with orange lines. Mainly, neglected influences of the laser field magnetic component and radiation pressure can lead to substantial electron energies (quiver– U_p and longitudinally kinetic $\mathcal{E}_{kin}^{radp}$) as well as excursion amplitudes (β_0) that violate the assumptions of the known quasi-static models described in this section. The mechanisms shown in the right column and the blue lines in the left panel are drawn with respect to the hydrogen atom (quantities in AU). The quasi-static picture contains (i) tunneling ionization and (ii) barrier suppression ionization. In the multi-photon picture, (iii) multi-photon ionization (MPI) and (iv) above-threshold ionization (ATI) are distinguished.

Field ionization is caused by the interaction of an external laser field with an electron and

the ionic potential it is bound to. In the established field ionization theory, two main ionization domains are usually distinguished, the *quasi-static* and the *multi-photon* domain. The earliest separation between the two was defined by the *Keldysh parameter*[78] $\gamma_K = \sqrt{\mathcal{E}_{ip}/2U_{pond}}$ (in atomic units) as the shaded areas above and below $\gamma_K = 1$ in Fig. 2.2) show. The ponderomotive potential $U_{pond} = I_L/2\omega_L$ describes the coupling of a plane-wave laser field to a charged particle and signifies the average *quiver energy* of this particle. If the oscillation period of the laser field is much longer than the typical atomic time scale of $T_{AU} = 150\,\text{as}/2\pi$, the field imposed on the atom is viewed as quasi-static. The atomic potential well is tilted such that the potential on one side is steepened while a barrier of finite height is created on the opposite side. Electrons can tunnel through this barrier quantum-mechanically. The ratio of a conceptual *tunneling time* τ_{tunnel} and laser period τ_L is another way to interpret γ_K. The early Keldysh tunneling ionization rate (Eq. 2.12)[78, 81] allowed to describe the time-dependent ionization of atoms in an external laser field already with reasonably good accuracy.

$$\Gamma_K = \frac{(6\pi)^{1/2}}{2^{5/4}}\mathcal{E}_{ip}\left(\frac{E}{(2\mathcal{E}_{ip})^{3/2}}\right)^{1/2}\exp\left(-\frac{2\left(2\mathcal{E}_{ip}\right)^{3/2}}{3E}\right) \tag{2.12}$$

The KFR–Theory (Keldysh [78]-Faisal [82]-Reiss [83]), also often called *strong-field approximation* (SFA)[84], is the basis for most strong-field ionization models today. The most prominent field ionization model for the tunneling regime, however, was formulated in 1986 by Ammosov, Delone and Krainov (ADK)[85] and has proven to be very accurate for the prediction of charge states.

$$\Gamma_{ADK} = \underbrace{\sqrt{\frac{3n^{*3}E}{\pi Z^3}}}_{\text{lin. pol.}}\frac{ED^2}{8\pi Z}\exp\left(-\frac{2Z^3}{3n^{*3}E}\right) \tag{2.13}$$

$$D \equiv \left(\frac{4eZ^3}{En^{*4}}\right)^{n^*} \qquad n^* \equiv \frac{Z}{\sqrt{2\mathcal{E}_{ip}}} \tag{2.14}$$

In the Keldysh description, tunneling ionization is dominant as long as $\gamma_K < 1$ and when the Keldysh parameter approaches unity, the physical picture changes to the multi-photon domain. Rather than tunneling through a quickly oscillating, externally perturbed potential well, the average binding energy of an electron is overcome via energy transfer of one or multiple mutually arriving laser photons of energy $\mathcal{E}_{ph} = \omega_L$ ($\hbar = 1$). The cross section $\sigma^{(n)}$ for n-photon ionization is strongly enhanced for $\mathcal{E}_{ip} = n\mathcal{E}_{ph}$ at which the process becomes resonant. If substantial initial momentum is transferred to the electron, the multi-photon process is often called *above-threshold ionization* (ATI). Generally, the multiphoton ionization (MPI) rate scales with $\Gamma^{(n)} = \sigma^{(n)}I_L^n$.

Within the last thirty years, the widely accepted quasi-static picture has repeatedly been questioned by Reiss [86] (one of the fathers of the KFR-theory) due to the increasing importance of the magnetic field component of the Lorentz force with relativistic laser pulses and especially at low frequencies[80, 87]. The criticism is based on the fact that the Lorentz invariant $\vec{E}^2 - \vec{B}^2$ (in Gaussian units) equals $\vec{E}^2$ for quasi-static electric fields but 0 for plane-wave laser fields, respectively[79]. Fig. 2.2 shows related limiting conditions past which the quasi-static picture becomes questionable. First, the dipole approximation[4] is clearly violated when the ponderomotive potential becomes as large as the electron rest energy $2U_p = m_e c^2$ (solid

[4] *Dipole approximation*: the laser field is treated as a purely electric field that does not vary spatially because the wavelength λ_L, e.g. 800 nm is large compared to the first Bohr radius $a_0 = 0.05\,\text{nm}$. The electron moves classically in it, instantaneously following the field change with no inertia.

orange line) because the interaction is already strongly relativistic. For reference, the usual laser-plasma quantity of $a_0 = 1$[5] (see Eq. 2.3) is shown (orange dashed line).

Additionally, radiation pressure can be non-negligible for ionization under these conditions. Modern attosecond spectroscopy methods are promising tools to research the sub-cycle ionization dynamics of UHI laser pulses. However, the additional influence of a highly overdense medium and a strongly damped laser field in the vicinity of the critical density remain unavailable to direct experimental observation.

Common modern UHI short-pulse laser systems used for laser-ion acceleration operate at central wavelengths of $\lambda_L = 800\,\mathrm{nm}$ (Ti:Sa), or 1053 nm (Yb:YAG) which have photon energies of roughly 1.6 to 1.2 eV. This is much lower than the lowest ground state ionization energies and the laser period is also much longer than typical atomic time scales. Therefore, the main ionization physics important for this thesis are situated in the quasi-static regime. In case that the laser intensity is very large, the corresponding electric field suppresses the potential barrier of the atom below the energy level of the bound electron, which can therefore leave classically.

$$E_{\mathrm{BSI}} = \frac{\mathcal{E}_{\mathrm{ip}}^2}{4Z^*} \tag{2.15}$$

This effect is called *barrier suppression ionization*[88] (BSI) and the analytical description in Eq. 2.15 provides a threshold field strength E_{BSI} that serves well as a useful, quick estimate for the highest abundant charge state, given that a laser pulse intensity surpasses the so-called *appearance intensity* $I_{\mathrm{app}} = E_{\mathrm{BSI}}^2$. The quantity Z^* is the residual charge state of the ion after the ionization process.

Collisional Ionization

In the early stages of laser-ion acceleration, field ionization dominates the creation of free electron density. An *ideal plasma* would contain intertial point particles that are exclusively dominated by long-range electromagnetic forces. In reality, only underdense laser-driven plasmas or astrophysical low-density plasmas fulfill these conditions very well since their long-range interaction time scales are much smaller than the typical two-body interaction times scales. The average collision frequency $\nu_{\mathrm{pp}} = n\sigma_{\mathrm{pp}}v_{\mathrm{rel}}$ between charged particles increases with the particle density n but decreases with higher average energies since the interaction cross section σ decreases for higher relative speeds v_{rel}. For the overdense targets studied in this thesis, however, collisions can play an essential role while the laser intensities are still moderate ($I_L = 10^{14} - 10^{17}\,\mathrm{W/cm^2}$) and then the ionization is quickly dominated by electron-ion collisions, especially where the laser cannot directly penetrate. The electron-ion collision frequency is calculated via[77]

$$\nu_{\mathrm{ei}} = \frac{4}{3}(2\pi)^{1/2}\left(\frac{Ze^2}{4\pi\varepsilon_0 m}\right)^2\left(\frac{m}{k_B T_e}\right)^{3/2} n_i \ln\Lambda \tag{2.16}$$

$$\nu_{ei}\left[\mathrm{s}^{-1}\right] = 3.6 \times \frac{Z n_e\left[\mathrm{cm}^{-3}\right]}{(T_e[\mathrm{K}])^{3/2}}\ln\Lambda \tag{2.17}$$

The term $\ln\Lambda$ in Eq. 2.16 is the *Coulomb-Logarithm* that is a measure for the average number of ions $N_i = N_D/Z^*$ in the so-called *Debye sphere* (see also Sec. 2.2.1)[89], that describes the volume in a plasma outside of which a point charge is shielded by surrounding charge carriers. It is usually identified with the ratio of the *Debye-length* λ_D (Debye sphere $V_D = \frac{4}{3}\pi\lambda_D^3$) and the

[5] Attention: not in atomic units! $a_0^{\mathrm{AU}} = a_0/\alpha \approx 137 a_0$, where α is the fine structure constant.

impact parameter for a 90° deflection, $b_\perp$, as

$$\ln \Lambda_{ei} \equiv \ln \frac{\lambda_D}{b_\perp} = \ln \frac{9N_D}{Z} \tag{2.18}$$

Typical values for $\ln \Lambda$ are between 2 to 5. A plasma is dominantly collisional if $\nu_{ei}/\omega_p \geq 1$ and collisionless for $\nu_{ei}/\omega_p \ll 1$. The majority of Coulomb collisions results in small angle deflections and τ_{ei} is the average time for a single electron to experience a net direction change by 90°.

One of the most simple ways to describe collisional ionization is via the *Thomas-Fermi* (TF) model which delivers a prediction for an average charge state $\langle Z^* \rangle$ based on the mass density of ions ρ and the electron temperature T_e of the ion species with base charge Z under LTE conditions.

$$\langle Z^* \rangle = \frac{Zx}{1 + x + \sqrt{1 + 2x}} \tag{2.19}$$

Here, the function x is a composite of the aforementioned input quantities and 9 fitting parameters to approximate the free electron density from the distribution of ion boundaries in a statistical average-atom model[90]. As such, the TF model neglects any inner structure of the atoms but provides a quick estimate of charge states based on local conditions. A better way to evaluate the frequency of impact ionization processes under conditions that have not yet thermalized is the treatment with a *binary collision* operator. But these calculations are usually performed numerically in kinetic plasma simulations of the *particle-in-cell* scheme that also account for the momentum change of individual representatives of the various particle distributions[53, 54, 91–94]. The approach requires to form pairs between colliding particle species and calculate the ionization probability $P = 1 - \exp(-v_{rel}\sigma n \Delta t)$ based on the relative particle velocity v_{rel} and ionization cross section during a discrete time step Δt.

2.2. Laser Ion Acceleration

The previous sections explained how a laser pulse creates free electrons through ionization, how particles move individually and collectively in the laser field and how laser energy is transferred to electrons via acceleration and heating. These are all prerequisites to the acceleration of ions from laser-driven plasmas. In this section, the most important acceleration mechanisms with a particular focus on ultrathin overdense targets are introduced.

2.2.1. Ion Acceleration Mechanisms

Target-Normal Sheath Acceleration

The most robust and well-researched ion acceleration mechanism to this day is called *Target-Normal Sheath Acceleration*[1–3] (TNSA). It depends on a capacitor-like field structure that is created during the expansion of plasma into vacuum. The laser pulse impinges onto the target, ionizing its surface and heating the up the freed electrons. Due to the large mass difference between electrons and ions, only the electrons are gaining a significant amount of energy in the laser field as was described in Sec. 2.1.3. This higher average kinetic energy $\langle \mathcal{E}_{kin} \rangle$ of the electrons cannot be transferred to the ions instantaneously. In case of $a_0 > 1$ electrons are pushed directly into the target which they traverse ballistically. As a result of both effects, electrons expand faster into the vacuum than the ions and an electron cloud forms at the surfaces of the target. This so-called *Debye sheath* is created at both the front (facing the laser / *upstream*) and the rear (facing away from the laser / *downstream*) surface. The front surface

sheath is usually suppressed compared to the rear surface sheath for as long as the laser pulse still exerts substantial light pressure on the plasma. The capacitor-like charge separation field at the rear surface ionizes the surface-near atoms. The strength of these fields can reach several 10 to 100 TV m^{-1} which is 6 orders of magnitude larger than the typical breakdown field strength in conventional accelerator machines, i.e. 100 MV m^{-1} [95, 96]. Due to symmetry reasons, i.e. the electrons and ions separating at approximately the cold target surface, the direction of the electric field is strongly aligned with the surface normal. As such, TNSA allows to accelerate short, well-collimated[97], highly charged bunches of ions to MeV energies over distances of only several microns which makes them very attractive for applications. Of all ions, the acceleration of protons is most effective because they have a normalized charge over mass ratio of $q/m = 1$.

TNSA was extensively researched in both experiment and simulation and as targets metal or plastic foils are often used. Research for modern target systems is ongoing, aiming to increase shot repetition rate and decrease debris that is harmful to the optical elements close to the interaction region. One of these recently devised and successfully tested targets produces a solid jet of cryogenic hydrogen [8, 98–102]. Even from metal foils, protons can be accelerated due to always present organic surface contaminants that will usually survive the vacuum conditions inside the target chamber.

Plasma Expansion into Vacuum A simple description of the expanding plasma sheath starts from the situation where a thermalized, i.e. Maxwell-Boltzmann-distributed, electron population of temperature T_e is in electro-static equilibrium with a step-like population of immobile ions[103, 104]. The electron density n_e and electro-static potential Φ are given by the following relations:

$$n_e = n_0 \exp(e\Phi/kT_e) \tag{2.20}$$

$$\varepsilon_0 \frac{\partial^2 \Phi}{\partial x^2} = e\,(n_e - Zn_i)\,. \tag{2.21}$$

For times $t > 0$, ions are allowed to move and numerical integration of the Poisson equation, assuming quasi-neutrality (overall $n_e = Zn_i$ is fulfilled), reveals that a self-similar solution describes the expansion well. The electron density evolves as $n_e = n_{e,0} \exp(-x/c_s t - 1)$ with the ion sound velocity $c_s = \sqrt{Zk_B T_e/m_i}$. The ion density front then moves with the velocity $v_{i,front} = 2c_s \ln(\omega_{pi} t)$ [43] where $\omega_{pi} = n_{e,0} Ze^2/m_i \varepsilon_0$ is the ion plasma frequency. These solutions are valid when the local Debye-length $\lambda_D = \lambda_{D,0}\sqrt{n_{e,0}/n_e} = \lambda_{D,0} \exp[(1 + x/c_s t)/2]$ is smaller than the density scale length $c_s t$ since this translates into $\omega_{pi} t > 1$ for which v_i is well-defined. Only a small region on the order of λ_D at the ion front exhibits local charge separation. While the original models assume a semi-infinite plasma reservoir with a sharp gradient on one side, the model works reasonably well for short (ps) and ultra-short (fs) pulses and thin (~µm) foils.

Plasma expansion model A very successful model for the prediction of final ion energies is the 1-dimensional, isothermal, self-similar, plasma expansion approach developed by Mora [43]. It was the first to reproduce the exponential shape of the spectrum with the characteristic energy cutoff that is observed in experiments. The model predicts

$$\mathcal{E}_{max} = 2\mathcal{E}_0 \left[\tau + \sqrt{\tau + 1}\right]^2 \tag{2.22}$$

where $\tau = \omega_{pi} t_{acc}/\sqrt{2e}$ (e is Euler's number) is the acceleration time normalized with the ion plasma frequency $\omega_{pi} = n_{e,0} Ze^2/m_i \varepsilon_0$ and $\mathcal{E}_0 = Zk_B T_e$. The result in Eq. 2.22 does not converge

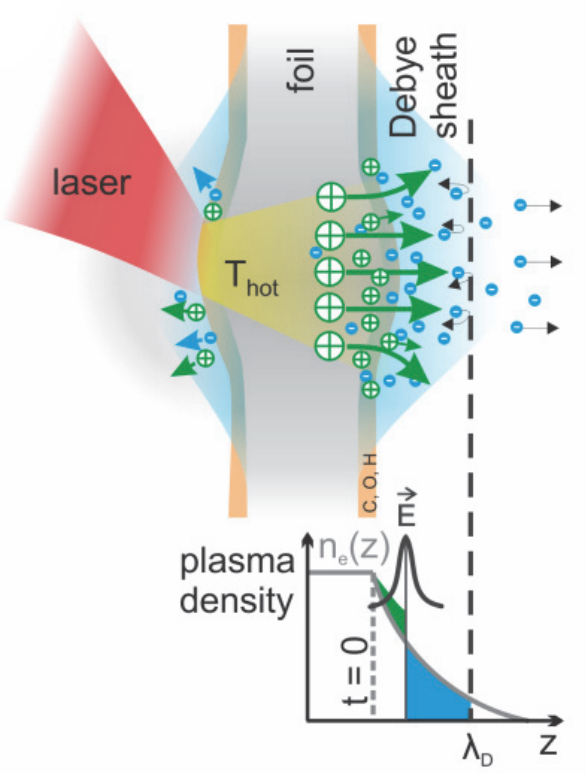

Figure 2.3.: Schematic image of the target normal sheath acceleration process from a foil coated with organic contaminants. Originally in [105])

for infinite times and has to be cut off with, e.g. the empirical acceleration time found by Fuchs et al. [106] of $t_{acc} \sim 1.3\,\tau_L$. The aforementioned semi-infinite reservoir of plasma with the constant electron temperature T_e is the reason for the infinite expansion and the energy divergence. While no laser pulse interaction is described in the model, its success lies in the simplicity of its equations and the good agreement with experimental observations of proton energies and 1D particle-in-cell simulations.

Mora [107] later refined the original model for targets of finite width d for which he assumed adiabatic cooling of the electron fluid. In the case of non-relativistic electron temperature ($T_e \ll 1$ MeV) the temporal dependence of the model reproduces the scaling $T_e \propto t^{-2}$ that was widely found by other works[108–113]. The ultra-relativistic case ($T_e \gg 1$ MeV) exhibits a scaling of $T_e \propto t^{-1}$ [107]. Additional considerations about the two phases of T_e-increase, then –decrease, two electron temperatures and final ion density gradients were made in [114, 115]. Reliable analytic models, accounting for 3D effects, and with predictive capability are unfortunately still not existing and numerical methods like particle-in-cell simulations (see 2.3) remain the best tool available to deliver quantitative trends *a priori* to laboratory experiments.

Quasi-static sheath model The model of Schreiber et al. [44] includes the laser-target interaction in a simple 1D analytic model for the prediction of final ion energies. As input parameters, only the laser pulse duration τ_L, power $P_L = \mathcal{E}_L/\tau_L$, spot radius r_L and the target foil parameters Z, n_e and thickness d are required. It describes how a fraction of electrons from the target front side $N_e = \eta\mathcal{E}_L/k_B T_e$ is accelerated as a bunch of length $L = c\tau_L$ through the target with a divergence θ_e, such that it spreads over a rear side source volume of $\pi B^2 \lambda_D$, where $B = r_0 + d\tan\theta_e$. The positive surface charge induced by the quasi-static electron sheath is then the driver for the ion expansion. No explicit distribution function for the electrons has

to be assumed which keeps the number of parameters low. However, the transverse spatial limitation B of the sheath keeps the potential finite and thus the final ion energies $\mathcal{E}_{i,\infty}$ finite, whereas they diverge in the model that was previously described.

$$\mathcal{E}_{i,\infty} = 2q_i m_i c^2 \sqrt{\eta \frac{P_L}{P_R}} \tag{2.23}$$

$$P_R = \frac{m_e c}{r_e^3} \tag{2.24}$$

Here, η is the total absorption efficiency into hot electrons and P_R is the relativistic power unit ($P_R = 8.71\,\mathrm{GW}$) with the classical electron radius r_e (see also Sec. 4.3.2). In contrast to the model of Mora [43], time is implicit and the evolution of ion energies is instead analytically accessible via the ion position as it travels down the accelerating potential.

The energy an ion gains after traveling down the accelerating potential partially is described by

$$\mathcal{E}(\xi) = \mathcal{E}_{i,\infty} s(\xi). \tag{2.25}$$

The proton energy evolves asymptotically towards $\mathcal{E}_{i,\infty}$ with the function $s(\xi) = 1 + \xi - \sqrt{1+\xi^2}$, where $\xi = z/B$ is the normalized, longitudinal coordinate along the electron sheath expansion direction. The maximum ion energies at infinite distance and after infinite acceleration time, $\mathcal{E}_{i,\infty}$, are further limited via the finite laser pulse duration that defines the length of the energetic electron cloud to be τ_L as well[44, 116].

$$\frac{\tau_L}{\tau_0} = X\left(1 + \frac{1}{2}\frac{1}{1-X^2}\right) + \frac{1}{4}\ln\frac{1+X}{1-X} \tag{2.26}$$

Eq. 2.26 is the solution to the non-relativistic equation of motion

$$\frac{\mathrm{d}\xi}{\mathrm{d}t} = \frac{v_i(\xi)}{B} = \sqrt{\frac{2\mathcal{E}_i(\xi)}{m_i}}. \tag{2.27}$$

and describes how a limited pulse duration τ_L limits the acceleration process with respect to real laser system parameters. There, $\tau_0 = B/v_i(\infty)$ and $X = (\mathcal{E}_m/\mathcal{E}_{i,\infty})^{1/2}$. Since the maximum pulse energy at any laser facility is limited but the maximum intensity and pulse duration are coupled but in theory variable, the follow-up work of Schreiber, Bell, and Najmudin [116] reported the existence of an optimum pulse duration for the TNSA process which achieves the maximum attainable $\mathcal{E}_m$.

Ultrathin Foils Typically, laboratory laser ion acceleration experiments as well as PIC simulation studies show an optimum target thickness for highest proton energies. This can readily be understood from simple equations following Schreiber et al. [44]. In TNSA, ions gain their kinetic energy from the charge separation fields between themselves and the expanding electrons in the plasma sheath. The more electrons are within that sheath and the smaller the its volume is, the higher the accelerating fields become that the ion front experiences. Since the targets are usually highly overdense, the laser field, which transfers its energy to free electrons at the target front side, can only efficiently interact with the first plasma layer. It is as wide as the spot size w_0 and as deep as a few skin depths $L_s = c/\omega_{pe}$ (electron plasma frequency $\omega_{pe} = \sqrt{n_e e^2/(\varepsilon_0 m_e)}$) over which the laser intensity decreases exponentially with $1/e^2$. This situation is schematically depicted in figure 2.4a. The skin depths for typical plastic or metal

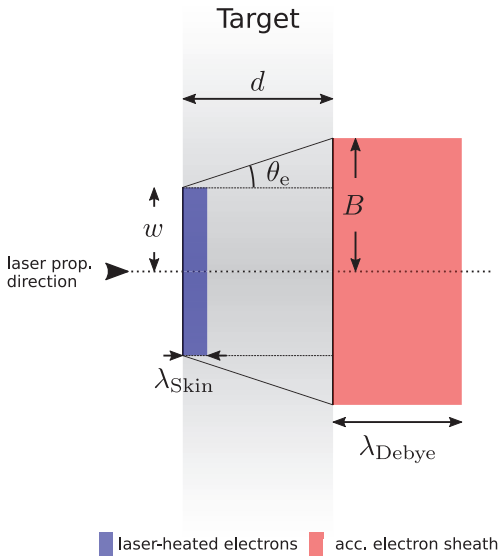

(a) Electron source and sheath geometry in target-normal sheath acceleration (TNSA).

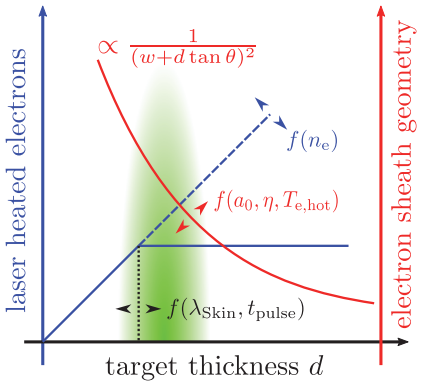

(b) Geometrical considerations implying the existence of an optimum target thickness d_{opt} for maximum ion energies in TNSA. The area shaded in green is where d_{opt} is suspected based on the maximized sheath density $n_{e,S} \propto d/\left(w_0 + d\tan\theta_e\right)^2$. For typical metal or plastic targets and current 10 J short-pulse ($\tau_L \sim 10$s of fs) laser systems, this means on the order of 10 to 100 nm.

Figure 2.4.: TNSA modeling schematic and geometry implications following [44].

targets with total electron densities of $n_{e,0}$ between 200 and upwards of 700 n_c are on the order of 10 down to 1 nm. Prompt electron bunches that the laser expels from the surface layer with a frequency of $2\omega_L$ travel through the target ballistically under an angle usually between $\theta_e = 15$ to 30° with respect to the surface normal. As the target thickness grows linearly with d, the sheath volume grows with $V_{sheath} = \pi\left(w_0 + d\tan\theta_m re\right)^2 \lambda_D$. The number of laser-accelerated electrons that can be spread over the sheath volume is furthermore limited by the laser pulse duration and follows $N_e \propto \eta I \tau_L$. While with increasing target thickness N_e first also increases, it is capped and with further increase of d, the possible electron density $n_{e,S}$ in the sheath rapidly decreases with growing sheath volume. Hence, the existence of an optimum thickness d_{opt} for TNSA logically follows. With existing short pulse laser systems ($\tau_L \sim 30$ fs, $I_{max} \sim 10^{21}$ W cm^{-2}), this optimal foil thickness is expected to be in the 10s of nanometer range[45, 117, 118].

Relativistically Induced Transparency The fact that the optimum thickness for TNSA is likely in the regime of ultrathin foils consequentially means that the target front- and rear-sides are not as disconnected as the ideal description would demand. Real laser pulses do not transfer their energy instantaneously (see later in Sec. 2.2.2) and early electron heating as well as evanescent penetration can lead to substantial rear-side ion gradients that have been shown to decrease the maximum attainable ion energy[119, 120]. Some earlier[121] and later works[122] couple the optimal target thickness d_{opt} and electron density n_e to an optimal areal electron density $\sigma_{e,opt} = d_{opt} \cdot n_e$. In the study of Mishra, Fiuza, and Glenzer [122], a relation is given to predict the optimum target parameters for an enhanced TNSA variant where the target becomes relativistically transparent just when the intensity maximum arrives:

$$\frac{n_{e,0}}{n_c} L_0 \approx c\tau_L \sqrt{\frac{a_0}{2^{3/2}} \frac{Zm_e}{m_I}} \tag{2.28}$$

As such, *relativistically induced transparency* (RIT) can supply the electron sheath with additionally heated electrons, leading to an increase in maximum proton energies[123, 124]. However,

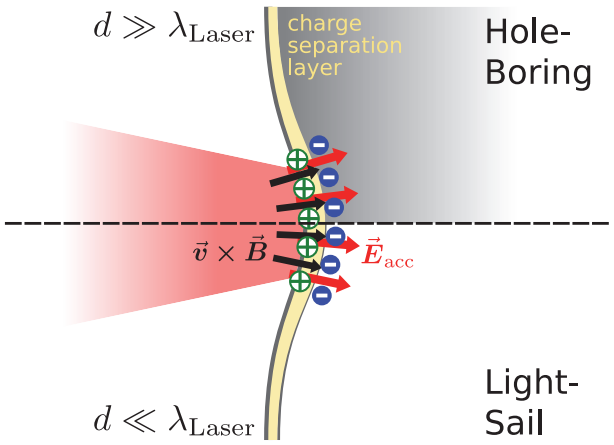

Figure 2.5.: Schematic figure of the radiation pressure acceleration (RPA) mechanism.

conditions like these are challenging to control with ultrathin targets where too early transparency could render the acceleration process completely ineffective. Additionally, other competing acceleration mechanisms complicate a possible diagnosis in laboratory experiments. These are briefly reviewed in the following part.

Radiation Pressure Acceleration

So far, only the acceleration of ions away from the target surface has been discussed. In a more direct manner, the laser pulse can also accelerate ions from the target front surface. If sufficiently high, the radiation pressure, $P_{\mathrm{rad}} = (1 + R - T)\frac{I}{c} = (2R + A)\frac{I}{c}$ (where $R + T + A = 1$ are the coefficients of reflection, transmission and absorption, respectively), depletes the front surface region of electrons[117]. A compression region of high electron density is created towards which ions from the electron-depleted region are accelerated. The electric field is directed antiparallel to the surface normal and ions are initially well-directed but increasing indentation of the critical surface upon longer interaction adds divergence to the accelerated beam. RPA strongly accelerates front-side ions but can occur in combination with the aforementioned TNSA. Which process dominates Depending on the target thickness, a further distinction between two variants of *radiation pressure acceleration* (RPA) is usually made.

Hole-Boring vs Light-Sail RPA The literature distinguishes between two different regimes radiation pressure acceleration (RPA), i.e. *hole-boring* (HB)[125–127] and *light-sail* (LS) RPA[17]. For the former, the target is usually only slightly overdense but of the order of or thicker than a laser wavelength. The surface at which the accelerating electric field is constantly created moves with a velocity v_{HB} as the laser pulse literally bores a hole through the target. In the light-sail regime the target is very thin such that the target moves as a whole entity[20]. Reportedly, quasi-monoenergetic ion bunches can be produced in such a scenario with higher energies than in the HB scenario since the there is no background plasma to shield the ions from the accelerating field. Many theoretical works have been performed under idealized, one-dimensional conditions, or been supported by 1D particle-in-cell (PIC) simulations. In higher dimensions, the occurrence of surface instabilities, such as the Rayleigh-Taylor instability, quickly broaden the proton spectrum and make this acceleration mechanism much more unstable.[21, 22, 128].

Conditions for RPA to be dominant Both previously discussed RPA mechanisms rely on the condition that the electrons at the front are pushed collectively as a dense slab. Therefore, the most efficient acceleration occurs with circularly polarized (CP) laser pulses for which the $\vec{j} \times \vec{B}$-force provides a constant push. As discussed in 2.1.1 for linear polarization, the oscillating part of the $\vec{j} \times \vec{B}$-force causes strong absorption and the creation of the high-energetic 2ω-bunches which can seed instabilities that break up the critical density surface. Furthermore, a 0° incidence ensures that no electric field components are perpendicular to the critical density surface which would also result in increased heating and expansion of the electrons. In 2D and 3D, the indenting of the target surface still causes electric field components pointing along the target surface normal and reduces the efficiency of the acceleration process again. It is reported that the ideal condition for LS RPA is at the threshold of transparency [20, 129] where $a_0 = \zeta$, when $\zeta = \pi \frac{n_e}{n_c} \frac{l}{\lambda}$.

Speed of the density front / light sail A relativistic treatment of the momentum balance of the ions and the laser pulse at the charge separation interface in its inertial rest frame [130] yields an expression for the hole-boring velocity.

$$\frac{v_{HB}}{c} = \frac{\sqrt{\Xi}}{1 + \sqrt{\Xi}} \tag{2.29}$$

Here in 2.29, Ξ describes the dimensionless *pistoning-parameter* $\Xi = I/m_i n_i c^3$. The final proton energy for a constant intensity is then

$$\mathcal{E} = m_i c^2 \left[\frac{2\Xi}{1 + 2\sqrt{\Xi}} \right] \tag{2.30}$$

Hole-boring RPA is similar to the so-called *Collisionless Shock Acceleration* (CSA)[19, 131] in the regard that the pistoning $\vec{j} \times \vec{B}$-force drives a sharp density gradient into the target. Here, the Mach number $M = v_{sho}/c_s$, i.e. the ratio between shock velocity v_{sho} and ion sound speed c_s (see 2.2.1), of the moving density front determines if an electromagnetic shock dominates ($M \leq 6.5$) or if pure pistoning ($M > 6.5$) drives the acceleration. In both cases, ions are viewed to be reflected from the moving density interface and gain double its velocity. While in principle the light-sail velocity can grow asymptotically towards c, the hole-boring velocity is due to the balance of radiation pressure and thermal pressure of the unperturbed plasma, i.e. the combined momentum change of the accelerated plasma portion[13].

Other Mechanisms

The most relevant acceleration mechanisms for the work performed in this thesis are TNSA and RPA. Other mechanisms exist and are also widely being researched as possible ways to constantly improve maximum ion energy and overall beam quality. The most promiment other mechanisms are briefly mentioned here. In *Directed Coulomb Explosion* (DCE), all the electrons are completely expelled from the (usually mass-limited) target, which can be a micro–/nanoscopic cluster or foil, and the remaining ions repel each other strongly due to their Coulomb-interactions, causing the target to explode. The ions are accelerated in all directions but with a preference for the direction of the removed electrons, yielding a broad energy distribution.

The aforementioned CSA, while sometimes confused with hole-boring RPA in literature, is more dominant in under– to near-critical targets. Another mechanism that is active in this density regime is called *Magnetic Vortex Acceleration* (MVA)[132]. It relies on a toroidal magnetic

field at the exit of a plasma channel that is created by escaping laser-accelerated electrons and returning less-energetic electrons. Its pressure expels low-energy electrons, thus forming a charge separation zone that accelerates and collimates ions from the ion density filament in the plasma channel. A mechanism that was proposed to still lead to very high ion energies after TNSA and RPA break down once the target has been turned transparent by the laser pulse is the so-called *break-out afterburner* (BOA) scheme[133–136]. It harvests the relativistic Buneman instability[137] which is believed to facilitate an efficient energy transfer between volumetrically heated electrons and expanding ions in the former TNSA sheath region.

2.2.2. Influence of Temporal Laser Pulse Shape on Ion Acceleration

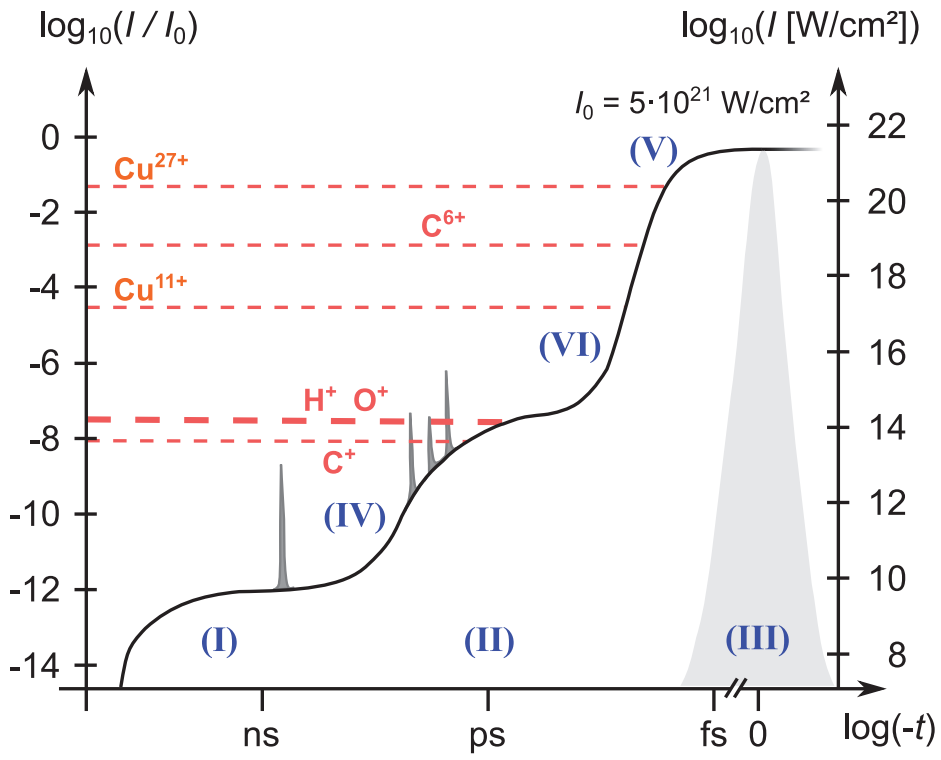

Figure 2.6.: Typical temporal contrast curve of a modern 10 J–class Ti:Sa short-pulse laser system, contrast-enhanced by a plasma mirror[138], such as DRACO at HZDR[G1]. Six different regimes and features have been marked in the schematic: Mainly there are the **(I)** nanosecond ASE region, the **(II)** picosecond "shoulder" / pedestal, and the **(III)** ~ 10 fs main pulse. Furthermore, distinct **(IV)** pre-pulses can already cross ionization thresholds 10s to few picoseconds prior to the main pulse arrival. The main pulse itself can show features such as **(V)** *skewness* and *chirps* on the "linear" intensity scale within the few FWHM around the peak. Prior to that, the **(VI)** intensity ramp from 10^{15} to 10^{18} W cm^{-2} on the "logarithmic" intensity scale can influence and shape the target during the early stage of the main pulse interaction. Several dashed lines mark the BSI appearance intensities (see Eq. 2.15) of selected atom species which mark the abundance of charge states in a classical field ionization treatment.

Ultra-high intensity, ultra-short laser pulses produced by current state-of-the-art 10 J class Ti:Sa laser-systems typically exhibit a temporal intensity contrast with three qualitatively different temporal regions and can have additional features which may all be of significance to the laser-ion acceleration process. These three regions along with three main feature types are displayed in a schematic laser contrast curve in Fig. 2.6. Not only the relative laser contrast but also the absolute intensity values and when, with respect to the arrival of the peak of the pulse, these intensities are reached is of vital importance to the dynamics of a laser-plasma

interaction. The work in this thesis (in particular Sec. 4.2) aims to shed light particularly on region VI, the last picosecond intensity ramp, and its ramifications for the laser-ion acceleration performance from ultra-thin foils.

Nanosecond ASE

Not only the laser contrast $I(t)/I_0$ in general is of importance but also the absolute intensities that are reached determine the character of the interaction. A laser pulse, such as the DRACO laser system at HZDR[G1] produces, consists of a nanosecond-duration base intensity **(I)** caused by *amplified spontaneous emission* (ASE). With $\lambda_L = 800\,\text{nm}$ (Ti:Sa) or $\geq 1054\,\text{nm}$ (Yb-doped diode-pumped) laser wavelengths the significant part of the ionization processes are predicted to occur in the quasi-static regime (see 2.2). However, at the ASE intensity of $I_{ASE} = 10^9 - 10^{11}\,\text{W/cm}^2$ is still too low to cause substantial tunneling ionization before the intensity increases further. Potentially, multi-photon ionization could free the lowest charge states but this process still requires on the order of 10 simultaneous photon interactions[6] and as such the MPI ionization rates are low. A study of how non-ionizing ASE can cause evaporation and early expansion of a target was conducted by Wharton et al. [139]. If the ASE does however surpass the ionization threshold, the target can be transformed into a near-critical density plasma by the time the laser maximum arrives[132, 140].

Experimental evidence obtained by Kaluza et al. [141] showed that a longer ASE duration increased the optimum thickness for maximum proton energies of aluminum foils irradiated with a $I_0 \geq 10^{19}\,\text{W/cm}^2$, $\tau_L = 150\,\text{fs}$ laser pulse and ASE durations varying between 150 ps and 6 ns. If the ASE creates a density shock that reaches the target rear well before the main-pulse-driven acceleration process starts, the sharp rear-side density gradient can degrade which was shown to be detrimental for TNSA[120]. Energy absorption, on the other hand, is likely increased at the target front side for metallic targets due to their partially quasi-free electrons in the interatomic band structure as opposed to insulator targets. Using analytical calculations, MHD and PIC simulations, Esirkepov et al. [142] already showed that an ASE and the "unclean" pre-(main)-pulse features of a finite contrast laser pulse can be beneficial in creating a pre-plasma from a "clean" µm–target that results in better laser energy absorption and possibly even self-focusing. When long pre-plasma gradients ($L_p \gg \lambda_L$) are created, vacuum heating loses significance in favor of resonance absorption and $\vec{j} \times \vec{B}$–acceleration, reducing the hot electron temperature T_e^{hot}[143].

Hadjisolomou et al. [144] recently examined ultrathin Mylar foils with a combined MHD and PIC study for 1 ns ASE and ultra-high main pulse laser intensities ($10^{20}\,\text{W/cm}^2$, $10^{21}\,\text{W/cm}^2$ and $10^{22}\,\text{W/cm}^2$). The authors found target density reduction prior to the main pulse, an exponential pre-plasma gradient shape and increased indenting towards thinner foils from the ASE interaction leading to patterns in the transverse proton profile and proton beam divergence increase. Initially the MHD simulations assumed a partly ionized plasma even though the intensities were $10^9\,\text{W/cm}^2$ and $10^{10}\,\text{W/cm}^2$. While melting and evaporation of the target can take place[142], with the advent of plasma-mirror techniques that increase the laser contrast prior to the PM trigger point by 1 to 2 orders of magnitude, the ASE is now often seen as non-ionizing.

[6] Taking hydrogen as an example, between 9 (800 nm) and 12 (1054 nm) photons are necessary to ionize the ground state.

Picosecond Shoulder

Uncompressed coherent light leaking around the optics in the laser's compressor chamber is usually the cause of a picosecond "shoulder" **(II)** in the laser contrast. With absolute intensities between 10^{12} to 10^{15} W/cm^2, this region is the cause of the first substantial cascade of tunneling and barrier suppression ionization processes. It can be cut away partly when contrast enhancement techniques like *Ultrafast Pockels Cells* and *Cross-Polarized Wave Generation* (XPW) are used which also reduce the ASE and possible pre-pulses[13, 143]. This part of the pulse can already significantly shape the target with ionization, heating and pre-expansion as described in Sec. 2.1.3 and above. With a duration between several 100 ps and a few 10 ps spanning several orders of magnitude in intensity, the ps-shoulder's exact influence on the ion acceleration process is still a topic of active research. This research is also complicated by the fact that it creates the state of *warm dense matter* (WDM) which is rich in atomic– and molecular physics where an exact *equation of state* (EoS) description would be required but is often unavailable. If the target was previously unperturbed, the ps pedestal can create a very short ($L_\mathrm{p} \ll \lambda_\mathrm{L}$) density gradient that was found have a beneficial effect on electron heating[145–147]. Flacco et al. [143] furthermore reported on the basis of MHD simulations that the rear-side plasma gradient of ultrathin targets can be degraded if energetic electrons and photons trigger an expansion earlier than 50 ps prior to the laser peak intensity. A completely missing picosecond shoulder, however, was shown to cause a similar expansion of target front– and rear sides, leading to comparable upstream and downstream proton expansion and final energies[148]. The latter is true for a long as the maximum laser amplitude does not significantly exceed $a_0 = 1$ and TNSA dominates since otherwise RPA effects would break this symmetry. Experimentally, the optimum target thickness and proton energy have been widely observed to shift with contrast settings that change the picosecond shoulder.

Ultrashort Main Pulse

The laser main pulse **(III)** is the shortest and most energetic part of the laser pulse contrast. Its duration is often given as the FWHM of intensity and for this thesis, ultrashort $\tau_\mathrm{L} < 100$ fs are the main focus. It has only become possible to produce such short pulses with intensities exceeding $I_0 \sim 10^{21}$ W/cm^2 with the advent of *chirped pulse amplification* (CPA)[16]. Nowadays, commercial tabletop CPA laser systems are common but the technique has continuously been developed over the years to push past contemporary intensity limits. Two representative systems are the double-CPA DRACO[G1] system at HZDR and the *optical parametric CPA* (OPCPA) J-Karen-P laser at the Kansai Photon Science Institute (KPSI)[149, 150]. While both are relatively similar Titan:sapphire laser systems, their characteristic laser contrast differs already due to the final amplification method, giving the latter an inherently better contrast on the ASE level while the former has a slightly steeper contrast on the last few hundred femtoseconds.

Most of the available theoretical, experimental and numerical work regarding acceleration performance and proton energy scalings (e.g. in [106]) is concerned with changes to the main pulse, namely τ_L and I_0, that is often assumed to be either Gaussian– or *sinc*–shaped in intensity. Naturally, said scalings exist for each of the aforementioned acceleration mechanisms (see Sec. 2.2.1) and for the sake of brevity here, the reader is referred to the excellent reviews of Daido, Nishiuchi, and Pirozhkov [13] and Macchi, Borghesi, and Passoni [117]. Fuchs et al. [106][151] have identified a robust scaling law for the TNSA mechanism which is based on the isothermal plasma expansion treatment explained before (see Par. 2.2.1) but that limits the acceleration time both due to energy transfer from electrons to ions in the sheath as well its divergence-related rarefaction. Vitally, very long pulses strongly influence the rear-side plasma expansion relatively directly as they constantly feed in laser-heated electrons which

relates the acceleration time t_{acc} linearly to τ_L. But for very short pulses ($\tau_L < 150\,fs$), the energy transfer time from electrons to ions seems to approach a constant value, $t_{transfer}$, such that $t_{acc} = \alpha(\tau_L + t_{transfer})$. The authors of [151] determined empirically that $t_{transfer} = 60\,fs$ and $\alpha = 1.3 .. 3$ varies linearly between $I_0 = 3 \cdot 10^{19}\,W/cm^2$ and $2 \cdot 10^{18}\,W/cm^2$. Concerning the maximum achievable proton energy $\mathcal{E}_{max}$, most works have found a dependence that scales with between the square root of the laser power P_L up to its linear proportion. . However, Zeil et al. [152] showed that even for 30 fs short, relativistic pulses, the intra-pulse phase plays a significant role in the accelerated proton spectrum and especially their spatial distribution as the sheath in a TNSA scenario is supplemented with highly directed, prompt electrons, thus breaking with a purely isothermal description of the process.

Pre– and Post-Pulses

A variety of other features in the laser temporal contrast can lead to plasma dynamics that substantially alter the acceleration process. One of these are well-separated *pre– and post–pulses* **(IV)** that originate in partial reflections of (e.g. within the amplifier chain) and possible subsequent non-linear coupling with the main pulse[153]. If pre-pulses ionize the target very early, already the ASE or picosecond shoulder can drive a shock through the target that leads to large rear-side gradients and a reduction of target density from early expansion[143]. This can even turn the priorly overcritical target into a near-critical one where the dominant acceleration mechanism was changed from TNSA to directed Coulomb explosion DCE.[154, 155]. On the other hand, nanosecond to 100 femtosecond pre-pulses were also reported as a means to shape the front-side density gradients to enhance absorption of micron-sized targets irradiated by picosecond kilojoule-class laser pulses[156, 157]. Currently, this degree of control has not yet been reported for nanometer-scale ultrathin foils and femtosecond main pulses where shot-to-shot fluctuations are all the more critical and time-resolved diagnostics of the state of the plasma are very challenging to come by. A special case of a pre-pulse scenario is one where one or multiple pre-pulses are on the same order of intensity magnitude as the main pulse. Such pulse-pairs or –trains were reported to produce distinct peaks in the proton spectra[158], increase conversion efficiency[159] and even increased proton energies[160, 161].

Main Pulse Asymmetries

Unlike the common simplification of a Gaussian- or Sinc-shape, the main pulse itself can exhibit temporal *asymmetries* **(VI)**.[7] Two categories that asymmetries can be classified into are *skewness* and *chirp*. Skewness describes a difference between the leading and trailing pulse envelopes while a temporal chirp introduces a varying central wavelength along the laser propagation direction. In the relevant literature both categories are still illuminated to very little extent.

Skewness In the TNSA regime of micron-sized targets with absolute laser intensities of $I_0 = 1 – 2\,W/cm^2$, previous works study the influence of skewness on the proton acceleration performance by means of 1D[40] and 2D[42] particle-in-cell simulations. Both studies treat overdense targets similar to cryogenic hydrogen with initial electron densities of $n_e = 12\,n_c$ and $20\,n_c$ and similar laser pulse duration $\tau_L = 34\,fs$ and 40 fs, respectively.

Souri, Amrollahi, and Sadighi-Bonabi [40] find that slow-leading[8] pulses have a beneficial

[7] Naturally, a real laser pulse has also spatial and angular irregularities. As before, this chapter mainly focuses on the temporal aspects of the laser contrast curve.

[8] *slow-leading* means that the leading pulse edge is shallower than in the symmetric case. Vice versa, *fast-leading* translates to a steeper edge.

effect on proton acceleration, as they observe an initial pile-up of front-side electrons without substantial rear-side expansion due to the lower penetration depth of the laser pulse. The later part of the pulse strongly heats the piled-up density to a larger depth and pushes electrons to the rear-side. An increased number of protons is accelerated to high energies. The fast-leading pulses led to a similar proton cutoff energy as the symmetric case but with a decrease in accelerated proton number. Different a_0 between 8 and 14 all show the same acceleration behavior.

Kumar, Gopal, and Gupta [42] performed 2D PIC simulations of asymmetric pulses interacting with a foil target and additionally varied between a steep density gradient and a long, underdense pre-plasma ramp. Contrary to the findings above, the best acceleration performance is observed for fast-leading pulses, whereas slow-leading pulses performed substantially worse. The results are explained with a stronger ponderomotive force from the strong intensity gradients of the fast-leading pulses that, especially in combination with a pre-plasma, accelerate more electrons to high energies which directly translates to the electron density in the sheath and accelerating electric field for rear-side protons. Fits to electron spectra show significantly larger electron temperature values as well as larger electron numbers for fast-leading pulses.

Both studies, however, do not show an evolution of proton energies over time and evaluate $\mathcal{E}_{max}$ relatively early after the peak interacts with the target. In reduced dimensions, the acceleration process takes longer than in 3D and since the probe time is well within 1 to 3 laser periods in both works, the acceleration process might still be qualitatively different before the real quantitative effect of the asymmetric pulses shows. However, it becomes clear the dynamics of ponderomotive steepening[72, 162] differ which can conceivably lead to altered proton acceleration with asymmetric pulses.

Chirp In their work, Souri, Amrollahi, and Sadighi-Bonabi [40] also investigated a chirp of the form $\omega_L(\xi) = \omega_0 + b\xi$ (with the retarded time $\xi = t - x/c$) and kept a symmetric electric field envelope. A chirp such as that results in an asymmetry of instantaneous a_0 between the rising and trailing flanks of the pulse. The authors observed an increased sheath field strength for positively chirped pulses where lower ω_L arrive earlier, leading to higher cutoff energies. Negatively chirped pulses were reported to turn the target transparent too early, which led to sub-optimal heating, a lower sheath field and less proton energy.

The same authors studied the influence of chirped, circularly-polarized pulses in numerical 1D3V PIC simulations[41]. Mainly, RPA was the dominant ion acceleration mechanism and both positively as well as negatively chirped pulses were found to be beneficial for the spectral quality of the proton beams. While the maximum energy did not change substantially, clear enhancement to the energy spread and number of accelerated protons was found especially for the negatively chirped pulse. These effects were attributed to better phase matching between electric field and protons, higher initial heating and suppressed self-induced transparency.

Tayyab et al. [163] found experimentally that with a constant laser pulse energy of 2.1 J there is an optimum pulse duration between 25 to 500 fs for the acceleration of protons from 0.4 to 1.5 µm metal foils. Even higher energies were observed when the pulse is positively chirped which manifested in a slow-leading edge. The enhancement is attributed to longer pre-plasma expansion, better absorption and an increased acceleration time.

Last Picosecond Ramp

A last and largely unexplored region of the laser contrast is the intensity ramp towards the main pulse maximum on the logarithmic contrast scale between 10^{16} W/cm^2 and 10^{19} W/cm^2. For an ultra-short UHI laser pulse this means the *last picosecond intensity ramp* **(V)**. The origin of

its shape is typically to be found in temporal chirps that have different wavelengths arriving at different times, thus broadening the pulse and possibly introducing distinct pre-pulses. Typically, the electric field of a laser pulse cannot be directly measured. Instead, the spectral phase $\Phi(\omega)$ and the spectral intensity $I(\omega)$ are recorded and with reconstruction methods the electric field $E(\omega) = A(\omega)\exp(-i\Phi(\omega))$ and via Fourier transform the temporal intensity contrast is recovered. The several existing methods deliver a dynamic range, temporal resolution and possibly artifacts that depend on the specific physical measurement process and often come at the expense of one another. As such, reconstructing the full laser contrast with a high degree of certainty and femtosecond resolution is very challenging.

A laser pulse of limited bandwidth $\Delta\omega$ and a so called *flat phase* (meaning $\Delta\varphi(\omega) = 0$ in the focal plane) creates a *Fourier-limited* temporal envelope that signifies the shortest possible pulse for a given spectrum. The analytical description of the spectral phase can be Taylor-expanded as follows

$$\varphi(\omega - \omega_0) = \varphi_0 + \varphi_1 \cdot (\omega - \omega_0) + \varphi_2 \cdot \frac{(\omega - \omega_0)^2}{2} + \varphi_3 \cdot \frac{(\omega - \omega_0)^3}{6} + \dots . \tag{2.31}$$

In order, the terms signify a constant phase offset φ_0, a temporal delay φ_1, the so-called *group delay dispersion* φ_2 or *group velocity dispersion*[9] (GVD), and the *third order dispersion* (TOD) φ_3. In particular the GVD and TOD terms can introduce significant alterations from the Fourier-limited pulse. The former, adding a *chirp* to the pulse, symmetrically stretches the pulse while the latter adds an asymmetry, effectively tilting the rising and falling ramps as Fig. 2.7 shows. The TOD effectively also stretches the pulse and reduces the maximum intensity but, additionally, changes features like the intensity ramp steepness while also introducing or reducing pre-pulses and shifting the main pulse away from a Gaussian or sinc-shaped form. The large-scale simulation campaign in chapter 4.2 of this thesis focuses especially on the the steepness of the leading intensity ramps based on measurements taken at the DRACO laser system. For

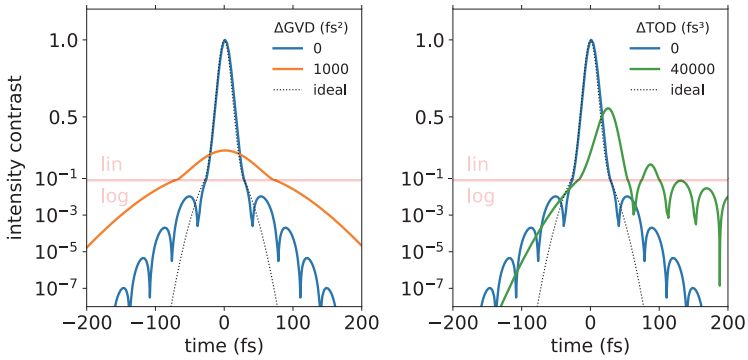

Figure 2.7.: Examples of spectral phase changes $\Delta\varphi(\omega)$ applied to a laser pulse with initially 30 fs Gaussian main pulse constructed from a Supergaussian spectral intensity. **Left:** Applying ΔGVD broadens the laser pulse symmetrically. This effect is independent of the sign of ΔGVD, as long as the initial spectral phase $\varphi(\omega)$ is constant (i.e. *flat*). **Right:** Applying positive ΔTOD also broadens the main pulse, reduces the intensity (steepens the edge) and pulsed envelope structure leading ahead of the main pulse and shifts it to the trailing pulse where more pronounced post-pulses are created. Reversing the sign of ΔTOD reverses the effect with respect to time, creating pronounced pre-pulses and a steeper trailing edge. In both cases, such phase terms lead to non-ideal compression of the pulse and reduction of peak intensity which instead shifts to leading and trailing flanks.

[9] since the φ_1–term is inversely proportional to the group velocity $v_g = \mathrm{d}\omega/\mathrm{d}k$

long pulse lasers and micron-sized targets, Schollmeier et al. [164] already studied the influence of an intensity ramp on the 10-picosecond scale experimentally and with simulations, concluding that especially this time scale defines the main pulse interactions and needs to be included in modeling as well since shot-to-shot proton acceleration performance fluctuations are a direct result of contrast fluctuations in the ramp. Furthermore, modeling of a realistic pre-pulse and pre-plasma scenario were found to be crucial to achieve predictability in kinetic plasma simulations where maximum energies were not only exceedingly high for the idealized Gaussian pulse but only the realistic pulse matched the experimental findings also with respect to the proton transverse profiles and electron spectra.

2.3. Numerical Plasma Modeling with Particle-in-Cell Codes

Numerical modeling of laser-particle acceleration experiments has been an extremely helpful tool for the understanding and prediction of the detailed dynamics involved in its real-life counterpart. Two main approaches are so far capable of reflecting the processes in sufficient detail while also allowing to simulate the full regions of interest on the computer resources that are available today. These are namely the *kinetic description* with particle-in-cell (PIC) codes and the *fluid description* with magneto-hydrodynamic MHD codes. Naively, two interaction types have to be modeled to describe a plasma numerically which are the *action of external electromagnetic fields* on every single particle in the plasma and the *particle-to-particle interaction* between all of its constituents. Following through with this thought, $N(N-1)/2$ interactions have to be computed. It becomes clear that very quickly extreme computational resources are necessary as the problem size increases since the number of computations grows with N^2. For this reason, so-called molecular dynamics (MD) codes are usually limited to setups of only a few thousand particles (millions on the most powerful machines). They do, however, provide understanding and very accurate predictions of complex materials since they perform *ab initio* calculations (i.e. from first principles) that cannot be covered to this degree of detail in MHD and PIC.

The naive *kinetic* description needs to be optimized to be able to tackle the problem sizes of relevance for laser-particle acceleration applications. These applications range from LWFA scenarios in a volume of $(0.1 \times 0.1 \times 10)$ mm with particle densities of $\gtrsim 10^{18}\,\mathrm{m}^{-3}$ to TNSA scenarios with micron-sized metal foils in an active volume of (15×30) µm and densities in excess of $10^{30}\,\mathrm{m}^{-3}$. This means one can easily have anywhere from 10^8 to 10^{15} particles directly involved in the interaction. The collective ensemble of all particles in a plasma is represented by the 7D distribution function $f(\vec{r}, \vec{v}, t)$ in phase space (position $\vec{r}$, velocity $\vec{v}$, time t). The *Boltzmann equation* (Eq. 2.33) expresses the time evolution of such a generalized distribution function.

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \vec{v}\frac{\partial f}{\partial \vec{r}} + \frac{\vec{F}}{m}\frac{\partial f}{\partial \vec{v}} = \left(\frac{\partial f}{\partial t}\right)_{\text{collisions}} \qquad \left|\omega_{\mathrm{pe}}/\langle\nu_{\mathrm{e-i}}\rangle \lesssim 1\right. \tag{2.32}$$

$$\approx 0 \qquad \left|\omega_{\mathrm{pe}}/\langle\nu_{\mathrm{e-i}}\rangle \gg 1\right. \tag{2.33}$$

When the force $\vec{F}$ is replaced with the Lorentz-force $\vec{F}_{\mathrm{L}} = \left(\vec{E} + \vec{v} \times \vec{B}\right)$, one receives the *Vlasov equation*. For sufficiently high laser intensities and low densities the plasma can be treated as *collisionless* and the right side of Eq. 2.33 becomes zero. In practice, this applies to laser-electron accelerators and their gas targets with very good approximation. For laser-solid interactions, long (> 100 fs) interaction times and especially for targets of micron-scale thickness or more, the plasmas become increasingly collisional and the right side of Eq. 2.32 cannot be neglected, anymore.

Fluid modeling Further analytical treatment of the first two moments of the Vlasov equation, i.e. the *average density* $n(\vec{r}) = \int f(\ldots)\mathrm{d}^3 r$ and the *average velocity* $\vec{u}(\vec{r}) = \int f(\ldots)\vec{v}\mathrm{d}^3 r$, leads to the full set of 14 MHD equations for the 14 unknowns $\rho, p, \vec{u}, \vec{j}, \vec{E}$, and $\vec{B}$. Of note here are the collective quantities *charge density* ρ, pressure p and current density $\vec{j}$. An important prerequisite for the fluid approach to be valid is the assumption of *local thermal equilibrium* (LTE) throughout the whole plasma. This implies "long" interaction time scales ($t \gtrsim$ ps), local transport models $\left(I_\mathrm{L} \lesssim 10^{14}\,\mathrm{W\,cm^{-2}}\right)$, quasi-neutrality, and linear laser-matter interaction $\left(I_\mathrm{L}\lambda_\mathrm{L}^2 < 10^{17}\,\mathrm{W\,cm^{-2}\,\mu m^2}\right)$. These restrictions limit the maximum laser intensity for the application of MHD codes.

The UHI laser pulses used for producing the currently highest ion energies are operating with peak intensities between 10^{21} and $10^{22}\,\mathrm{W\,cm^{-2}}$. They are, however, not perfectly Gaussian-shaped in time but have a non-trivial pulse shape that transitions through many different intensity-regimes on picosecond and nanosecond ramp-up phases for which fluid models are applicable. The very steep ramps and overall highest intensities around the peak ionize localized regions of the target in a matter of femtoseconds and the free electrons that are produced become strictly non-thermal. That is because these processes do not leave the time necessary for inter-particle collisions to drive the distribution to a more Boltzmann-like state. This means that averaged quantities of pressure, charge- and current density cannot be calculated directly, anymore and a fully kinetic treatment becomes necessary. Naturally, this requires a larger compute effort since the momentum coordinates of individual representatives of the particle distribution are needed, instead of purely spatial information. Particle-in-cell simulations are able to describe plasmas in a fully kinetical way and in the following section its key concepts are introduced.

It remains to be noted here that usually only the few 100 fs around the peak of the main pulse can be treated with PIC simulations. This is, on the one hand, due to the extreme computational cost but also because the system quickly leaves the region of applicability again. Additionally, there is a gap between the conditions where a plasma can be fully described in the fluid approach and the region of fully kinetic treatment where neither of these methods applies to the full system. So-called *Vlasov-Fokker-Planck* (VFP) codes are promising to bridge this gap[89]. Regions where N-body collisional processes at solid density dominate and begin to matter on distances within the Debye-sphere ($\omega_\mathrm{pe}/\nu_\mathrm{e-i} \ll 1$) the plasma is treated by the Fokker-Planck part, whereas the regions of already expanding, close to fully-ionized, HED plasma ($\omega_\mathrm{pe}/\nu_\mathrm{e-i} \gg 1$) is treated by the Vlasov part. For almost 40 years, these methods have been studied and the field is still very fertile with new developments[165, 166] Until now, however, no single code can universally describe a short-pulse laser-plasma system with few assumptions while covering all relevant physical aspects and so the VFP method remains mostly reserved to specialized problems in, e.g. inertial confinement fusion.

2.3.1. Kinetic Plasma Modeling: PIC Method

As described earlier, the sheer number of particle-particle interactions rules out the treatment of each individual electron and ion for realistically sized laser-particle acceleration experiments. To alleviate this problem, so-called *macroparticles* (also super-particles) are introduced. They each represent a patch $dN = f(\vec{r}, \vec{v}, t)\,\mathrm{d}^3 r\,\mathrm{d}^3 v$ of the 6D particle phase space, i.e. within the *discrete* element $\left[\vec{r}, \vec{r} + \mathrm{d}\vec{r}\right] \times \left[\vec{v}, \vec{v} + \mathrm{d}\vec{v}\right]$. This reduction in particle number is possible since the Lorentz-force acts the same way on all particles that have the same charge to mass ratio q/m.

Furthermore, the electromagnetic fields are discretized on a grid that defines a *cell* as its smallest structure. Macroparticles move within these cells and all interactions between the external electromagnetic fields and the fields of the particles are mediated by grid-operations

(the *particle-mesh* method). Each macroparticle represents also a certain number of real particles, i.e. its *weighting*. This number typically ranges between several 100 to 10 000s. A macroparticle is a marker within the phase space distribution and has only one momentum $\vec{p}$ but spans over a finite spatial volume. While computational effort is reduced, the artificially increased granularity can cause cause numerical noise. A *weighting function* (or also *macroparticle shape*) of higher degree can reduce this noise but spreading the macroparticle to more cells also requires more computational resources.

The basic particle-in-cell algorithm consists of four distinct steps that are performed in cycles and each completed iteration advances time in the simulation for one finite *time step* Δt_{PIC}. The governing set of equations is in principle the full set of Maxwell's equations.

$$\vec{\nabla} \cdot \vec{E} = \frac{1}{\varepsilon_0} \sum_{\mathrm{s}} \rho_{\mathrm{s}} \tag{2.34}$$

$$\vec{\nabla} \cdot \vec{B} = 0 \tag{2.35}$$

$$\vec{\nabla} \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{2.36}$$

$$\vec{\nabla} \times \vec{B} = \mu_0 \left(\sum_{\mathrm{s}} \vec{J}_{\mathrm{s}} + \varepsilon_0 \frac{\partial \vec{E}}{\partial t} \right) \tag{2.37}$$

However, Gauss' laws for electric (Eq. 2.34) and magnetic fields (Eq. 2.35), $\vec{E}$ and $\vec{B}$, are not explicitly enforced. Specifically, close-range Coulomb interactions between particles are neglected this way on the inter-cell level. Particle-in-cell codes are valid for plasmas where all collisions are accounting for only small-angle momentum changes and their effect results in the collective compensation of external fields and dynamic shielding of single particle charges on a scale the Debye length λ_{D}. Faraday's law of induction (Eq. 2.36) and the extended circuital law of Ampère (Eq. 2.37) combined with the Boltzmann equation make the set of *Vlasov-Maxwell* equations which are computed within PIC.

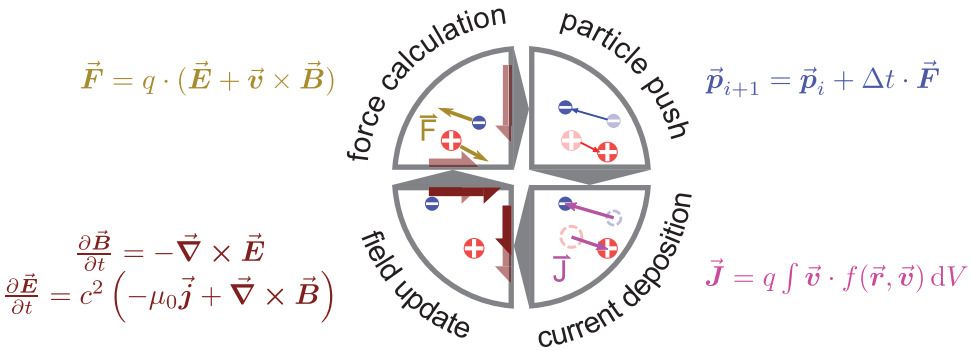

Figure 2.8.: Chart of the explicit particle-in-cell cycle for electro-magnetic PIC codes. Each iteration starts with the interpolation of the Lorentz-force to the macroparticle positions. Its effect in terms of position and momentum changes for each particle is then calculated in the so-called *particle push*. Afterwards, the resulting electrical current of each particle is deposited onto the grid. As a result of these currents, the electric and magnetical fields are finally updated.

Figure 2.8 shows the explicit PIC-loop and the essential equations for each step.

1. Calculation of forces The Eulerian representation of electromagnetic fields on a mesh stores values for E and B on specific points of a cell (corners, sides, ...). At first, the electric

and magnetic fields have to be interpolated to the positions of the individual particles they are acting upon. Particles, while represented by their continuous position and relativistic momentum coordinates in a Lagrangian frame of reference, are assigned their fields and forces under consideration of their macroparticle– or *cloud shape*[167]. It describes the degree by which the macroparticles are spread out spatially. In the cloud assignment scheme, the particles are weighted for each cell by their overlap, i.e. the amount of charge that is contained by that cell. In a different approach, a so-called *assignment function* is evaluated at the position of the center of a cell.

$$\vec{F}_{\mathrm{L}} = q \cdot \left(\vec{E} + \vec{v} \times \vec{B}\right) \tag{2.38}$$

From the electromagnetic fields, weighted with the correct assignment for each cell, the Lorentz-force (Eq. 2.38) that moves the particles is calculated.

2. Moving the particles In a second step, the forces move the particles and their new position and changed momentum values are obtained by a discrete integration of the equations of motion (Eqs. 2.39, 2.40).

$$\ddot{\vec{r}} = \frac{1}{m_{\mathrm{S}}}\dot{\vec{p}} = \frac{q}{m_{\mathrm{S}}} \cdot \left(\vec{E} + \vec{v} \times \vec{B}\right) \tag{2.39}$$

$$\dot{\vec{r}} = \frac{\vec{p}}{\gamma m_{\mathrm{S}}} \tag{2.40}$$

Here, γ is the relativistic Lorentz-factor $\left(1 - v^2/c^2\right)^{-1/2}$ and m_{S} denotes the mass of the particle species. Very common is the *leap-frog* method[168] which uses finite-difference calculations to advance position and momentum with second order accuracy, but shifted by half a time step against each other, since these quantities cannot be known at the exact same time. Force and position are *time-centered* to full iterations $t = 1, 2, 3 \ldots \Delta t_{\mathrm{PIC}}$ and velocities are calculated in between at $t = \frac{-1}{2}, \frac{1}{2}, \frac{3}{2} \ldots \Delta t_{\mathrm{PIC}}$ from forces adjusted to these times.

3. Current Deposition Then, the electrical current from each particle motion is calculated and deposited onto adjacent grid points. Again, the extended macroparticle shape $\Delta\rho_{\mathrm{S}}$ demands that the charge density needs to be integrated for each overlapped cell boundary and each macroparticle, which usually makes this step the most expensive and time-consuming one in the PIC-loop.

$$\underbrace{\dot{\rho}_{\mathrm{S}}}_{\substack{\text{depends on} \\ \text{particle shape}}} = - \underbrace{\vec{\nabla} \cdot j_{\mathrm{S}}}_{\substack{\text{depends on} \\ \text{field discretization}}} \tag{2.41}$$

This step implements the continuity equation and implicitly also fulfills Gauss' law2.34 self-consistently on a local cell level. In addition to that, initial conditions and boundary treatment have still to be chosen to explicitly enforce a scenario where both Gauss' laws (2.34, and 2.35) are fulfilled.

4. Electromagnetic field update One full iteration concludes with the update of electric and magnetic field from the accumulated particle current. As the step before, the field update depends on the specific discretization scheme. Typically, the method of Yee[169] is chosen, where next-neighbor finite differences are employed to explicitly solve the laws of Faraday (2.36) and Ampère-Maxwell (2.37). There, E and B are shifted by half a time step, $\Delta t_{\mathrm{PIC}}/2$, and the discrete positions for their components are staggered by half a cell. This choice allows for second-order accuracy in the *finite-difference time domain* field-solver methods.

Stability Conditions and Rules of Thumb

Solving the equations of motion with the *finite difference time domain* (FDTD) solvers like Yee's method is affected by numerical errors which scale quadratically with the length of the time step Δt_{pic} and the grid resolution $\Delta\vec{r}_{\text{PIC}} = (\Delta x, \Delta y, \Delta z)_{\text{PIC}}$. To fulfill Gauss' law implicitly and avoid having to solve a global Poisson equation, it is necessary to limit the size of each cell such that particles can only reach neighboring cells within one PIC iteration. But also the explicit solutions for the curls of the electromagnetic fields, Eqs. 2.36 and 2.37, place a restriction on the time and grid resolution to achieve numerical stability [170]. This is the *Courant-Friedrichs-Levy* (CFL) condition which takes a different form depending on the discretization scheme and Maxwell solvers that are used. In the Yee scheme it reads

$$c_0 \Delta t \leq \left(\sum_i \frac{1}{x_i^2} \right)^{-0.5} \tag{2.42}$$

With Eq. 2.42 the cell size of the spatial grid becomes coupled with the temporal resolution via the length of the time step.

Furthermore, it is easily comprehensible that if the PIC-simulation is meant to predict the detailed trajectories of particles inside the electric field of a laser pulse, the laser wave has to be sampled sufficiently well enough. This concept can be generalized to the statement that the overall fastest oscillation of a plasma system has to be resolved for the physics to be properly modeled. In underdense laser-matter interaction setups the laser wavelength mainly determines the spatio-temporal resolution. Especially in laser-propagation direction, the grid spacing needs to be particularly fine to resolve the central laser wavelength λ_{L}. At least 25 to 35 sample points are usually recommended to avoid significant numerical dispersion that would build up over many iterations. Perpendicular to the laser-propagation direction this requirement is more relaxed and cells can be made longer. This sacrifice in transverse resolution gains the advantage of less compute memory that can either be used for more simulations in parallel or a larger, more realistic simulation box.

Section 2.1.3 showed that one very important mechanism of energy transfer from the laser pulse to a solid-density plasma can be via electron plasma waves. Hence, the frequency $\omega_{pe} = \omega_{\text{L}} \sqrt{n_{\text{e}}/n_{\text{c}}}$ begins to dominate the resolution requirements as soon as $n_{\text{e}} > n_{\text{c}}$. Strongly dependent on plasma density, plasma waves can propagate in every spatial direction based on the geometry of the target. Therefore, a grid of cubic cells is often the best choice. In that case, the CFL criterion (2.42) simplifies to $c_0 \Delta t \leq \Delta x$ (1D), $c_0 \Delta t \leq \Delta x/\sqrt{2}$ (2D), and $c_0 \Delta t \leq \Delta x/\sqrt{3}$ (3D), respective to the dimensionality of the simulation.

Following Tskhakaya et al. [171] (or [172]), a linear harmonic oscillator can be considered in the leap-frog scheme where the differential equation $\mathrm{d}x/\mathrm{d}t^2 = -\omega_0^2 x$ turns into

$$\frac{x_{k+1} - 2x_k - x_{k-1}}{\Delta t^2} = -\omega_0^2 x_k \tag{2.43}$$

and plugging in the solution $x_k = a \exp(-i\omega t)$ leads to $\sin(\omega_\Delta t/2) = \omega_0 \Delta t/2$. The cumulative phase error after N_t steps should then be approximately $\Delta(\omega\Delta t) \approx \left(N_t(\omega_0 \Delta t)^3\right)/24$, where N_t denotes the number of PIC-iterations. Assuming $\Delta(\omega\Delta t) < \pi$ and $N_t = 100\,000$, the often-used criterion

$$\omega_{\text{p}} \mathrm{d}t \leq 0.1 \tag{2.44}$$

follows. As section 3.2 will show later, the combination of these restrictions can severely limit the extent of the possible simulation volume for a full 3D3V PIC simulation in a solid-density scenario, even on the most powerful supercomputers today.

2.3.2. Extensions to the PIC-Cycle

To establish control over the intricate processes involved in laser ion acceleration it is necessary to understand the dynamics at the femtosecond-nanometer level. Being the most used and trusted tool to study high energy density plasmas numerically, the particle-in-cell method is able to resolve these time and length scales. However, the PIC method by itself describes only collisionless plasmas of constant phase space volume and already the physics of how free electrons are created are not included. Therefore, it has become commonplace to include physics extensions that calculate ionization–, collision–, and even quantum electrodynamic processes. This section summarizes the modeling extensions to the highly scalable particle-in-cell code PIConGPU, which was the main tool that was used to create the results in this thesis. Additional details about parallel ionization implementations of the author can be found in Sec. 3.1.1. It is important to note that state-of-the-art atomic physics modeling in PIC codes does not necessarily imply the most recent and exact modeling of an atomic system. Rather, the immense computational cost of PIC simulations and the statistical representation of particles through marker-particles demands that the physics extensions are efficient, ideally low-costed in terms of memory consumption and computation time, and well-tested to ensure a versatile PIC framework for various physics setups.

Field Ionization

In underdense gas targets, the laser can interact very directly with the atoms and due to the low density, each atom can be viewed in isolation from each other when it comes to ionization. Since the invention of the laser 60 years ago[173], many ionization models have been developed for such a scenario and are constantly being refined (see, e.g. [81] for a brief overview). For gas targets and laser intensities up to $10^{18}\,\mathrm{W\,cm^{-2}}$ it was also possible to validate these models with specifically designed experiments[174–176]. Recently, it has also become possible to deconvolve the sub-femtosecond multi-ionization dynamics of gases like helium[177, 178] and neon using the attoclock technique[179–181] and few-cycle optical laser pulses. These works suggest that for the extreme cases of ultra-short, super-intense laser pulses simple rate models are not fully reliable, as sequential ionization events are influenced by the orientation of electron holes in the valence shells of atoms, created by prior ionization processes. Nevertheless, simple rate models are widely employed for modeling ionization in PIC codes since they greatly reduce the resources needed to describe complex atomic physics while still delivering reasonably accurate results. For highly *overdense* targets that are used in laser-ion acceleration, the same ionization models are used but here experimental validation is basically non-existent. Under such conditions atoms and ions are surrounded closely by others and the application of models that assume single atoms in a vacuum implies an error that is usually treated with parameter scans over different initial conditions. In addition to that, the hope is usually that the ramp phase is so short and steep that any energy distributions imprinted on the earlier plasma are quickly *overwritten* by the following laser periods as the intensity increases exponentially. This happens out of the reasonable presumption that the interaction modeled with the PIC never reaches thermal equilibrium and also for lack of more profound methods and diagnostics.

The particle-in-cell codes that were employed in this thesis, PIConGPU[52] and PICLS[54, 182], both model field ionization. In the former, a combination of models can be chained together to cover the tunneling and the barrier-suppression regime. The early Keldysh model (Eq. 2.12 , the ADK model as well as BSI variants are available for modular use. In the latter, just the ADK model is implemented.

Particle Collisions

Collisional ionization in plasmas occurs mainly due to electron-ion impacts facilitated by *ionization potential depression* (IPD)[183], the lowering of the energy gap between bound and free states due to the overlapping of atomic potentials under high pressures. The *Thomas-Fermi* (TF) model[90] is one of the simplest approaches as it predicts an average charge state $\langle Z \rangle$ that depends on the atomic species, its ionization potentials, the ion density n_i and the electron temperature T_e. As such, it assumes a local thermal equilibrium (LTE) which is typically not the predominant state for a strongly laser-heated solid target. To improve its applicability for relativistic laser interactions, the typical $2\omega_\mathrm{L}$-electron-bunches are typically excluded via an upper energy cutoff[182]. These electrons could otherwise skew the local temperature calculation since they exhibit a broad high-energy spectrum but extremely low interaction probability as they propagate ballistically through the target with close to light speed. The TF model is implemented in both PICLS and PIConGPU, the latter of which was done by the author of this thesis. Further empirical cutoffs, specifically tailored to physical setups, have been introduced in this work and will be explained in Sec. 4.1.2. In the Thomas-Fermi model, atoms of a single or averaged species are treated as point-like and structure-less, surrounded by an electron gas. Thereby neglecting a more complex atomic structure and atomic excitation effects, its low computational complexity on the other hand allows for use in large-scale PIC simulations while providing good accuracy when compared to external equation-of-state databases[94] or atomic physics tools like SCFLY and FLYCHK[184, 185] (for a comparison, see also Sec. 4.2.3).

Another LTE model to calculate the relative charge state populations due to impact ionization is the Saha ionization equation that is based on the Saha-Boltzmann equation[186]. It is reportedly more computationally expensive than the TF model as it uses iteration cycles to determine the final ionization distributions while using the same input parameters as the Thomas-Fermi model and both models being in reasonably well agreement with each other[94].

In a computationally more demanding approach, the collision frequency between electrons and ions is calculated inside each simulation cell which constitutes to a certain number of collision processes per time step. In the *binary collision model*, an according number of electrons and ions is then chosen randomly to perform these collisions[187]. The ion charge states are increased and the electrons lose energy equivalent to the ionization potential. Mishra et al. [94] introduced performance improvements by randomly ionizing atoms in a cell as long as the local kinetic electron energy is still larger than the required ionization energy. Electron screening of the nuclei in partially ionized plasmas was also considered.

A binary collision approach can also be chosen to model atomic excitations within the plasma as was recently shown for PIConGPU and published in a master's thesis by Marré [188]. Keeping track of every macroparticle's atomic state, even in a reduced form, is however extremely memory intensive and the current implementation does not yet allow for the modeling full-sized laser-ion acceleration setups.

The initialization of cold solid-density plasmas can have numerical consequences due to the mesh not resolving the local Debye-length ($T_\mathrm{e} \to 0 \Rightarrow \lambda_\mathrm{D} \to 0$, see Eq. 2.8) that lead to artificial heating[189] or unwanted cold plasma instabilities [167]. However, a high resolution and high number of particles per cell, randomized particle initialization, field ionization and collisional modeling can quickly reduce this behavior drastically.

Bremsstrahlung

A selection of the simulations presented in Ch. 4.2 were configured to produce bremsstrahlung from the interaction of laser-accelerated electrons with the target bulk material. Bremsstrahlung from laser-driven solid-density targets has recently received increased attention as

it provides a diagnostic capability unaffected by the strong electromagnetic fields in the vicinity of the target[G2, 190–192]. Especially for ultra-relativistic electrons, emitting X-ray photons as a consequence of scattering off atomic potentials can be the dominant process of energy-loss in matter[193]. Thus, bremsstrahlung diagnostics are a promising avenue for immediate access to the electron energy– and target density distribution during the short phase (mostly the the temporal FWHM) of highest laser intensity when such electrons are produced.

PIConGPU supports dynamic photon creation for bremsstrahlung and synchrotron radiation with classical and QED models. The implementation is described in full detail, including tests and benchmarks, in Burau [194] and is based on the *Furry*-picture[195] and its application for particle-in-cell by Gonoskov et al. [196]. Here, the main assumptions and ideas are summarized from these works to better judge the results in the later parts of this thesis. The highly-parallel photon creation was adapted from the implementation described in Sec. 3.1.1 and in [G5]. Therefore, photons in PIConGPU propagate as PIC-macro-particles representing multiple real photons. This approach sets up the possibility for future in-situ modeling of radiation transport. The average number of interactions of an electron in the simulation is given by $\overline{N}_{\text{int}} = n_{\text{ion}} \cdot v_{\text{el}} \Delta t_{\text{PIC}} \cdot \sigma$. To avoid the differential cross-section $\mathrm{d}\sigma/\mathrm{d}\omega$ from diverging for $\omega \to 0$, a soft-photon cutoff W_{soft} separates low-energy from high-energy photons. Only the latter are explicitly modeled while the former result in an average energy reduction for the electron. The angular emission characteristic follows

$$\frac{\partial^2 \sigma}{\partial W \partial \Omega_{\text{ph}}} = \frac{\mathrm{d}\sigma}{\mathrm{d}W} \frac{1}{2\pi} p(\mathcal{E}_{\text{kin}}, W, \cos\theta) \tag{2.45}$$

where W is the energy above the soft photon cut-off W_{soft}, p is the probability density which depends on the azimuthal part θ of the solid angle Ω_{ph}. The full expression for the angularly integrated Bethe-Heitler cross section $\mathrm{d}\sigma/\mathrm{d}W$ of an electron in the vicinity of a Wentzel-model[197] atom can be found in [194]. The polar angle φ is determined randomly under the assumption that the nuclear potentials have no multipole moments. In the ultra-relativistic approximation, the dipole model for azimuthal emission reads

$$p(\mathcal{E}_{\text{kin}}, W, \cos\theta) \approx p(\gamma, \theta) = \frac{3\gamma^2}{2\pi} \cdot \frac{1 + \gamma^4\theta^4}{(1 + \gamma^2\theta^2)^4}, \tag{2.46}$$

showing that the signal is strongly peaked into the direction of electron motion. This part is precalculated on the CPU-(host-)side for given combinations of electron energy and soft photon cut-off ($\mathcal{E}_{\text{kin}}, W_{\text{soft}}$), written into lookup-tables and later bilinearly interpolated between during the kernel execution. The electron recoil is directed against its propagation direction, not the emitted photon's to keep the numerical effort manageable. It is modeled as elastic scattering in the Rutherford picture as[193]

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega_{\text{e}}} = \left(\frac{e^2}{4\pi\varepsilon_0} \frac{Z}{4\mathcal{E}_{\text{kin}}}\right)^2 \frac{1}{\sin^4\left(\frac{\theta}{2}\right)} \approx \left(\frac{2e^2}{4\pi\varepsilon_0} \frac{Z}{\mathcal{E}_{\text{kin}}}\right)^2 \frac{1}{\theta^4}, \qquad \theta \geq \theta_{\text{min}}. \tag{2.47}$$

To achieve a compromise between resources and statistics, there is a weighting ratio between macro electrons and macro photons that allows to customize the number of emitted photons for the specific simulation setup.

2.3.3. Influence of the Dimensionality

In the early days of numerical treatment of plasmas with particle-in-cell methods, the available computational power did usually not suffice to simulate full systems in multiple dimensions and at full density. However, many problem setups, like plasma expansion from a planar

surface, exhibit enough symmetry for them to be reduced to a single dimension with good accuracy. Even today, where computational resources are exponentially more powerful, low-dimensional simulations are still an important tool to inexpensively observe qualitative trends for the existence of optimum configurations in a wide parameter space.

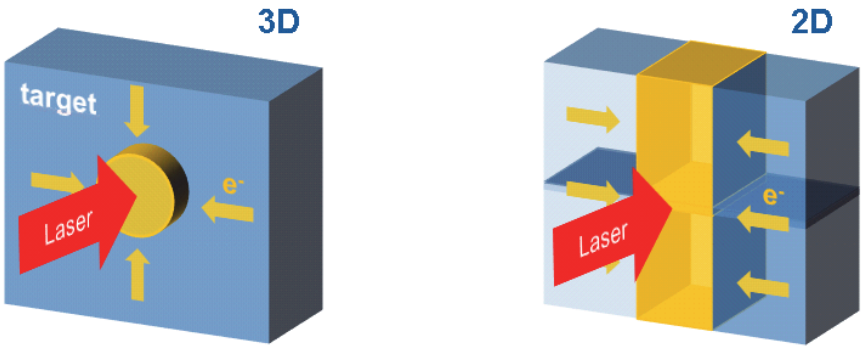

Figure 2.9.: Geometric consequences from 3D and 2D simulations with respect to laser spot and electron currents that compensate the removal of particles from the foil by the laser.

2D vs. 3D Preparatory simulations in 2D are a valuable tool to gain insight about how the plasma dynamics are likely going to be in 3D. However, 2D simulations do not make good quantitative predictors due to an effectively different geometry that results from the periodic boundary in the missing dimension. As Fig. 2.9 illustrates, a laser focus of circular shape becomes the equivalent of a line focus in 2D. For the same reason, the electron sheath that will evolve behind the target is closer to cylindrical than spherical (or Gaussian) symmetry. Thus, the charge separation fields and electric potentials Φ scale with $\vec{E} \propto r^{-1}$ and $\Phi \propto \ln r$ in 2D, instead of $\vec{E} \propto r^{-2}$ and $\Phi \propto r^{-1}$. For the acceleration of ions in these fields this has two consequences. Ions are accelerated more strongly over longer distances and the time to saturation of the kinetic energies is also longer.

Babaei et al. [198] (2017) quantified this effect with an empirical law based on the model by Schreiber et al. [44] that they benchmarked against 2D and 3D PIC simulations with ALaDyn[199]. As the two main parameters, the scaling time t^* and the asymptotic maximum proton energy $\mathcal{E}_\infty$ are identified. These take different values between 2D and 3D and the temporal evolution of the proton energy is defined by them as follows:

$$\mathcal{E}_{\max}(t) = \begin{cases} \mathcal{E}_\infty \log \frac{ct}{ct^*}, & \text{(2D)} \\ \mathcal{E}_\infty \left(1 - \frac{ct^*}{ct}\right)^2, & \text{(3D)} \end{cases} \tag{2.48}$$

The scaling laws given in Eq. 2.48 are valid for times $t \geq t^*$. The 2D part, however, does not converge and the asymptotic maximum energy still remains difficult to determine. In reality, the acceleration process in both 2D and 3D simulations is cut short or altered by particles leaving the simulation box, be it the electrons that provide the accelerating field or the ions themselves. From case to case, the necessary box size needs to be carefully configured to avoid pre-mature termination of the acceleration before any trends have solidified.

3. Methods and Development for Highly Scalable Plasma Simulations with PIConGPU

For the main part of this thesis, the author performed a large-scale simulation campaign investigating the influence of the laser intensity ramp on the ion acceleration performance (results in Secs. 4.2, 4.3). This campaign was granted computation time in the 15th PRACE Call for Proposals[1] as the 2nd largest allocation with 109 million core hours ($1.6 \cdot 10^6$ GPU node hours) on Piz Daint[201], hosted by CSCS[202][2], Switzerland. In this chapter, the main "work horse" for this project, the fully 3D, open-source particle-in-cell code PIConGPU will be introduced. Feature characteristics will be briefly explained, in particular what makes the code scalable to the world's largest supercomputers with a focus on how dynamic particle creation, for e.g. ionization and bremsstrahlung processes, is handled to maintain this efficiency. Following this, the typical conundrum is presented that a computational plasma scientist is faced with when weighing the choice of simulation parameters against the available resources. A real-world-example, based on a common experimental laser-ion acceleration scenario will be used to show how limited even nowadays most powerful supercomputers still are in terms of the capacity to resolve the fundamental solid-density laser-plasma dynamics in a meaningfully large spatial volume. This example is the basis for the results presented in Sec. 4.2 where the adjustments and compromises are briefly discussed that allowed the scenario to fit the compute resources on Piz Daint. Running explorative large-scale PIC simulations on a supercomputer implies creating raw big data. Therefore, section 3.3 focuses on the management and life-cycle of PetaByte-scale data as it was produced in the campaign described Sec. 4.2. Ensuing the explanation of the methods used for the creation, the transporting and the later archiving of scientific data, section 3.4 will be dedicated to big data analysis and reduction.

3.1. PIConGPU

PIConGPU is a fully-relativistic, 3D3V particle-in-cell code developed for many-core architectures[15, 52]. It is mainly developed and maintained by the *Computational Radiation Physics* group at HZDR in a complete open-source and open-access fashion on GitHub[204]. Since

[1] `https://prace-ri.eu/call/call15-for-proposals-for-project-access/` 200.
[2] At the time ranked 3rd highest in the Top500 list of supercomputers.[203]

it was published under a GPLv3+ open-source license, the philosophy has always been to transparently develop the code in a collaborative way with regular DOI-assigned releases. The authors of PIConGPU believe that only such practice serves the open-science paradigm sustainably and predictive capabilities can truly be achieved or improved (For more details on this, see also [15, 205]). When the non-linear plasma response of laser-particle accelerators is simulated with particle-in-cell codes, the detailed dynamics are subject to change not only due to the variation of physical input parameters. Also the choice of numerical solvers and algorithms, the specific additional physics models (ionization, etc.), simulation dimensionality (see 2.3.3) and most importantly, the sampling resolution of every physical process can all either influence the quantitative prediction or even change the qualitative behavior and physical interpretation of the virtual experiment. In the past, PIC simulations could take weeks to complete and the high computational cost forced plasma physicists to put their trust in rather small sets of runs. PIConGPU focuses on achieving reliable predictive capabilities by providing fast, highly scalable simulations at high resolution and full geometry. A time-to-solution of days to only hours was made possible with the use of graphics processing units (GPUs) for scientific computation.

Tightly coupled online diagnostics and post-processing modules Many forms of analysis that can be performed on the particle– and field data of PIC codes are standard procedure, regardless of the physical scenario. Unfortunately, the gap between compute performance and input/output (I/O) performance of modern supercomputers is also perpetually growing. Additionally, there exists a memory bandwidth disparity between where data is generated (on CPUs & GPUs), local fast memory RAM and the disk to which the data is finally stored, amounting to a total bandwidth loss of currently about $1 \cdot 10^4$. The larger the data becomes, the more tedious the analysis is in post-processing and, increasingly, the necessary operations cannot be performed on a single notebook or cluster compute node anymore, because the datasets are simply too large for the memory of a single machine. It is therefore reasonable and feasible to do as many as possible of the foreseeable computations directly coupled to the simulation while the relevant data is still in memory. TeraByte-scale data can be reduced to physically meaningful output of energy histograms, phase space arrays and slice images of only kiloByte or MegaByte size, achieving data reduction factors of $1 \cdot 10^6$ to $1 \cdot 10^9$. A full list of all plugins available in PIConGPU can be found in the respective chapter of its online documentation[206]. PIConGPU also ships with several Python modules to quickly analyze and plot its diagnostic output or judge the computational resources for a new simulation setup.

Programming model An extensive description of the programming model employed in PIConGPU can be found in the dissertation of Huebl [15], chapter 3. PIConGPU is based on C++14 (as of 2020) and modern programming paradigms like the zero-overhead abstraction technique of *template meta-programming*[207–209]. As such, PIConGPU follows a single-source approach and therefore retains full functionality on all different architectures and under all physics scenarios without loss of features due to a change of platforms or simulated physics. The code base is made performance-portable by the `alpaka`[210–212] library which abstracts the underlying levels of parallelism of each supported compute architecture (x86, ARM, NVIDIA GPU, ...) to fit PIConGPU's kernels to each accelerator backend. Of the implemented backends, CUDA is the most prominent as it is the parallel computing platform for NVIDIA GPUs. Despite tendencies towards increasing numbers of AMD graphics processors in supercomputers, most notably the next generation exascale machine Frontier[213] (projected

peak performance > 1.5 ExaFlops), the HPC market is still mostly populated by NVIDIA.[3].

PIConGPU has already shown outstanding performance and scalability in 2013 and 2015 on Titan, the formerly most powerful supercomputer in the world, which also led to a nomination as a Gordon Bell[4] finalist. In 2016 and again 2017/18 with the work of this thesis, the code was again testing the limits of the most powerful supercomputer with GPU accelerators, Piz Daint[214]. The next generation promises to deliver the first *ExaScale* machines and as part of the Frontier Center for Accelerated Application Readiness (Frontier CAAR) project[215], PIConGPU is currently being prepared for Frontier[216, 217], the fully AMD-powered upcoming supercomputer at OLCF. To achieve this, the code, via the `alpaka` library, has adopted an additional backend to enable programming for AMD GPUs, called HIP. PIConGPU, by now, has

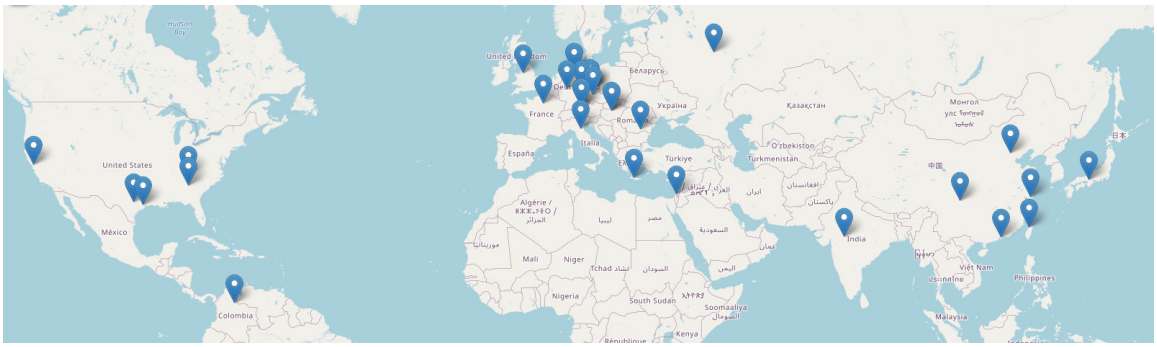

Figure 3.1.: PIConGPU community map [218] displaying contributors, collaborators and users worldwide. Map tiles ©OpenStreetMap contributors under CC BY-SA license.

users, developers, contributors and collaborators in many groups all over the world as the Community Map in Fig. 3.1 shows. A wide range of research subjects has been studied in the past, ranging from laser-electron and -ion acceleration, laboratory HED plasmas, astrophysical plasmas and radiation signatures as well as secondary radiation sources. Two workshops for future users and developers among theoreticians and experimental physicists alike were held already at HZDR. The second of these workshops was organized in February 2019 by the author and was held under the banner of the WHELMI collaboration[219] between HZDR and the Weizmann Institute of Science in Rehovot, Israel.

3.1.1. Compute Performance and Consequences

The defining property of PIConGPU that sets it apart from other particle-in-cell codes is its unparalleled speed and scalability to the largest supercomputers on earth. This is enabled by the hardware-aware design of the underlying data structures and compute kernels.

Particle species Particles in PIConGPU are created from a generic species description. In this concept, two different categories of particle properties are distinguished. Particle *attributes* (e.g. position, momentum, weighting) are objects that can change at runtime and are stored for each PIC particle. Particle *flags* are immutable and define the species as a whole in terms of the algorithms and physical processes that these particles take part in during a PIC cycle. Each particle species can be tailored to the specific needs of the the simulation and is extensible by combined observables (for instance the product of velocity and electric field at the particle

[3] a combined 89 % of the current Top500, according to the June 2020 systems share comparison of *Accelerator / CP Family* at `https://www.top500.org/statistics/overtime/`

[4] The Gordon Bell prize is a awarded each year in the wake of the international Supercomputing Conference (SC) for outstanding achievements in high-performance computing.

position $\vec{v} \cdot \vec{E}$) that help study the relevant dynamics in more detail. In particular the results described in section 4.1.1 were obtained by heavily relying on the adaptable species concept. There, it was employed to determine at which position and by which ionization mechanism the electrons in a laser-wakefield accelerated (LWFA) electron bunch were created. Thus, PIC particles in PIConGPU are not only markers in a 6D phase space distribution but can also be markers for different physical processes to increase the insight into femtosecond-submicron plasma dynamics. In order to execute algorithms efficiently on the particles and take advan-

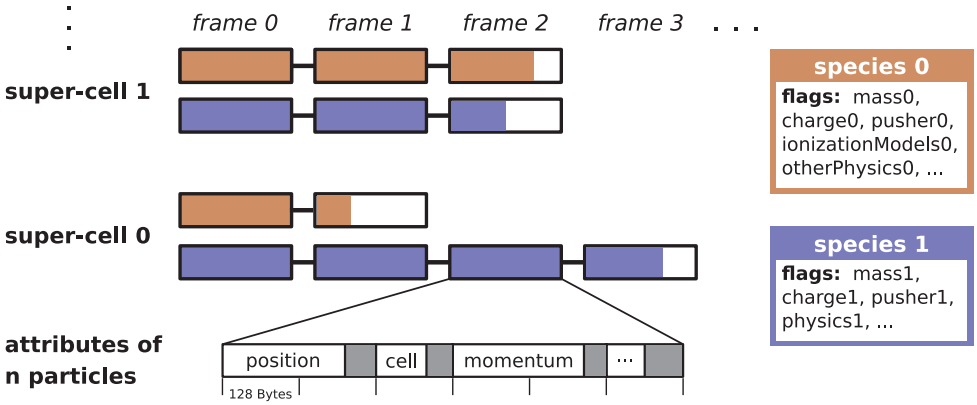

Figure 3.2.: Frame lists for adaptable particle species in PIConGPU. Per super-cell in the simulation volume, the individual particle data of each species is represented by a list of *frames* (i.e. the storage structure). Different colors represent different generic particle species. Each frame is of constant size and they are kept contiguously filled, except the last one (per species and supercell) whose fill-level is indicated by the partial coloring[5]. Image adapted from [205].

tage of the parallel thread execution in modern accelerator hardware, PIConGPU organizes particles in so-called *frame lists*. As figure 3.2 shows, particle attribute lists are stored inside a frame in what is called a *struct-of-arrays*. However, frame lists in PIConGPU do not reflect particles per cell but rather particles per *supercell*. A fixed number of cells are grouped together for optimized SIMD[6] processing and caching of field data. As such, per particle species, a list of particle frames is assigned to each supercell. The number of cells in a supercell and also particles in a frame is chosen to reflect the number of threads in a *thread block* that operates in parallel on modern manycore (GPU) hardware. The supercell length and particle frame size are fixed and determined before compiling the code. This additional level of abstraction maximizes the utilization of the architecture's compute performance and avoids global sorting operations.

A very challenging operation for optimized, highly parallel PIC codes, is the creation and removal of particles at run-time. This is especially true for PIConGPU because it requires to carefully manage the available memory, in particular on GPU accelerator architectures (where memory is very limited), to retain the performance advantage that it exhibits over other, CPU-based codes. Both the *ionization* implemented by the author and the *bremsstrahlung* module in PIConGPU require this to happen regularly. Figure 3.3 shows a schematic view of the necessary steps at the kernel[7] level on the example of ionization. Here, different colors describe different types of operation regarding the memory access. The displayed flow of operation is executed by each thread (in a thread block) in parallel. First, super-cell-local electromagnetic field information is cached and the kernel begins to iterate sequentially through the list of ion frames. Each thread then calculates the number of electrons to create for an ion in this frame

[5]The padding in figure 3.2 which is used to align all data to 128 B blocks was recently changed to a smaller size through automatic alignment performed by the compiler. This leads to a very small performance increase.
[6]Single Instruction Multiple Data
[7]A device function that is executed across a set of parallel threads is called *kernel*.

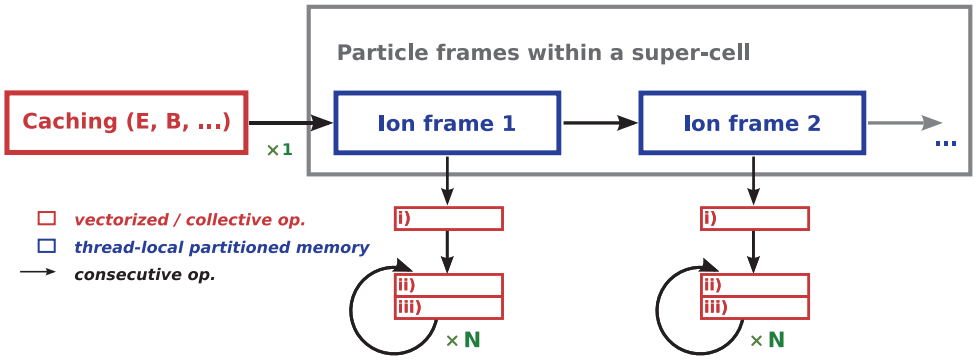

Figure 3.3.: Local field data required for ionization is cached. The kernel iterates sequentially over the list of ion frames and executes the ionization model prediction and following particle creation in parallel (see Fig. 3.4 below) with vectorized SIMD-like operations for each *frame*.

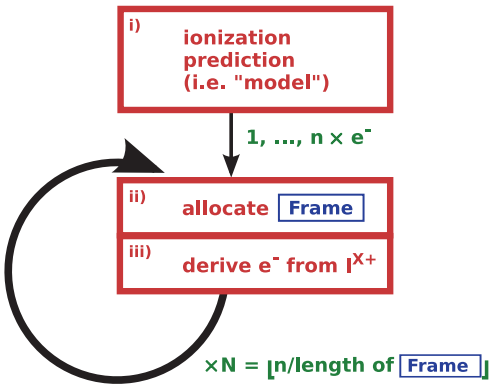

Figure 3.4.: A *block* of *threads* calculates the prediction for the number of ionization processes in parallel with the local cached field data using the configured ionization model (i). If there is a non-zero number *n* of electrons to create, a chunk of thread-local partitioned memory, i.e. an empty frame, is allocated. The *threads* derive free electrons from the ions that are ionized and fill the empty frame in parallel (iii). The process is then repeated *N* times as long as new electrons still need to be created.

according to the ionization model and local field data. Figure 3.4 shows this process in detail. The predicted electron number is accumulated and, if non-zero, a new empty electron frame is created by the coordinating *Primary* thread. Threads *derive*[8] new electrons from their parent ions until the electron frame is completely filled and an additional frame has to be allocated or until no electrons need to be spawned anymore within this time step. Upon completion of all electron creation operations the kernel moves on towards the ionization prediction of the next ion frame (compare Fig. 3.3). In extreme cases, as many as 10 million new particles per time step may need to be created per GPU, requiring fast, well-coordinated memory allocation. The example of a solid-density laser-matter interaction scenario later in Sec. 3.2.1 shows that GPU memory is the most limiting resource for large-scale PIC simulations.[9] Only with dynamic allocation and freeing of memory, with the memory manager `mallocMC`[220], which extends and adapts the concepts of the parallel memory allocator `ScatterAlloc`[221], can the available memory be efficiently used, thus avoiding sacrifices in code performance. Less obvious is the fact that if a particle enters (or is created in) a supercell that previously did not contain any particles of that species, the memory for a full new frame (usually for 256 particles) is al-

[8] Electrons are given the same position, velocity (if not otherwise specified) and weighting as the parent ion.
[9] Host-sided RAM is more abundant than the faster memory of GPUs but using it for this purpose would severely reduce the performance of the code, eliminating almost all benefits from calculating on the accelerated hardware.

located. In physical scenarios like Laser-Wakefield Acceleration (LWFA, see Sec. 2.1.2), where the full simulation volume is often homogeneously filled with particles from the very beginning, the consequences are barely noticeable. However, in laser-ion acceleration simulations, which are prone to high load imbalances due to the geometry of localized foil targets, this side-effect of PIConGPU's efficient data structures can lead to unforeseeable crashes of the simulation even though on average the GPU memory load is far from its limit. Particles are initially highly concentrated to a small region but upon laser irradiation many new particles are produced via ionization processes. When the resulting plasma is expanding, further particle concentration occurs in a ring-like structure (see Fig. 3.5) around the laser beam axis due to the ponderomotive force (given that $a_0 > 0$). Both of these effects can overload the memory capacity of a single

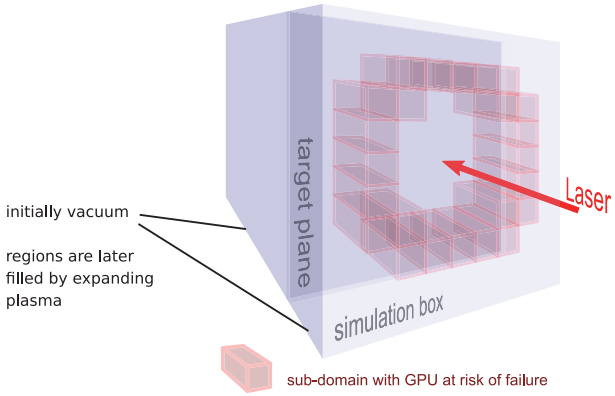

Figure 3.5.: Schematic of GPU subdomains at risk of failure during plasma expansion. The laser propagates in positive y-direction, displaces particles from the center of the laser spot on target and causes accumulation in a ring or ellipse around. Additionally, the expanding electron and ion clouds also possess a divergence, further populating these sub-domains. This has a high risk of overloading the local subdomain of a GPU causing the failure of the whole simulation.

GPU, causing the crash of the whole simulation. This behavior is very challenging to predict and with the implemented static load balancing methods, it can only be averted by reducing the sub-domain size around the initial target. This inevitably means that less overall volume can be computed, the grid resolution has to be lowered or when other GPUs are compensating by calculating larger regions, they become more prone to crashes if the plasma behaves differently than expected. Modern dynamic load-balancing methods exist[222–225] but PIConGPU currently tries to minimize communication overhead by maximizing its subdomain volumes (i.e. a cuboid for each GPU) and benefits strongly from its fixed buffer structures without the requirement of reallocation. Implementing dynamic load balancing requires deep infrastructural changes since each accelerator's subdomain would be subdivided further to account for the local load imbalances. Changing to dynamic buffer structures and keeping track of all the communication is a very challenging task within PIConGPU for which dedicated development is necessary and planned for the future. A more detailed discussion on the consequences of the current implementation on the calculation of numerical resolution and simulation size of the scenarios used in this thesis can be found in section 3.2.2.

3.2. Planning a Large-Scale PIC Simulation

For particle-in-cell simulations to accurately portray the detailed kinetics of charged particles under the influence of high-intensity laser fields, a certain set of fundamental conditions have to be fulfilled. In section 2.3, these conditions were physically motivated and introduced. This section shows how these rules are employed in practice for a real-life laser-ion acceleration

scenario. Using one of the simulation setups from the main study of this thesis (see Ch. 4.2) the staggering computational expenses of simulating real-world-scenarios in full-resolution 3D are calculated. Since the necessary resources for such simulations are not commonly available, solid-density PIC simulations are often performed in 2D. A brief discussion highlights the implicit assumptions and possible effects that reducing the dimensionality can have on the physical outcomes. The exploratory nature of the study described in section 4.2 of the next chapter required as few initial assumptions as possible and gaining high-quality 3D data was a valuable goal in itself for this regime where literature is still scarce and contradictory (compare 2.2.2). Finally, the trade-offs between available computing resources, physical resolution and limitation of the simulation volume are discussed in Sec. 3.2.2.

3.2.1. Real-world Example

The main contribution of this thesis is concerned with answering the question how the intensity ramp of the optical drive laser on the last picosecond prior to the pulse peak shapes the ensuing acceleration of ions from ultrathin solid-density foil targets. The most extensively simulated setup contained copper foils of thicknesses ranging from 10 to 300 nm. First, the necessary resolution on a cell- and time step-level needs to be calculated. At solid density and room temperature, copper has a mass density of $\rho_{\mathrm{Cu}} = 8.96\,\mathrm{g\,cm^{-3}}$. Fully ionized, this corresponds to $n_{\mathrm{e},0} = Z_{\mathrm{Cu}} N_{\mathrm{A}} \rho_{\mathrm{Cu}} / M_{\mathrm{Cu}} = 1414\, n_{\mathrm{c}}$. where, $Z_{\mathrm{Cu}} = 29$, $M_{\mathrm{Cu}} = 63.546\,\mathrm{g\,mol^{-1}}$ are the atomic number and molar mass of copper, and $n_{\mathrm{c}} = 1.74 \cdot 10^{21}\,\mathrm{cm^{-3}}$ is the critical electron density for laser light of $\lambda_{\mathrm{L}} = 800\,\mathrm{nm}$ wavelength, respectively. While the oscillation of the laser wave has to be resolved in underdense setups, this is usually implicit for optical drive lasers in overdense setups where the electron plasma wave determines the numerical resolution. Using Eq. 2.44 and Eq. 2.6 the simulation time step ideally measures

$$\begin{aligned} \Delta t_{\mathrm{PIC}} \leq \frac{0.1}{\omega_{\mathrm{pe}}} &= \frac{0.1}{2\pi} T_{\mathrm{pe}} = \frac{0.1}{2\pi} T_{\mathrm{L}} \sqrt{\frac{n_{\mathrm{c}}}{n_{\mathrm{e}}}} \\ &= \frac{0.1}{2\pi} \cdot 2.67\,\mathrm{fs} \cdot \sqrt{\frac{1}{1414}} \\ &= \underline{1.13\,\mathrm{as}}. \end{aligned} \tag{3.1}$$

Here, T_{pe} and T_{pe} are the periods of the laser– and the electron plasma wave, respectively. As explained in Sec. 2.3, temporal and spatial resolution are coupled via the CFL criterion. For a 3D3V solid-density PIC simulation with a Yee scheme Maxwell solver, cubic cells are employed. Following Eq. 2.42, they ideally measure $\Delta x_i \gtrapprox \sqrt{3} c_0 \Delta t_{\mathrm{PIC}} = 5.87\,\text{Å}$. In fact, with attosecond-Ångström grid resolution, the simulation is also already pushing far into atomic time- and length scales, leaving the validity of the often LTE-based or *quasi-static* atomic physics modules implemented in particle-in-cell codes questionable[G6]. The simulation box needs to comfortably fit the laser pulse spot size as well as have enough vacuum in front and especially behind the foil target to allow for the accelerated ions' energies to stabilize. If the first condition was not satisfied, significant effects from laser radiation reflecting off the absorbing transverse boundaries could turn the whole simulation dynamics unphysical. So-called *Perfectly Matching Layer* (PML) boundary conditions[226, 227] can reduce unwanted boundary effects significantly and thereby relax the transverse simulation size requirements. While such PMLs were not yet implemented during the time of the simulation campaigns described in this thesis, PIConGPU's most recent release 0.5.0[G7] now offers this feature.

For laser-ion acceleration with a short-pulse laser, the lowest estimate on the acceleration time is at least the laser pulse duration τ_L but well adapted is the empirically obtained value of $1.33\,\tau_{\mathrm{L}}$ by Fuchs et al. [106] (intensity FWHM). Assuming that the main acceleration starts with

the peak of the laser pulse on target when the largest electric fields are at the center of the laser spot, it is extremely undesirable that fields reflected off this point travel to the boundaries and back within this acceleration time. In the campaign on Piz Daint, 30 fs laser pulses with a_0 values ranging between 20 and 63 are simulated. The main pulse length is 10 µm and so at least 5 µm outside of the laser spot should be added to each side in transverse direction. Besides the reflection, the normalized amplitude a_0 is still equal to unity at 3.67 µm away from the laser axis if the spot size is $w_0 = 3$ µm. Ideally, the particle dynamics should be fairly weak in the border region to avoid too much artificial transition radiation from particles running into absorbing layers. As a result, it is best to make the transverse size large enough such that $a_0 < 1$ is reached.

The full length of the laser pulse does not need to fit into the simulation box, since the pre-vacuum can be stacked with more absorbing layers to smoothly compensate the back-reflection. If the rear-side vacuum region does not extend far enough, however, the acceleration process is cut short by a significant part, possibly resulting into misleading trends of the actual acceleration performance. The above conditions are comfortably met by a box of, e.g. $(15 \times 30 \times 15)$ µm with $25000 \times 50000 \times 25000$ cells. To model laser-target interaction over the last picosecond preceding the laser-pulse maximum as well as up to 300 fs afterward, 1.15 million steps are required.

During this thesis, the author contributed a Python `memory_calculator` utility to the packages shipped with PIConGPU to help determine the necessary computing resources[10]. Estimating the necessary total memory required for the fields, $(2 \times 3 + 1 \times 3)$ scalar values need to be reserved for the components of $\vec{E}$, $\vec{B}$ and $\vec{j}$ (or $\Delta\rho$), respectively. Furthermore, if there is any field data output desired, the memory for a temporarily stored field of one scalar value per cell has to be reserved. Another temporary field is necessary to compute the collisional ionization prediction within the Thomas-Fermi model by having electron energy density and ion density available at the same time (see 2.3.2). With 32 bit precision per scalar value, the fields alone require 1377 TB of memory. Not only fields but also particles have to be stored in the directly accessible memory (RAM / VRAM) of the computing devices. A representative of a basic particle species in PIConGPU requires 7 scalar values for $\vec{x}$, $\vec{p}$ and w, the position, momentum and particle weighting, respectively. In the simplest case, electrons and photons can be expressed with just these, while ions need additional charge state information that is stored in another scalar value. In the simulation campaign, copper foils ranging from nanometer to micron scale thickness are studied. Taking $d_{\mathrm{Cu}} = 1$ µm as the most extreme example – all thinner foils should then easily fit the resources – and distributing only one copper ion alongside its 29 electrons per cell, the target would be represented by $3.2 \cdot 10^{14}$ particle markers, accounting for 994 TB of fast memory.

A very important feature of the study of the last picosecond intensity ramps was the inclusion of dynamic ionization of the target. The probabilistic nature of the ionization processes requires that also the state of a random number generator (RNG) is stored at all times. With an initial seed and the RNG state, a new random sample can be generated and subsequently the state is changed. In PIConGPU, there are several available random number generators implemented and the one with the smallest state size, `XorMin`, requires 6×4 bytes per simulation cell, adding 750 TB to the memory requirement. In total, the required memory amounts to 3121 TB.

As of June 2017, the five most powerful supercomputing systems were hosted by China, Switzerland and the United States[228]. Table 3.1 shows these systems and their rating based on the High Performance Linpack (HPL) benchmark[229]. Notably, the Top500 list purely

[10] A usage example can be found in the online code documentation at `https://picongpu.readthedocs.io/en/0.5.0/usage/workflows/memoryPerDevice.html#usage-workflows-memoryperdevice`

	NAME	SITE	COUNTRY	CORES	R_{MAX} (PFLOPS/S)	POWER (MW)
1	**Sunway TaihuLight** Shenwei SW26010 (260C 1.45 GHz) Custom Interconnect)	NSCC in Wuxi	China	10,649,600	**93.0**	15.4
2	**Tianhe-2 (Milkyway-2)** Intel Ivy Bridge (12C 2.2 GHz) & Xeon Phi (57C 1.1 GHz), Custom interconnect	NSCC in Guangzhou	China	3,120,000	**33.9**	17.8
3	**Piz Daint** Cray XC50, Xeon E5-2690v3 (12C 2.6GHz), Aries interconnect, NVIDIA Tesla P100	CSCS	Switzerland	361,760	**19.6**	2.27
4	**Titan** Cray XK7, Opteron 6274 (16C 2.2 GHz) + Nvidia Kepler GPU, Custom interconnect	DOE/SC/ORNL	USA	560,640	**17.6**	8.2
5	**Sequoia** IBM BlueGene/Q, Power BQC (16C 1.60 GHz), Custom interconnect	DOE/NNSA/LLNL	USA	1,572,864	**17.6**	7.9

Table 3.1.: World's five most powerful high-performance computing systems as of the Top500 list in June 2017[228]. The first publicly available system is the Swiss hosted Piz Daint supercomputer at CSCS where the large-scale simulation campaign described in 4.2 was performed.

shows a ranking according to the maximum compute performance R_{max} which does not reveal anything about the available memory. Particle-in-cell codes are mainly memory transfer-bound, however, and while landing in the top 10 fastest supercomputers is a tremendous incentive for investments, increasing the available memory accordingly is not. The most powerful publicly available system at the time was Piz Daint at CSCS with 5704 Cray XC50 compute nodes. Each of these nodes has 64 GB of RAM and contains an NVIDIA Tesla P100 GPU with 16 GB of VRAM. To distribute the simulation setup described above to these resources would require about 200 000 GPUs or 50 000 CPU nodes. While PIConGPU can technically run on both hardware systems, the required resources to barely fit the problem exceed the available capacity easily by 1 to 1.5 orders of magnitude.

In reality, the mismatch is even more severe due to the planar geometry of the target. Most particles are initially concentrated on a comparably thin region with respect to the vacuum around it. In the 1 µm copper foil example, the target would only be contained by 1700 of the 50 000 cells in longitudinal (y in PIConGPU) direction. The extremely high spatiotemporal resolution is only needed for the most dense regions of the simulated plasma and could theoretically be reduced in the outer regions. Approaches to better distribute available compute resources to the most compute-heavy regions, combined with adaptive mesh refinement (AMR)[230, 231] exist, but are increasingly challenging to implement for GPUs and come with their own numerical caveats. These methods of *dynamic load balancing* are interesting and worth to be studied as they have the potential to further exceed the capabilities of PIConGPU where static load balancing is implemented. Unfortunately, the very concept that makes PIConGPU scale so well on any hardware with `alpaka`[211], i.e. the execution hierarchy of *grid, block, thread, and element* that distribute the work– and memory load to architecture-specific chunks, as well as the many compile-time optimizations of the code, are antagonal to the idea of dynamic, heterogenous load balancing. Neighboring subdomains of the simulation volume are, furthermore, close together in memory to reduce the communication and boundary-treatment overhead. A major overhaul of the current memory management would very likely be necessary to evolve past static load-balancing[232] in PIConGPU. Within the community, dynamic load balancing for GPU-accelerated PIC codes is being actively researched and progress is being made[233].

The estimate for particle and field memory from before left out another factor that increases the overall memory requirements, i.e. the resources required for communication between

nodes or GPUs. It increases with the overall bordering surface between the sub-domains contained by the compute units. At least one supercell layer is reserved for the *guard region* (sometimes *ghost region*) around each sub-domain which stores duplicates of the particle- and field data of the neighboring sub-domains. All things considered, the currently most powerful supercomputers of the *Petascale* era still do not possess the fast memory capacity to handle laser-solid PIC simulations of micron-sized copper (or similar density) metals in full-size, full 3D setups. The capacities to deal with these setups without further restrictions or simplifications will likely be available in the future generations of top 5 supercomputers belonging to the *Exascale* era.

Nevertheless, nowadays compute facilities can still be used to perform novel, physically reasonable 3D PIC simulations of solid-density laser-plasma interactions. In the following it is described by which additional considerations this was made possible, eventually enabling the simulations described in chapter 4.2.

3.2.2. Simulation Volume vs Resolution vs Available Resources

How can the (always) limited computational resources be used to simulate as much of a physical region of interest while sacrificing as little as possible of the spatiotemporal resolution that is vital to capture all the relevant plasma dynamics?

This is a dilemma that researchers in computational plasma physics have always been faced with. Historically, 1D and 2D simulations were employed because computers were much less powerful than nowadays and, on the other hand, some numerical algorithms take a simpler form or are free of numerical artifacts in lower dimensions, which is not necessarily true for their full 3D description[234, 235] (e.g. very prominently *Numerical Cherenkov* radiation[236]). Some of these methods, like fully dispersion-free Maxwell solvers and the directional splitting method[168, 237] for 3D are still an active field of research. Later it became feasible to simulate Laser-Wakefield accelerators in 3D due to their low density and thus, lower resolution requirements. In the case of LWFA it has also been shown that the cylindrical symmetry of the problem can be taken advantage of by employing quasi-3D PIC codes with rotational symmetry, such as `FBPIC`[238], that allow for extremely wide-ranged parameter scans at simultaneously low cost. These codes work remarkably well within their suitable class of setups and only for quantitative predictions of some parameters, such as electron bunch charge, or dynamics like the hosing instability a higher dimensionality is required.

In laser-solid interactions, the underlying plasma dynamics are far less understood and symmetry arguments often do not apply. More and more, full 3D3V simulations show that the plasma evolution and final outcome of a laser-particle acceleration simulation can reach completely different results, when compared to their 2D counterparts. This can be seen in recent works at near-critical plasma density[155, 239], that are less restrictive in terms of resolution while still requiring some of the largest supercomputers. Concerning the simulation campaign of Sec. 4.2, several adjustments to the setup described before in Sec. 3.2.1 were possible that allowed to run fully 3D simulations even with solid copper targets, while still resolving the relevant plasma physics.

An especially strict requirement is the criterion for long-term stability of the second-order linear Maxwell solvers $\omega_p \Delta t_{PIC} \leq 0.1$ (as shown before for Eq. 2.44, this stems from a phase error of π after 100 000 steps). This means, that the smallest occurring wave of the system, i.e. the electron plasma wave is temporally resolved with $2\pi \cdot 10 \approx 63$ sample points. While the heating of the target depends strongly on resolving the electron plasma oscillation, the actual regions in the simulation volume where densities are close to or in excess of $1414\, n_c$ are actually very small. Since the simulated targets are ultra-thin (between 10 and 300 nm), the

plasma waves cannot travel very far. Furthermore, the temporal overlap between these high densities and the locations where protons are accelerated is also not extensive since the extreme electric fields at fully ionized copper plasma edges are quickly removing protons from those regions. Therefore it is reasonable to relax the resolution and, as chosen in Sec. 4.2, only resolve the plasma wavelength spatially by $\Delta x = \lambda_{pe}/8$ which translates to a larger time step according to the CFL criterion. The new cell size is larger by a factor of $2\pi/8\sqrt{3} \approx 4.5$, which, when applied to each dimension, leads to an overall reduction of cells by nearly two orders of magnitude. This is the largest adjustment since it not only reduces the number of cells but also lowers the total number of particles, if the number of particles per cell is kept constant. Before every large-scale 3D simulation campaign, 2D simulations are still a very important tool to estimate the maximum particle energies and determine the temporal evolution of the plasma dynamics. Especially knowing how the plasma evolves and expands can help to prepare for events that threaten the stability of the full-scale simulation, e.g. from local load imbalances (see Fig. 3.5). An overview of the 2D simulation results can be found in section 4.2.4. Another advantage of 2D simulations is that different grid resolutions can be tested before to see if the overall dynamics are robust against their change. However, a full-resolution 2D laser-solid simulation might already exhaust the computing capacities of a university scale cluster. From the full-scale simulation pre-runs it was assessed that reducing the simulation box to $(10 \times 10 \times 10)$ microns is very likely enough to have ions be accelerated up to 90 % of their final energy. It was observed that the transverse box size is still large enough to exclude possible boundary effects of the reflected laser. As will be motivated more extensively in section 4.2.3, reasonable pre-ionization of the copper ions allows to free up additional memory by representing multiple initially free electrons with a single higher-weight macro-particle.

3.3. Management and Sustainability of Data

Large-scale Dataset Creation A large part of the analysis of particle-in-cell simulations involves workflows that are similar for every simulation run. Among these are the creation of particle counts, particle histograms and the continuous checks for total particle and field energies. Such common operations on particle and field data have already been incorporated into the plugin package of PIConGPU (compare Sec. 3.1). Nevertheless, when performing explorative simulations, not all types of analyses can be prepared beforehand and therefore the intuitive perception pertains that: *Large-scale simulations create large-scale data*.

Between a laptop, a university-scale cluster and one of *Top 5* supercomputers on earth (Tab. 3.1), the attention required towards how to perform I/O increases manifoldly. There are several pitfalls to avoid when doing high-performance I/O. On HPC systems, persistent storage space is a shared resource and writing out too much data might result in exceeding the system quota which puts a stop to the work of all the other users of the system, if there is no user quota system in place for the specific partition to automatically prevent this. The *$SCRATCH* file system at CSCS's Piz Daint imposes a limited lifetime on its data, rather than a fixed user quota. It is for fast data production and procession, while their *$PROJECT* partition is intended for longer storage. Equivalently, writing out too many small data files in the range of kB to MB with deeply nested directory structures may cause the index server to overload and can slow down the whole system and I/O for all other running jobs and users. As a best practice on parallel HPC file systems (FS), aggregating the data in only very few file handles and writing larger files is generally recommended. HPC systems are optimized to increase data bandwidth, not read-write operations per second which are always required when new files are created. It is also generally better to write out binary files instead of ASCII files and avoid global re-ordering and synchronization operations.

The PRACE simulation project described in chapter 4.2 was one of the first large-scale projects on the Piz Daint supercomputer at CSCS after it received its major hardware upgrade in November 2016. With this upgrade, it became the number one most powerful GPU-accelerated supercomputer in the world. Closely in collaboration with the technical staff of CSCS, the scalable I/O of PIConGPU was optimized for Piz Daint's Lustre[240, 241] FS, HPC workflows were implemented that were vital for the success of this thesis and will also benefit future HPC projects at CSCS. Unlike the computer at home, the parallel FS of a supercomputer not only manages one

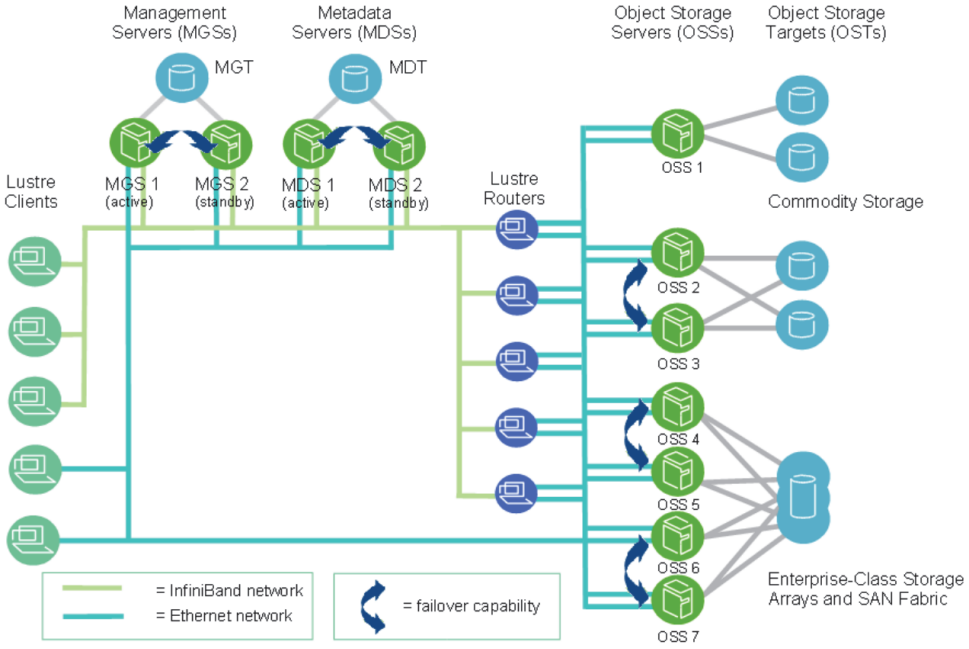

Figure 3.6.: Schematic of the Lustre file system, as present on the Piz Daint supercomputer at CSCS. The shared fast storage of the latter attributes to one Cray Sonexion 3000 and one Sonexion 1600 system, offering 6.2 PB and 2.5 PB memory capacity, respectively. Figure originally from *Lustre Software Release 2.x - Operations Manual* [242].

or two hard disks, but has several servers, which in turn contain several hard disks. As Fig. 3.6 portrays, in a Lustre file system these servers are called *Object Storage Servers* (OSS) and their disks are called *Object Storage Targets* (OST). Via management and metadata servers the OSSs are interconnected to the Lustre clients, i.e. the compute nodes. Even though these clients make the total of Piz Daint's *$SCRATCH*, `snx3000` and `snx1600`, Sonexion storage partitions look homogeneous, it is important to be aware of the underlying hardware structure.

3.3.1. ADIOS and Compression

Together with the data science team at ORNL and shortly prior to this thesis, the ADIOS[11] library[243] was implemented as an MPI-I/O backend into PIConGPU next to the already available HDF5. ADIOS, and by now its successor ADIOS2[244], are being developed with dedicated focus towards parallel, asynchronous I/O, automatically employing the best performing strategy depending on the underlying file system parameters. A schematic overview of efficient ADIOS I/O on a Lustre FS is displayed in Fig. 3.7. Processes can concurrently read from and write data into binary sub-files, which are created by so-called ADIOS *aggregators*. The metadata of the complete output iteration is stored in an XML-file and can be created offline

[11] *Adaptable Input Output (I/O) System*

after the simulation has completed. Creating the metadata file post-mortem is also generally recommended for maximum optimization of a highly parallel simulation with large datasets, many sub-files and rather high output frequency, since this is a serial operation that would otherwise stall the system unnecessarily. These ADIOS-BP files[12] are self-descriptive and allow local buffering which increases I/O speeds by streaming large chunks of data. In addition

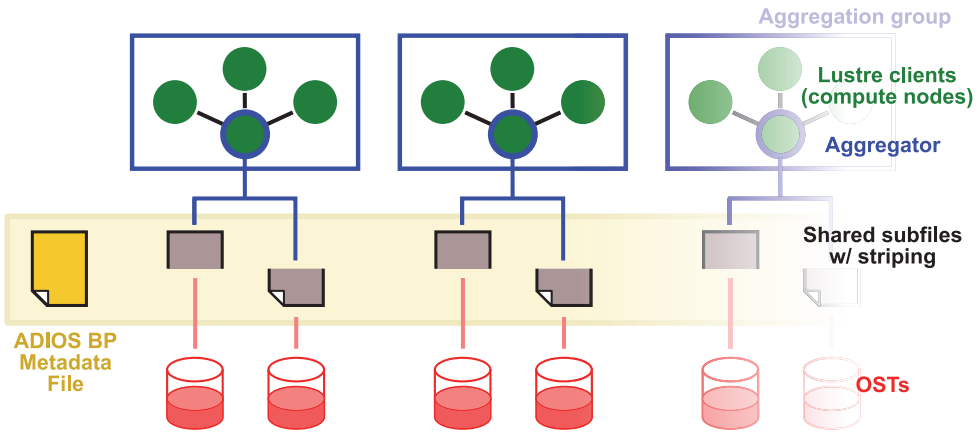

Figure 3.7.: Schematic view of efficient parallel I/O (of a single data-set written by a single simulation) as is possible with the ADIOS library on a Lustre file system. To harness the full I/O bandwidth of an HPC file system, ideally all or most OSTs are engaged. This can be achieved by file *striping* which splits them between multiple OSTs (hard drives). Keeping all OSTs at approximately the same fill level helps to maintain the performance of the whole file system. Bottlenecks like conflicting *write locks* are largely averted with ADIOS' subfiles (gray). These are created by so-called *aggregators* (blue-bordered green circles) that accumulate the data within their aggregation groups that consist of multiple compute nodes (green circles). The data is then chunked and locally buffered, thus enabling faster, stream-like writing. For full optimization, the serial operation of creating the self-descriptive ADIOS metadata BP-file (yellow) can be performed post-mortem.

to that, the Lustre file system also allows for the *striping* of files which essentially splits the file between several OSTs. In doing so, the occupation of each OST can be kept to a similar level. Highly asymmetric distribution of data can cause single OSTs to be filled close to their limit which harmfully affects the performance of the whole system. Only if all (or most) OSTs of a file system are engaged in parallel can the maximum I/O bandwidth be reached. For the fast scratch partitions of Piz Daint the manufacturer specified this maximum with $\mathcal{T}_{\text{max}}^{\text{total}} > 100\,\text{GB/s}$.

Another knob by which the file output performance of large-scale I/O can be optimized is *compression*. The work of Huebl et al. [245] quantifies the concept that data reduction in the form of compression prior to its writing to file is the most feasible if the additional compression step does not slow down the effective I/O throughput past the uncompressed output performance. In essence, any data reduction algorithm must fulfill the following inequality.

$$\frac{\mathcal{T}_R \times (1 - f_R)}{1 - \mathcal{T}_R} > \mathcal{T}_{\text{out}} \tag{3.2}$$

Here, $\mathcal{T}_R$, $\mathcal{T}_{\text{out}}$ are the per-node data throughput with reduction, and without reduction, respectively. The reduction ratio f_{out} describes the efficiency of compression and takes values of $0 < f_R < 1$. Compression in PIConGPU is efficiently enabled by the high-performance compressor `blosc`[246] which is optimized for binary data. Since PIConGPU mainly uses the GPUs

[12] File extension `*.bp` for *binary-packed*

for computation and usually one CPU for scheduling, the spare CPU cores (11 on the Piz Daint nodes) are tasked with the multi-threaded chunking, bit-shuffling[13] and compression of data.

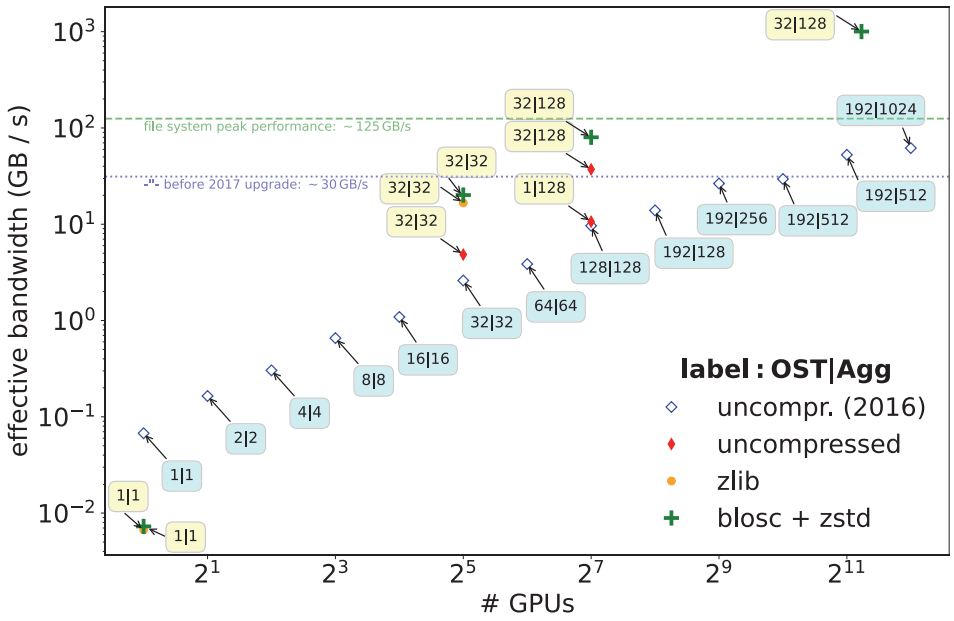

Figure 3.8.: PIConGPU ADIOS I/O weak scaling runs performed on Piz Daint with varying compression and number of OSTs and aggregators. The efficient throughput was measured as the ratio of uncompressed output size and the combined duration for compression and writing of the whole data dump. The blue-labeled markers show a dedicated weak scaling run using the Kelvin-Helmholtz Instability (KHI) setup in preparation for the PRACE application one year prior to the start of the project described in this thesis. In mid-2017 Piz Daint received a file system upgrade that reportedly quadrupled the performance due to the use of burst-buffers. The update reduced the number of available OSTs from 200 to 40. The horizontal dotted and dashed lines mark the peak performance before and after the upgrade as given by the manufacturer. The yellow-labeled data points mark weak scaling tests performed by the author in December 2017. In contrast to the KHI scenario, the production run setup contained the foil target and, initially, the majority of space was filled with vacuum. Therefore, the compression with blosc+zstd was particularly effective and fast, resulting in an efficient throughput that is much larger than the system specifications.

In a study on Piz Daint prior to its file system update, the effective I/O performance was measured and is shown here in Fig. 3.8. Most of the displayed data was taken in 2016, in preparation for the application to the PRACE project that started one year later. With an intermediate layer of burst-buffers[247] between compute nodes and OSTs, the peak performance of the file system was reportedly quadrupled[214] from $\sim$ 30 GB/s to $\sim$ 125 GB/s. Due to the file system changes and a reduction of the number of OSTs from 200 to 40, the author sparsely tested the weak scaling again during the first computation quarter of the project. Together with the I/O team of CSCS an excellent configuration was quickly identified. As a rule of thumb, the file system was observed to perform best if the following relation is satisfied:

$$\#Agg = \#OST \cdot 4 \cdot x, \qquad x \geq 1. \tag{3.3}$$

There, x is an arbitrary positive scaling factor such that the number of MPI ranks ideally becomes a multiple of the number of aggregators. However, the number of aggregators should

[13] random rearrangement of typed data into data blocks that enable more efficient compression

be equal or less than the number of MPI ranks. In practice, 32 (not fixed but varying) of 40 available OSTs were used to write 128 or 256 collectively aggregated sub-files per raw data output from the 2400-GPU jobs on Piz Daint. Using this technique and the `blosc` library with the `zstd`[248] compressor, it was possible to regularly size down the raw data outputs from 13 TB per iteration to 1.2 to 2.1 TB and write them to disk within 30 to 45 s, thus reaching effective bandwidths exceeding $\sim$ 400 GB/s.

3.3.2. Computation Time and Data Management

In total, but not persistent at one time during the simulation campaign, an estimated 4 PB of data were produced. The largest amount present on one partition were 2.2 PB at the end of the first simulation quarter. The initial data management plan, however, did not foresee the production of such a large amount of raw data at the same time. This section explains which circumstances made changes of the plan necessary and which new challenges had to be solved in its wake.

To fully make use of the 1.6 million GPU node hours that were awarded by PRACE for use within one year, the whole system needed to run for a total equivalent of 11.7 days with all nodes. The maximum resources per simulation run that CSCS allowed were 2400 of the 5704 compute nodes for 24 consecutive hours (the "*large*" queue offered 4400 nodes, but for 12 hours and by arrangement only). The computing year at CSCS is furthermore split into four quarters. All of the compute budget that is allotted for one quarter has to be completely used or it is irretrievably lost. As explained before, the massive scale of the simulations and the unfavorable target shape made the setup a balancing act between a crash and a successful completion of a run. Therefore, it was originally planned to split the compute budget irregularly over the year, especially with less node hours during the first quarter. This is reasonable because the first quarter is usually needed to install the software, run stability checks and the preparatory 2D simulations to judge the temporal evolution of the target and identify the most interesting points in time for large-scale raw data output.

However, several factors influenced this plan unfavorably. While the application for the computation time was prepared and sent in one year in advance, the successful outcome of it became known to the group only one week before the start of the first quarter. At the same time, the yearly budget was split evenly between the compute quarters and 400 737 node hours were allotted to the first quarter. During this first quarter, the simulation code had to be compiled and run successfully on the updated system, first 2D simulations had to be performed and the necessary data analysis pipeline had to be brought into place. The accompanying challenges and possible solutions are also discussed and general conclusions for HPC systems are drawn later in this chapter. A detailed view on the possibilities of analyzing large-scale data will subsequently be given in Sec. 3.4. This section also focuses on data reduction strategies (see 3.4.2) that can help lower the amount of data that has to be transferred or stored permanently. While it was still possible to usefully employ all of this compute time to the first parameter scan on copper foils, roughly 2.2 PB of raw data were also created at the same time.

Especially, structures for efficient analysis and local transport of data were not initially present and could not be implemented during the first quarter. The situation became more critical due to the 30 day grace period of data on the fast `$SCRATCH` partitions. Upon expiry of these 30 days, data was to be automatically deleted. The project resources also included 400 TB of persistent project storage which only expire 90 days after conclusion of the project. This `$PROJECT` partition was, however, not reachable from the compute nodes and data had to be on `$SCRATCH` for analysis. Additionally, the internal data mover was a SLURM queue with one node that only allowed for sequential file transfer. Not only did this make it highly overbooked

by other users as well, but it also limited the transfer bandwidth severely. Only with assistance of the CSCS support groups and the unconventional use of the external GridFTP data-mover, as well as local data reduction jobs could the situation be handled.

3.3.3. Open Standards

One of the major principles of good scientific practice states that from the data and methods that are described in a publication, the described experiment must be reproducible for other researchers. This principle also counts for *virtual experiments*, and gradually the awareness for open access and open science seems to increase throughout the laser-plasma community. More and more open-source particle-in-cell codes are being advertised at conferences[249], experimental results are collaboratively and openly compared[250] and open research policies are increasingly supported, for instance with by the Horizon 2020[251] and Horizon Europe[252] initiatives of the European Commission.

There still remains an issue of comparisons between different particle-in-cell (and other) codes that is almost considered a taboo topic because at the moment it is likely both a titanic but also unrewarding effort for any single group of researchers. However, with the *Particle-in-Cell Modeling Interface* (`PICMI`[253]) and *Accelerator Modeling Interface* (`AMI`[254]) as well as the *Open Standard for Particle-Mesh Data Files* (`openPMD`[255] and `openPMD-api`[256]), the first steps have been taken to facilitate such efforts. These standards aim to establish conventions for input files of PIC and particle accelerator simulations, and their various output data, respectively. As these are increasingly adopted by the community[257], regularly scheduled benchmarks between codes based on standard plasma physics problems, as well as code-interchangeability based on physical scale, available hardware resources, and used numerical algorithms will become easy best-practices of computational plasma physicists.

So far, and especially with respect to laser-ion acceleration simulations, real predictive capabilites are not only lacking due to missing analytical solutions of the calculated problems but also because no single numerical code can typically describe the whole range of physics in laser-matter interactions. Also with this issue the aforementioned standards are of great benefit, as the example of `simex_platform` in chapter 4.3.2 shows. Different photon science research groups concerned with laser pulse creation, light-matter interaction, beam propagation and detector technology have already joined forces within the EUCALL[258] (now PaNOSC[259]) project to combine simulation codes into full start-to-end simulations. Thereby, the work in this thesis has also contributed that virtual experiments can now be performed with less assumptions and instead more realistic modeling of laboratory experiments. Even though file formats might differ and all of the involved codes are underlying constant development, agreeing on a common standard about how particle- and/or mesh data is stored significantly facilitates scientific exchange.

openPMD

The *Open Standard for Particle-Mesh Data Files* - `openPMD`[255] is one of the first standards that aims to unify meta data and naming schemes for increased data portability between applications. Apart from HZDR-born codes like PIConGPU and ParaTAXIS other PIC codes like FBPIC[238, 260], Smilei[224], WarpX[261, 262] and OSIRIS[263] have adopted openPMD for their raw data outputs, already. Any hierarchically organized, self-describing data formats are suitable to be expressed with openPMD flavor. Two among these, namely HDF5[264] and ADIOS BP[243] are produced by PIConGPU. The main concepts of the standard include that any physical quantity that describes a particle or a field on a mesh is stored as a *record*. Multi-

dimensional *records* have *components* (e.g. x, y, z of the record *position*). A minimal set of *attributes* is required for each record, consisting of a *conversion factor* to the SI system, the *dimensionality* in SI units, stored as a length-7 array of exponents in a fixed order, and the *iteration* (time step) within the simulation.

By now, several other projects have evolved around the openPMD standard[257]. Among these are tools to explore PIC datasets using Jupyter (`openPMD-viewer`, see 3.4.1), a plugin for the scientific visualization suite `VisIt` (`openPMD-visit-plugin`), and the `openPMD-api` which is replacing the legacy library `libSplash` in PIConGPU this year, to unify the data output completely.

3.3.4. Transport of Big Data

During the first quarter of the one-year-long PRACE computation project, the massive amount of 2.2 PB of data was produced. Due to the 30-day purge-period on the fast `$SCRATCH` partition the lifetime of the raw large-scale data was very limited. This meant that only data that was actively used in analysis could stay on the partition. At first, the output iterations had to be pre-screened quickly to identify important points in time that could then be marked for keeping. For this purpose the synthetic diagnostics output of the PIConGPU plugins was vital. Each simulation run's parameter set as well as all the output except the raw particle and field data dumps and checkpoints were also immediately duplicated on the `$PROJECT` directory. The pre-screened raw data outputs that were judged less important were deleted from disk. Even with the thinning-out of the data, massive amounts of them had to be moved. The internal file transfer queue `xfer` only allowed for sequential transfer via `rsync`. This was sufficient for the small diagnostics and parameter sets but not the large datasets. Even though they were compressed, as described earlier in Sec. 3.3.1, each output iteration still measured between 1.2 and 2.1 TB. For the same reason that smaller numbers of large files are more efficient during the initial I/O, they are also more efficient during transport. Still, parallel data transfer is necessary to increase the data rates enough to avoid automatic deletion via the 30-day protocol. Together with the data transfer team of CSCS, several MPI-based tools were tested. Among these were `dcp` (distributed file copy)[265], bbcp (a peer-to-peer network file copy spawned within the BaBar collaboration at SLAC)[266], and `parsyncfp` (a parallel wrapper for the usual `rsync` utility). Unfortunately, with these methods the transfer rates could also not be improved sufficiently.

A short calculation example illustrates the challenges of moving the data. Without the reduction techniques that are later described in 3.4.2, the data transfer via the `xfer` queue at 130 MB/s would have taken 210 days. On average, `bbcp` performed with 250 MB/s. To get the data to HZDR directly would have used the 1 Gbit/s connection, also taking 214 days. Fortunately, the Technical University of Dresden is connected to the internet with 20 Gbit/s, which, if it could be used at full bandwidth, would have brought down the duration to 11 days. Using the entirety of the university bandwidth was, of course, not a feasible solution since it is also a shared resource. Eventually, the simultaneous internal transfer to the `$PROJECT` partition and external transfer via `GridFTP`[267] to the university were the best available methods at the time by which the data was successfully moved.

For high-speed remote data transfer, the CSCS offered a GridFTP[14] service[269] that is part of the recently retired *Globus toolkit*[270]. By now, it has been mostly succeeded by the *Globus research data management cloud* which continues to offer a data transfer API[271]. GridFTP is an efficient file transfer protocol that uses a control channel opened by a Globus GridFTP Server to govern the data exchange between two endpoints. These endpoints are specified with a

[14] *FTP – File Transfer Protocol*

Figure 3.9.: Amazon Web Services (AWS) Snow Family. **Left:** AWS *Snowball*, a "PetaByte scale" data migration or edge computing device that can either be optimized for computing performance or data storage with up to 80 TB capacity. **Right:** AWS *Snowmobile*, an "ExaByte scale" data transfer service that delivers up to 100 PB on an 18-wheeler semi-trailer truck. The container is almost 14 meters long and fully powered the truck has a power consumption of ∼ 350 kW. It is reportedly designed to offer a transfer speed of 1 Tbit/s on-site which would fill the 100 PB in under 10 days, if such transfer capabilities were available. For comparison, transferring 2.2 PB of simulation data using these services would result in a data transfer rate of 66 GB/s (considering a 9h15min drive between CSCS and HZDR). Image source: AWS[268].

Uniform Resource Identifier (URI), at the TU Dresden and the CSCS, for instance. The control channel opens several data channels, thus enabling multiple, encrypted TCP[15] streams via the nodes of a global grid. GridFTP was widely adopted as the main protocol for large-scale remote data transfer employed by various international research facilities, such as the Large Hadron Collider at CERN or the Oak Ridge National Laboratory. Before it could be used for this thesis work, it had to be set up at both the TU Dresden as well as HZDR to establish a chain between these data centers. This task was performed in close collaboration with the IT support of all three involved institutions. Once the GridFTP connection was established, the data transfer speed was first boosted to 400 MB/s and after more optimization with the parallel streams and an upgrade of the connection between HZDR and TU, even 1.5 GB/s were observed. During the critical phase between first and second computation quarter, many different solutions for data reduction and transport were considered. For comparison, *Amazon Web Services* (AWS) offers two data transport solutions that involve the physical moving of data storage devices, both of which are shown in Fig. 3.9. With the *AWS Snowmobile*, an 18-wheeler semi-trailer truck carrying a 45 feet container, up to 100 PB of data can be moved with a single trip between data centers. Transporting 2.2 PB between CSCS and HZDR by car (travel time est. 9 h) would have been the equivalent of 68 GB/s transfer bandwidth. For an interesting comparison, cutting-edge research set new transfer speed records via existing optical fiber connections in 2020, achieving data transfer rates between 5.53 TB/s (44.2 Tbit s^{-1})[272] and staggering 22.26 TB/s (178.08 Tbit s^{-1})[273], respectively.

Unfortunately, the global exploit of system vulnerabilities via Meltdown and Spectre in late 2017 affected HPC systems all around the world. Therefore, shutdowns and slow and careful reopening further delayed the data transfers. Thanks to another grace period of 30 days given by the CSCS User Engagement & Support team, significant parts of the data were saved from automatic purging.

Large-scale I/O is expected to grow massively over the next years as more applications are adapting to new and more powerful compute hardware, whereas the persistent storage space is not growing proportionally.[245] Avoiding to write out large-scale data wherever possible

[15] *Transfer Control Protocol*

thus becomes increasingly desirable, since dealing with petabyte-scale data will likely remain cumbersome as it exceeds single-node fast memory capacities. PIConGPU's plugins take over a significant part of the frequently recurring analysis operations *in-situ*, without ever having to write the raw source data to disk. In the case that it still has to be written, the next section focuses on how the analysis of such data can be performed quickly and efficiently using the Jupyter framework[274]. Furthermore, basic techniques are described that reduce the output of PIC simulations without losing physically relevant information.

3.4. Analysis of Big Data and Avoiding of Transport

As HPC systems are addressing the physical limits of chip density by introducing multiprocessing and heterogeneous architectures[275], they are surpassing the scaling of *Moore's law*[276] by 1-2 orders of magnitude in the last 30 years[277]. Memory capacity and I/O bandwidth increase at a much smaller rate, creating an increasingly severe disparity[245]. This is a main incentive to couple analysis and diagnostic workflows very closely to the simulation. A variety of plugins in PIConGPU offers to do this *in-situ* while the simulation data is in memory. For exploratory simulations, however, it is still necessary to write raw particle- and field data to disk. With nowadays supercomputers a single simulation iteration can easily land in the range of several 1 to 10 TB. Data of this size is naturally challenging to deal with because even cluster nodes may not have enough RAM to fit in the total of one particle species or field component. Data management policies that are rightfully aimed at enforcing good conduct and fairness in using the shared resources on a supercomputer can additionally complicate these matters for users who are running full cluster-size jobs.

3.4.1. Jupyter: Exploring and Prototyping

A critical prerequisite to explorative simulation analysis is the ability to quickly look at the produced data and make the first decision if to keep an output iteration or delete it right away. The Jupyter project[274] allows for such a workflow of easy data access with powerful libraries like `NumPy`[278] and easy plotting with `matplotlib`[279] using the production or analysis nodes of the supercomputer with the data close to where it was produced. Code is written in Python and interpreted as well as executed cell-wise. With the Markdown syntax, code can be documented in a tidy manner and prototype scripts for parallel analysis are easily created and then ported to pure Python files. Jupyter notebooks are viewed in a browser, either in *tree* mode as separate tabs or with the lightweight *jupyterlab* GUI. The exploratory workflow resembles a scientific lab book and has established itself as the main workhorse in the first analysis of PIConGPU simulations. Jupyter notebooks need a server instance that can be started by the cluster user, preferably on a compute node. It is necessary that the location of the data is available to the machine the notebook server is running on. With SSH-tunneling and port-forwarding a connection can then be established between the user client computer that runs the browser and the compute node with the Jupyter server. By now, many research facilities have also established a so-called `JupyterHub` which allows the user to easily run notebooks on an HPC machine via a web interface, secured with their login credentials. At the beginning of the compute project described here and in 4.2, such a service was not yet available on Piz Daint but was delivered via the Cray Urika XC package in mid-2018[280]. Since the Python or Anaconda environment that contains Jupyter can become quite large with user-specific software (several 100 000 files and between 5 to 10 GB of file space), it is important that the user quota on the file systems take this workflow into account.

To tackle the issue of large datasets not fitting into the memory of compute nodes, there are several solutions that work well with Python and Jupyter. *Dask* and *Apache Spark* are two of these solutions that can handle parallel collections of large arrays distributed onto multiple nodes. This is organized by dynamic task schedulers, giving the appearance that the separate memory patches of the nodes make a homogeneous entity. Recently, an impressive campaign[281] has used the resources of ORNL's Summit to accelerate the search for COVID-19 therapeutics by using NVIDIA RAPIDS[282] with BlazingSQL[283] on GPUs. With the latter, not only are query-times of large databases reduced from days or hours to minutes or seconds but currently there are endeavors to also use the framework for openPMD datasets. Together with the ADIOS2[244] streaming backend, which allows to loosely couple a PIConGPU simulation to an analysis program by directly sharing its data, promising new avenues towards fast analysis and avoiding of large-scale output become possible.

3.4.2. Data Reduction Strategies

Good scientific practice dictates that primary data used in publications should be archived alongside the analysis results. For PB-sized file outputs this rule can mean a significant monetary expense to always keep enough tape drives for this purpose. However, especially solid density PIC simulations are often highly over-resolving large parts of the simulation volume to ensure stability of the algorithms and avoid the accumulation of numerical errors. For the interpretation of physics results, this resolution is usually not needed since the structures and relevant dynamics evolve on scales that are often larger by a factor of 10 to 100. The following collection of simple data reduction techniques was used to significantly reduce the data footprint of output iterations.

In-situ plugins created angle- and energy-resolved histograms with high temporal frequency. Field data was spatially averaged over neighboring sets of cells ($2 \times 2 \times 2$). Volumes of interest, spanning one quarter of the focal plane, were extracted from the fields, preserving the original resolution. Particle data was filtered in-situ for their angle θ between the target normal and their momentum direction, creating subsets of particles which correspond to the particle yield that would be accepted by the aperture of a virtual pin-hole ($\theta = 4.5°$) and Thomson-Parabola spectrometer (TPS), or a virtual radiochromic film (RCF) stack detector.

3.4.3. Dedicated Resources

The tasks in a data-centered scientific workflow are manifold and can best be described by the example of the data life cycle. Particle-in-cell simulation data is produced on dedicated compute nodes with a high-bandwidth I/O infrastructure. Ideally, the system First analyses are performed on the data in-place, for which a dedicated (sub-)cluster with so-called *fat-nodes* is beneficial. These nodes ideally have a large amount of RAM (e.g. 0.5 to 2 TB) which allows not only the original data set to be kept in memory but also its computation results. Simple, parallelizable analysis operations benefit from many fast CPU cores. The dedicated analysis cluster itself ensures that data analysis does not interfere with production runs of users on a time-limited budget. Access to fast scratch memory should be enabled for both the production as well as the analysis cluster to avoid unnecessary transfers and file system load. On a dedicated and backed-up partition, more persistent data is stored and organized by projects with access rights for the users belonging to the same compute project. Personal user memory besides the `$HOME` directory should be available for additional software. The file system there should be able to handle large amounts of small files which are usually created during the compile processes of software. In the case of longer (one- or multi-year) projects, an effi-

cient “warm” archiving solution can be used to allow users to store data from the beginning of their compute project and compare it again with datasets produced at the end of the project. As such, some facilities like OLCF offer temporary archiving of simulation data with internal transfer speeds of up to 100 TB/h via fast intermediate buffers. This also acts as a safety net for hardware maintenance periods that leave only part of the file system active or allow users to quickly remove data from the fast scratch space to avoid reaching the global quota. The final step of the data life cycle can be long-term archiving after it was used for publication in a scientific journal. The spirit of open science suggests that this data is then also made available upon request or just for download and uniquely addressed with a Digital Object Identifier (DOI). Unfortunately, quick access to the tape drive storage system at Piz Daint was not available at the time. However, after transferring the data from the project described in Sec. 4.2 to HZDR via the TU Dresden’s Taurus cluster storage, it was successively archived to the HZDR tape drive library. As of April 2019, the HZDR is running an open-source powered publication platform for datasets which assigns a DOI using the HZDR-specific prefix `10.14278` to each of its entries.

4. Results

While the understanding about physics of laser-driven underdense plasmas is already quite substantial, the intricate details of its overdense counterpart are much less understood. This lack of understanding stems from the shortage of probing methods that are able to directly resolve the femtosecond-nanometer dynamics within the target at the time of interaction with the most intense part of the laser pulse. Furthermore, detailed knowledge about the exact state of the target as well as the earlier conditions that lead up to this state are typically only accessible up to a few picoseconds pre- and post-arrival of the laser peak. Particle-in-cell simulations are commonly employed to model the physics of the few 100 femtoseconds around $t_{\mathrm{rel}} = 0$ but as the previous chapter showed, the enormous computational cost of laser-solid PIC simulations severely limits the physical resolution and the overall size of the physical volume, or often rule out a fully 3-dimensional treatment of the problem, at all. However, to gain predictive capabilities, rigorous scanning of the parameter space is necessary and even then initial conditions like pre-plasma gradients or electron temperature distributions are decisive for the evolution of the simulated plasma. With increasing control over the temporal pulse contrast in laboratory experiments, reproducible pulse shapes on the linear intensity scale are reliably generated but a high shot-to-shot fluctuation of the ion acceleration performance remains. In addition to the known influence of temporal pulse shape features reviewed in section 2.2.2, this hints at the increasing importance of the still largely unexplored intensity ramps during the last picosecond. This work aims to shed light on the influence such intensity ramps have on the target conditions at the time the pulse maximum arrives and how they shape the ion beams that are generated from thin foil targets.

This chapter presents the results obtained in the large-scale simulation campaign that was dedicated to answer this question and performed on the Swiss supercomputer Piz Daint. Since the simulations were explorative in nature and massive in size, certain prerequisites had to be met first to make them possible and ensure that the relevant physics could be addressed.

Efficient absorption of laser energy into free electrons is crucial for later conversion into ion energy. The absorption efficiency η still remains one of the most challenging parameters to measure in experiment and the existence of pre-plasmas can significantly alter it. Since the temporal structure of the laser determines when and how the target material is ionized, i.e. free electrons are produced, the ionization physics need to be carefully modeled. Expanding on the author's previous works [G5, G6], the ionization framework in the highly scalable particle-in-cell code PIConGPU was improved to meet these requirements. The results of these improvements and their application to two other separate studies performed at HZDR, partly in collaboration with other institutions, are presented in the beginning of this chapter.

The first application was a study of laser-wakefield acceleration (LWFA) via self-truncated ionization injection (STII) supporting an experimental campaign at HZDR that set new standards for very stable, high-current tabletop laser-electron accelerators. The intricate dynamics that lead to such stable operation could only be revealed due to synthetic diagnostics in the simulation that relied on being able to distinguish separate electron species by their origin in different atomic shells. The findings of this campaign were published by Couperus et al. [G3] and Pausch [62].

With the implementation of the Thomas-Fermi (TF) collisional ionization model in the course of this thesis, PIConGPU is now also applicable to ion acceleration scenarios using solid density targets. The second application studied laser-generated structural changes in colloidal silicon crystals. The collaborating partners at DESY, Hamburg, performed the supported experiment that used time-resolved X-ray diffraction to probe the disintegration of the crystal structure. Here, the success of the simulation was first challenged by characteristic artifacts of the Thomas-Fermi model, specifically low temperature plasma regimes, some of which have been the subject of intense discussions in recent literature[284–286]. These were overcome by additional extensions, resulting in more physical behavior of the TF model, which were also implemented and tested by the author. More detailed results are published in the manuscript of Mukharamova et al. [G4].

In the second section, the results of the campaign regarding the influence of different leading edge intensity ramp shapes of the UHI drive laser pulse are presented. First, the general effect on the pre-plasma development, laser absorption and the ion acceleration process is evaluated for different target types. Then, the main finding of optimum proton energies from metal targets for non-"ideal" pulses is presented. It was found that laser pulses that deviate from a perfect, Fourier transform limited (FTL) Gaussian shape can inject protons into the electron sheath behind the target at the location and instance of largest accelerating fields. Furthermore, an extension to the analytic ion acceleration model by Schreiber et al. [44] was found that describes the ion energy evolution with respect to their origin. These results are discussed in the light of recent experimental results at the DRACO laser-ion acceleration facility at HZDR.

In a final section, novel avenues for time-resolved probing and diagnostics of the acceleration process are explored. Radiation signatures observed in the large-scale simulation campaign of the previous section are discussed as promising candidates to gain information about the spatial and spectral conditions of the electron population during the ultrashort interaction window with the laser maximum. Particularly the highest energy part of the bremsstrahlung spectrum is created only during this time and new detector developments at HZDR promise to give insight into this previously inaccessible but most important part of the interaction on a single-shot basis.

Active, optical probing methods cannot penetrate past the surface of solid-density plasmas and most existing methods that characterize the particles or radiation emitted from these plasmas, like electron calorimetry or Kα diagnostics, are still largely time-integrated and indirect. A new class of pump-probe experiments promises to give a more direct access to the very fast and very small time– and length scales that are vital for the acceleration of ions from solids. These experiments involve ultrashort X-ray pulses of unprecedented brightness produced in novel X-ray free electron laser facilities (XFELs). Using the method of Small-Angle X-ray Scattering (SAXS), it was predicted that the simultaneously high temporal and spatial resolution, that other plasma diagnostics are falling short of, can now be utilized to deliver new insights into short-pulse laser-matter interactions[50]. Unfortunately, the interpretation of the scattering images is not straightforward.

Therefore, 2D PIC simulations of silicon grating targets are presented that preceded and

prepared the first direct measurement of plasma expansion on a femtosecond-nanometer scale[G8, G9]. These targets were designed with the intent for a clear scattering signal that is more accessible to interpretation. Once the surface grating structure is heated by the laser, the scattering signal changes as the sharp grating edges soften. Thanks to the knowledge gained from the simulations, a parametric model (described in extensive detail in the thesis of Rödel [287]) could be built that allows the direct deduction of the plasma gradient from the scattering image. In addition to that, an unexpected expansion signature was found that explains why a simple one-grating model is not sufficient for all experimental shots, especially at later interaction times and one shot could be identified that shows a strong indication that the signature was directly measured.

While ultrathin targets are promising to deliver the highest proton energies, their interaction dynamics are often unstable. Thicker targets above the micrometer range are less sensitive to laser pulse parameters and show smaller shot-to-shot fluctuations which make them better suited for reproducible experiments with SAXS. However, effects like the time structure of the probe pulse, the temporal evolution of the probed structure and absorption processes make predictions with simple Fourier transforms challenging. Moreover, multiple scattering cannot be incorporated at all. This section concludes with a proof-of-concept simulation study of the aforementioned grating targets with ParaTAXIS, the PIC-like photon tracing code developed at HZDR. The opportunities and feasibility of such tools are discussed as they were presented in the context of full start-to-end simulations with `simex_platform`, the simulation suite of the international collaboration of advanced laser light sources, EUCALL[G10–G12, 288–290].

4.1. Application and Validation of Ionization Methods

Reliable modeling of the spatiotemporal electron density evolution is crucial in the pursuit of answering the question what the influence of leading laser pulse intensity ramps is on the acceleration performance from ultrathin foils. Since the intricate laser-particle acceleration dynamics in solid-density plasmas are much less understood than in the underdense case, the latter can be used as a testbed of the additional physics models that are employed in PIC simulations of the former where experimental validation is all the more challenging. Laser-gas interactions historically provide a valuable experimental validation to field ionization models (see section 2.3.2). This is due to the laser's direct interaction with the atoms without virtually any other interfering processes. However, for laser-solid interactions the electric field of the UHI laser penetrates only the surface and exponentially decreases with the target depth. There, field ionization effects are much more challenging to validate because the important scales shrink down to femtseconds and nanometers while collisional effects also start to play an increasingly dominant role. To tackle this issue gain more insight into the origin of electrons created in ionization processes, the author has developed a modular, adaptable ionization framework for the particle-in-cell code PIConGPU.

Rigorous testing on the verification of model functionality and validation of physical correctness preceded the two real-life physics applications described in this section. As such, the studies presented here involve complex dynamics and therefore serve as additional indirect validation since the processes heavily rely on ionization physics. They furthermore show the necessary capabilities that were developed during this thesis to enable the pursuit of the original thesis goal.

4.1.1. Field Ionization: Laser-Wakefield Acceleration of Ionization Injected Electrons

The first successful large-scale simulation study that utilized the advanced ionization framework described in chapter 3 was mainly performed by Pausch [62] and aimed at explaining the extraordinarily stable production of high-charge electron bunches in an LWFA STII setup that was experimentally demonstrated at HZDR by Couperus et al. [G3]. The basic mechanics of electron acceleration from a gas target in the wake of an ultrashort superintense laser pulse have been explained in chapter 2.1.2. Nowadays, novel accelerators operating with the LWFA mechanism can stably produce beam currents at a magnitude of several 1 to 10 kiloamperes[291, 292]. To qualify as sources of free-electron lasers and other novel light sources it is highly desirable to increase these currents by at least a factor of 100. This entails increasing the electron bunch charge captured in the plasma cavities from the picocoulomb to the nanocoulomb level. Unfortunately, at such bunch charges, the self-generated field of the accelerated bunch is strongly modifying the accelerating field, potentially deteriorating the quality of the electron beam. This effect is called *beam-loading*[293, 294].

Self-Truncated Ionization-Injection

Couperus et al.[G3] demonstrated the stable operation of a nanocoulomb-class laser wakefield accelerator using the technique of *self-truncated ionization injection*[55, 295, 296]. The main idea that allows to meet the optimal loading condition while keeping a high beam quality is to take control over the exact location and amount of injected charge within the accelerator. This is done by doping the low-Z main constituent of a gas jet target with a high-Z species whose atoms not only carry many more electrons but show a steep jump in ionization potentials between charge states. The low-Z component acts as the main carrier of the plasma wave while in the high-Z component a specific atomic shell is targeted to become the main source of electrons that will be accelerated to high energies with a low spread. As the relativistic laser self-focuses inside the plasma, only at the center of the pulse the intensity surpasses the ionization threshold of the chosen shell (typically the K-shell).

There are then three main experimental parameters that allow for precise control of when and how many electrons are injected into the bubble. Those are *laser intensity*, *spatial gas density profile* and the *dopant concentration*. The first two determine the general acceleration regime, the degree of self-focusing, and the plasma bubble size. The dopant concentration influences the amount of injected charge during the time the conditions for successful injection are met. By controlling these parameters, electron injection is allowed to happen only in a well-defined interaction region and then stopped (*truncated*) due to the tailored gas density profile, thus avoiding too much charge inside the bubble which would eventually lead to beam breakup and filamentation.

Ionization Setup for LWFA STII

The experimental campaign conducted at HZDR is described in full detail in [G3, 297]. This section focuses on the application of the ionization framework that allowed to perform this simulation study and clearly distinguish between the plasma background electrons and the nitrogen K-shell electrons that were targeted by the ionization injection mechanism. The contribution of the author lies in the implementation of the methods and the ionization setup for the simulation campaign. In the following, the simulation strategy is presented after which the results are also reviewed briefly.

Level	Transition	$\mathcal{E}_{\text{ion}}$ [eV]	$\lvert\vec{E}_{\text{ion}}\rvert$ [a_0]
He I	$He \rightarrow He^{1+} + e^-$	24.6	$2.6 \cdot 10^{-2}$
He II	$He^{1+} \rightarrow He^{2+} + e^-$	54.4	$6.4 \cdot 10^{-2}$
N I	$N \rightarrow N^{1+} + e^-$	14.5	$9.1 \cdot 10^{-3}$
N II	$N^{1+} \rightarrow N^{2+} + e^-$	29.6	$1.9 \cdot 10^{-2}$
N III	$N^{2+} \rightarrow N^{3+} + e^-$	47.4	$3.3 \cdot 10^{-2}$
N IV	$N^{3+} \rightarrow N^{4+} + e^-$	77.5	$6.5 \cdot 10^{-2}$
N V	$N^{4+} \rightarrow N^{5+} + e^-$	97.9	$8.3 \cdot 10^{-2}$
N VI	$N^{5+} \rightarrow N^{6+} + e^-$	552.1	2.20
N VII	$N^{6+} \rightarrow N^{7+} + e^-$	667.0	2.75

Figure 4.1.: Helium and nitrogen ionization states with their ionization energies $\mathcal{E}_{\text{ion}}$ and the respective threshold fields $\lvert\vec{E}_{\text{ion}}\rvert$ in units of normalized laser amplitude a_0. The threshold fields are calculated from the BSI model following equation 2.15. Originally in [62].

The gas mixture that was used in the experiments at HZDR contained mainly helium and nitrogen which was doped in with between 0.5 % and 3.0 % partial pressure. This low ratio ensures that the overall plasma dynamics are solely determined by the helium density and are not affected by the amount of nitrogen dopant. Table 4.1 shows the complete list of atomic transitions of both gases and their respective ionization energies $\mathcal{E}_{\text{ion}}$. In the column on the far right the corresponding electric field thresholds $\lvert\vec{E}_{\text{ion}}\rvert = \mathcal{E}_{\text{ion}}^2/4Z^*$ for abundance of the charge state were calculated, using the BSI model (see 2.1.4). The helium `He I-II` as well as nitrogen charge states `N I-V` (marked in orange) are completely ionized by the laser pre-pulse. Only the nitrogen K-shell electrons require a laser pulse of relativistic strength. An amplitude of $a_0 \geq 2.2$ is only reached within the center of the laser spot as it undergoes self-focusing in the plasma. With the first simulations featuring initially neutral atoms of both atom species it was confirmed that the lower charge states are really ionized early on the rising flank of the laser pulse. Since it was found that the ionization dynamics of these states do not influence the plasma dynamics during the interaction with the main pulse maximum both atoms can safely be assumed as pre-ionized. This allowed to combine the electrons from these states into a single macro-electron species which reduced the overall number of particles and freed computational resources. These resources were then employed to increase the transverse size of the simulation box and model a more sophisticated laser structure, better matching the measured profile in the real laser pulse. Both species' atomic nuclei barely contribute to the plasma dynamics in the laser blow-out region due to their high mass. For all relevant interaction times, the helium atom is fully ionized and its nucleus can thus be entirely neglected from the simulation. The solution of Maxwell's equations automatically creates stationary mirror charges complementing the electrons that are placed on startup. Of prime importance is now that PIConGPU was configured to produce a *separate species* of electrons during the ionization of the nitrogen K-shells. These electrons behave physically in the same way as the plasma background electrons but remain clearly distinguishable within the simulation and conveniently produce synthetic diagnostics output separately.

In addition to this, three different ionization models and their combinations were tested for their influence on the dynamics to ensure that the results are robust against the choice of model.

The Keldysh[78] and ADK[298] model both describe ionization in the tunneling regime (compare 2.2) while the BSI model treats the regime where the field strength completely suppresses the atomic potential barrier and frees the electron classically. Currently, most ionization models that qualify for PIC do not cover both regimes sufficiently well[G5, G6]. Thus, a combination

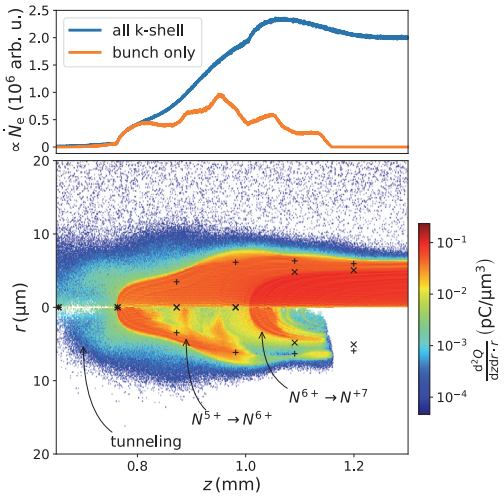

(a) Ionization dynamics from a LWFA STII simulation with combined ADK and BSI model. **Upper panel** Number of all newly created K-shell electrons versus the ones meeting the right conditions to be accelerated inside the high-energy electron bunch. **Lower panel** Electron charge density at their point of creation with respect to longitudinal position along the plasma channel and distance from the beam axis (normalized the finite ring volume of $2\pi r \Delta r$). The upper half shows all K-shell electrons while the lower half just shows the electrons in the accelerated bunch. For chosen time steps, markers denote the radial extent of the electric field where it exceeds the BSI threshold of **+**: N7+ and **x**: N8+ ($r = 0$ means *nowhere*).

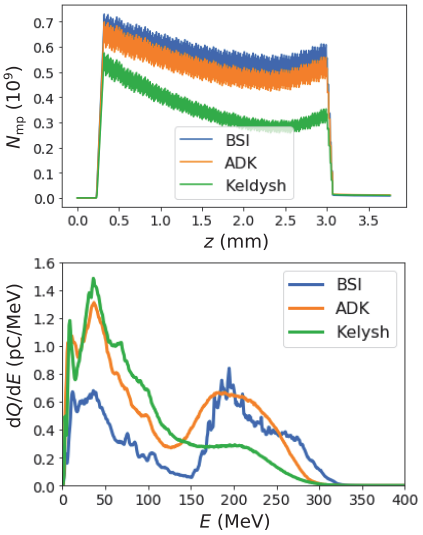

(b) **Upper panel** Nitrogen macro-electron number within the moving simulation box in three PIConGPU simulation runs of the LWFA STII setup that only differ in the choice of ionization model: ADK, BSI, and Keldysh. **Lower panel** Final electron spectra.

Figure 4.2.: Ionization dynamics in the LWFA STII configuration, with varying ionization models.

of the ADK and the BSI model were employed first to cover the ionization dynamics as best as possible. Figure 4.2a shows results from this simulation. The upper panel shows the number of newly created nitrogen K-shell electrons, whose dynamics were tracked[1] from the points of their creation throughout the rest of the simulation. It becomes apparent that of all K-shell electrons that are created (blue line) in the focal region of the laser, only a limited number fulfill the trapping condition for the LWFA process and end up in the accelerated electron bunch (orange line). Therefore, one can see very clearly the self-truncation of the injection mechanism as the number of injected electrons stops at a distance of $z = 1.2\,\text{mm}$. In the lower panel, a 2D histogram is shown that depicts the charge density of newly created electrons versus the position along the beam propagation axis and the respective radial distance away from it. In the upper half, again all electrons are shown while in the lower half only the electrons were counted that ended up in the accelerated bunch. With this filtering and a comparison to the local electric field at a few chosen time steps (see + and x markers), an otherwise hidden sub-structure in the ionization and trapping dynamics becomes apparent. It emphasizes two regions of high charge density originating from the two ionization levels of the nitrogen K-shell on a background of electrons produced during tunneling at lower intensities. Within these regions, smaller lobes appear where the conditions of electron momenta and in-bubble position are just right to be trapped and accelerated with the bunch.

Additionally, each ionization model was also run by itself and figure 4.2b compares the amount of nitrogen electron macroparticles produced in three otherwise identical PIConGPU

[1] A subset of particles that were issued an ID upon creation for more detailed information on single-particle trajectories.

simulation runs. In the run with the Keldysh model, significantly less electrons were produced which did not result in the acceleration of a well-formed quasi-monoenergetic bunch but rather a broad distribution of energies. Both the bunch-shape, injected charge and peak electron energy measured in the experimental campaign these simulations supported were represented much better with either just the ADK or BSI model. Only the shape of the electron energy spectrum was smoother in the setup that used the ADK model due to its probabilistic nature. The BSI model, however, is computationally much less demanding since only a threshold field strength is considered for the ionization process. While it is generally possible that the exact spatial and temporal ionization dynamics influence the acceleration process's outcome even more significantly, the BSI threshold model was sufficiently accurate in the simulated STII scenario. This, in turn, allowed for simplifications in the modeling that freed more resources, which could be employed to widen the parameter scan to include other parameters like the laser focus position.

Results

Experimentally, the new STII scheme was demonstrated to have a very high shot-to-shot reproducibility and to yield electron beams with several 100 pC of charge at simultaneously low energy spread. Together with the PIC simulations it could be confirmed that focusing the laser towards the end of the gas jet led to a long interaction phase during which the laser pulse was focused to a spot size smaller than the vacuum spot would have been. This self-focusing increased the a_0 and upon surpassing the nitrogen K-shell ionization threshold, injection ensued and continued for about 1 mm while the pulse underwent self-guiding in the plasma. The electron injection was not only delayed but also prematurely terminated by the rapid decrease in gas density due to the tailored laser profile, causing the pulse to defocus and the bubble size to increase. Careful optimization of the focus position and gas profile resulted in the precise control over the charge of electron bunch and the possibility to optimize it towards the theoretical limit[293]. Despite 1.5 years of simulations with varying parameters, an exact match for the experimental results could never be reached. This was due to an oversight in the focus position measurement where the defocusing effect of a damping filter led to a later focus than in the actual experiment where the filter was not used. However, the extended simulation campaign provided much insight into non-linear laser focusing in plasmas. The disagreement between

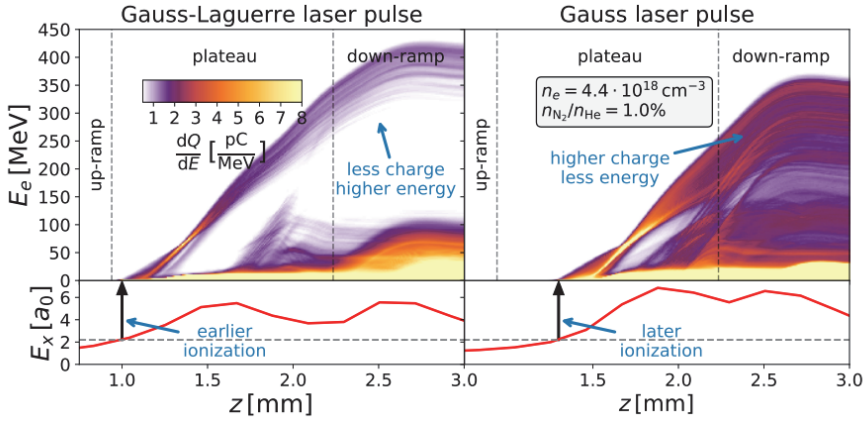

Figure 4.3.: Electron energy evolution for laser pulses with higher spatial Gauss-Laguerre modes (left) and a Gauss-only mode (right). In the upper panels, the electron energy is depicted while the lower panel shows the laser electric field E_x (red line) versus the propagation distance (eqivalent to time). The dashed gray lines illustrate the ionization threshold for the transition of N^{5+} to N^{6+}. The black arrow points to the start of injection which coincides with the time the nitrogen K-shell is ionized. Originally in [62].

experiment and simulation led to more careful laser modeling and the newly implemented *Gauss-Laguerre* profiles showed a much earlier focusing of the pulse inside the plasma. Hence, injection was observed at a much earlier time and since the bubble shape was similar, the acceleration gradient was left unchanged. As a result, the final energy of the electron bunch was higher than in the Gauss-only case. The direct relation between focus position inside the plasma and final energy, as well as the initial misconception of the experimental measurement, revealed the increased need for a reliable focus diagnostic that could facilitate the comparison between simulation and experiment. Such a diagnostic was developed by Pausch [62] from a unique radiation signature of the plasma electrons.

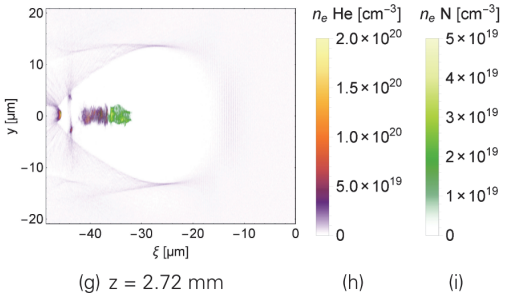

(a) Electron density slice through the bubble in an LWFA STII simulation near the end of the accelerator. Areas shaded in green mark the nitrogen K-shell electrons while the plasma background consisting mostly of helium electrons is marked in violet and red. Towards the left end of the bubble, background electrons have entered via the downramp injection mechanism while the wakefield structure is leaving the gas jet and plasma density decreases. Originally in [297].

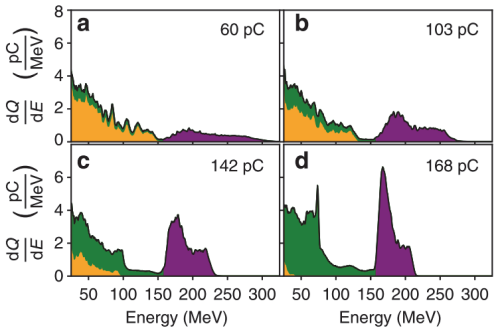

(b) Electron histograms of selected simulations (**a**–**d**) with increasing bunch charge. *Green-* and *purple-shaded areas* are the contribution of nitrogen K-shell electrons. *Orange areas* mark the (mostly) helium and (vanishingly few) nitrogen L-shell *background electrons*, respectively. Originally in [62].

With about 5 days per full simulation, it becomes apparent that the virtual experiment is becoming already as much of an effort in execution as well as analysis, since the amount of data from a single 3D LWFA STII run on a university cluster easily exceeds 10s of GigaBytes. Each of the solid-density simulation runs described in section 4.2 created 10s of TeraBytes instead, showing that it is all the more vital to increase the number of in-situ methods that allow on-the-fly analysis for explorative simulations.

4.1.2. Collisional Ionization: Structural Changes in Laser-Heated Colloidal Crystals

In a collaboration with the Coherent X-Ray Scattering and Imaging Group[2] of the *Deutsches Elektronen-Synchrotron* (DESY) the author contributed to the simulation modeling and interpretation of a pump-probe experiment performed at LCLS at SLAC[299] by providing simulation setup and support. In the experiment, colloidal crystal samples were pumped by an infrared laser ($\lambda_L = 800\,\text{nm}$, $I_L \sim 10^{14}\,\text{W cm}^{-2}$) and after a delay of up to 1000 ps the created plasma was probed by X-ray pulses (photon energy $E_{ph} = 8\,\text{keV}$) to deduce structural changes from the temporal evolution of the Bragg scattering signal. The experimental methodology of the experiment (schematic setup in Fig. 4.5) performed by Mukharamova et al. was published in [300]. The subsequent analysis of the experimental results, supported by the aforementioned simulation work can be found in [G4]. During this campaign, the author extended the freshly

[2] https://photon-science.desy.de/research/research_teams/x_ray_crystallography_and_imaging/research_areas/coherent_x_ray_scattering_and_imaging/index_eng.html

implemented Thomas-Fermi collisional ionization model to counteract unphysical model artifacts that obstructed the success of this study.

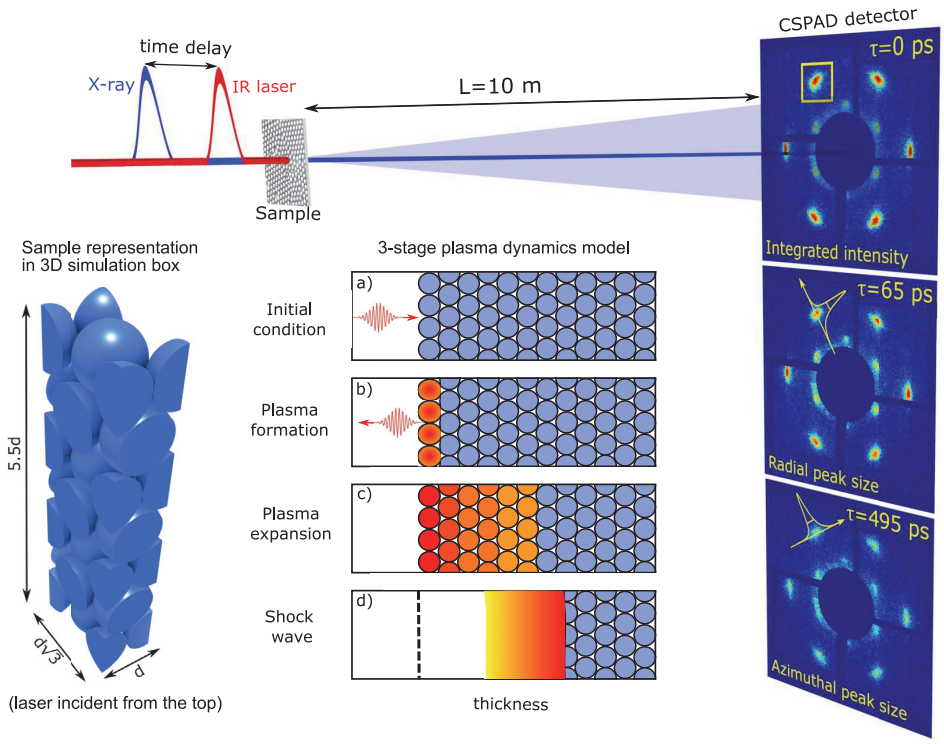

Figure 4.5.: **Top:** Schematic view of the IR-pump-X-ray-probe setup used to study the structural changes of the polystyrene sample. Selected scattering images for different pump-probe delays τ stacked on the right. **Bottom left:** Representation of the colloidal crystal sample used for 3D simulations with PIConGPU (laser direction here from top to bottom). **Middle panels *b) – d)*:** Three-stage model developed with results from the combined PIConGPU + HELIOS simulation study. Adapted from [G4].

Target and Ionization Setup

To explore the dynamics of the IR-pumped colloidal polystyrene (mass density $\rho_{PS} = 1.05\,\mathrm{g\,cm^{-3}}$), a combined simulation study was performed. The 3D particle-in-cell code PIConGPU (see 3.1) was employed for the direct, short timescale ($\lesssim$ 1 ps), laser-matter interaction. The longer evolution ($\lesssim$ 1000 ps) of the created plasma was modeled using the 1D radiation MHD code HELIOS-CR[301].

The target material consisted of closely-packed spheres (diameter: 163 nm of polystyrene $(C_8H_8)_n$ which is an electrical insulator and translucent for optical light. PIConGPU simulations were performed at the Hypnos supercomputer at HZDR. Given the structure of the target (see Fig. 4.5 bottom left 3D illustration) a reduction of dimensionality to 2D was inadequate as it would have imposed a cylindrical symmetric on the colloidal crystals. Given that 3D PIC simulations are computationally costly, the simulation volume had to be reduced to a small region inside the laser focus with the longest edge along the laser propagation direction. Such a quasi-1D setup allowed for the simplification to a plane-wave laser and periodic simulation box boundaries since the lateral size of the box was small compared to the focal spot size. When modeling a target material for high-power laser matter interaction, the effect of the intensity ramp leading to the main pulse is usually modeled by pre-plasma and pre-ionization

of the target. Initially free electrons are also a good approximation for the electron gas in metals. For the comparably low intensity of the laser in this experiment and the polystyrene target, pre-ionization was not a suitable initial condition.

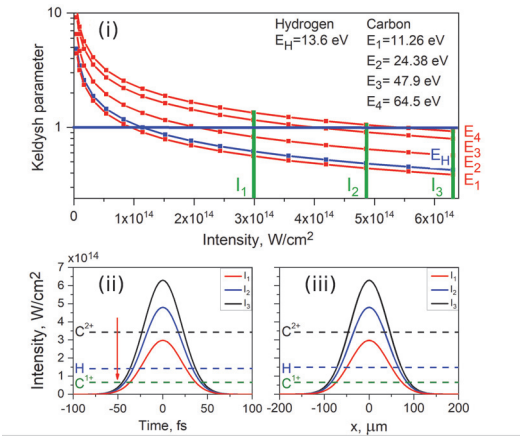

(a) Keldysh parameter vs laser intensity (i) for different ground state ionization potentials of hydrogen (E_H) and carbon (E_1 – E_4). Temporal (ii) and spatial (iii)laser intensity profiles in the PIConGPU simulations (I_1 – I_3). Dashed lines represent the appearance intensities for selected lower charge states. The red arrow marks the beginning of the simulation. Originally in supplementary material of [G4].

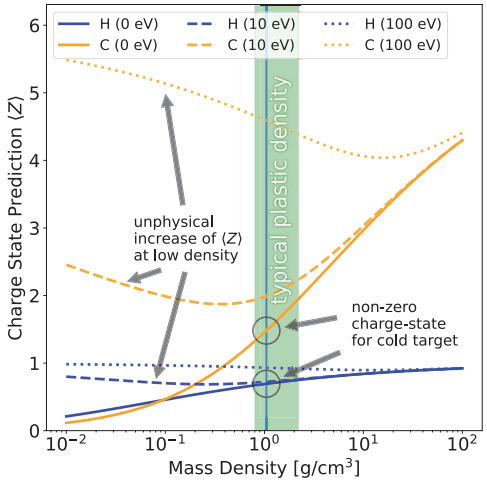

(b) Thomas-Fermi average charge state prediction following [90] for hydrogen and carbon atoms. Regions of questionable physical correctness exist for low densities and cold target conditions, which are addressed by cutoffs that are discussed in this chapter.

Figure 4.6.: Ionization prediction of atomic species in polystyrene.

Ionization in the particle-in-cell simulations was treated in two ways: a combined ADK (Eq. 2.14) and BSI (Eq. 2.15) model for field ionization and the Thomas-Fermi (TF, see Eq. 2.19) model for collisional ionization. Three different peak intensities I_1, I_2 and I_3 on the order of $10^{14}\,\mathrm{W\,cm^{-2}}$ were used for the simulations. Figure 4.6a shows the Keldysh parameter $\gamma_K = \omega\sqrt{2E/I}$ (in atomic units) for these intensities and the lowest carbon and hydrogen charge states of ionization energy E. It can be seen that for most of the laser duration γ_K is smaller than 1 and using models from the quasi-static ionization regime is justified. Comparing the appearance intensities (see 2.1.4) of these charge states against the intensity envelope shows that only charge states up to C^{2+} and H^+ were expected to be saturated. It is important to start simulating the laser pulse already at intensities before it reaches a level where it becomes large enough to ionize the lowest bound states. This ensures more natural ionization and electron motion dynamics and avoids that a quick succession of ionization processes creates a sharp electron density boundary at which the laser is reflected in an unphysical way.

As figure 4.6b illustrates, the Thomas-Fermi prediction of the initial charge states of both Hydrogen and Carbon at plastic density and room temperature (even at 0 K) are non-zero. Double-checking with the atomic population kinetics tool FLYCHK[184, 185] yielded the same prediction since the molecular structure of polystyrene is not represented in either. This is just one of the cases where the applicability of the Thomas-Fermi model is questionable. Better models like e.g. direct impact ionization (see section 2.3.2) exist but many codes still rely on the Thomas-Fermi implementation for its simplicity and general performance. In the specific case of PIConGPU, more advanced models require finding actual collision partners between the simulated particles which severely challenges the scalability and compute speed of the code that is one of its key features. Development regarding these methods is currently ongoing at HZDR.

Results

First simulations without any TF adjustments showed that plasma was already present at the start of the simulation with charge states of H^+ and C^{2+}. Therefore, the target was immediately overcritical to the relatively moderate intensities of the laser pulse. Heating of the target did not occur in a physical way which affected the later expansion of the target that did not agree with the measured data.

The shortcomings of the Thomas-Fermi model are known within the literature but their treatment for more physical results is a rich source of discussion[284–286, 302, 303]. Within the scope of this thesis, three distinct, user-configurable thresholds were implemented in PIConGPU to counteract the unphysical behavior in regions where the model is not applicable.

1. A **low-temperature cutoff** ensures that below a certain free electron average energy $T^{\text{low}}_{\text{cutoff}}$ no collisional ionization is enforced.

2. A **low-density cutoff** $n^{\text{low}}_{\text{cutoff}}$ is motivated by an average charge state $\langle Z^* \rangle$ increase that follows from the model upon decrease of the ion density at electron temperatures of already 10 eV. This is unphysical behavior since a lower ion density results in both less overlap of the ion spheres and less likely electron ion collisions.

3. A **high-energy cutoff** $T^{\text{high}}_{\text{cutoff}}$ was introduced to exclude super-thermal electrons that are commonly produced within the $2\omega_L$-bunches by laser pulses of $a_0 > 1$. These electrons have a much lower interaction time and cross section and move ballistically through the target.

These three thresholds were successfully implemented and employed for the PIConGPU simulations of the laser-colloidal crystal interaction. The value of $T^{\text{low}}_{\text{cutoff}} = 1$ eV was chosen to leave the first ionization levels to field ionization processes at the target surface and only allow collisional ionization once free electrons of significant average energy (that is still below the lowest ground state binding energy value) are present. The density threshold was set to 1 n_c below which the ionization is dominated by the high-power laser that can directly penetrate the target. Finally, $T^{\text{high}}_{\text{cutoff}}$ was set to 50 keV as is also discussed in [284].

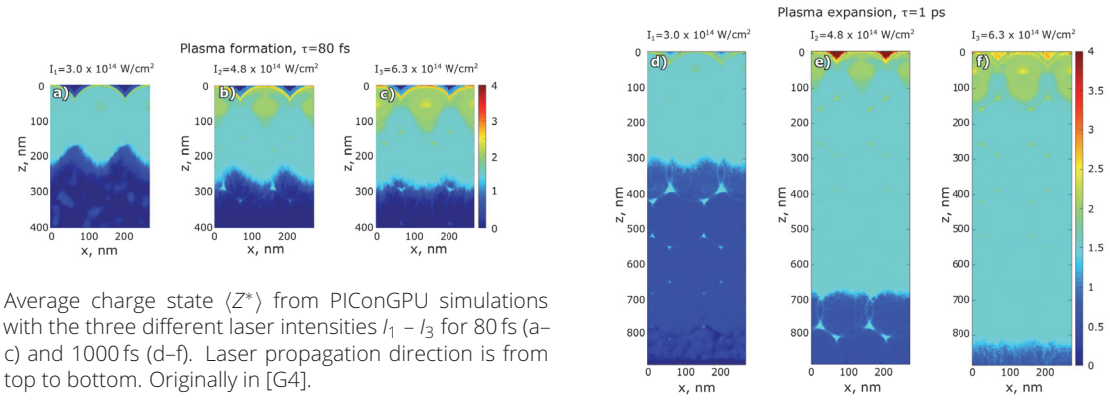

Average charge state $\langle Z^* \rangle$ from PIConGPU simulations with the three different laser intensities I_1 – I_3 for 80 fs (a–c) and 1000 fs (d–f). Laser propagation direction is from top to bottom. Originally in [G4].

Figure 4.7.: PIConGPU simulations of the colloidal crystals being pumped by the IR laser.

After the maximum simulation time of 1000 fs the charge state distribution, ion density and electron energy density were transversely averaged and extracted. From the latter, the electron temperature for polystyrene was determined using PROPACEOS (PRism OPACity and Equation Of State code)[301] and SESAME[304] data tables. The temperature was then used

as an input for the HELIOS-CR hydrodynamic code that simulated the long-term evolution of the plasma. This combined simulation was compared to two sets of HELIOS-only simulations where 44 meV, the melting temperature of polystyrene, was chosen as the initial condition. In one of the two sets, the initial 1D mass density was chosen to be the transverse average of the unperturbed colloidal crystals and in the other set homogeneous polystyrene was assumed.

As the main result of these sets of simulations a three-stage model, depicted in the middle panel of Fig. 4.5, was proposed. The Bragg peak analysis of the experiment described in [300] resulted in two distinct time scales. At first, on a *short scale* of $\sim$ 5 ps the Bragg peak scattering signal decreased sharply and the radial and azimuthal width increased significantly. On a second *longer timescale* of $\sim$ 300 ps the signal decrease and peak width increase had slowed down. In the proposed model, the IR laser first ($T_1 < 1$ ps) creates a hot, dense plasma in the surface layer of the colloidal crystal which then propagates downstream. Since the highly ionized layer width only measures a few 10 nm and the moderate laser pressure at $\sim 10^{14}$ W cm^{-2} is not strong enough to dent in the surface, the scattering signal from the structure would not change noticeably at first. Upon the thermalization of ions and electrons after the direct laser irradiation, hydrodynamics dominate the ensuing evolution of the target. The highest temperatures and highest pressures, exceeding 100 GPa, are still at the target front where the plasma is ablated into vacuum. However, there is also a lower pressure shock front that already moves into the target. During $T_2 \sim 5$ ps the high pressure front quickly ablates the target further and causes the quick drop in scattering signal. It moves along the beam propagation axis until it catches up with the lower pressure shock front and increases its speed. At this point in time the energy of the shock has already been depleted such that the disruption of the colloidal crystal structure is not as severe, anymore. For $T_3 \sim 300$ ps the shock keeps on propagating and exponentially loses energy until it cannot surpass the resistance of the bulk material any longer and is reflected. With higher laser intensities the ionization degree, plasma temperatures and pressures increase. This results in more ablation, longer shock wave propagation and greater stopping distance. The shock propagation speed and maximum mass propagation speed were found to be on the order of 5 km s^{-1} and 2 km s^{-1} and the resulting changes in the projected scattering signal are in good agreement in both simulation and experiment.

4.2. Explorative Simulations of Enhanced Ion Acceleration from Ultrathin Foils with Optimized, Realistic Laser Pulses

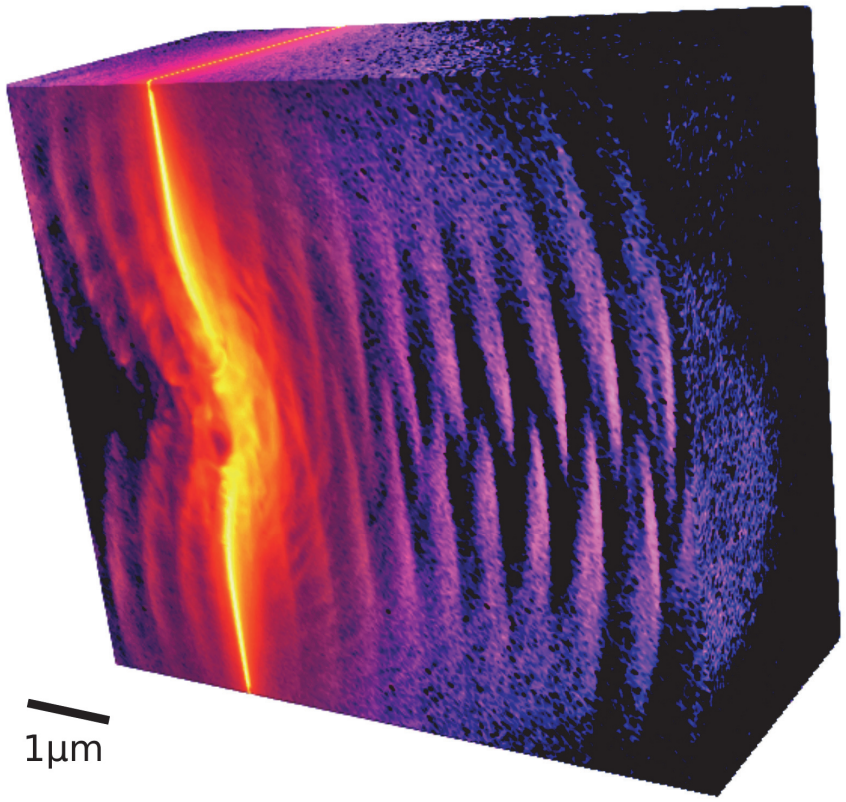

Figure 4.8.: Volumetric rendering of electron density (in arbitrary units) from a Gaussian laser pulse of $a_0 = 63.25$ interacting with a 32 nm copper foil with rear-side organic contaminant at $t = 30$ fs after the laser maximum arrived at the target front surface. Additional parameters the campaign surrounding this simulation encompassed are summarized in Tabs. 4.1, and 4.2.

This section contains the main results of this thesis on the influence of leading edge intensity ramps of the UHI laser pulse on the acceleration of ions from ultrathin foils. These were obtained in a large-scale simulation campaign on the Swiss supercomputer Piz Daint hosted by CSCS. These simulations were enabled by the highly scalable numerical and I/O methods described in chapter 3, the extension of PIConGPU to solid-density plasmas with collisional ionization that also lead to the results described in the previous section and, of course, preparatory work and support from the maintainers of PIConGPU at HZDR.

Starting with a brief reminder why the contrast dynamics of the last picosecond are especially interesting for laser ion acceleration, the simulated targets are described and the campaign parameters are presented. The general influence of the intensity ramps is discussed for different target materials and the results of preliminary simulations in 2D are shown. During the main campaign, plastic targets behaved rather like expected whereas especially metal targets showed an unexpected increase of proton energies by 50 % due to injection dynamics that transport protons from the target front side to the point and time of highest accelerating fields in the rear-side sheaths. This mechanism is extensively explained and supported with synthetic diagnostic output of PIConGPU. The well-known maximum energy scaling model of Schreiber et al. [44] is then adapted to better represent proton energies with respect to their point of origin in the sheath. Recent experimental results obtained at the DRACO laser at HZDR

are also discussed in the light of the results obtained in this study and conclusions are drawn about how to proceed with this new understanding of the last picosecond intensity ramp.

4.2.1. Motivation

In laser-ion acceleration from foil targets, already the earliest models have shown that the highest overall ion energies are to be expected from especially thin targets. Priorly in 2.2.1, it was motivated that the compromise between the number of directly accelerated electrons at the front and the rear-side sheath geometry strongly indicates that this optimum emerges for target thicknesses close to the skin depth L_s or low multiples of it.

Thanks to recent advances in plasma mirror PM technology[305–307], it has become feasible to shoot such targets without destroying them long before the main pulse maximum arrives. Employing a single plasma mirror typically improves the laser contrast by 1 to 2 orders of magnitude[138]. As laser development in recent years has inexorably pushed the maximum laser intensities past their previous limits, the nanometer-scale targets are suspected to become sensitive to other features of the temporal laser contrast. Even with the deliberate elimination of defined nano- and picosecond prepulses, a characteristic temporal contrast shape remains which originates in spatiotemporal couplings along the full laser– and amplifier chain and is partly permitted onto the target despite PM operation. For nanometer scale targets, the last picosecond intensity ramp conceivably gains significant importance towards the overall acceleration performance as the leading ramp that already surpassed the ionization threshold is expected to shape the target pre-plasma conditions for the main pulse arrival. So far, the influence of this part of the laser on the whole acceleration process has never been systematically studied before in the regime of 10 J-class PW short-pulse laser-solid interactions.

Widely used are initial conditions for particle-in-cell simulations of the main pulse interaction, that approximate the history of the nanosecond and picosecond laser-target interaction by setting an initial temperature– ($T_{e,0}$), and a pre-plasma (sometimes also post-plasma) density distribution often characterized by one or two pre-plasma scale-lengths.While the region of validity for PIC-simulations lies in the 100s of femtoseconds around the main peak maximum (2.3.1), the detailed example in 3.2.1 clearly shows the computational cost involved with a 3D or even 2D laser-solid simulation that fully resolves the plasma dynamics. Initial temperatures and pre-plasma scalelengths can be obtained from, e.g. MHD simulations or laboratory measurements. However, the regime of 10^{16} to 10^{18} W cm^{-2} is questionable to treat with either MHD or PIC due to the increased importance of collisional effects while at the same time non-local transport mechanisms are already at play. For solid targets and these intensities the local equilibrium and local transport assumptions of hydrodynamic simulations do not hold. At the same time, this regime belongs to the field of *warm dense matter* where a multitude of collisional and atomic effects are occurring which are not inherently covered by the particle-in-cell method. Therefore, the accuracy of initial assumptions about plasma temperature and density close to the pulse maximum in PIC-simulations remains problematic, but is often largely inaccessible by any other means.

In section 2.2.2, the known influences the temporal laser pulse shape can have on ion acceleration performance were already briefly reviewed. It became apparent that the rather steeply increasing intensity ramp on the last picosecond prior to the arrival of the pulse maximum is still poorly explored; experimentally – due to the lack of precise control and knowledge over the real laser contrast at full intensity, and in simulation – due to the increased time-to-solution at already immensely expensive computational demand. Existing 2D simulation work reflects the whole parameter space only sparsely and suggests contradicting trends[40, 42]. Therefore, the author performed the first ever large-scale, fully 3D simulation campaign that includes con-

trast features of the last picosecond of laser-solid interaction into the full laser-ion acceleration process.

4.2.2. Last Picosecond Intensity Ramp

To study the influence of the intensity ramp of the laser pulse on the last picosecond before the intensity maximum arrives at the cold target front surface position (from now $t_{rel} = 0$), the relevant intensity scales and commonly occurring ramp shapes had to be identified. Contrast measurements at the DRACO short-pulse laser system[G1] in 2016 revealed two characteristic ramp regions. Figure 4.9 displays idealizations of a variety of ramp shapes that were observed. Typically, the shortest UHI pulses are ideally created from a laser spectrum with a *flat phase*.

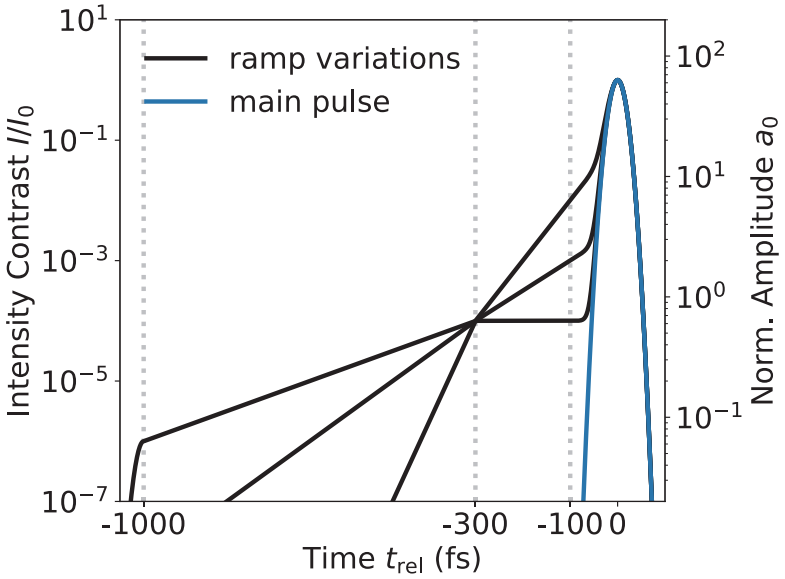

Figure 4.9.: Idealized pulse shapes for preliminary 2D simulations preceding the 3D large-scale campaign on Piz Daint. The study combined intensity ramp segments in two stages with 3 contrast values each and in all permutations. At $t_{rel} = -1000$ fs the curves start at a contrast of $I/I_0 = 10^{-16}, 10^{-8} and 10^{-6}$. All meet at 10^{-4} at –300 fs and continue to contrast levels of $I/I_0 = 10^{-4}, 10^{-3} and 10^{-2}$ where they merge into an FTL Gaussian main pulse with 30 fs FWHM. For the large-scale 3D simulations only the second stage was used in an updated variant (See Fig. 4.10b).

This means a top-hat-like distribution of the phase φ with respect to each spectral component ω within the measurable window. The resulting temporal shape (i.e. the Fourier transform) of such a distribution has distinct pre– and post-pulses. However, since the decisive quantity for plasma heating and the following acceleration is the energy flux on average, and to not introduce additional features into the studies of the ramp, these pulses have been replaced by plain exponential functions. Modifications of the spectral phase terms, namely GVD and TOD, with Acousto-Optic Programmable Dispersive Filters (AOPDF change the laser contrast shape and average energy flux by introducing chirp and asymmetries. Modifying GVD and TOD numerically on an idealized spectrum shows that the distinct pulses can also be suppressed on either the rising or falling ramp of the pulse while making them more pronounced on the other side. It is important to note that these simple manipulations *do not directly translate* to the temporal pulse evolution in laboratory experiments where the phase term changes are coupled, may have higher-order effects and be subject to change throughout the rest of the

laser chain. Determining the exact pulse shape in the focal plane of a fully amplified UHI laser pulse remains a highly difficult undertaking to this day.

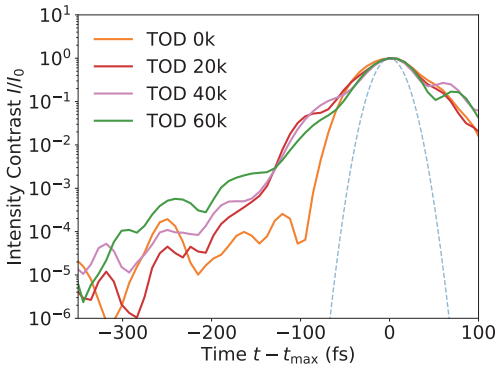

(a) Examples of measured intensity contrasts of selected, consecutive shots at the DRACO laser[G1] using the SRSI-ETE technique[308]. The dashed blue line marks an ideal 30 fs FWHM Gaussian. With changes to the 2nd and 3rd order phase terms, GVD and TOD, applied via Acousto-Optic Programmable Dispersive Filters (AOPDF, here *Mazzler* and *Dazzler* from Fastlite[309]) the temporal evolution of the laser pulse can be shaped. The broadening of the peak is introduced both by the changes to the TOD, which can partly be corrected for with careful GVD adjustments, as well as the phase reconstruction from the SRSI-ETE measurement which had a temporal resolution of $\Delta t_{res} = 50$ fs.

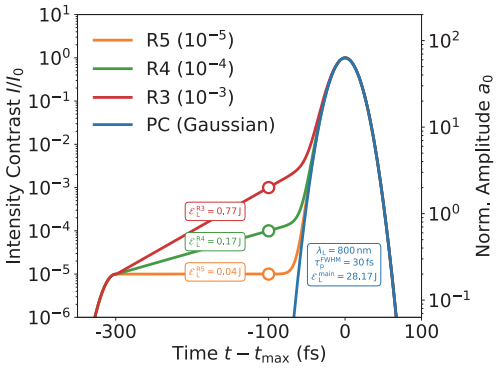

(b) Temporal pulse shapes used for the simulations in this study. Identified from commonly occurring intensity ramps, they were simplified and used as input for PIConGPU. The different pulse shapes reach relativistic intensities at different points in time: –51.82 fs (PC), –53.20 fs (R5), –60.76 fs (R4) and –159.52 fs (R3). The amount of energy, calculated from the intensity envelope, that the ramps add to the 28.17 J main pulse are 0.04 J (R5), 0.17 J (R4), and 0.77 J (R3), respectively. Since the pulse broadening in Fig. 4.10a is likely exaggerated due to the relatively low temporal resolution of the reconstruction, the main pulse is still modeled as a 30 fs Gaussian. Evidence for better compression in reality is given with the SPIDER measurements ($\Delta t_{res} = 5$ fs) in [10].

Figure 4.10.: Last picosecond intensity contrast for 3D simulations on Piz Daint.

In the results of the 2D pre-runs (see Sec. 4.2.4) it became apparent that the first 700 fs of the contrast curves shown in Fig. 4.9 were much less significant than the next 300 fs. After a scheduled maintenance on the DRACO laser where degraded optical components were exchanged and the total laser energy output saw an increase of almost 100 % new contrast measurements with SRSI-ETE technique[308] directly before the last *off-axis parabolic mirror* (OAP) led to the updated intensity ramps shown in figure 4.10.

Three different temporal contrast settings in addition to the purely Gaussian pulse were investigated. Figure 4.10b illustrates the simulated settings that all share the same Gaussian main pulse but the three settings of imperfect contrast feature an intensity ramp starting at $t_{rel} = -300$ fs. Each intensity ramp exponentially interpolates between the starting contrast of 10^{-5} and the contrast levels 10^{-5}, 10^{-4} and 10^{-3}. For the remainder of this manuscript, the following shorthands will be used for the laser pulse settings from best (no ramp, Gaussian pulse) contrast to most energetic ramp case: *PC, R5, R4, and R3*. The main pulse is $\tau_L = 30$ fs long, has a $w_L = 3$ µm spot size (both intensity FWHM), at the central wavelength of $\lambda_L = 800$ nm. The a_0 was varied between 20, 36.5 and 63.25 which also covers the (ideal) range of laser intensities delivered by the two arms of the DRACO laser that offer 100 TW and 1 PW operation. As such, the results obtained in this simulation campaign describe the physics occurring with 3 J and 30 J class short-pulse laser systems. Although characteristic representatives of consecutive shots with TOD changes were chosen for Fig. 4.10a, the exact shape of the real intensity ramps underlies a strong shot-to-shot fluctuation.

4.2.3. Foil Target Setup

Figure 4.11.: Schematic representation of target types. *Left:* copper foil with only rear-side contaminant. *Middle:* copper foil with organic contaminant on both sides. *Right:* "mixed" target like the LCT

Parameter	Copper	LCT	Formvar
Composition	Cu + organic contaminant $(C_nH_{2n})_x$	8CB – $(C_{21}H_{25}N)_x$	Polyvinyl (formal + alcohol + acetate)
Total n_e (n_c)	1414	192	229
Laser a_0	20, 63.25	63.25	36.5
Target thickness d (nm)	8 to 301	10 to 221 (736)	37 to 309 (617)

Table 4.1.: Physical properties for the three different target setups that were run in the PRACE campaign at Piz Daint. The majority of simulations focused on the the copper targets where the beneficial effect of section 4.2.5 was observed. The liquid crystal targets (LCT) as well as the Formvar foils have already been successfully used in experiments performed at HZDR (see e.g. [124]). Target thicknesses in parentheses did not run to completion and only results from early times are available (see 3.2.2 for discussion of simulation stability at extreme scales). Reduced data and selected raw data was archived on the HZDR tape storage with IDs listed in Tabs. C.1 and C.2 of the appendix.

Copper Targets Figure 4.11 shows the general simulation setup of the simulations performed on Piz Daint. The laser pulses are incident under $\theta_L = 0°$ and irradiate metal and hydrocarbon foils of varying thickness d. Table 4.1 shows an overview of the physical parameters that were investigated with the three target types, copper, Formvar and LCT. The considerations in 3.2.1 showed that a reduction in simulation size and resolution had to be made to fit the problem to the Piz Daint supercomputer. Since the optimum target thickness d_{opt} for maximum proton energies was expected to be on the order of the skin depth L_s, only ultrathin, submicron foils were simulated.

At a drive laser wavelength of 800 nm, Copper has a total electron density of $n_e = 1414\, n_c$ and was chosen to represent high-Z metal targets. With its heavy ions and high density it was expected to withstand the early radiation pressure and also reach relativistic transparency only very late in the interaction, if at all. As is illustrated in 4.11, even a single-element metal foil is in reality a compound target where the copper serves as the main source of electrons and carrier of the proton source layers. For the organic contaminant a CH-compound similar to Paraffins was configured. The latter have a mass density of 0.88 to 0.92 g cm^{-3} at room temperature, which is close to the density of liquid water. In many organic compounds, like e.g. alkanes (C_nH_{2n+2}), hydrogen is about twice as abundant as carbon atoms. With the following simple

calculation

$$n_{\text{ions,cont}} = 3n(\text{C}) = \frac{\rho_{\text{cont}}}{M_{\text{cont}}} N_{\text{A}} = \frac{0.9\,\text{g cm}^{-3}}{(12+2)\,\text{g mol}^{-1}} \cdot 6.022 \cdot 10^{23}\,\text{mol}^{-1} = 22.22\,n_{\text{c}} \tag{4.1}$$

$$\implies n(\text{C}) = 7.41\,n_{\text{c}}, \quad n(H) = 14.82\,n_{\text{c}} \tag{4.2}$$

The final electron density for the contaminant upon full ionization is then $6 \cdot n(\text{C}) + 1 \cdot n(\text{H}) = 59.28\,n_{\text{c}}$. TPS traces in experiments hint at the presence of oxygen in the contaminant (e.g. from water) but was neglected due to its similarity in charge-to-mass ratio and ionization dynamics compared to carbon.

A recent work by Sommer et al. [310] substantiates the findings of previous authors[311, 312] that the hydro-carbon contaminant layer only measures very few nanometers. When now the thickness d_{transp} is considered upon which the target turns transparent (see 2.2.1) it becomes apparent that d_{transp} as well as the proton source layer are on the same order of magnitude. Using the relation $d_{\text{transp}} = \lambda_{\text{L}} a_0 n_{\text{c}} / \pi n_{\text{e}}$, the copper target is predicted to become transparent at 3.60 nm (11.39 nm) for a laser a_0 of 20 (63.25). If the critical electron density n_{c} is replaced by the relativistic critical density $\gamma_{\text{e}} n_{\text{c}}$ and the electron energy scaling of Kluge et al. [74] (Eq. 2.9), $\gamma_{\text{e,hot}} = T_{\text{e}}^{\text{hot}} = \frac{\pi}{2\mathcal{K}(-a_0^2)} - 1$, is employed, d_{transp} becomes 25.93 nm (205.03 nm), respectively. Since the interaction is mainly in the TNSA regime, the accelerated protons are expected to originate in the rear-side contaminant but with decreasing foil thickness the 3 nm thin contaminant might increasingly influence the overall plasma dynamics.

In reality however, it is unclear if the nanosecond ASE, the picosecond pedestal or distinct pre-pulses already ablate the laser-facing contaminant layer before it can contribute to the main interaction. Therefore, two scenarios were extensively tested: one *with* and one *without* the front-side hydro-carbon layer. Later on, the front-side and rear-side originating protons were configured to be in two separate ion *species*. This allowed for separate in-situ diagnostic outputs and an immediate way to account the proton dynamics to their point of origin. As described in 3.1, using this technique sacrifices more GPU memory and with that, possibly endangers simulation stability. Hence, the separation in H_{front} and H_{rear} was only used in the latest simulations to exactly pinpoint the physical dynamics around the optimum observed target thickness d_{opt}.

Ultrathin copper nanofoils are also in reach of recent material science production capabilities[313]. Furthermore, past experience with SAXS and resonant SAXS, where the X-ray photon energy of LCLS was tuned to an inner-atomic transition, in recent pump-probe experiments[G13] make them a promising candidate for future studies concerning their heating and proton-acceleration performance.

Hydrocarbon Targets In addition to copper, two lower density targets were simulated to study the behavior of materials that likely become transparent earlier and are homogeneously abundant in hydrogen. The first is the liquid crystal 4-octyl-4'-cyanobiphenyl (also called *8CB*) which is a material that has performed very well in the past as both a renewable plasma mirror[306] as well as a a target for laser-ion acceleration[124, 314]. The liquid crystal target (LCT) molecules are composed of chains represented by the chemical formula $(C_{21}H_{25}N)_x$. Upon full ionization, the target has a density of $192\,n_{\text{c}}$ which is about one seventh of the copper density. It offers a solution to one of the major challenges in laser-ion acceleration, i.e. achieving a high repetition rate of shots as is required by some of the future applications. The material can be dispensed from a valve and drawn out as a thin film across an aperture by a movable knife-edge or wiper. By adjusting the flow and wiper speed as well as material temperature, a precise and repeatable selection of film thickness is possible which allows for the adjustment of the target in the vacuum chamber without breaking the vacuum first.

Secondly, a plastic target was simulated in the interaction with a 10 J laser pulse. The material with $229\,n_c$ is made up of *Formvar*, a composition of three organic substances: polyvinyl formal (PV1), polyvinyl alcohol (PV2) and polyvinyl acetate (PV3). Their relative molecule abundances are $(PV1\,|\,PV2\,|\,PV3) = (412\,|\,63\,|\,113)$. From this results an average atomic content of $(H\,|\,C\,|\,O) = (7.19\,|\,4.49\,|\,1.89)$ per molecule and number density ratios of (14.59 %|9.11 %|3.84 %) with respect to total electron density.

In contrast to the LCT above, the handling of such film targets requires no intricate target apparatus and the production process is cheap and fast[315]. While they lack the repetition rate of LCT, they have become a regularly used target in laser ion experiments at HZDR, especially in studies aiming to obtain the optimum temporal pulse shape by variation of second (GVD) and third order dispersion (TOD)[10]. As such, full 3D PIC simulations for all of the above three target types have very active relevance for ongoing research in laser-ion acceleration.

Technical realization of the setup Table 4.2 displays the extensive technical details of the simulation campaign. Of the 5704 compute nodes with each a single NVIDIA Testa P100 GPU available on Piz Daint only about half was available for use at the same time in one job. This cut the possible simulation volume for distribution to the compute resources in half.

As was described in detail in 3.2.2, foil targets are unfavorable in terms of simulation load balancing and especially with several thousand nodes, the risk of crashes is high. In the following the considerations for *domain decomposition* of the 3D volume are presented on the example of the copper target cases. All the particles are located on a very thin plane compared to the rest of the simulation volume and initially, the compute problem is field-dominated. As the laser pulse starts to ionize the foil, more macro-particles are produced and the load imbalance increases.

As per traditional convention in PIConGPU, the laser propagates along the y-axis and target foil extends in an $x-z$-plane. With the amount of available total GPU memory of $2400 \times 16\,\text{GB} = 38\,400\,\text{GB}$, a simulation volume of $(10 \times 11 \times 10)$ µm was chosen as a compromise between available ion acceleration distance and transverse space for resolving the laser spot of $w_0 = 3\,\mu\text{m}$ without running the risk of possible boundary effects. To lower the load for the GPUs that carry the target initially, it is placed near a GPU layer boundary. Furthermore, the simulation domain is decomposed as such, that as many GPUs as possible hold the initial target. In the beginning, a full *"spaghetti"-setup* was considered (meaning, each GPU extends along the full y-direction. Unfortunately, the *surface-to-volume* ratio grows for such a distribution and increasing amounts of device-memory are contributed only to communication purposes between the GPUs. This would have reduced the possible physical simulation volume. Since the laser-heated plasma expands mostly along the target normal directions, in first order approximation, the number of particles per transverse patch should stay constant. A division of the laser propagation and primary target expansion direction into three GPU layers was selected.

When looking at the combined memory requirements of particles and cells in tab. 4.2, the largest setup would only require about 20 % of the available memory. However, that number is still misleading due to a much heavier load close to the initial target and also due to the issue of memory being reserved for a *full frame per species and supercell* whenever a new species of particles enters a supercell (see Section 3.1.1). So as particles begin to spread heavily through the simulation volume, the GPU memory is being used up in a way that is challenging to plan for but the most critical regions are still the ones being handled by the GPUs in the ellipse around the laser propagation axis where plasma accumulates due to ponderomotive displacement and radiation pressure.

For these reasons, the simulations of the overall thickest targets of micrometer scale (compare with values in parenthesis in tab. 4.1) could not run to completion and crashed in-between.

Parameter	Copper	LCT	Formvar
Box size (µm)	$(10 \times 11 \times 10)$	$(10 \times 10 \times 10)$	$(10 \times 20 \times 10)$
Cell size (cubic, nm)	2.667	2.667	4.402
No. of cells ($x \mid y \mid z$)	(3840 | 4272 | 3840)	(3840 | 3744 | 3840)	(2400 | 4544 | 2240)
No. of GPUs ($x \mid y \mid z$)	(20 | 3 | 40)	(20 | 3 | 40)	(20 | 6 | 20)
Macro-particles per cell	26 (e), 10 (H), 5 (C), 1 (Cu)	23 (e), 2 (H), 2 (C), 1 (N)	32 (e), 4 (H), 4 (C), 4 (O)
Simulation time (fs)	310, 620	310, 620	300, 510
Time step (as)	5.134	5.134	8.474
No. of time steps	$6 \cdot 10^4$, $1.2 \cdot 10^5$	$6 \cdot 10^4$, $1.2 \cdot 10^5$	$3.5 \cdot 10^4$, $6 \cdot 10^4$
No. of particles	$2 \cdot 10^9$ to $5 \cdot 10^{10}$	$2 \cdot 10^9$ to $1 \cdot 10^{11}$	$3 \cdot 10^9$ to $5 \cdot 10^{10}$
Total memory: particles (GB)	650 to 2400	1000 to 3600	700 to 2000
Total memory: cells (GB)	4600	4000	2000
Run time	4 to 12 h	6 to 12 h	6 to 18 h

Table 4.2.: Simulation parameters for the three different target setups that were run in the PRACE campaign at Piz Daint.

Even much thinner targets were observed to become unstable once the foils were exploded by the laser and more particles were pushed outside of the central focus region. Nevertheless, the large majority of simulations ran up to a point where the ion acceleration process was largely completed and all the physics dynamics could be extracted.

Choice of Initial Conditions

Based on the 2D pre-studies it was determined to focus the fully 3D simulations on the last 300 fs intensity ramps. To study the interaction following a natural ionization and pre-plasma

formation, no initial pre-plasma gradient or electron temperature was assumed. This is reasonable given that there are neither pre-pulses nor a plateau phase in the laser contrast that deliver intensities above the ionization threshold causing expansion over a multi-picosecond duration. Additionally, two-color probe measurements of the expansion dynamics at the much less dense (30 n_c) cryogenic hydrogen jet target at HZDR confirmed that with a plasma mirror no ionization occurs prior to $t_{rel} = -700$ fs[316, G14].

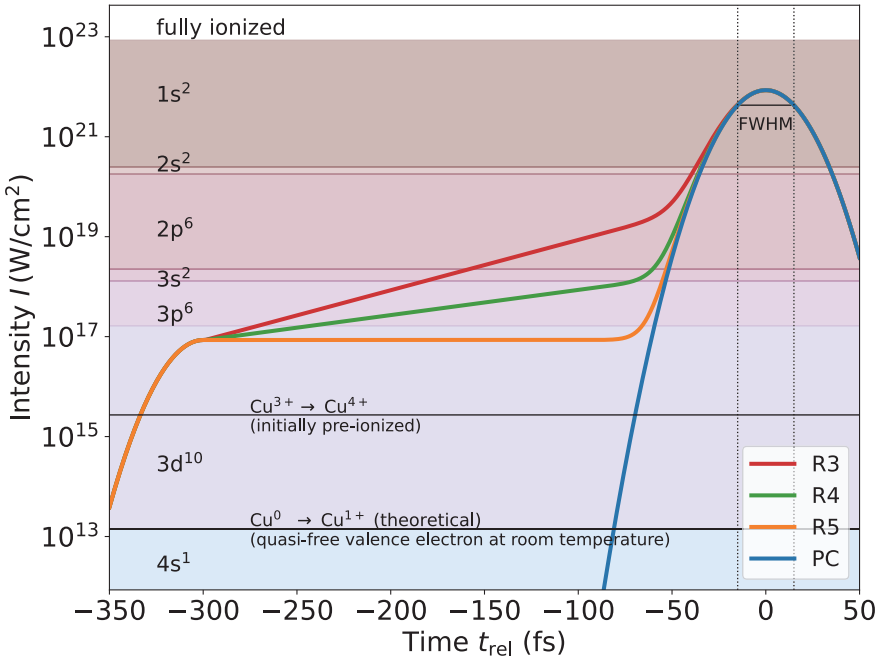

Figure 4.12.: Prediction for the ionization of copper following the BSI model. The four different pulse shape setups are depicted with laser intensity versus time. The colored areas represent the shells of the ground state configuration of copper [Ar]$3d^{10}4s$. Transitions between them are at the appearance intensities (see 2.1.4) of the last bound electron in that respective shell.

With the start of the simulation the laser intensity quickly increases to $10^{-5}\, I_0$. At this intensity level the BSI model predicts that most of the copper 3*d*-shell charge states are saturated. The ionization prediction for respectively the first and last electrons in each shell is shown in Fig. 4.12. By the same reasoning the lowest charge states of hydrogen, carbon, nitrogen and oxygen in the other target setups are also pre-ionized. Additionally, the ionization prediction of the Thomas-Fermi model as well as the 0D atomic physics simulation suite FLYCHK[184] are depicted by figure 4.13. Even at low temperatures, an average charge state of $\langle Z \rangle = 4.41$ results from this model. Copper at room temperature has one quasi-free valence electron within the conduction band, but both the temporal intensity contrast as well as the TF model, after surpassing the low-energy threshold introduced in section 4.1.2 is surpassed, justify a pre-ionization to Cu^{4+}. This saves additional resources for particle memory since multiple, initially free, electrons need to start at the same position to ensure quasi-neutrality and can hence be combined to a single macro-particle with higher weighting.

Subsequent field ionization is modeled via a combination of both the tunneling ionization ADK model with the classical BSI model, to ensure a natural ionization progression.

Another important consideration to take is always that the laser pulses are initialized with consideration of the temporal ionization prediction of the target (compare Fig. 4.12). There-

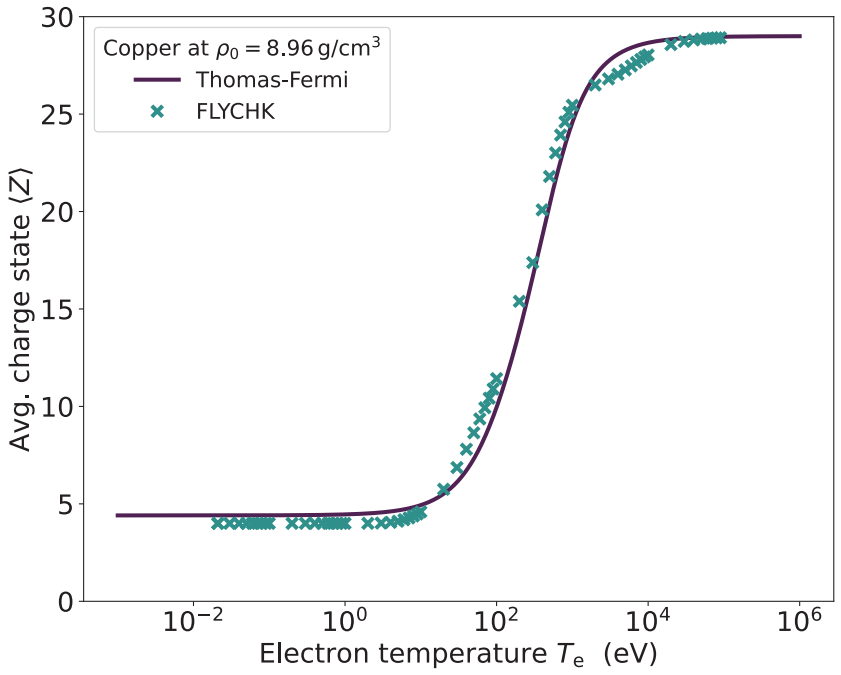

Figure 4.13.: Average charge state prediction for copper ions at solid density of $\rho_0 = 8.96\,\mathrm{g\,cm^{-3}}$ using the Thomas-Fermi collisional ionization model[90] (solid line) that is implemented in PIConGPU. Even for an electron temperature of $T_e = 25\,\mathrm{meV}$ (room temperature), the predicted average charge state is already $\langle Z \rangle = 4.41$. The 0D atomic physics code FLYCHK[184] shows very similar predictions (x-markers). While the prediction is technically unphysical, quasi-free electrons in metals behave plasma-like and the initial state is valid for the PIC simulations.

fore, all laser pulses have been initialized before they reach the intensity that is large enough to make the pre-ionized charge states abundant, thus ensuring a physically valid evolution of free electron density.

4.2.4. Influence of the Intensity Ramp

The laser pulses in this campaign are configured such that all share the same main pulse and reach the same maximum intensity. Real changes to the laser shape based on manipulations of the phase terms broaden or shift the pulse which affects the maximum a_0. However, the ramps added here only deliver more energy to the target at earlier times at the same central wavelength. Based on the existing literature that was reviewed in chapter 2.2.2, more energy during earlier stages of the interaction usually means earlier ionization, expansion and often lower maximum energies. With increasing energy prior to the last picosecond, existing proton energy maxima have been observed to shift to larger target thicknesses (e.g. [141]).

2D Pre-study at HZDR

The early 2D studies of pulse ramps with the parameters given in figure 4.9 showed that the time interval from –1000 to –300 fs only affects the maximum energies to a few percent. Therefore, the simulations were repeated with updated ramp shapes from more recent measurements, with a focus on the last 300 fs.

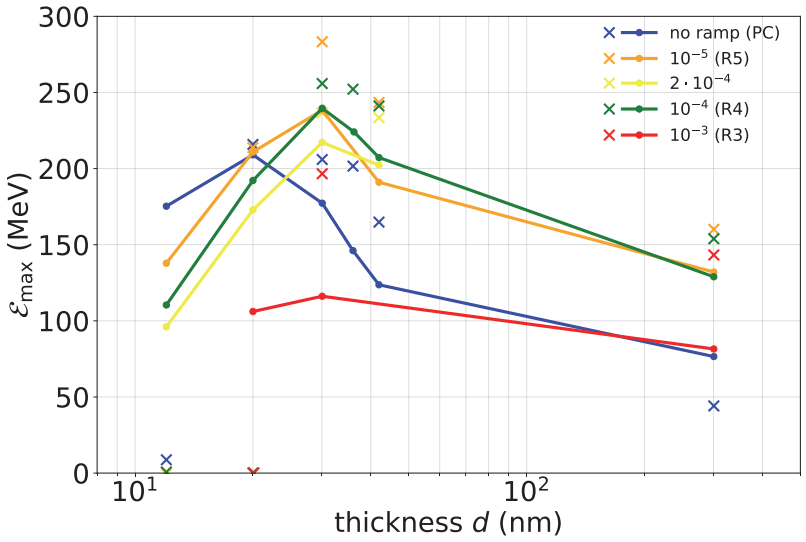

Figure 4.14.: Maximum proton energies from 2D PIConGPU simulations of copper foil targets irradiated by a 30 J, 30 fs laser pulse in different leading intensity ramp configurations. Protons from front (× markers) and rear (dots connected by lines) side contaminant layers are represented by separate particle species. Extremely low energies for front-side protons at lower thicknesses mean that the most energetic front-side protons were accelerated in backwards direction and left the box already.

Figure 4.14 shows the maximum proton energies at $t_{rel} = 180$ fs from 2D PIConGPU simulations of copper targets with various intensity ramps, starting at $10^{-5}\, I_0$ and growing exponentially to 10^{-5}, 10^{-4}, $2 \cdot 10^{-4}$ and $10^{-3} I_0$. The lines connect maximum energy values of rear-surface protons. Narrow maxima at $d = 20$ nm and 30 nm appear at around 200 to 250 MeV for all cases except the R3 case where the maximum reaches only 115 MeV. The expected shift to larger thicknesses with increased ramp energy is already visible and coincides with a decrease in maximum proton energy of ramps shallower than R4. The R4 and R5 case both show the overall optimum with a 30 nm foil at almost 240 MeV, which is 30 MeV more than the PC case. Since 2D simulations usually show increased energies by a factor of about 3/2[198], the maximum energies from the copper targets in 3D are expected at $\mathcal{E}_{max} \approx 150$ MeV. Cross markers show maximum energies of front-side originating protons in these simulations. At 30 nm, the R5 case shows the overall maximum with 289 MeV, followed by the other configurations in the order of increasing ramp energy with the exception of R3, where the perfect Gaussian contrast still produces higher energies. Front-side originating protons are reaching overall higher maximum energies than the rear-side originating ones. Only for thicknesses of 20 nm and below, as well as the PC case at 300 nm, this was not observed. For the thinnest targets, the maximum front-side proton energies were recorded in backwards direction and protons quickly left the simulation box there because the vacuum region before the target was set up to be much smaller than the region behind.

Hydro-Carbon Targets

The first results from the 3D PIConGPU campaign presented here consider the effect of the intensity ramps on the observed proton acceleration performance from hydrocarbon targets. The initial hypothesis stated that possibly more energy inside the ramp could lead to higher front-side expansion, higher absorption and thus higher proton energies. It was formed based on early SPIDER measurements of the laser pulse shape during the last 300 fs taken at the

DRACO PW laser system at HZDR, Dresden. For the chosen shape of intensity ramps, however, already the R5 case with the least energy inside the ramp (i.e., a plateau of 10^{-5} contrast) showed that increased energy in the leading edge ramps leads to lower final proton energies due to actually strongly decreased laser absorption.

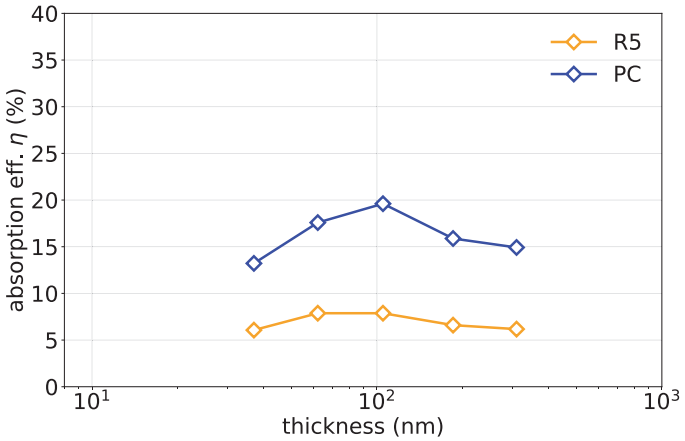

Figure 4.15.: Absorption efficiency of Formvar targets irradiated by laser pulses with a maximum a_0 of 36.5 and shapes according to Fig. 4.10b.

Formvar Targets Figure 4.15 compares the absorption efficiency of Formvar targets irradiated with laser pulses of the R5 and PC variant. The thickness was varied between 37 nm and 309 nm. It becomes immediately apparent that in the cases with perfect Gaussian (PC) contrast the laser couples into the electrons with more than twice the absorption efficiency than the ramp cases showed. The optimum of $\eta = 19.6\,\%$ was observed at $d = 105\,\text{nm}$. Already the best contrast ramp case R5 shows only between 6 to 8 % conversion. An absorption efficiency below ten percent as well as clear signatures of the laser electric field in electron phase spaces at the time of maximum intensity indicate that the targets became transparent to the laser as a result of the leading ramp.

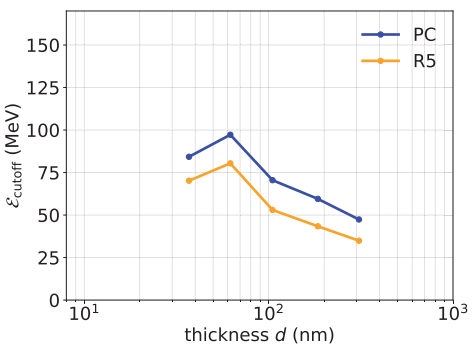

(a) Proton cutoff energies obtained at $t_{\text{rel}} = 120\,\text{fs}$.

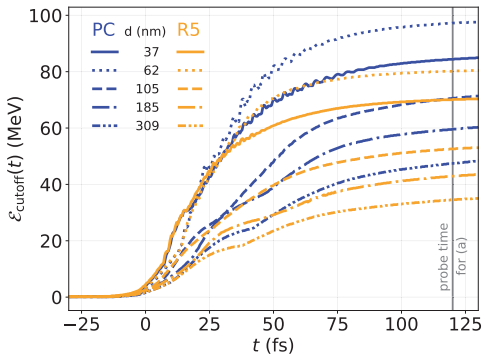

(b) Evolution of proton cutoff energies with respect to the time of laser maximum on target.

Figure 4.16.: Proton cutoff energies in the contrast cases PC and R5 from PIConGPU simulations of formvar foils irradiated by a 10 J, 30 fs ($a_0 = 36.5$) laser pulse.

Becoming transparent too early resulted in a directly adverse effect on the proton acceleration performance, as figure 4.16a shows. All simulations show an energy optimum for 62 nm

targets with all ramp cases coming out 20 MeV lower than the ones with a Gaussian drive laser pulse. The overall maximum proton energy was 96 MeV and at 309 nm it dropped to 37 MeV.

Liquid Crystal Targets The simulations with liquid crystal targets featured the same absolute grid resolution as the ones run with copper that are presented hereafter. Due to more violent plasma dynamics caused by earlier expansion of the target in the case of leading intensity ramps, several simulations did not make it past $t_{\mathrm{rel}} = 20\,\mathrm{fs}$ as the memory capacity of single GPUs was exhausted. At this point in time the total energy transferred to all particles has not peaked yet. Additionally, the proton energy is still strongly increasing and only at about 40 % or below of its final asymptotic value. A reliable comparison between the different intensity ramps and their influence on ion acceleration from LCTs was not possible within this campaign. However, in Sec. 4.3.1, bremsstrahlung emission from these low-Z targets is compared to the radiation signatures from copper foils for several runs.

4.2.5. Optimum Proton Energies from Imperfect Contrast

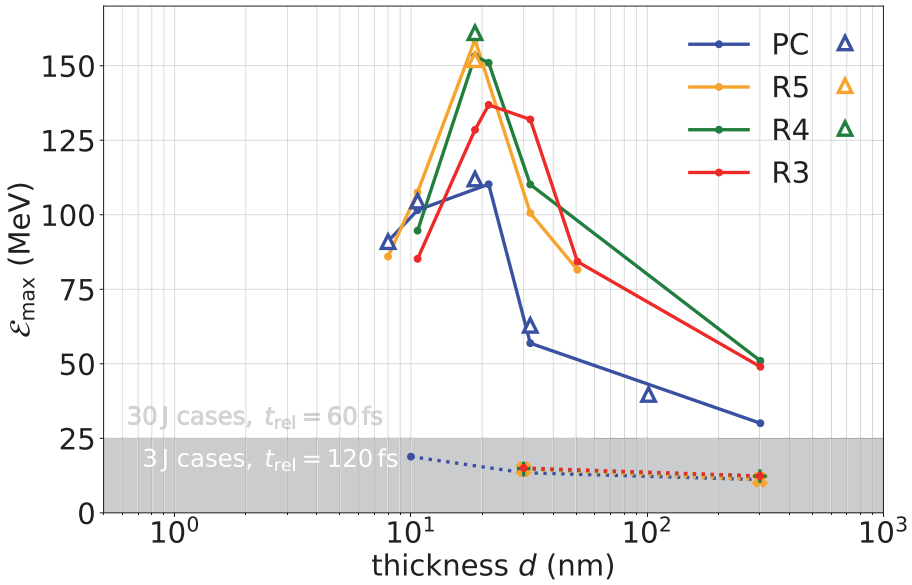

Figure 4.17.: Maximum proton energies from 3D simulations of copper foils. Simulated laser parameters were: $\lambda_{\mathrm{L}} = 800\,\mathrm{nm}$, $\tau_{\mathrm{p}} = 30\,\mathrm{fs}$, $w = 3\,\mu\mathrm{m}$ and laser maximum amplitude $a_0 = 63.25$ (20 for data points in the lower, gray shaded area). The laser contrast setting was varied according to Fig. 4.10b. Solid lines connect cases where the proton source layer was put on the rear side only. For the triangle markers both target sides carried the contaminant layer. The proton energies for all 30 J laser pulse cases were measured 60 fs after the laser maximum hit the target front side. The lower energy, 3 J, cases were measured at $t_{\mathrm{rel}} = 120\,\mathrm{fs}$.

In a thickness scan performed on copper targets with 3 J and 30 J laser pulses, a pronounced maximum in the final proton energy[3] was found for each contrast setting of the 30 J case in Fig. 4.10b. The results of the scan are shown in Fig. 4.17. Remarkably, every maximum in the ramp cases R3–R5 is both more pronounced and *higher* than in the perfect contrast case, PC. With more energy in the intensity ramp, unlike the observation on Formvar before, the

[3] Criterion: $N(\mathcal{E} > \mathcal{E}_{\mathrm{max}}) \approxeq 500$; the highest energy bin where only about 500 particles have higher energy marks the maximum proton energy $\mathcal{E}_{\mathrm{max}}$.

maximum proton energy optimum shifts to larger d_{opt} as initially expected. This shift coincides with a decrease in peak proton energy. Between the best result of the PC case and the highest overall proton energy of all the ramp cases, an increase of $\sim$ 50 MeV is gained, amounting to a factor of about 1.5.

All contrast settings show a similar target thickness dependence. From 300 to 30 nm the maximum proton energy grows roughly by a factor of 2. Then, a very narrow energy peak is observed where $\mathcal{E}_{max}$ varies strongly within the range of ten nanometers below and above d_{opt}. Towards even thinner targets the energy drops rapidly again.

Temporal evolution of the acceleration process

Figure 4.18 shows an example of the total energy in electromagnetic fields and particles in absolute units evolving over time on the example cases of $d_{opt} = 18.67$ nm copper in the PC and R4 scenarios as well as Formvar with $d_{opt} = 62$ nm in PC and R5 configuration. The EM-field data clearly shows the energy influx via the intensity ramp from –300 to –100 fs, followed by the Gaussian main pulse. In these cases, electrons gain energy very earl, roughly following the field curve. The energy stored in copper ions becomes visible in the upper row of Fig. 4.18 before hydrogen and carbon ions due to their relative abundance and comparably larger mass.

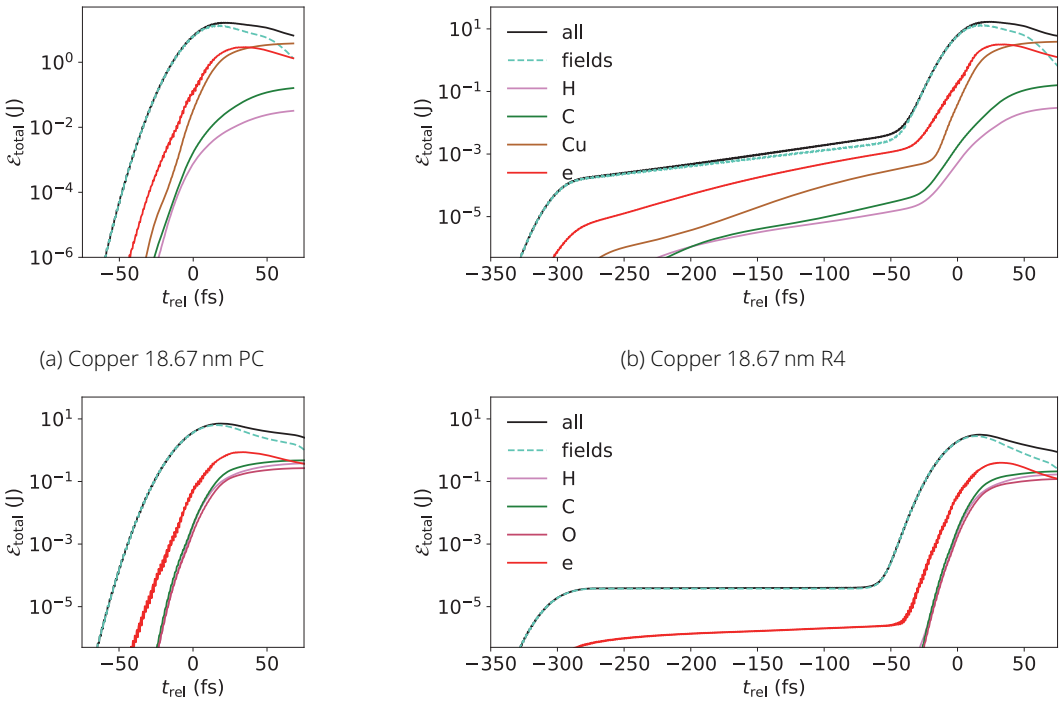

(a) Copper 18.67 nm PC

(b) Copper 18.67 nm R4

(c) Formvar 62 nm PC

(d) Formvar 62 nm R5

Figure 4.18.: Total energies of all particles and electromagnetic fields from 3D PIConGPU simulations of copper (upper row) and formvar targets (lower row) irradiated by 30 J and 10 J laser pulses, respectively. Particles are separated by species and energy values are absolute in units of Joule.

In comparison to the PC case, all simulations with intensity ramps showed the expected earlier ionization and expansion of the target. For copper targets with thicknesses of 32 nm and lower, the flat target surface was not fully intact anymore when the laser maximum arrived. Nevertheless, the majority of rear side contaminant protons always detached as a layer to be accelerated in the electron sheath. Increasing energy in the ramp or decreasing thickness

accelerated this effect as figure 4.19 shows.

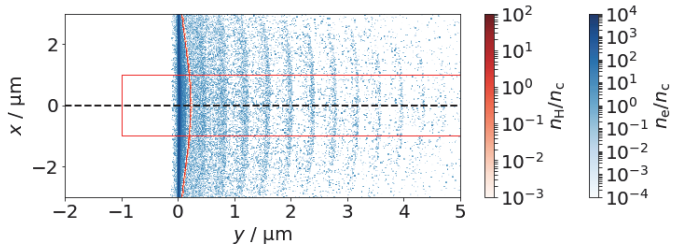

(a) PC case

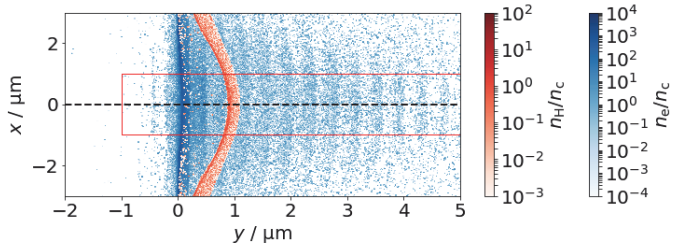

(b) R4 case

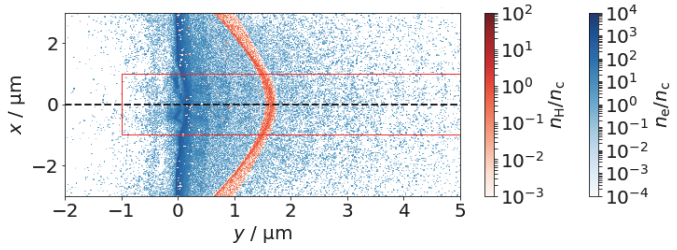

(c) R3 case

Figure 4.19.: Electron and proton x–y density slices close ($T_{rel} = T_{max} \pm 2$ fs) to the instance of maximum laser intensity on target (32 nm Cu-foil, $a_0 = 63$) for different contrast settings. Electron $2\omega_L$ are visible in all cases but their spatial separation softens with increasing ramp energy. Especially the R3 case shows disturbed foil surface integrity. The red boxes from $y = -1$ µm and between -1 µm $< x < -1$ µm mark the volume over which densities and fields are averaged later in Fig. 4.22.

Prompt electron bunches accelerated by the $\vec{j} \times \vec{B}$-force of the laser pulse are decreasingly well-defined in space as the energy content in the intensity ramp increases. Initially formed pre-plasma is swept up by the main pulse, leading to a pile-up of electron density on the laser-axis shortly prior to the arrival of the maximum intensity.

4.2.6. Discussion and Explanation of Results

Spectral shape

Figure 4.20 displays the proton spectra at the observed optimum thickness d_{opt} in the PC and R4 cases. The former shows an exponentially decreasing particle number for increasing energy and a distinct cutoff at 112 MeV where the proton number per MeV drops off by several orders of magnitude. These features are both typical for a TNSA spectrum. The ramp case R4 shows a similar exponential decrease and then also a cutoff, but at reduced energy (91 MeV) when compared to the PC case. However, the spectrum does not drop off entirely but instead continues with a tail of higher energy protons. In this tail, the particle number is also exponentially decreasing with energy. The rate of decrease is slightly steeper than in the

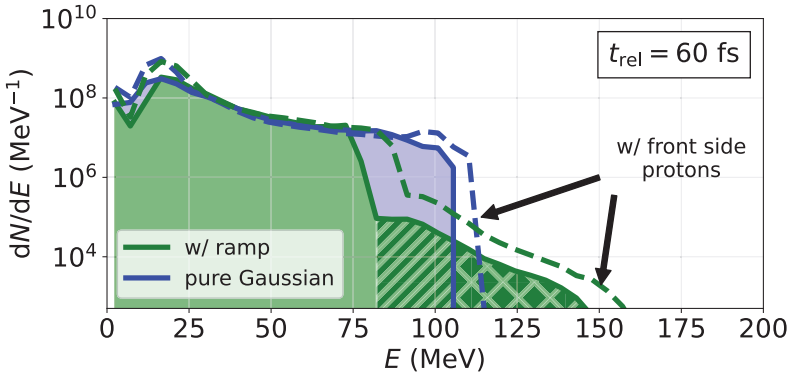

Figure 4.20.: Proton spectra from cases R4 (green) and PC (blue) at d_{opt} = 17 nm. Solid (dashed) lines represent cases with proton source layer on only the rear side (on both sides) of the target. Hatched and cross-hatched areas correspond to proton populations identified in Fig. 4.22.

lower energy part of the spectrum. All ramp cases over all thicknesses consistently show this feature with particles above the first TNSA cutoff, while the PC cases have only a negligible amount of particles above the cutoff, if at all.

Absorption efficiency Since a front-side plasma gradient has been shown to increase the distance over which the laser pulse is able to efficiently heat electrons, comparing the absorption efficiency η for all the simulations is important to solidify the picture of the acceleration physics at play. Figure 4.21 displays the absorption efficiency η versus target thickness for the

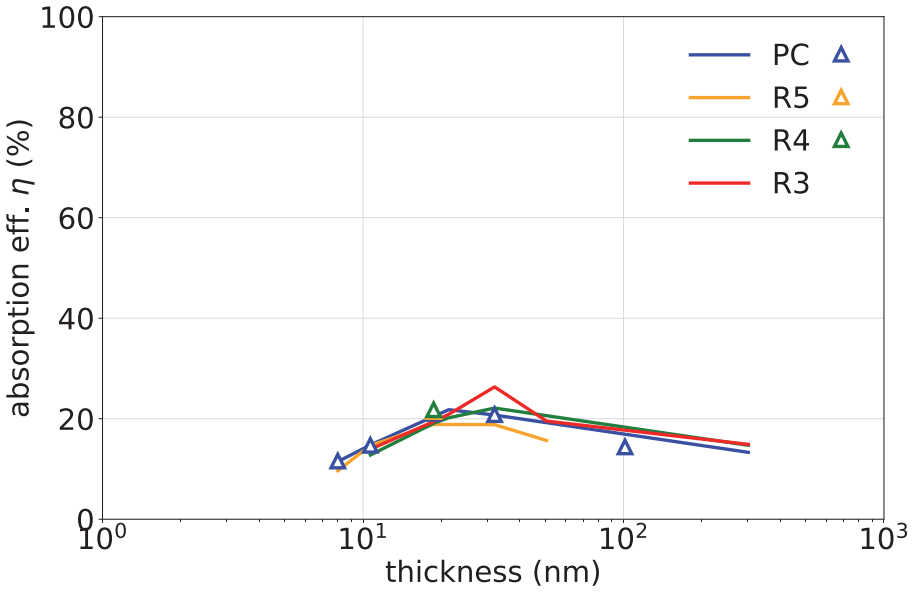

Figure 4.21.: Absorption efficiency η vs. target thickness. Similar to Fig. 4.17, solid lines connect cases with organic contaminant at the rear, while triangle markers show cases with proton source layers on both target sides.

30 J copper simulation series. The values for η were determined from the maximum total kinetic energy of all particle species in the simulation divided by the total laser pulse energy (also accounting for the ramps, see Fig. 4.10b). This maximum occurs roughly at $t_{rel} = 1.3 .. 1.5\,\tau_p$. At this point in time the laser pulse has been mostly reflected and traveled out of the simu-

lation box again. Almost all of the particles, however, are still inside the simulation volume and their combined kinetic energy represents the energy transferred by the laser. Except the single outlier of the 32 nm foil in contrast setup R3, the absorption efficiency of all ramp cases lie within an absolute range of $\Delta\eta = 2$ to 4 % from the perfect Gaussian contrast case. Overall, the absorption of laser energy into particles ranges between 9 and 27 % which is within the range of expected values for highly overdense, ultrathin foils (see 2.1.3). The optimum of η lies between 20 nm and 30 nm foil thickness which coincides well with d_{opt} for maximum proton energy $\mathcal{E}_{\mathrm{max}}$ in the PC case. However, the significant differences between the contrast settings do not show a direct dependence of η and indicate a different rationale.

Density profiles To study the origin of the differences in the high energy part of the proton spectra, the optimum thickness cases of PC and R4 concerning their accelerating field structure and evolution over time are compared. Especially the condition of maximum laser intensity on target is of interest to investigate if the increase in ion energy is related to TNSA enhanced by transparency of the target. The fields and particle charge densities, transversely averaged over a 4 µm^2 extent around the laser axis (red boxes in Fig. 4.19), are plotted against the longitudinal position in figure 4.22 for the time $t_{\mathrm{rel}} = 0$. The region is large enough to include enough protons even in less populated regions but small enough that the surface and layer curvature resulting from target expansion are still negligible.

In addition to the charge densities, the accelerating electric field component parallel to the target normal direction is displayed. Most prominently, the two cases differ in the degree of target expansion. The very steep increase in intensity in the PC case did not leave sufficient time for the target to significantly heat prior to $t_{\mathrm{rel}} = 0$, and subsequently expand into vacuum. The only brief exposure to $2\,\omega_L$-bunches of electrons created by laser amplitudes of $a_0 > 1$ at the front side caused the protons at the rear to still remain in a thin layer that has pre-expanded only slightly ($\sim$ 301 nm) by the time the laser maximum hits. In contrast to this, with realistic temporal contrast R4 the target was continuously heated during the ramp phase of the pulse and expansion is more prominent. As long as the a_0 is sufficient for the earliest ionization levels but still below unity, the target expands more or less uniformly on both sides. The earlier but steady expansion causes the charge density distribution of both electrons and ions to broaden, and thermal motion within the target leads to a mixing of previously layered target and contaminant species. As seen in Fig. 4.22, the majority of protons (solid blue/orange areas) in the ramp case has expanded significantly further into vacuum than in the PC case. It is broadened and, vitally, protons have populated the whole depth of the target. At the instant of maximum laser intensity on target, the displacement of electrons is strongest, and the resulting $2\omega_{\mathrm{L}}$ *prompt* electron bunches can be seen clearly in the electron density modulations. The electric sheath field, however, exhibits a similar shape in both cases.

One can readily see that the acceleration of the main proton population of the rear side contaminant layer is exclusively caused by the electron sheath in front of it. Hence, the accelerating field is lower for the the *main proton population* in the ramp case, since they have already expanded far at the moment of the most substantial contribution of electrons to the sheath. Consequently, they arrive at lower final energies. From now on, these protons will be referred to as *regular TNSA protons* for the remainder of this manuscript.

However, the more interesting proton population is the one spread out throughout the whole target region. Accelerating fields are strongest at the very instant when the electrons pushed by the main pulse maximum exit the target rear, surpass the ion bulk and add to the origin of the Debye sheath.

In the R4 case protons are spread throughout the whole target within the laser spot region. Therefore, protons are also present at the position and time the largest fields are created

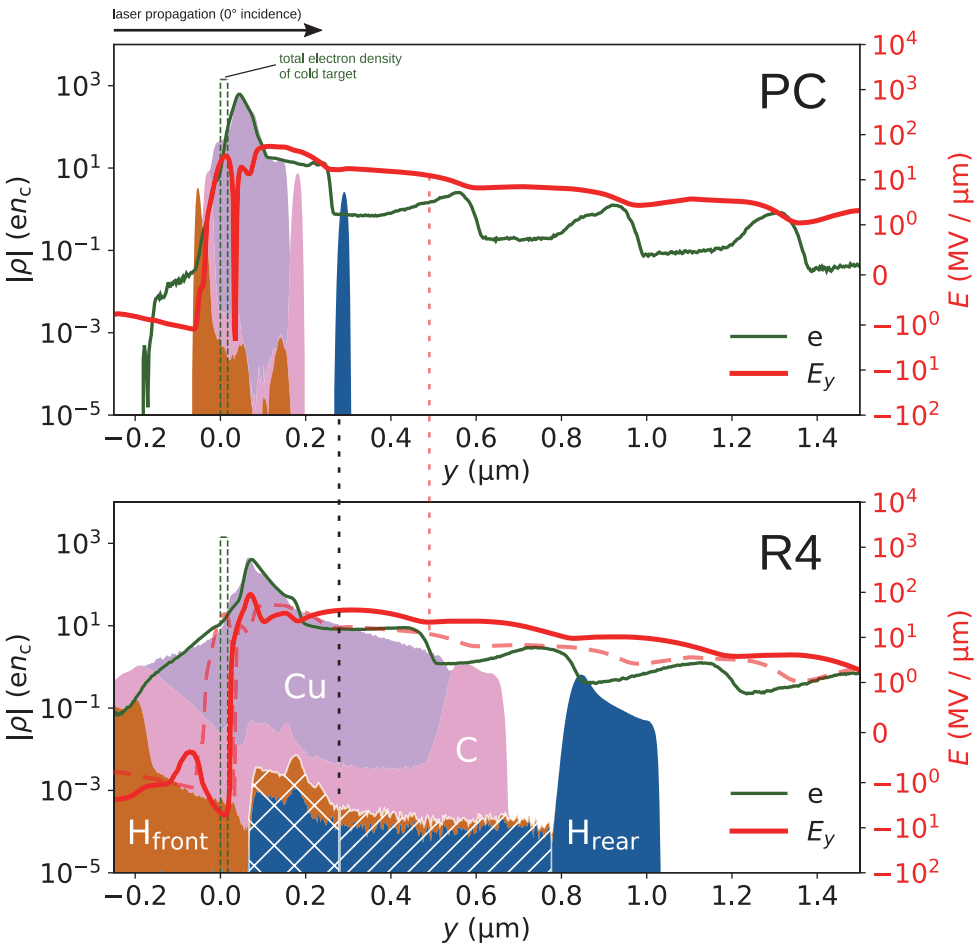

Figure 4.22.: Charge densities and accelerating field for the moment of arrival of the laser intensity maximum on the cold target front surface position, i.e., $t_{rel} = 0$ for the case of perfect Gaussian pulse (upper panel) versus best intensity ramp case, R4, (lower panel) at optimum foil thickness $d_{opt} = 18.67\,\text{nm}$. The ion charge densities are represented by the stacked, colored areas. The electron density is depicted by the dark green line. All densities and fields have been transversely averaged over a small region ($r < r_{spot}$) around the laser propagation axis. The hatched population of protons in the ramp case can be identified as the protons that are later found in the high energy tail of the spectrum. The position of the rear side proton layer in the PC case is marked in the panel of the R4 case. Protons between the origin of the sheath, identified by the maximum accelerating field, and this position (cross-hatched area) can later be found at energies higher than the maximum energy of protons in the PC case. The solid red lines show the accelerating field E_y in units of MV/µm in longitudinal direction y. For reference, the electric field of the PC case is shown as a pale dashed red line in the lower panel.

and these are then able to see and travel down the full acceleration potential. With the pre-expansion in the PC case the protons there are already behind the region of largest fields, but the capability to probe these fully in the R4 case is what leads to the higher energies

Thus, the fastest protons are the ones which are injected into the electron sheath at the right place and the right time, which will hereafter be identified by *injected TNSA protons*. At the moment displayed in Fig. 4.22, all protons in the on-laser-axis region have gained only a maximum of a few MeV of energy. Hence, almost all the energy the injected protons (cross-hatched population in Fig. 4.22) gain over the regular TNSA protons in the PC case stems from the difference in their *starting point* and the higher accelerating field there when they travel down the acceleration potential.

Accelerating Fields

The TNSA model of Schreiber et al. [44] uses the assumption of a quasi-static field structure in the Debye-sheath that the fastest ions traverse before it is subjected to significant change. For the *injected protons* to reach the overall largest energies, they need to experience sustained high accelerating fields as the field structure, opposed to the assumption before, may evolve over time. The overall size of the volumetric electric field data made it unfeasible to write these files to disk with the required high temporal frequency. Instead, the accelerating field can also be reconstructed from representatives of the proton populations of interest by taking the discrete derivative of kinetic energy with respect to position.

$$E_y = \frac{\partial \mathcal{E}_{\mathrm{kin}}}{\partial y} \approx \frac{\Delta \mathcal{E}}{\Delta y} \tag{4.3}$$

Therefore, proton particle data outputs for both the species $\mathrm{H}_{\mathrm{front}}$ and $\mathrm{H}_{\mathrm{rear}}$ in the contrast cases of R4 and PC have been created with increased temporal resolution, every 257 as.

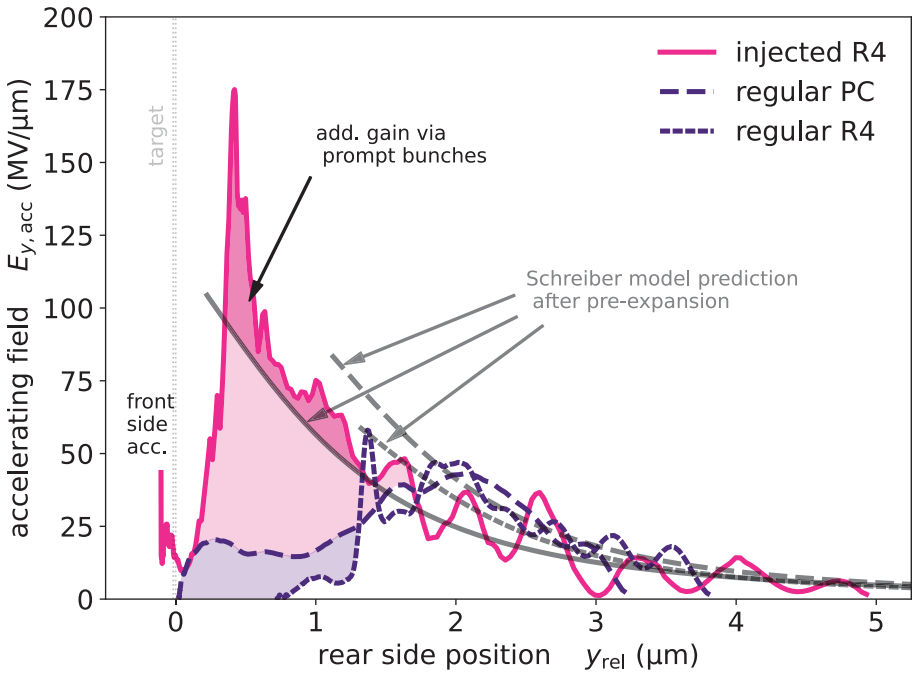

Figure 4.23.: Spatial evolution of average accelerating field $\langle E_{y,\mathrm{acc}} \rangle_y$ with respect to the position of the fastest protons, separated by their classification as *regular* (indigo) or *injected* (magenta) as defined before in section 4.2.6. The position y_{rel} is defined relative to the cold target rear surface. The magenta-colored area signifies the energy gain the injected protons of case R4 (solid line) gain over the regular protons of the PC case (dashed line) . Vice versa, the indigo-colored area amounts to the energy the regular PC protons gain over the regular R4 protons (dotted line).

The evolution of the accelerating electric field over relative position y_{rel}, with respect to the cold target front-surface, is shown in Fig. 4.23. Like mentioned before, it was obtained from the average energy gain $\Delta\mathcal{E} = \int E(y)\mathrm{d}y \approx \langle E_y \rangle_y \Delta y$ the protons of the injected or regular population experience. The highest-energy protons travel close to the laser axis that coincides with the target-normal direction. The isothermal ion energy scaling models need to be compared to these protons since the region in the center of the ion beam is the most similar to the 1D approximations from the analytical models.

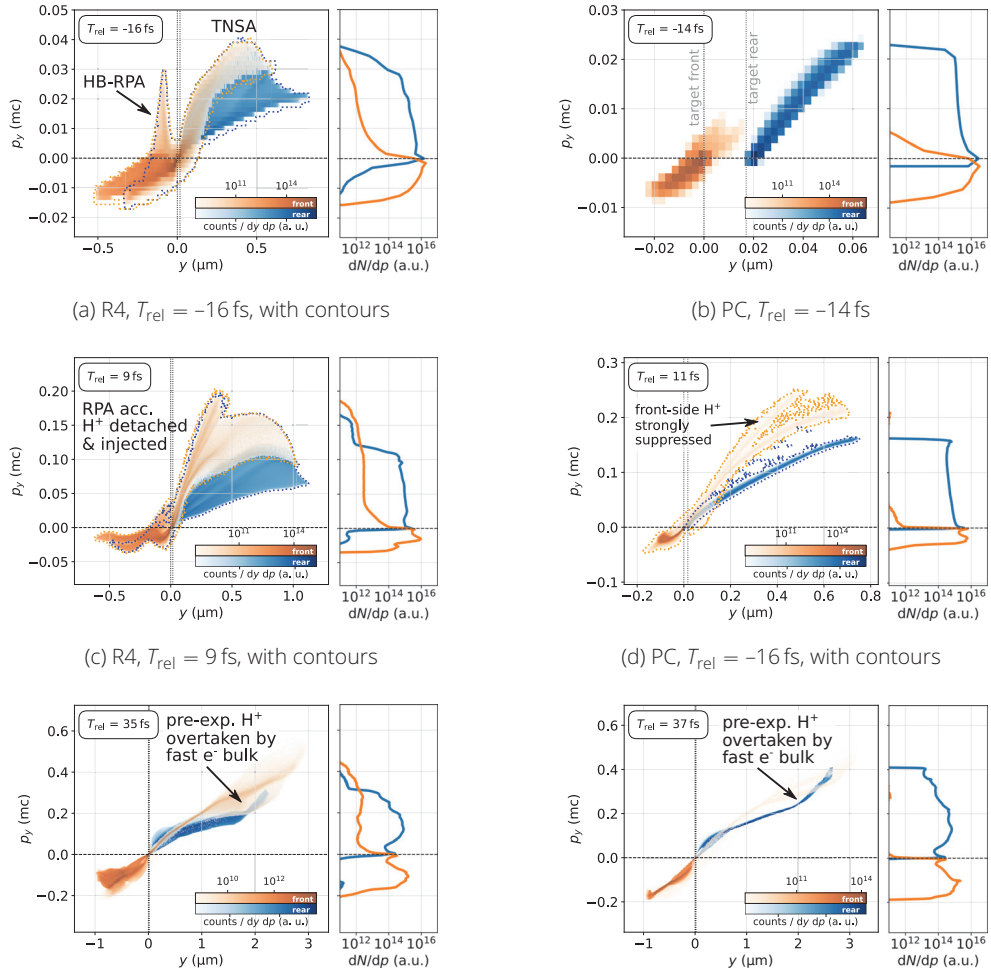

(a) R4, $T_{rel} = -16$ fs, with contours

(b) PC, $T_{rel} = -14$ fs

(c) R4, $T_{rel} = 9$ fs, with contours

(d) PC, $T_{rel} = -16$ fs, with contours

(e) R4, $T_{rel} = 35$ fs

(f) PC, $T_{rel} = 37$ fs

Figure 4.24.: **(a,c,e):** Longitudinal ($y - p_y$) phase space images of the H_{front} (*orange*) and H_{rear} (*blue*) proton species in the $d_{opt} = 18.67$ nm R4 contrast case. **(b,d,f):** PC case, respectively. Phase spaces showing the stages of acceleration are displayed on top of each other to illustrate the evolution of both species, but the partly translucent H_{front} is obscuring H_{rear} to a degree. Therefore, figures (a),(c) and (e) also show contours as guides to the eye in dotted lines around the phase spaces, separating occupied from empty phase space volume. The right side vertical figures show momentum spectra, respectively.

However, the energy spectra alone do not qualify to distinguish the different populations at all times. For this purpose, the phase space timeline in figure 4.24 shows a more detailed view of the acceleration in the R4 contrast case at $d_{opt} = 18.67$ nm. Three phases of the laser-proton acceleration process can be distinguished. Before $t_{rel} = 0$, the laser pulse compresses the pre-plasma that formed during the ramp-phase. At the surface of critical density, the $\vec{j} \times \vec{B}$-force displaces a large number of electrons against the protons populating this region. Consequently, these protons are accelerated by the hole-boring process, which transports them into the highly overcritical region of the still-intact copper foil where they detach from the laser. This concludes the *first stage* of acceleration. A sharp spike in the electric field at this time signals the hole-boring which rapidly accelerates the protons to single-MeV energies. Around $t_{rel} = 0$, the target front side has reached maximum compression. Due to that, the interaction length of the laser with the overcritical plasma is minimized, and its resistance against the radiation pressure reaches a maximum. Once the region around the critical density surface, i.e. the region the laser can most effectively interact with, is depleted, the supply of injected protons

diminishes. When the electrons that were accelerated by the peak of the laser pulse reach the rear side of the target, they create the strongest accelerating fields. Around 0.5 µm the accelerating field reaches a maximum for the injected protons (Fig. 4.23 solid magenta line). The majority of electrons accelerated by the peak of the laser pulse overtakes the injected protons in this *second phase* of acceleration. Only when these electrons also overtake the pre-expanded regular protons (dashed/dotted indigo lines), the latter experience their maximum in the accelerating field. In this *third stage* of acceleration the shoulder-like TNSA cutoff signature in the final proton spectra is created. After about $t_{\mathrm{rel}} = 2\,\tau_{\mathrm{p}}$, the protons have reached approximately 85 to 90 % of their final energy.

Contribution per Mechanism

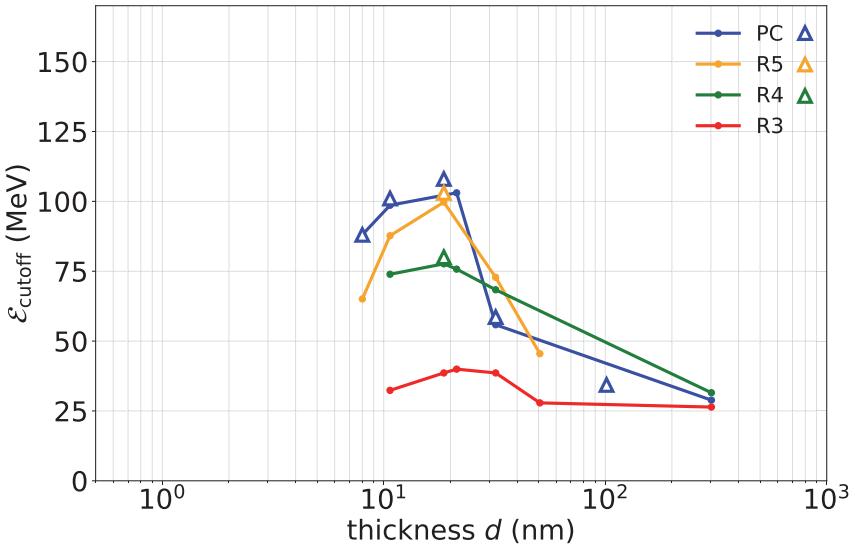

Figure 4.25.: *Regular* TNSA cutoff proton energies from 3D simulations of copper foils with an organic contaminant only at the rear (dots and solid lines) or on both sides (triangles) at laser maximum a_0 of 63.25. The cutoff energies $\mathcal{E}_{\mathrm{cutoff}}$ have been identified from the proton spectra at $t_{\mathrm{rel}} = 60$ fs.

After $t_{\mathrm{rel}} = 60\,\mathrm{fs} = 2\tau_{\mathrm{p}}$ the proton energies are largely stabilized and with about 90 % of their asymptotic maximum. From the shape of the spectra the values for $\mathcal{E}_{\mathrm{max}}$ and $\mathcal{E}^{\mathrm{TNSA}}_{\mathrm{cutoff}}$ were obtained. On the other hand, the proton and electron phase spaces revealed the locations of the critical density (see 4.2.5). Figure 4.24 showed the detachment of the protons from the target front side and, finally, the end of the hole-boring process. The momentum gained from hole-boring for on-axis protons is mostly directed in the target normal direction. That is because the target is still intact at that point in time and in the inner spot area electrons are displaced against the ions forming two almost parallel fronts. If the hole-boring were to continue for prolonged times, the indent of the foil would add more divergence to the proton momenta and also a preferential shift towards the laser propagation direction emerge. However, since the hole-boring stage for front-side protons stops before the peak of the pulse, all the momentum is directed along the y-direction. After identifying which ion acceleration mechanisms are at play at which times, the question remains how much energy they contribute and how this ratio scales with the parameters that were varied in the simulation campaign. Figure 4.26 shows this comparison for the R4 case with only protons at the target rear side . The target thickness of 18.67 nm is the optimum for both the regular TNSA population (dotted lines) as well as the

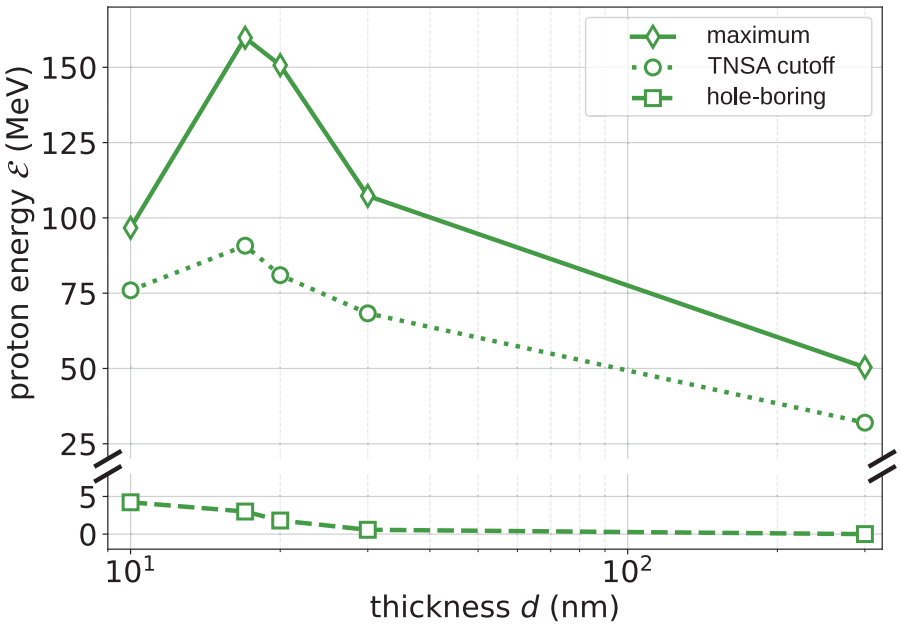

Figure 4.26.: Proton energy contribution per acceleration mechanism for the ramp case R4 and a proton source layer only on the rear side of the foil. The maximum energy and TNSA cutoff energy were determined at $t_{\mathrm{rel}} = 60$ fs from the proton energy spectra. The hole-boring contribution was identified from proton phase spaces at the instance in time when front-side accelerated protons detach from the region of direct laser influence and enter the highly overdense target region before they are injected into the sheath field.

injected protons (solid lines), which reach an overall maximum of 160 MeV. The overall contribution of hole-boring to the final energy was found to be below 5 MeV in all cases. The fact that there seems to be a dependence of $\mathcal{E}_{\mathrm{HB}}$ on the target thickness is counter-intuitive at first. This dependence can be explained with the fact that in the cases shown in figure 4.26 the proton source layer was only on the target rear side. These protons need to travel to the target front side first during the phase of relatively homogeneous pre-expansion while the laser pulse is still non-relativistic, i.e. $a_0 \ll 1$. With an increase in target thickness, the rear surface is heated more slowly and later in time and protons are less likely to arrive at the front surface until the hole-boring stage begins.

Even though the hole-boring stage of the acceleration process at the target front did not contribute significantly to the maximum proton energy, the number of protons at energies higher than the first visible cutoff can be increased if a hydrocarbon contaminant layer is also placed on the target front side. This initial condition relies on the assumption that the nanosecond and picosecond pedestals of the laser pulse do not slowly ablate the front side contaminant. Solid lines in figure 4.20 represent this case. Dashed lines are the simulations with only a rear-side contaminant. Furthermore, the blue and magenta shaded areas mark ranges in the spectrum where the ramp case lacks, or exceeds, protons compared to the perfect contrast case, respectively. The exact moment of injection into the TNSA sheath fields is crucial for reaching the highest energies.

4.2.7. Modeling of Final Proton Energies

Adapted Schreiber Model

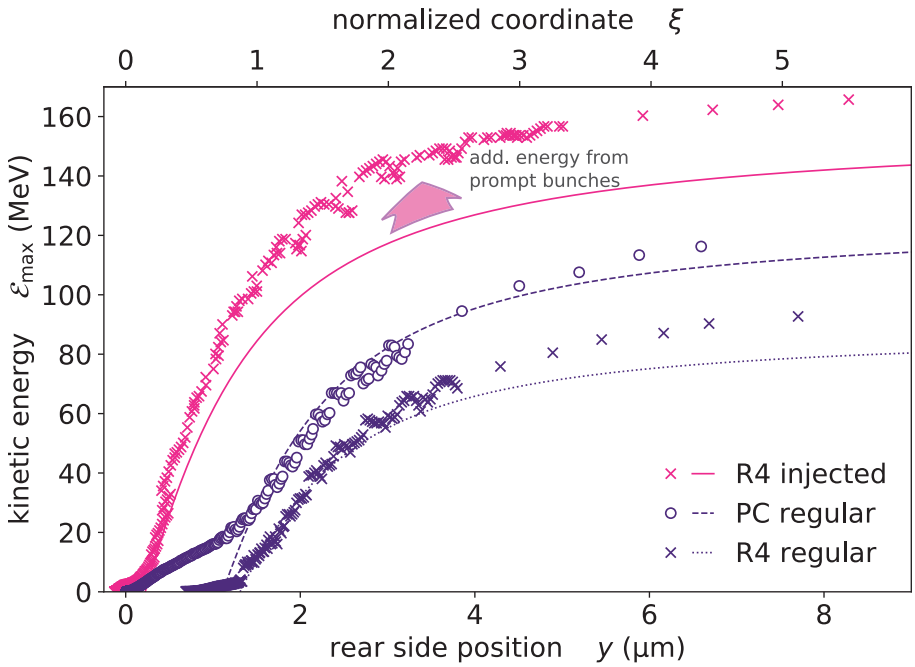

Figure 4.27.: Kinetic energy evolution of the fastest representatives of the *regular* and *injected* proton populations in the $d = 17$ nm R4 and PC cases. Markers represent measured data points from the simulation. The continuous lines are predictions following a modified model based on Schreiber et al. [44]. Additional energy gained by R4 injected proton population stems from increased electric field due to prompt electron bunches at the beginning of the acceleration process close to the sheath origin (compare Fig. 4.23).

To include the evolution of the proton energy based on the detailed dynamics of the proton populations into an analytic description, the well-known static sheath model by Schreiber et al.[44] is used as a basis. In the original model, the accelerating potential does not change with time. It is an isothermal model which depends on time only implicitly. Knowing the laser power P_L and energy absorption efficiency η, it allows to predict a value for the maximum possible ion energy obtained by traveling down the full length of the potential.

$$\mathcal{E}_{\mathrm{i},\infty} = 2 q_\mathrm{i} m_\mathrm{e} c^2 \sqrt{\eta \frac{P_\mathrm{L}}{P_\mathrm{R}}} \tag{4.4}$$

$$\mathcal{E}_\mathrm{i}(\xi) = \mathcal{E}_{\mathrm{i},\infty} s(\xi) \tag{4.5}$$

The authors also give equation (4.5) for the energy at a normalized distance $\xi = y/B$ from the rear side sheath origin. Here, y is again the longitudinal propagation direction along the laser axis and $B = r_\mathrm{L} + d \tan\theta$ is the sheath radius at the rear side foil surface, taking into account the electron divergence θ and the foil thickness d. The laser power $P_\mathrm{L} = \mathcal{E}_\mathrm{L}/\tau_\mathrm{L}$ is readily available. In laboratory experiments, usually only an approximate value for the absorption efficiency into hot electrons η can be given and for solid density targets it ranges between 5 to 50 %. One way to determine the total energy absorbed by the target in the simulation would be to integrate over the reflected and transmitted parts of the laser light and subtract their sum from the

total laser pulse energy. Due to the enormous computational cost of the simulations, the full length of the laser pulse $L = c\tau_L \approx 10\,\mu m$ did not fit into the part of the simulation box that was allocated for the vacuum in front of the target. Instead, the sum of all kinetic energies of all particle species in the simulation as well as the total energy in electromagnetic fields were separately monitored. As soon as laser energy is transferred into hot electrons, they continuously pass on that energy to the several ion species. Within $t_{rel} = 1\,..\,1.5\,\tau_L$ the total kinetic energy of all particles peaks across all runs in the parameter study. Since the time evolution of the acceleration process varies with target thickness and contrast setting, the overall maximum kinetic energy sum $\mathcal{E}_{kin}^{total} = \max\left(\sum_{species,k} \mathcal{E}_{kin,k}\right)$ was calculated separately for each simulation (here k is the single-particle index) and chosen as the characteristic quantity to compare each simulation with. While the detailed plasma kinetics might differ from case to case, the particle kinetic energy always peaks and as such is the feature that all simulations have in common and can be compared by. The absorption efficiency is then calculated with respect to the total laser pulse energy $\eta = \mathcal{E}_{kin}^{total}/\mathcal{E}_L$. Now it becomes possible to calculate the theoretical maximum energy for the R4 and PC case, respectively. With $\eta_{R4} = 21.77\,\%$ and $\eta_{PC} = 21.05\,\%$, the resulting final energies are $\mathcal{E}_{i,\infty}^{R4} = 157.04\,MeV$ and $\mathcal{E}_{i,\infty}^{PC} = 153.96\,MeV$, respectively. From Fig. 4.23 it becomes apparent, that only the injected protons of the R4 case experience a sheath field distribution as theorized by Schreiber et al. [44], i.e. a field of the form

$$E(\xi) = \frac{\mathcal{E}_{i,\infty}}{eB}\left[1 - \frac{\xi}{\sqrt{1+\xi^2}}\right]. \tag{4.6}$$

In Fig. 4.22 the situation is depicted at $t_{rel} = 0$, when the electron sheath is about to be supplemented by the most energetic electrons accelerated by the main pulse maximum. The different proton populations are already at different points in the sheath but have not yet gained a significant percentage of their final kinetic energy. With a slight adjustment to the equation, the energy that these pre-expanded protons can gain, can be modeled again, however. With the sheath field shape as calculated from Eq. (4.6) assume that protons move along the field starting from the measured origin positions $y_{o,inj}^{R4} \overset{!}{=} 0\,nm$, $y_{o,reg}^{R4} = 900\,nm$ and $y_{o,reg}^{PC} = 300\,nm$. Only the injected protons start at the exact origin of the sheath since they originate from inside the target after pre-expansion or the target front side by a brief hole-boring phase. Since the sheath within the Schreiber model is constant in time, the resulting maximum proton energies can easily be calculated from the new origin positions by subtracting the amount of energy they would have gained until the position of origin from the result obtained with eq. 4.4.

$$\tilde{\mathcal{E}}_{i,\infty}^{red} = \mathcal{E}_{i,\infty}\,(1 - s(\xi_o))\,, \quad \xi_o = y_o/B \tag{4.7}$$

Based on these adjustments, the new final energies follow Eq. 4.7, where $s(\xi) = 1 + \xi - \sqrt{1+\xi^2}$ is the same as in the original publication. The adjusted energies from non-zero origin positions are $\tilde{\mathcal{E}}_{reg,\infty}^{PC} = 126.29\,MeV$, $\tilde{\mathcal{E}}_{reg,\infty}^{R4} = 89.06\,MeV$ and $\tilde{\mathcal{E}}_{inj,\infty}^{R4} = \mathcal{E}_{i,\infty}^{R4} = 157.04\,MeV$. The sheath is, however, not static and the protons are probing an electron sheath whose origin and strength is determined by the charge separation of the expanding copper plasma. At the same time, the large electron cloud that is accelerated by the peak of the laser pulse first needs to overtake the already expanded proton populations before the sheath field they experience reaches its full strength. In Fig. 4.23, the regular proton population of the PC case show this effect plainly. Using the position where the electric field first reaches half of its peak value, the *new* origin of the *reduced* sheath in Fig. 4.27 is shifted. With that, a good agreement of the energy evolution with respect to the proton front position for the regular protons in both the R4 and the PC

case is found. The measured proton energy from the simulation is depicted as single markers, while the modified Schreiber prediction is shown with line plots. The adjusted model fits best for the regular proton populations and the energy increase with propagation distance follows the prediction very closely for the first 4 µm. For the injected protons of the R4 case, the adjusted Schreiber prediction underestimates the maximum energy by about 25 MeV. Two factors could possibly be the cause for this mismatch. In the structure of the accelerating field at the injected proton front position in Fig. 4.23 regularly spaced spikes hint at the influence of prompt electron bunches. The peaks are more visible for the injected protons since they are faster than the others and also further down the sheath where the relative density perturbation of the bunches is larger. At the same time, their increased speed causes them to co-propagate longer with the fast $2\omega_L$-bunches. However, especially when the injected protons are still close to their sheath origin, the bunches are much more condensed and they accelerate the protons more strongly. The latter fact in particular could be the reason for the sharp peak of the electric field that cannot be captured by the Schreiber model which does not cover $2\omega_L$-bunches.

However, the original publication of the model[44] and its follow-up[116] also predict that ion energies are further limited by a finite acceleration time that comes from the limited duration of the electron cloud that is passing over the proton front to build up the sheath. The next section briefly discusses the effect for the adapted modeling presented here.

4.2.8. Discussion of the Model

Figure 4.28 shows a prediction of final ion energies $\tilde{\mathcal{E}}_{i,\infty}^{\mathrm{red}}$ in the adapted Schreiber model (Eq. 4.7) with consideration of the starting position inside the Debye sheath. The solid curve shows the energy normalized to the value $\mathcal{E}_{i,\infty}$ that would be reached if a proton were to use the full acceleration potential for an infinite acceleration time. This ratio just reduces to the form $\sqrt{1+\xi^2}-\xi$ according to Eq. 4.7. Absorption efficiency η, laser energy $\mathcal{E}_L$ and pulse duration τ_L affect $\tilde{\mathcal{E}}_{i,\infty}^{\mathrm{red}}$ via the the original model prediction for the ion energy at infinity $\mathcal{E}_{i,\infty} = q_i m_e c^2 \left(\eta \mathcal{E}_L / \tau_L P_R\right)^{1/2}$. According to the result of [44], the expectation for the real cutoff energies would still be limited again by a finite acceleration time that stems from the longitudinal extent of the electron cloud $c\tau_L$ which is accelerated by the laser main pulse. With Eq. 2.26, i.e. the solution of the equation of motion $\mathrm{d}\xi/\mathrm{d}t$, a relation for $\tau_L/\tau_0(X)$ is given, where $\tau_0 = B/v_{i,\infty} = B/(2\mathcal{E}_{i,\infty}/m_i)^{1/2}$ and $X = (\mathcal{E}_m/\mathcal{E}_{i,\infty})^{1/2} = \sqrt{s}$, which predicts a lower final energy $\mathcal{E}_m$. This prediction was calculated for every single simulation and its ratio versus $\mathcal{E}_{i,\infty}$ is the intercept of the vertical axis in Fig. 4.28. All adapted predictions due to $\tau_L = 30$ fs are represented by the grey-shaded area.

Data points represent the TNSA cutoff energies from the simulations of copper targets presented before, normalized to $\mathcal{E}_{i,\infty}$. To compare the model prediction to the simulations, the positions of the proton front, regarding the *regular protons* that expand in the early-formed TNSA sheath, at approximately $t = 0$ were extracted. Assuming that these positions are the origin positions inside a model sheath, normalized to the rear-side sheath radius $B = r_L + d\tan\theta_e$, the measured proton TNSA cutoff energies at $t = 60$ fs were normalized to $\mathcal{E}_{i,\infty}$. Different colors mark the laser contrast configurations as before and different markers signify the different target thicknesses. Most simulations only featured the organic contaminant at the target rear side while the empty, paler markers represent simulations that also had contaminant at the front.

Ideally, all data points would lie within or close to the shaded area and below the solid curve of the model. Most data points are indeed below the solid curve of $\tilde{\mathcal{E}}_{i,\infty}^{\mathrm{red}}(\xi_0)$. The better the laser contrast, the more ideal the pulse and the less pre-expansion was observed. Further-

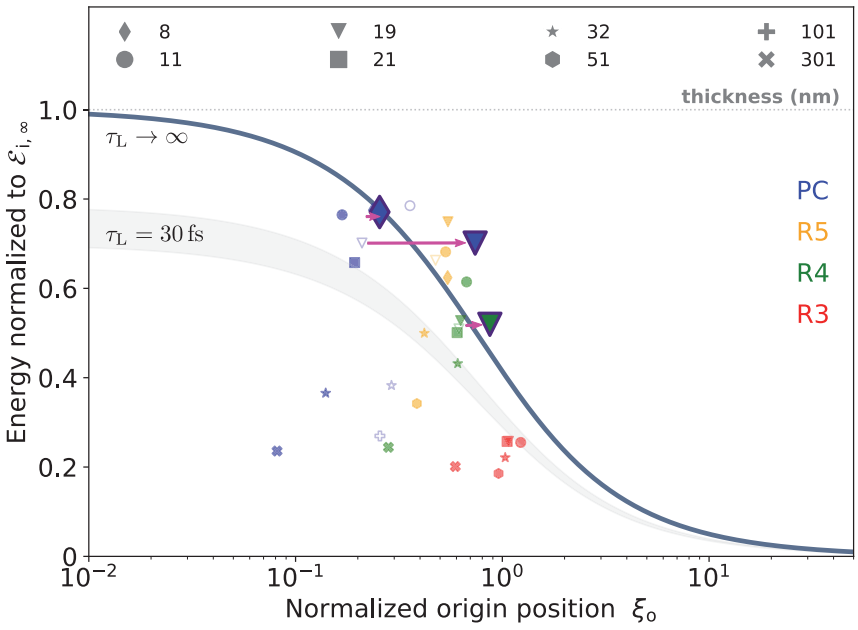

Figure 4.28.: Prediction for the reduction of final ion energies $\tilde{\mathcal{E}}_{i,\infty}^{red}$ due to pre-expansion of the *regular proton* front during the ramp phase of the laser pulse. The curve follows Eq. 4.7 without the leading factors, reducing it to $\sqrt{1+\xi^2}-\xi$. Proton TNSA cutoff energies obtained from simulations of copper targets at 60 fs after main pulse irradiation were normalized to their maximum expected value $\mathcal{E}_{i,\infty}$ within the model of Schreiber et al. [44], taking Eq. 4.4 and the absorption efficiency η measured in each simulation. The coordinates on the position axis were measured as close as possible to the instance of maximum laser intensity at the target front side ($t = 0$) and normalized to B, the rear-side sheath radius that is a function of target thickness and electron angle that was chosen to be 15°. Different target thicknesses are encoded with different markers and the pulse ramp configuration is encoded as before in color, according to Fig. 4.10. The three enlarged markers (PC: 8 and 18.67 nm; R4: 18.67 nm) identify simulations where proton data of higher output frequency allowed to see when the electron cloud accelerated at $t = 0$ actually overtook the pre-expanded ions, resulting in additional pre-expansion and thus updated starting position.

more, with more ideal laser pulse contrast the final proton energies vary stronger with target thickness variation. However, there is no good agreement with the shaded area that indicates the limited acceleration time of $\tau_L = 30\,\mathrm{fs}$ and instead cases close to the respective optimum thickness are closer to the model prediction of $\tau_L \to \infty$ while especially thicker targets deviate far. This can be understood as the pre-expansion is reduced with increasing target thickness and the delay until the majority of laser-maximum-accelerated electrons reach the proton front also increases. However, the laser absorption does not decrease substantially enough so that the predicted final energies from a low expansion and moderately high absorption ($\sim$ 15 %, compare Fig. 4.21) are still high.

In reality, protons expand further until the sheath is actually supplied with more electrons and this distance also increases the longitudinal spread of $2\omega_L$–bunches, reducing their effectiveness for acceleration. The effect of these bunches is, however, not reflected within any of the 1D isothermal models. Also, while the number of laser-accelerated electrons is limited and the sheath is re-supplied only for a limited amount of time, it is upheld for much longer than τ_L. Any form of pre-expansion also causes the protons the co-propagate with the evolving sheath, which increases the temporal overlap and therefore the interaction time with the sheath. Especially thickness scans around the optimal thickness are challenging to describe with only a single 1D model. Necessarily, the nature and contributions of different acceleration mecha-

nisms changes there which is especially substantial if the laser intensities are strong enough to also trigger radiation pressure acceleration. In that case, the detailed dynamics of the acceleration process for each parameter combination have to be taken into account to match parts of the acceleration with the model. Such a detailed treatment beats the purpose of such models, however, which retain their use for the prediction of general trends of ion energies over a wide range of laser intensities, given that TNSA stays the dominating process.

The three enlarged markers in Fig. 4.28 mark simulations where additional proton data with higher output frequency was written. With this data, it was possible to identify the delay between $t = 0$, the moment at which the strongest sheath is forming, and the instance when the pre-expanded proton front would experience the field of the electron cloud once it overtakes the front (see Fig. 4.23). This delay accounts for additional pre-expansion and thus the positions in Fig. 4.28 are updated, as the pink arrows indicate. In the model, the sheath is static but in reality the origin moves forward as also heavier ions like copper expand and the electron cloud accelerated at $t = 0$ moves as well. However, its center of mass is slower than the speed of light since only the most energetic electrons move at approximately c and escape the sheath whereas less energetic electrons recirculate.

In effect, parts of the acceleration of ions can still be described very well with the adjusted model of Schreiber et al. [44], despite the fact that a 1D model is used to describe a 3D process, but only if the acceleration stages can be identified in the data. Unfortunately, this was not possible for all simulations of the campaign described here since the available output frequency only allowed for a very rough estimate of when the pre-expanded proton front actually saw the freshly supplied electron sheath. The increasingly detailed picture such simulations deliver also emphasizes how the complex spatio-temporal coupling of a realistically shaped laser pulse is challenging to encapsulate in simple model parameters. Especially if the leading intensity ramp were to have an even higher fraction of energy compared to the main pulse it is questionable how to define the pulse length, and consequently, the acceleration time. For ultrathin targets, the energy content of the electrons varies strongly over 10s of femtoseconds around $t = 0$ which makes the choice for a single absorption efficiency η, or alternatively an electron temperature T_e (interchangeable since $\eta \propto n_e T_e$), difficult. Lastly, different target thicknesses around d_{opt} change the timing of the acceleration process enough that the description with a single model requires adjustment and detailed insight.

4.2.9. Transverse Proton Profiles

The results so far showed that changes in laser contrast on the last few hundred femtoseconds before the arrival of the ultra-high intensity maximum can already have a large impact on the proton acceleration performance from ultrathin targets. But since this acceleration performance is a result of the spatio-temporal coupling of laser energy into the (pre-)expanding plasma it is conceivable that the transverse profile of the accelerated particle beams are likewise affected.

Figure 4.29 shows an extrapolation for the total proton number deposited on a virtual detector in 55 mm distance from the point of laser-matter interaction. Proton data has been recorded with the `ParticleCalorimeter` diagnostic of PIConGPU in a 3D histogram binning the dimensions of *energy*, *yaw*– and *pitch angle* (see [206] for documentation). The detector measures $(50 \times 50)\,\mathrm{mm}^2$ such that the full opening angle is about 40° in both angular directions and the laser propagation axis is normal to the detector surface and directed at its center. Fig. 4.29 compares the three different contrast settings of PC, R4 and R3 at a 32 nm copper target with organic contaminant at the rear side. It becomes apparent that the accelerated proton beam has two separate centers aligned in the plane of laser polarization. With de-

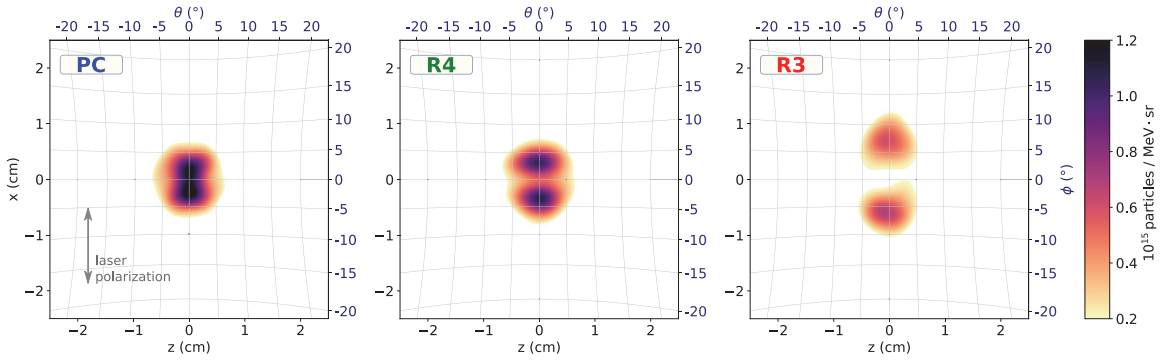

Figure 4.29.: Extrapolation for total proton number deposited on a virtual planar detector with its center on the laser beam axis at a distance of 55 mm and extent of $(50 \times 50)\,\mathrm{mm}^2$. The simulated target was a 32 nm copper foil with organic contaminants on the rear side that was irradiated by $a_0 = 63.25$ laser pulses of decreasing contrast quality (*from left to right*) in the configurations PC, R4 and R3 as shown in Fig. 4.10b. The warped grid shows lines of equal yaw (φ) and pitch (θ) angles.

creasing contrast quality, i.e. with increasing energy inside the exponential ramp preceding the Gaussian main pulse, these two centers separate more in the direction of the yaw angle φ. For the PC case, they are about 5° apart, for R4 it is 7.5°, and 14.5° in the R3 case. Thereby, the temporal laser contrast not only affects the maximum proton energies but also the general pointing of the beams. The virtual detector dimensions are chosen such that it corresponds to radiochromic film (RCF) stacks (like in [10]) which are often used for measuring proton energy in laboratory experiments. It is customary to perforate these stacks with pinholes for guiding a part of the beam to a Thomson parabola spectrometer (TPS) which can deliver shot-to-shot energy measurements while RCF stacks need to be read out manually after extraction from the experimental chamber. Even with such small changes in laser contrast as they were varied here, the majority of the proton beam might completely miss the pinhole (often around 2°–5° opening angle) when ultrathin targets of the same thickness are irradiated.

Figure 4.30 display a set of virtual RCF layers from 3D simulations of laser-irradiated 18.67 nm copper foils with organic contaminants on both surfaces. Hydrogen atoms on both surfaces could be clearly distinguished, since they were configured as two physically identical but separated ion species in PIConGPU. These are the simulations that showed the overall highest maximum proton energies (compare Fig. 4.17) following the interaction with laser pulses of maximum $a_0 = 63.25$ and $\tau_\mathrm{L} = 30\,\mathrm{fs}$ Gaussian main pulse (PC), only – as well as with added exponential leading ramp (R4) (see Fig. 4.10b).

The virtual RCF layers just show intervals of $\Delta\mathcal{E} = 0.146\,\mathrm{MeV}$ and do not account for the effect that higher energy particles also partially deposit dose in lower energy layers. Neither do the virtual RCFs account for the increase in proton beam divergence which comes from different layer distances that are separated by absorbing layers in reality. In other words, the virtual RCF stack here is infinitesimally thin while a real stack can be several centimeters long. In Figure 4.30, each layer shows an automatically scaled proton number that is normalized between 0 and the maximum value in the image which is written in every bottom right corner, given in $(\mathrm{MeV\ and\ steradian})^{-1}$, respectively.

Both contrast cases in Fig. 4.30 show very similar transverse proton profiles throughout the different energies. A very prominent feature is the ring of higher proton number and depleted center region which was observed in experiments and simulations before[310, 317]. At the center of the sheath its curvature is strongest and with increasing distance from the beam axis the normal vectors point increasingly away. At some distance from the axis, the transverse component of the normal vectors begins to decrease again until the target is fully flat, again. Therefore, the expanding protons are focused into a ring structure. As described earlier, the

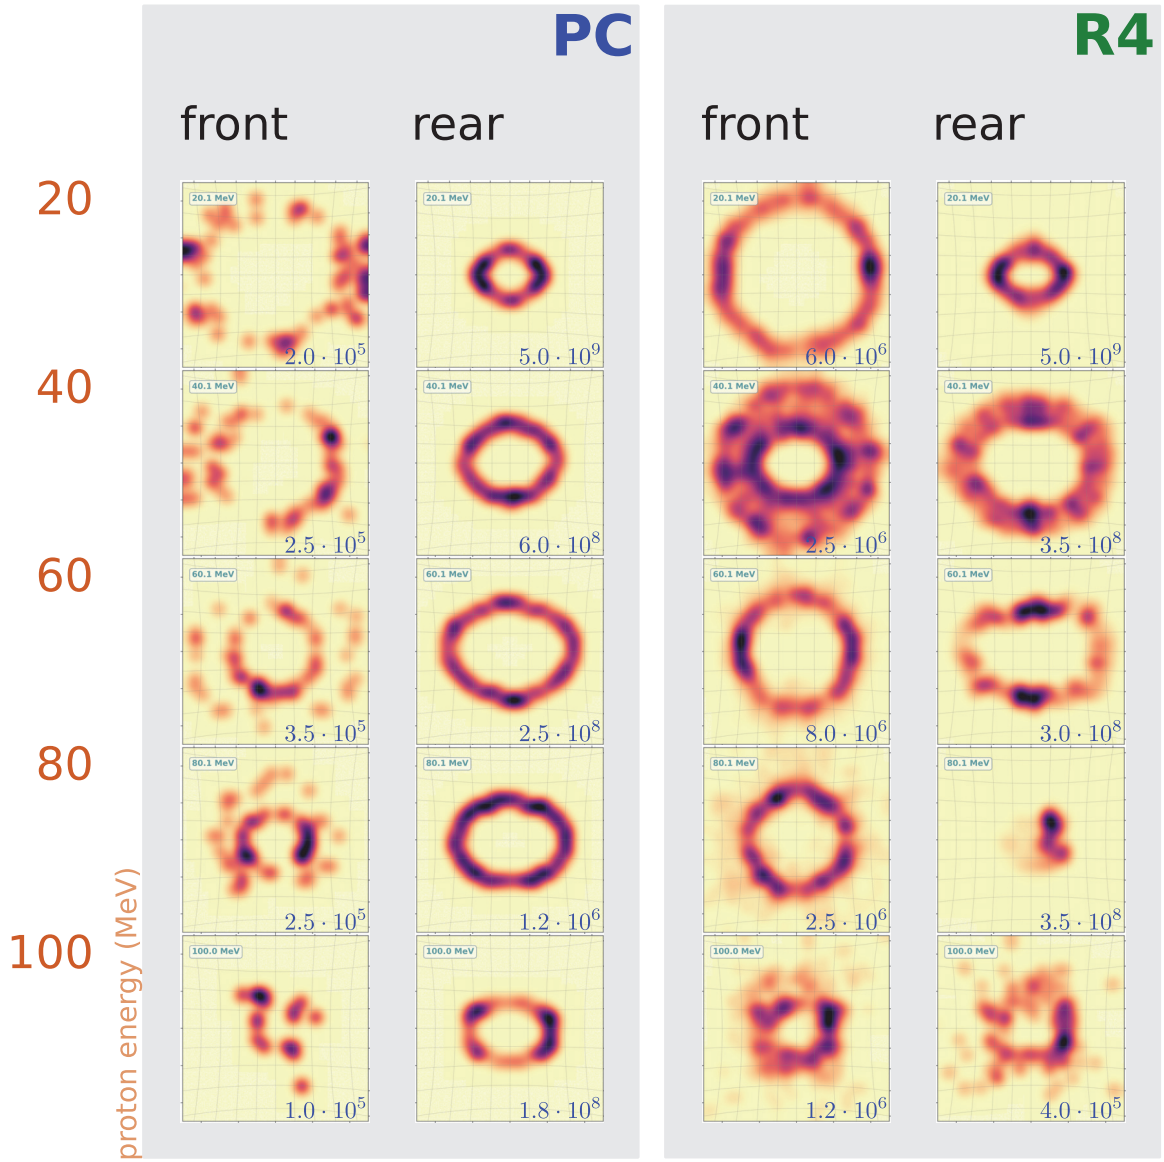

Figure 4.30.: Set of chosen virtual RCF layers from simulations of copper foils irradiated by $a_0 = 63.25$ laser pulses at the optimum thickness of $d_{\text{opt}} = 18.67\,\text{nm}$. In the left (right) panel the stack layers of the PC (R4) contrast case are displayed. Virtual detector dimensions are the same as in Fig. 4.29. Each layer's energy width is $\Delta\mathcal{E} = 0.1\,\text{MeV}$ and is normalized linearly to its maximum value given in protons per MeV and steradian which is shown in each bottom right corner, respectively. Warped grid lines are drawn every 5° for both angles. The laser is polarized vertically here.

PC case shows a much lower number of front-side protons accelerated in forward direction than the R4 case with less ideal contrast, where there are about an order of magnitude more in each layer $\Delta\mathcal{E}$. The ring feature of the rear-side protons is much more pronounced in the ideal contrast case whereas the profiles are more blurred out with the initial R4 intensity ramp. The fact that they are still very clean in both cases originates in the very limited contaminant layer thickness of only 3 nm. The profiles of the Formvar cases hereafter in Fig. 4.31 show other features as well since the complete expanding target is abundant with protons. Some layers of the R4 case exhibit a slightly larger ring radius than the PC case (except the rear-side protons at 80 MeV that are very collimated), indicated stronger pre-expansion and sheath curvature. In both cases, the rear-side originating protons show a ring of increasing radius with increasing energy until between 60 and 80 MeV the radius decreases again. The exact forward direction

is also populated for higher energies than 100 MeV but the layers are so thinly spaced with $\Delta\mathcal{E} = 0.146$ MeV that no features but single-proton speckles remain.

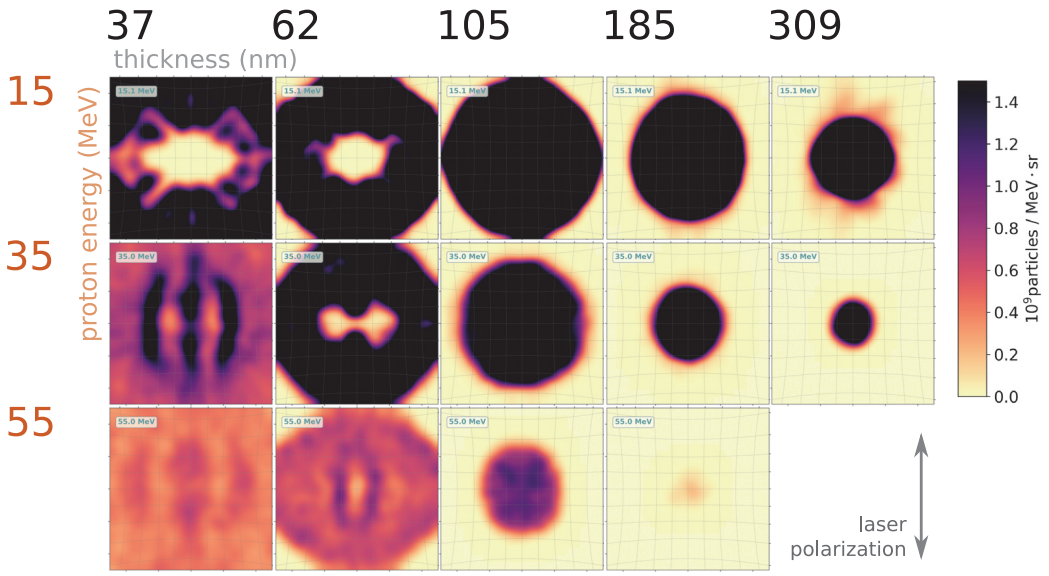

(a) PC (Perfect Gaussian contrast) case.

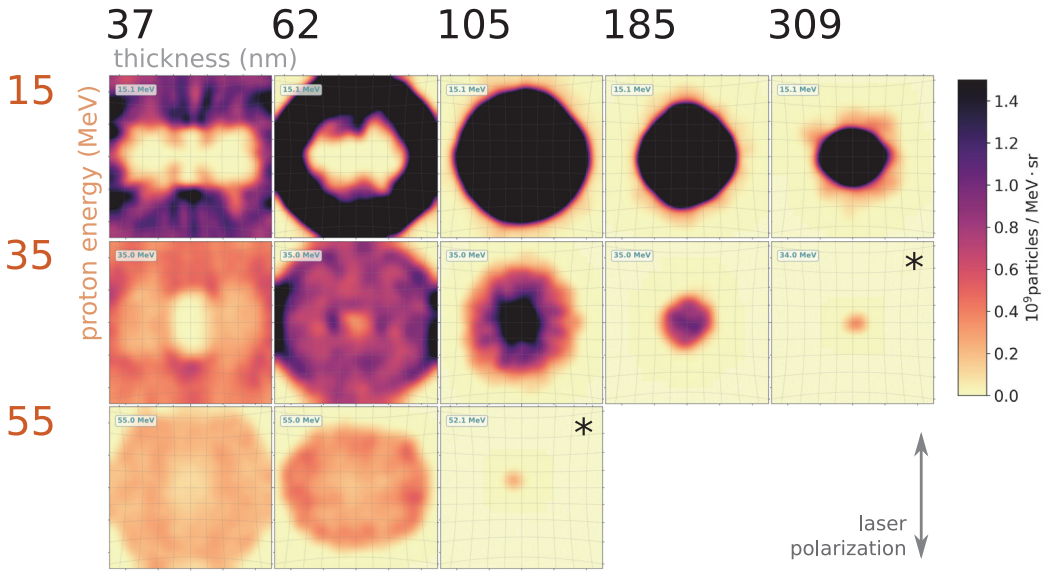

(b) R5 ramp case (plateau of 10^{-5} I_0 from –300 to –100 fs before the main pulse). Images marked with an asterisk are layers where the last visible signal was found at slightly lower energy (34 MeV and 52 MeV, respectively).

Figure 4.31.: Virtual RCF stack images from PIConGPU `ParticleCalorimeter` output of 3D simulations modeling Formvar targets of varying thickness irradiated by laser pulses of $a_0 = 36.5$ in two different contrast configurations. RCF distance and spatial scales are the same as in Fig. 4.29. Warped grid lines show equal yaw and pitch angles every 5°.

Figure 4.31 shows example layers from virtual RCF stacks created from `ParticleCalorimeter` data of the thickness scan performed on Formvar foils that were irradiated by laser pulses of maximum $a_0 = 36.5$. The detector dimensions are chosen the same as in Fig. 4.29 and all images in Fig. 4.31 share the same color scale.

Three energies (15, 35, and 55 MeV) are shown to compare the five target thicknesses among each other as well as with the two simulated contrast settings of PC and R5 (according to Fig. 4.10b). Under both contrast settings the optimum observed target thickness was found at

62 nm (compare Fig. 4.16a). It becomes immediately apparent that for target thicknesses above the optimum thickness the spot on the virtual detector is circular and the beam is well-collimated. With increasing target thickness the beam spot diameter decreases which hints at less pronounced target expansion such that most normal vectors of the target rear side are well aligned. In the case of an intensity plateau of $10^{-5}\, I_0$ for 200 fs preceding the Gaussian main pulse (case R5), less particles reach higher energies and the edges of the beam spot are softening and fraying more, like the layer of 35 MeV at $d = 105$ nm shows. While the beam spot is saturated in many of the lower energy layers[4], almost all transverse proton profiles show the aforementioned ring feature. This ring structure is most prominent for the lower energy part of the spectrum since these protons experienced the accelerating potential of a pre-formed sheath.

Beam spots at a target thickness of 37 nm, i.e. less than d_{opt}, are extending beyond the limits of the virtual RCF and are highly structured. Whereas in the lower energy layers the center region is clearly depleted of protons, the higher energy layers show the protons in a three-lobed structure, also at lower divergence angles, again, but with a preferential direction along the laser polarization axis. This is because the $2\,\omega_L$ electron bunches that were accelerated by the $a_0 > 1$ half-waves of the laser pulse increasingly contribute to the acceleration of protons. They are traveling through the target under a preference angle of alternating direction, depending on where the electric field pointed during their creation, and are divergent themselves. Both their energy spread as well as their angular divergence causes them to be smeared out along the polarization axis.

4.2.10. Campaign Summary and Discussion

The simulations in this chapter revealed a mechanism that showed how proton energies can be boosted to 150 % of their final energy if instead of a perfect Gaussian laser pulse a pulse with a leading intensity ramp is used. This energy increase is due to better timing and placement of protons in the expanding Debye sheath caused by the main pulse that was not synchronized optimally in the PC (perfect contrast – Gaussian) cases. Only when they are present at the origin of the Debye sheath, protons can subsequently gain kinetic energy throughout the full acceleration potential and reach the final final energies predicted by Schreiber et al. [44] and other models. If proton layers pre-expanded away from the sheath origin, they reach a lower asymptotic energy and an adapted model was introduced to account for this behavior. The dependence on pre-expansion implies that the model only applies to multi-layer targets like metal foils coated with organic contaminants. Plastic targets are homogeneously abundant with hydrogen and are also populating the sheath origin. Additional pre-expansion effects like lower accelerating fields due to softened plasma gradients as described by Grismayer and Mora [120] can still occur at the same time and overlay the beneficial effect described above. To avoid this and get increased ion energies from intensity ramps, heavy, high-Z ion species are beneficial since they provide more resistance against radiation pressure even at lower target thicknesses while being less prone to complete depletion of electrons which would result in Coulomb explosion effects.

The large-scale 3D simulations of plastic targets irradiated by a 10 J laser showed that the maximum proton energies decreased with increasing energy in the intensity ramp, already when just a constant contrast plateau was added in the R5 case. This resulted from consistently lower absorption efficiencies over all scanned target thicknesses. Even though plastic targets showed ion species demixing in later stages of the acceleration process, it is not promi-

[4] Owing to the choice of the colorbar range to ensure that the proton signals of both very thin and very thick targets are clearly visible.

nent during the time of strongest accelerating fields. However, as expected, the heavier ion species separate clearly in the energy spectrum due to their lower charge-to-mass ratio. Recently, a model has been proposed by Huebl et al. [G15] to predict maximum energies from homogeneous two-species planar targets. There, hydrogen and deuterium were chosen for their simple ionization physics but the model could in principle be extended to other ion species and adjustments for different charge states could be made. While not in the scope of this thesis, the 3D simulation data that was obtained will be subject to further study after the conclusion of this work.

The large-scale simulation campaign described in this section was the first study that produced high-resolution, high-quality data of the full 3D laser-ion acceleration process from foils under the consideration of realistic laser intensity ramps within the last picosecond prior to laser peak arrival. In the wake of this one-year computation project, all three involved institutions of HZDR, ZIH and CSCS built up new capabilities for data-intensive simulation workflows (e.g. JupyterHub) and transport (e.g. GridFTP) that future projects and users will be able to benefit from. But the simulations also delivered valuable information of the divergence and spectra of ions and electrons that can be used to maximize the high-energy ion yield with magnetic focusing optics like demonstrated recently at HZDR by Brack et al. [318].

The simulation campaign showed another peculiarity that was not studied in much detail but will give rise to further investigations. Different laser pulse shapes influence the evolution of target ionization and expansion. Subsequently, the target front and rear surfaces are indented or bulging out, respectively, at different times with respect to the arrival of the laser maximum which geometrically affects the acceleration of ions.

In laboratory experiments, diagnostics for particles and electromagnetic radiation usually only cover a very limited part of the full solid angle. Proton diagnostics like Thomson parabola spectrometers (TPS), for instance, will admit particles only through a tiny pinhole. Under such circumstances, the measurements become very sensitive to the transverse proton emission characteristics which can only be appropriately modeled in 3D simulations. Furthermore, the comparison between 3D simulation results and experimental measurements can be further complicated by the detection thresholds of available detector systems such that the real maximum energies cannot be observed.

However, being able to accelerate protons from thin foil targets has often been an unofficial quality indicator for the laser contrast and standard targets can give a general clue about the laser performance on the beginning of a given experiment day. In this sense, an extensive characterization of transverse proton profiles accelerated from a standard thickness ultrathin foil target and measured on a scintillating screen after different absorbing layers at the same time (to account for different proton energies) could as well be an indicator for the last picosecond contrast of the system on that day.

Concurrent Experimental Campaigns In parallel to the simulations of the author, experimental work at the DRACO laser system studied the influence that ΔTOD and ΔGVD dispersion adjustments have on the ion acceleration performance from various targets.

Figure 4.32 shows the results obtained in these experiments[10]. Depicted is the maximum proton energy from changes of ΔTOD, *after* a stable, best-compression (*initial*) setting was found, using a Wizzler-Dazzler feedback loop and SPIDER measurements. It can be seen that in all target types positive values of ΔTOD lead to increased proton energies. Especially in the case of Formvar targets the energy gain surpasses 50 %. Further changes to ΔGVD (not shown here) return the pulse to better compression from the TOD-caused broadened shape, increasing the gain even further. Ziegler et al. [10] conclude that short, asymmetric laser pulse with a shallow intensity ramp can effectively double the maximum proton energy while also

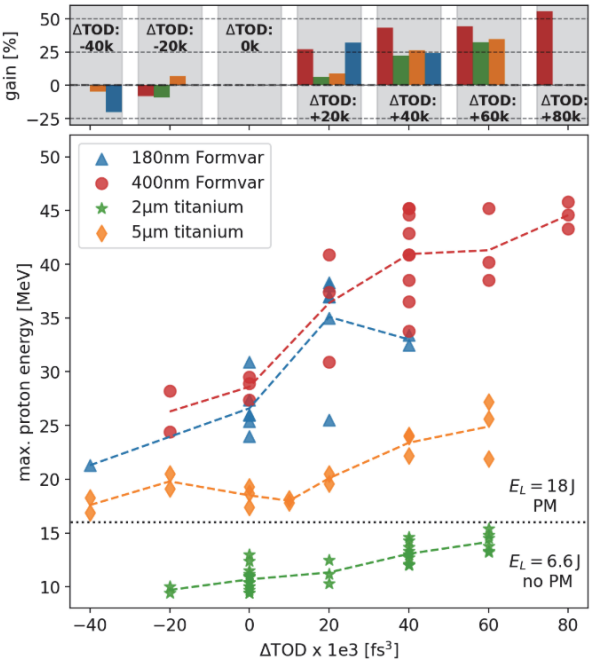

Figure 4.32.: Experimentally obtained maximum proton energies from various targets with respect to ΔTOD changes (ΔGVD= 0). Different laser energy (E_L) and contrast settings (with PM / without PM) were also employed. The upper panel shows the relative energy gain compared to the initial settings of best laser pulse compression. Originally in [10].

increasing the proton flux significantly.

While the intensity ramps varied in the simulations are common, averaged representatives of the measured intensity profiles at DRACO, the trend of increasing proton energies with shallower ramps was not observed for plastic targets. In the simulations maximum energies are decreasing due to much lower laser absorption in the ramp cases and the early onset of transparency. However, a direct comparison between the idealized conditions and the much more complex experimental dynamics remains challenging.

4.2.11. Conclusions

The 3D simulations with realistic intensity ramp shapes performed in this thesis work revealed that commonly used initial conditions for PIC simulations do not sufficiently describe the state of the target following the interaction with an intensity ramp within the last hundred femtoseconds prior to the main pulse arrival. Non-trivial density– and electron energy distributions as well as particle drifts caused by the leading edge can significantly enhance the maximum proton energy from metal targets over the performance achieved with a perfect Gaussian laser pulse in the short-pulse ~ 10 J regime. If the available resources do not prohibit such a study, a more realistic leading edge should always be included into the parameter scans for a laboratory experiment to ensure that possible effects from the interaction are not missed. On the other hand, more suitable initial conditions can possibly be identified that allow for a later initialization point with respect to the pulse maximum.

The spectral feature of the high-energy tail observed in simulations with copper targets can be found in several other works that were published before this thesis which also used metal foils (e.g. [21, 23, 319, 320]). Revisiting these target and laser setups could provide new insight on how to control and enhance the number of protons and proton maximum energies by

exploitation of the process described above. Especially with metal foils of varying thickness and possibly a proton-rich layer on the laser-facing side, experiments are conceivable that may allow to distinguish protons of front-side and rear-side origin based on their location in the energy spectrum or their spatial profile. Protons injected into the sheath from the front side would be accelerated to high energies but would likely remain very collimated. Lower energy protons could exhibit a transverse profile that is spatially structured due to instabilities or disturbances inside the Debye sheath.

Comparing the results of recent experiments at HZDR and the simulation results with plastic targets from this study, opposite trends for maximum proton energies were observed. This indicates that the simulations still do not reflect all significant features of the laser pulse provides cause for further, careful study. Changes to GVD and TOD in the experiment do not directly translate to the fully-compressed, fully-amplified laser pulse in the focal plane and neither do the measured intensity shapes reflect the full spatiotemporal structure of the laser pulse. More diagnostics in both simulation and experiment as well as close collaboration between both are necessary to shed light on the complex acceleration dynamics to ultimately gain precise control over the ion beam parameters.

4.3. Novel Diagnostics Accessing the Nanoscopic Ultrafast Processes in Laser-Driven Solid-Density Plasmas

Most of the currently available diagnostics during a laser-ion acceleration experiment suffer from a lack of either spatial or temporal resolution. But the precise state of the target during the short period of time including and around the arrival of the main pulse with its peak intensity needs to be known to fully characterize the acceleration process. Posterior electron and ion detection methods as well as most of the radiation diagnostics are largely time-integrating and, therefore, cannot give direct insight into the few ten to hundred femtoseconds around the main pulse interaction. Probing methods using optical light contain valuable spatial information but fail to penetrate the overdense plasma regions. Characteristic X-ray radiation from atomic transitions, e.g. Kα, inside the target has been shown to yield spatial information about ionization states and the plasma temperature[321, 322], but again only in a time-integrated fashion with large contributions after the main pulse interaction has already happened. Additionally, limited capacities in experimental vacuum chambers only allow for detectors to be placed in fixed locations, subsequently covering only very limited regions of the full solid angle.

During the large-scale simulation campaign described in the previous section, a subset of simulation runs featured self-consistent bremsstrahlung radiation generation. These X-rays are produced when laser-accelerated electrons traverse the target and scatter off atomic nuclei (scaling with Z^2) and shell-electrons (scaling with Z). In PIConGPU, the module was implemented in the course of the diploma thesis of H. Burau [194]. Recently, bremsstrahlung directly generated from the laser-driven target has been proposed to give insight into the dynamics of the laser-ion acceleration process[190]. While radiation signatures in the integrated measurements of ejected electrons, ions and photons can all be used to deduce specific conditions of the acceleration process, bremsstrahlung X-rays are especially of interest. Due to their low interaction probability with the target, they hold the potential to decode very direct information about the state of the target during the time of their creation, whereas electrons and ions can be deflected in the strong fields that are inherent in laser-ion acceleration.

This section shows only a selection of the results that were obtained in the simulation study mentioned before. Unfortunately, the creation of bremsstrahlung macro-photons that propagated through the target volume was an additional source of instability for the already compu-

tationally very expensive simulations. As several simulations did not run to completion, parameter scans were left fragmented and incomplete. Nevertheless, to the knowledge of the author, these are the first fully 3-dimensional simulations of laser-ion acceleration from metal and plastic targets with realistic laser contrast conditions and unprecedented resolution featuring in-situ bremsstrahlung generation. Therefore, the results presented here are of general interest and may serve as a motivation to continue the research of the option to use bremsstrahlung as a novel diagnostic to finally access the femtosecond timescales of the main-pulse-target interaction.

4.3.1. Radiation Signatures from the Target During Interaction

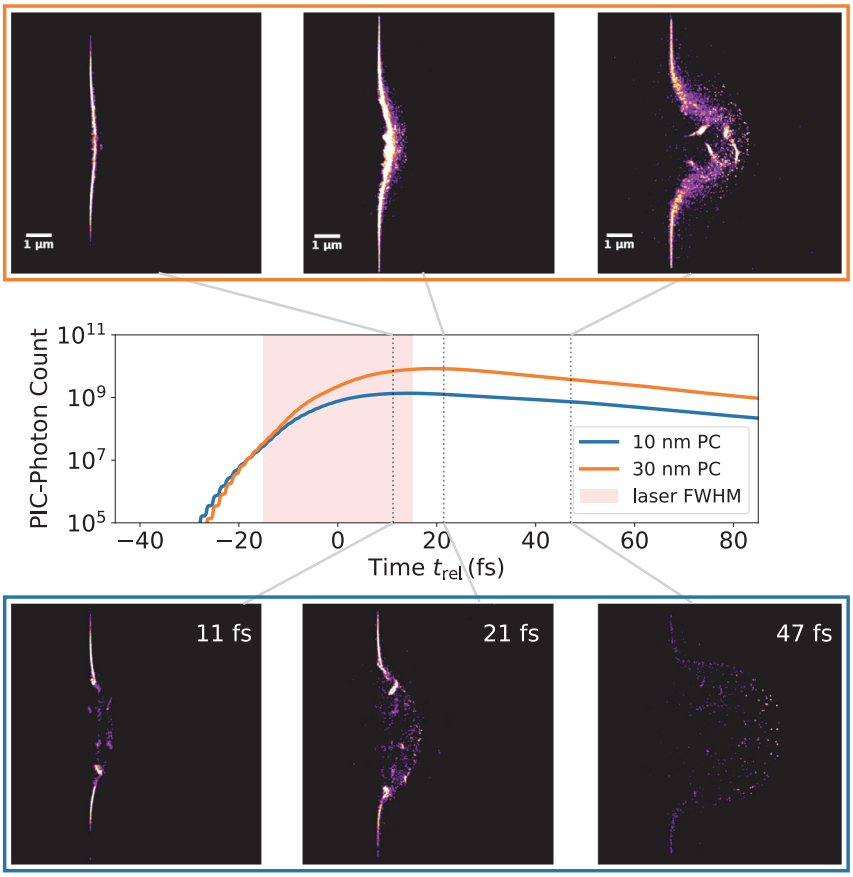

Figure 4.33.: *Middle panel:* Photon macro-particle count of two copper simulations with perfect Gaussian laser contrast (PC) and identical setup, apart from their target thicknesses of 11 nm (blue line, lower panels) and 32 nm (orange line, upper panels). The maximum a_0 was 63.25. The lower and upper panels show example photon macro-particle density slices from the plane spanned between laser propagation and polarization direction. While the density for each image is in arbitrary units, the different degrees of target indentation as well as differences in bremsstrahlung creation hot spots become already apparent.

Figure 4.33 shows photon density slices obtained from the simulated interaction of a Gaussian laser pulse ($a_0 = 63.25$) with copper foil targets with thicknesses of 32 nm (top row) and 11 nm (bottom row), respectively. Each image is normalized to its maximum value and while the three images in both time series cannot be compared quantitatively, they show very different photon distributions, qualitatively. This hints at different states of the target, with respect to expansion and transparency. For the first snapshot, at $t_{rel} = 11$ fs after the arrival of the

main pulse maximum, the 32 nm foil is still structurally intact and so bremsstrahlung photons are created in a very confined layer. Target indentation becomes already visible at 21 fs, close to the time at which the photon count inside the simulation is maximal (see middle panel) and 47 fs after the main pulse interaction, photons are created in a spatially bunched structure inside a bulk plasma that has expanded at least 2 µm past the original rear surface of the target. The thinner foil on the other hand already shows strongly localized photon creation at the earliest shown time. This hints at a strong disturbance of the foil surface and bulk where the target was already (or is being) burned through in its center. By the latest shown time of 47 fs, no strongly featured structure can be identified, anywhere, indicating that the target expansion is so pronounced that bremsstrahlung photons are not efficiently created, anymore. These early observations already show that the bremsstrahlung emission depends strongly on the target conditions over the very short time that the laser is strongest until only very shortly after, which is also the period that is of highest interest and importance for the ion acceleration process. In the following, selected angular photon distributions and energy histograms are shown and examined for their connection about the target conditions.

Energy Spectra

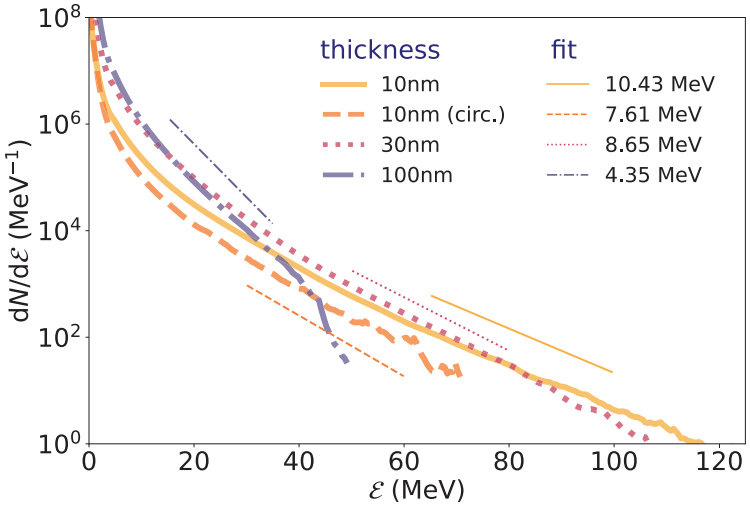

Figure 4.34.: Energy histograms of **bremsstrahlung photons** from four simulations of Copper targets irradiated with a linearly (*circularly) polarized 30 fs Gaussian laser pulse of $a_0 = 63.25$ (*PC* in Fig. 4.10b). Target thicknesses are 11 nm, 32 nm, and 101 nm. Thin lines represent temperature-like fits to the high-energy part of the bremsstrahlung distributions.

Figure 4.34 shows total bremsstrahlung energy histograms of four different simulations of Copper foils at 170 fs after the peak of the laser pulse interacts with the target front surface. The respective high-energy parts were fitted with exponential functions of the form $f(\mathcal{E}) = N_{\mathrm{ph},0}/k_B T_{\mathrm{ph}}^{\mathrm{tail}} \cdot \exp\left(-\mathcal{E}/k_B T_{\mathrm{ph}}^{\mathrm{tail}}\right)$. Among the simulations with linear laser polarization, the fitted photon tail "temperatures"[5] $T_{\mathrm{ph}}^{\mathrm{tail}}$ decrease with increasing target thickness d whereas the overall photon number increases. The 11 nm case with circular polarization exhibits a

[5]The author would like to emphasize, again, that the fits to photon, electron and proton spectra in this section do not represent real *temperatures* in a thermodynamical sense. While the high-energy tails in the spectra are approximated with a Boltzmann law, the particle energy distributions originate in highly non-equilibrated physics.

lower photon temperature than even the 32 nm case and shows the lowest overall photon number. Since electrons experience a rather constant driving force in this case, high-energy $2\omega_L$–bunches are not created which translates to a lower-energy bremsstrahlung signal.

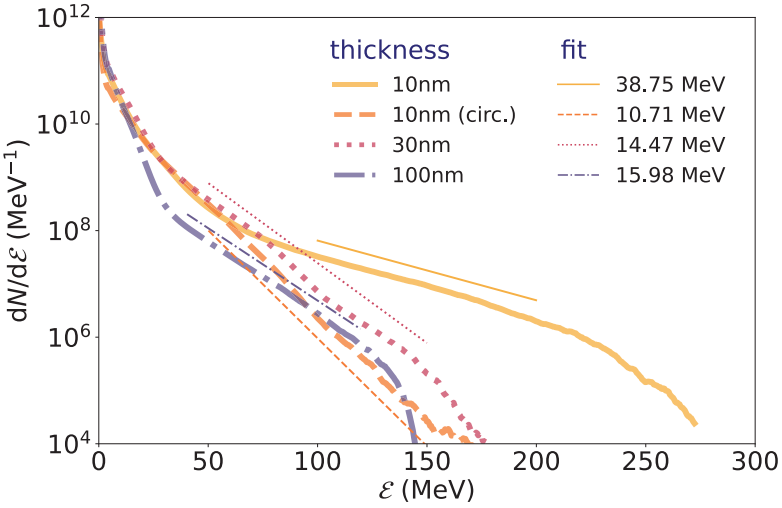

Figure 4.35.: Total energy histograms of **electrons** emitted in four simulations of Copper targets irradiated with a linearly (*circularly) polarized 30 fs Gaussian laser pulse of $a_0 = 63.25$ (*PC* in Fig. 4.10b). Target thicknesses are 11 nm, 32 nm, and 101 nm. Thin lines represent temperature fits to the high-energy part of the distributions.

Figure 4.35 shows the accompanying electron energy histograms of the simulations displayed in Fig. 4.34 and Fig. 4.36 shows the spectra of all protons. The trends observed in the photon spectra can also be seen in the electron spectra, i.e. temperature fits showing decreasing values of T_e^{tail} for thicker targets. The high-energy part of the electron spectra consists of the prompt electrons created at the target front which are energetic enough to overcome the restoring forces in the charge separation zones at the target expansion surfaces. For reference, the scaling of Kluge et al. [74] gives $T_e = 8.6$ MeV for $a_0 = 63.25$ (see Eq. 2.9) whereas the ponderomotive scaling yields 22.3 MeV. However, these values are not directly comparable because they describe the characteristic "temperature" of electrons directly after the interaction with the laser field but prior to the interaction with the target. According to Kluge et al. [73], the prompt electron bunches exhibit a spectral structure that is proportional to $1/\gamma_e$ and peaks again around their maximum energy. Only in convolution with the temporal intensity distribution of the laser pulse, the collective of all prompt $2\omega_L$–bunches has a spectral distribution that can look similar to a thermal distribution with a temperature. As the interaction with the laser comes to an end, the electrons remaining in the larger target area equilibrate and develop a real temperature. The multi-10 MeV tail of the prompt bunch electrons has a low interaction probability with the bulk of typical micron-sized targets and is also able to escape the sheath field. As such, its spectral structure is mostly preserved. It creates highly energetic bremsstrahlung during a short period in time and the following considerations will show how target state is imprinted into this radiation which is promising for allowing to reconstruct more details about the proton acceleration mechanism (see Figs. 4.37, 4.38, and 4.39). The lower energy part of the bremsstrahlung, on the other hand, is also of interest but should give more general information about the total absorption efficiency, for instance.

Finally, the proton spectra in Fig. 4.36 show exponential distributions with clearly identifiable maximum energies. The single circularly polarized case has the overall highest proton energies that were detected in this simulation study with close to $\mathcal{E}_{max} = 200$ MeV. This is because

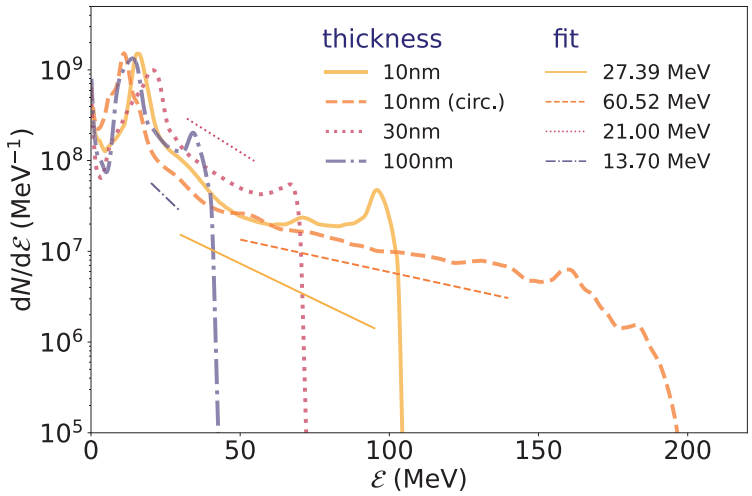

Figure 4.36.: Total energy histograms of **protons** emitted in four simulations of Copper targets irradiated with a linearly (*circularly) polarized 30 fs Gaussian laser pulse of $a_0 = 63.25$ (*PC* in Fig. 4.10b). Target thicknesses are 11 nm, 32 nm, and 101 nm. Thin lines represent partial fits to the distributions.

the interaction process strongly exhibits Light-Sail RPA. For the linearly polarized cases, the maximum energies decrease from 105 MeV (11 nm), over 70 MeV (32 nm) to 40 MeV (101 nm). The temperature fits to the spectra follow the same trend as before with the exception of the circularly polarized case. It has the the highest value since the electrons there are displaced as a single layer with a relatively low energy spread which leads to strong proton acceleration in LS-RPA. The dip in proton number at the very beginning of all the shown spectra accounts for the fact that the proton source are the organic contaminant layers of the targets. Therefore, the particle number is very limited and a significant amount of all available protons gets accelerated.

The information that can be gained from integrated spectra is still very limited. In the following, angularly resolved spectra are shown which yield more information about the target structure during the time of maximum laser intensity.

Angular Emission Characteristics and Energy Spectra

Figure 4.37 displays angular distributions of bremsstrahlung photons, electrons and ions for the four simulations from Fig. 4.34. The images were obtained from the `ParticleCalorimeter` diagnostic of PIConGPU[6] and show relative particle numbers in the respective energy interval distributed over the *yaw* angle, i.e. the angle in the $y-x$ plane[7] which coincides with the polarization plane in the cases a)-c). The distributions were integrated for the *pitch* angle outside of the polarization plane. Only a single circularly polarized simulation was performed with a target thickness of 11 nm and the results are shown in the top row (a*) of Fig. 4.37.

In all cases, the highest energy bremsstrahlung is directed in the general forward direction because high energy electrons accelerated by the laser radiation pressure and $\vec{v} \times \vec{B}$–force are driven through the target bulk where the bremsstrahlung is created. Also, the overall

[6]The particle calorimeter diagnostic maps all occurring particle momenta to an angular bin, regardless if the particle has already left the simulation volume. For all particles still inside the simulation this is done on-the-fly. Once a particle leaves, its emission direction as well as energy contribution to an energy bin is frozen in time.
[7]In PIConGPU the laser propagates along the y-axis

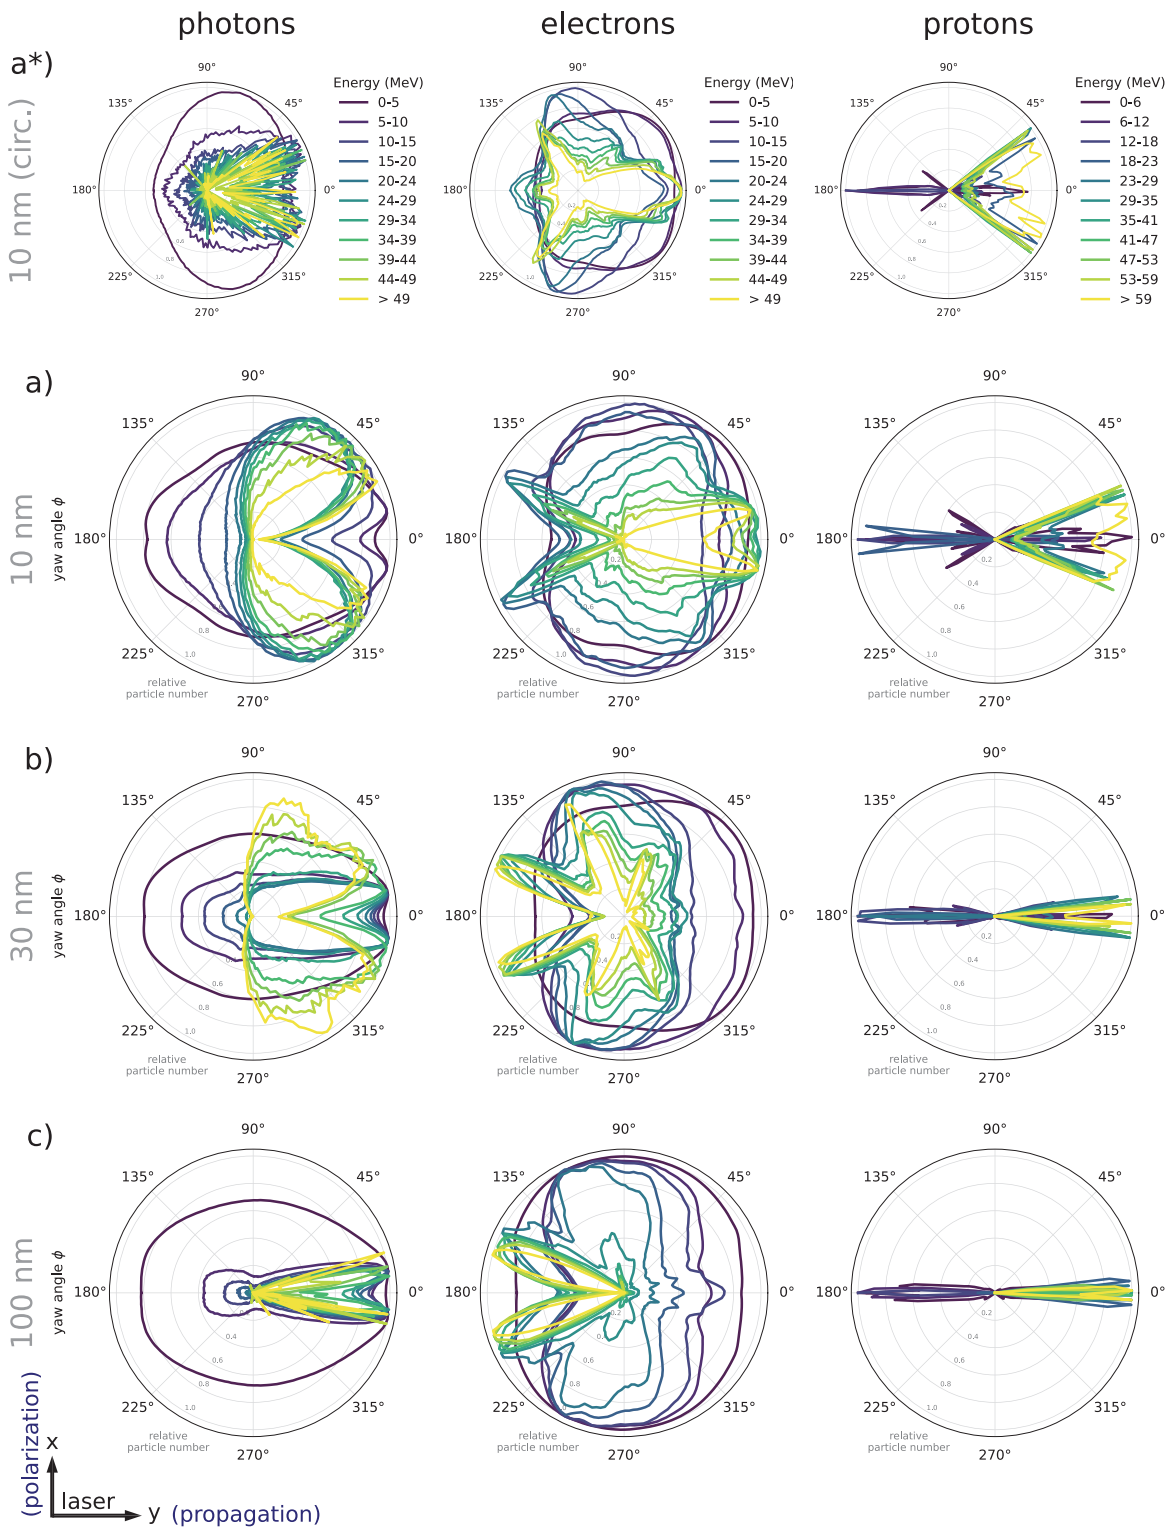

Figure 4.37.: Angular distributions of bremsstrahlung photons (left column), electrons (middle column) and protons (right column) in the the plane of polarization (*"yaw"*, i.e. $y - x$), integrated over the angle around the orthogonal transverse axis (*"pitch"*) at 170 fs after main-pulse-target interaction. Shown are relative particle numbers, normalized to the maximum value in the respective energy interval. Copper targets were irradiated with a linearly (*circularly) polarized 30 fs Gaussian laser pulse of $a_0 = 63.25$ (*PC* in Fig. 4.10b). Target thicknesses are **a)** 11 nm, **b)** 32 nm, and **c)** 101 nm. The legends shown next to the angular distributions of the circularly polarized case in the top row describe the energy intervals for the entire column, respectively.

emission characteristics show that different angles of preference exist for the each of the energy intervals. In both the circular polarization case a*) as well as the linearly polarized cases a)-c) the laser propagation direction under 0° exhibits a dip in particle numbers over almost the complete range of energies. Furthermore, the angular distribution of photons with $\mathcal{E}_\gamma > 44\,\mathrm{MeV}$ were observed to not change anymore after about $t_{\mathrm{rel}} = 42\,\mathrm{fs}$, or $\sim 1.5\,\tau_{\mathrm{L}}$ in every case. The time span between τ_{L} and $1.5\,\tau_{\mathrm{L}}$ was also found to exhibit the highest average electron energies $\langle \mathcal{E}_{\mathrm{e}} \rangle_{\mathrm{crop}} = \int_{\mathcal{E}(\beta\gamma>1)}^{\infty} \mathcal{E} \cdot \mathrm{d}N/\mathrm{d}\mathcal{E}\ \mathrm{d}\mathcal{E}\ \ / \int_{\mathcal{E}(\beta\gamma>1)}^{\infty} \mathrm{d}N/\mathrm{d}\mathcal{E}\ \mathrm{d}\mathcal{E}$ for electrons with momenta above $\beta\gamma = 1$. This shows that the highest energy bremsstrahlung photons are only created from when the laser pulse is strongest, until electron energies peak and then emission quickly ceases, again. The emission probability also scales quadratically with the electron energy, $p(\mathcal{E}_{\mathrm{kin}}, W, \cos\theta) \propto \gamma_{\mathrm{e}}^2$, and so the contribution to the overall bremsstrahlung signal is very strong during that time. As such, these photons provide a strong signal with virtually undiluted information about the most important time period of laser-ion acceleration.

In linear polarization, the bremsstrahlung signal is forked into two lobes, symmetrically pointing away from the axis of laser propagation. This pattern is caused by the $2\omega_L$-bunches that are accelerated at the target front when the normalized laser amplitude satisfies $a_0 > 1$. Similar to the 2D simulation study of Vyskočil, Klimo, and Weber [190] on micron-sized aluminum targets, the emission angles from sub-micron copper targets also generally decrease with increasing foil thickness d. Especially very thin foils become indented earlier and together with electrons already gaining some radial momentum away from the laser axis due to the spatially Gaussian intensity shape, the bremsstrahlung opening angle in the circularly polarized case can be explained as well. The aforementioned study[190] showed the photon emission characteristics only up to 10 MeV for targets irradiated by a laser of $a_0 = 86$. For higher photon energies it becomes apparent that there is a clear energy cutoff above which no bremsstrahlung photons are emitted in laser-backwards direction. That the backwards emission direction (180°) is populated at all for lower energies stems from the fact that electrons can recirculate in the TNSA sheath behind the target and return to again interact with the bulk and produce breaking radiation. Above $\mathcal{E}_\gamma = 25\,\mathrm{MeV}$, almost no counter-propagating photons are found. It is conceivable that a closer study of the backward-directed bremsstrahlung and its cutoff over a more extended parameter scan can give insight that would allow reconstructing the strength of the TNSA sheath and its lifetime, which decides for how long and below which kinetic energy $\mathcal{E}_{\mathrm{e}}$ can recirculate.

The second column of Fig. 4.37 displays angular normalized electron number distributions. A striking feature is again that different energy intervals show very different emission patterns. For the thinnest targets of 11 nm the preferential direction where the most high-energetic electrons are detected is generally forward. For 32 nm and 101 nm, most high-energetic electrons of $\mathcal{E}_{\mathrm{e}}$ are found to be propagating in two very narrow angular lobes between 20 to 25° away from the laser backward direction. An analysis of the time evolution of the detector signal reveals that these lobes only begin become so pronounced once the laser has been mainly reflected. As long as the laser pulse still applies radiation pressure to the target, a front-side sheath is strongly suppressed and electrons expand towards the outer regions of the laser spot where the pressure is smaller. When the laser pressure decreases, the strongly heated target also expands more in laser backwards direction. In laser-forward direction however, the sheath field not only accelerates ions but also keeps electrons back unless their energy is too great to be retained. For this reason, electron numbers in the higher energy intervals are reduced in forward-direction .

The angular distributions of total proton number in all cases, displayed in the third column of Fig. 4.37, show a behavior similar to the photons throughout the different energy intervals. Protons are accelerated in both forward and backward directions but to visibly higher ener-

gies in forward direction. For the lowest energies below 20 MeV, protons are well collimated towards 0° and 180°. The emission angle increases with increasing energy up to 59 MeV. Protons with higher energies are emitted to slightly smaller divergence angles, again. In general, the proton emission half-angles decrease with increasing target thickness from up to almost 40° ($d = 11$ nm circ. pol.) down to 5° (101 nm).

Temporal Analysis of Acceleration Dynamic Encoding Figure 4.38 shows a temporally resolved analysis of the bremsstrahlung distribution shown in Fig. 4.37 in the case of the 11 nm foil. In the upper part (I), distributions for newly created bremsstrahlung photons are shown for chosen time intervals between consecutive outputs of the `ParticleCalorimeter` diagnostic (placed every 5.13 fs) which originally records the accumulated energy deposited over time. The lower part (II) shows the temporal evolution of total photon number N_{ph} (blue line) and its time derivative $\mathrm{d}N_{\mathrm{ph}}/\mathrm{d}t$ (orange line) for photons currently contained in the simulation box. The change in photon number is displayed in a spectral representation $\mathrm{d}^2N_{\mathrm{ph}}/\mathrm{d}\mathcal{E}\,\mathrm{d}t$ in the panel below. This data arises from consecutive differences of the `EnergyHistogram` output which is available every 0.25 fs. The pink area marks bins where the number of particles that leave the box is larger than the number of newly created particles.

In Fig. 4.37 it was already visible that the divergence angle of the high-energy photons coincides well with the opening angle of the highest energy protons for the 11 nm foil. The temporally resolved analysis shows now that between 5 and 25 fs after the main pulse maximum hit the target front bremsstrahlung with energies larger than 49 MeV is produced in these opening angles. From Fig. 4.24 it is already known that between $t_{\mathrm{rel}} = 0$ and ≈ 40 fs protons gain already most of their energy. Therefore, the electrons that produce bremsstrahlung during this time also contribute majorly to the acceleration of protons. However, seeing that medium-high energy photons (around 30 MeV) show a structure with two pronounced lobes of wider opening angles (not unlike butterfly wings) in Fig. 4.38 I b), already hints at larger target indentation and beginning transparency. In this case, electrons accelerated transversely in the laser field already oscillate in relatively dense pre-plasma before being pushed into the overdense region where they detach from the laser and produce bremsstrahlung in more narrow forward direction. Note that "spikier" structures represent less particles than smooth curves in an energy interval. That is because the numbers in each energy band are normalized to the band's maximum number per angle bin. In the subfigures d) and e), i.e. after 35 to 45 fs, high energetic photons are produced less and less. After 98 fs, the production of medium-energy photons has ceased as well. The general emission pattern is also more homogeneous, indicating that only (re-)circulating electrons contribute to the bremsstrahlung signal. Asymmetries with respect to the 0° axis in the bremsstrahlung emission at earlier times are the result of changing laser field strength due to the Gaussian temporal profile and the output interval length of $2\,T_{\mathrm{L}}$, which allows to see pronounced emission during specific laser half-waves. The sharp gradient in $\mathrm{d}N_{\mathrm{ph}}/\mathrm{d}t$ in the upper panel of 4.38 II) can be explained with the fact that many low-energy photons that were produced earlier are leaving the simulation box. Since they are always created in much larger numbers than the high-energy photons (due to the electron bunch spectrum) they dominate this quantity. The $2\,\omega_L$–bunch structure is well reproduced in the creation of bremsstrahlung photons for which there are roughly 6 local maxima every 10 fs, i.e. every 3 laser periods, as Fig. 4.38 II) shows.

Figure 4.39 shows the same analysis as Fig. 4.38 but for the 32 nm foil. Analogous to the 11 nm case, the production of the highest energy bremsstrahlung is limited to the intervals b) through e), i.e. roughly within 1 to 1.5 laser pulse durations after the arrival of the maximum intensity of the laser on target. However, in the 32 nm case the bremsstrahlung is directed more towards the forward direction which indicates that the target plane is still mostly intact

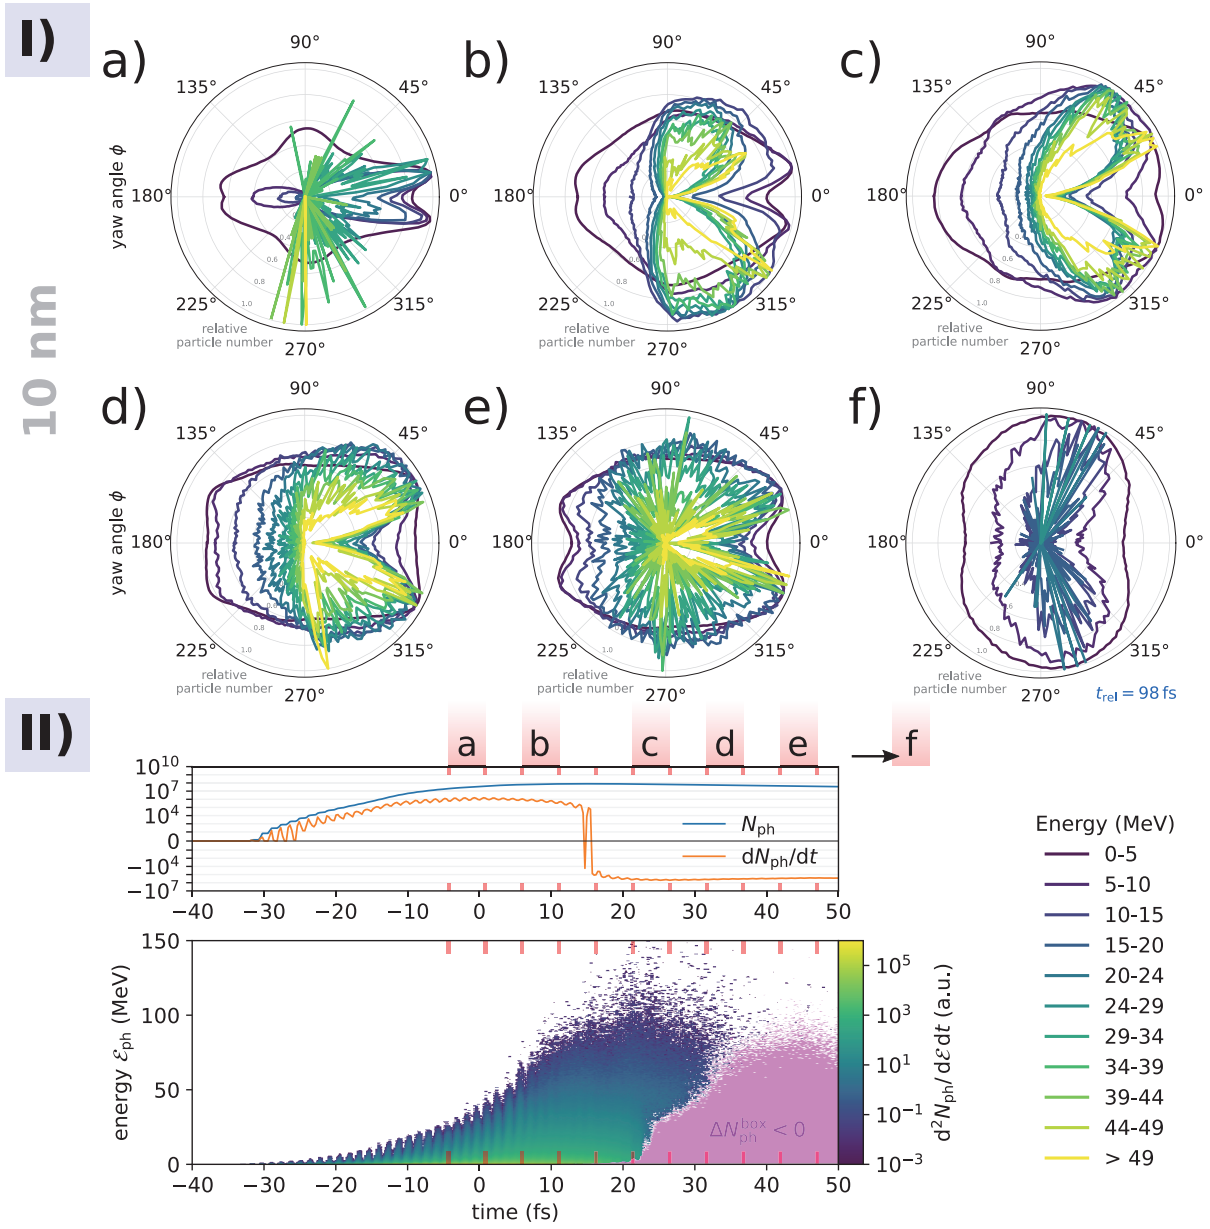

Figure 4.38.: **I)** Angular distribution of newly created bremsstrahlung photons during chosen *time intervals a) – f)* in the 10 nm copper foil PC case. Each polar plot shows all photons that are created between the surrounding consecutive output iterations (red ticks in (II)) which are 5.13 fs apart. Photons are sorted into energy bands (legend below, next to the panels in (II)) and the angular distribution is normalized to the maximum signal within each energy band. **II)** *Lower panel:* Temporal evolution of the time derivative of photon energy spectra that are available every 0.25 fs. *Upper panel:* Temporal evolution of the change in total photon number dN_{ph}/dt (orange) and total photon number N_{ph} (blue) within the simulation box, at every 0.25 fs. The pink area marks bins where more photons leave the simulation box than are newly created. The UHI drive laser is incident under 0° from the left (like in Fig. 4.37).

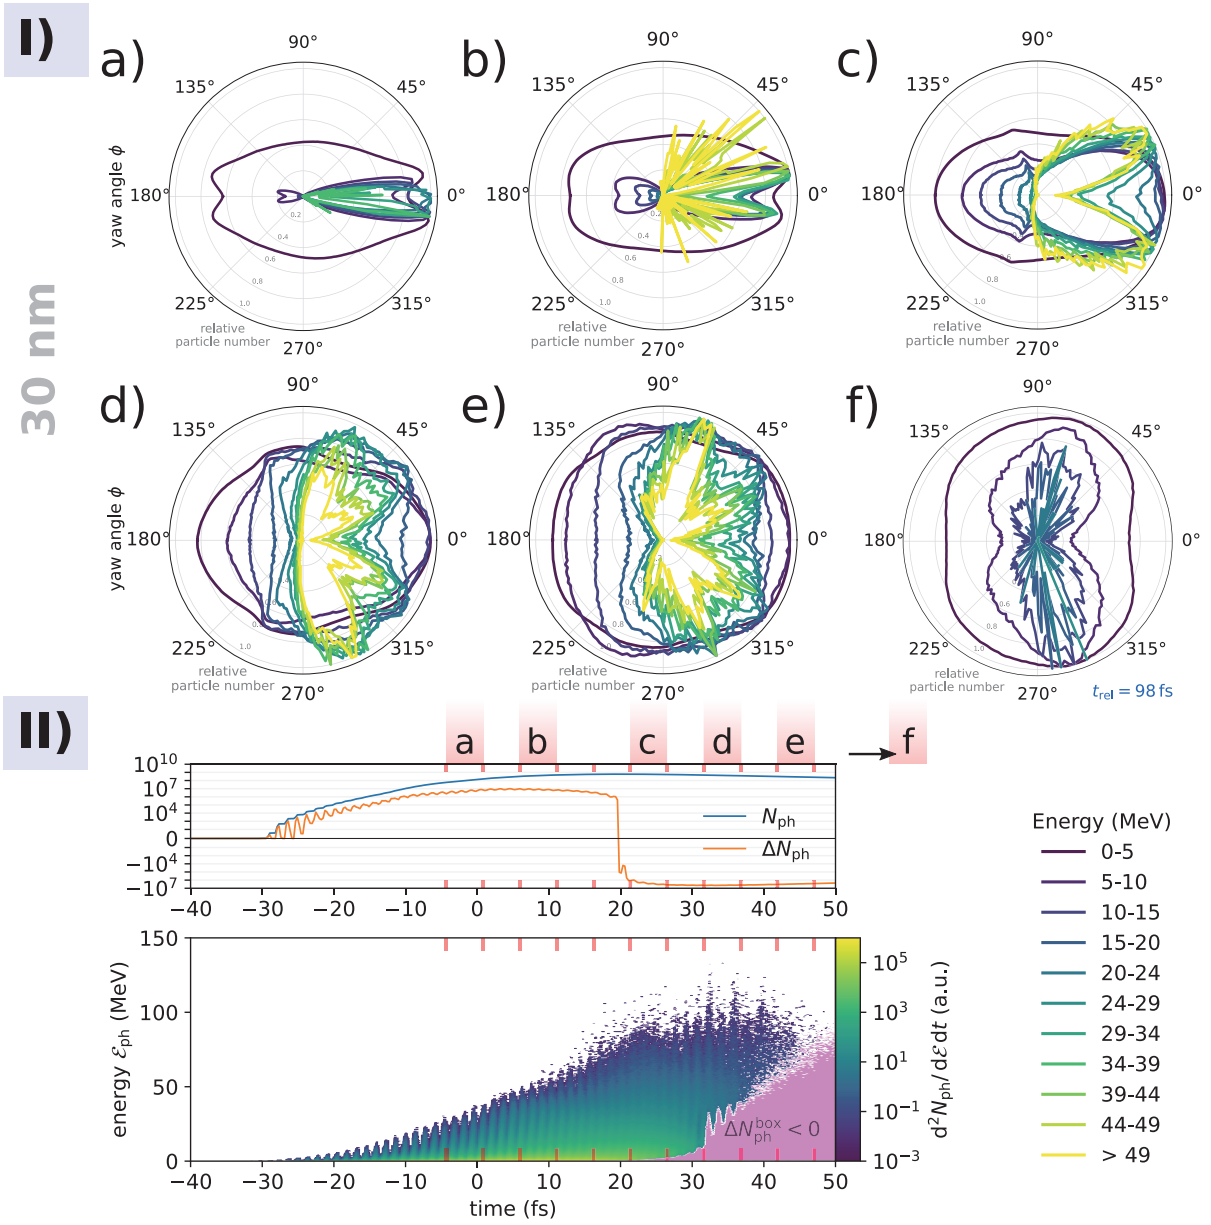

Figure 4.39.: **I)** Angular distribution of newly created bremsstrahlung photons during chosen *time intervals a) – f)* in the 30 nm copper foil PC case. Each polar plot shows all photons that are created between the surrounding consecutive output iterations (red ticks in (II)) which are 5.13 fs apart. Photons are sorted into energy bands (legend below, next to the panels in (II)) and the angular distribution is normalized to the maximum signal within each energy band. **II)** *Lower panel:* Temporal evolution of the time derivative of photon energy spectra that are available every 0.25 fs. *Upper panel:* Temporal evolution of the change in total photon number dN_{ph}/dt (orange) and total photon number N_{ph} (blue) within the simulation box, at every 0.25 fs. The pink area marks bins where more photons leave the simulation box than are newly created. The UHI drive laser is incident under 0° from the left (like in Fig. 4.37).

and not as indented as with the 11 nm. Only in subfigures c) and d) the "butterfly wing" structure from before starts to become visible, speaking again for the onset of target transparency. During the time between b) and c) that is most crucial for the energy gain of accelerated protons, it can be seen that a majority of electrons produce medium-energy bremsstrahlung in more narrow forward direction than in Fig. 4.38 I b), c). These preference angles are the ones that the fastest accelerated protons imprint on and which do not change much later since the protons will already be moving too inertly to be influenced by later, lower density, regions of the electron sheath. This analysis confirms the general impression about the laser-solid interaction dynamics that already the photon density slices in Fig. 4.33 for the 11 and 32 nm cases created.

In the 101 nm case of Fig. 4.37, any bremsstrahlung but the lowest shown energy bands is directed in forward direction with only very narrow opening angles. Also, there are no "butterfly wings" to be seen which indicates that the target was intact for all of the time that energetic electrons probed the bulk density. Therefore, the opening angles of the high energy proton emission and the highest energy photons mainly coincide, as well.

Conclusion

This section has shown how in-situ generated bremsstrahlung could possibly be used in the future as a *bremsstrahlung tomography* diagnostic of the target during the previously inaccessible time period shortly after the laser peak interaction with the target. When both the bremsstrahlung and proton emission are known, there is promising indication that the state of the target during maximum laser intensity irradiation (and shortly after) could be reconstructed. This diagnostic, supplemented by existing on-shot diagnostics will hopefully lead to an even more accurate picture of the interaction dynamics that would in turn allow for better control over accelerated proton beam parameters in the future.

Unfortunately, the detection of multi-10 MeV photons is not a simple feat. Either, scintillation detectors would require a large amount of absorber material to shield away the lower-energy photons or detection techniques like they are used in high-energy physics require large detection volumes. Both factors make it likely unfeasible to place an on-shot detector into the target chamber while the low particle number of high-energy photons make the detection further away outside the chamber very challenging. On-shot detectors developed at HZDR are a work-in-progress diagnostic tool and variants currently reaching up to ~ 10 MeV have been recently employed in laser-matter interaction experiments at DRACO and other facilities[192]. Since realistic target chamber geometries usually only allow for very specific angles to be covered for X-ray detection, accompanying simulation support will continue to be important for the selection, planning and interpretation of experiments.

With the advent of X-ray free electron lasers (XFELs) in recent years, ultra-brilliant X-ray sources have become available that produce coherent X-ray beams that are short enough ($\sim$ fs) and carry enough photons ($N_{\mathrm{ph}} \sim 10^{12}$) to feasibly conduct pump-probe experiments of laser-driven solid density plasmas. Like Sec. 4.3.2, the next section shows simulation work performed within this thesis to explore these novel diagnostic methods for the probing of transient laser-driven plasma processes on nanometer scales.

4.3.2. Probing Plasma Expansion Dynamics of Structured Foil Targets with SAXS

Ion acceleration from laser-driven solid density plasmas includes transient phenomena and non-equilibrium physics that occur at femtosecond and nanometer scales, such as the target ionization at the critical density surface or the creation of plasma waves. After being inaccessi-

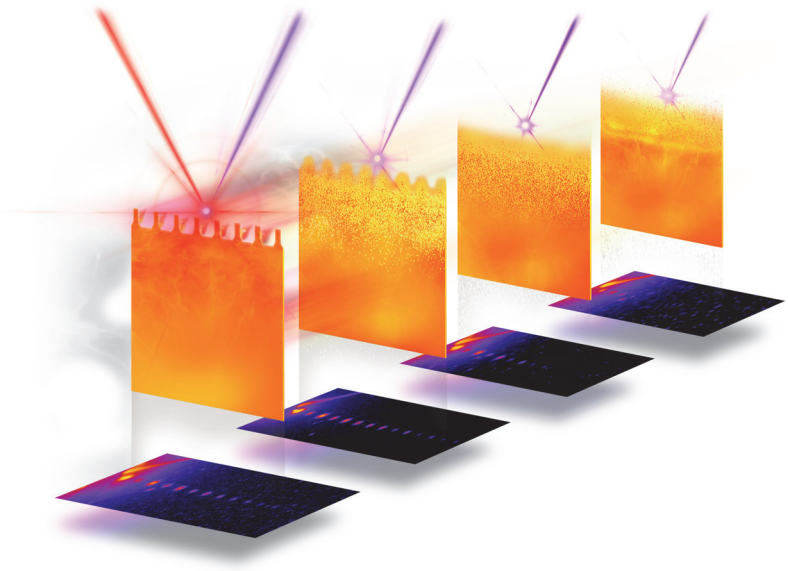

Figure 4.40.: Artist impression of Small-Angle X-ray Scattering (SAXS) experiment performed at SLAC shown later in Fig. 4.42. A silicon foil with pre-inscribed grating surface is driven by a high-intensity laser pulse (red) and subsequently the evolving plasma structure is probed by a perpendicularly impinging X-ray pulse (blue, different delays expressed by the several slices). Simulation images created by the author. Scattering patterns below were obtained in the experiment and indicate the gradual disintegration of the sharp grating surface structure. Image originally in HZDR press release [323].

ble to experimental diagnostics for a long time, these scales have now simultaneously become available for study due to the unparalleled peak brilliance of X-ray free electron lasers (XFELs). In order to gain deeper understanding about the complex process of laser-ion acceleration, it is sensible to study a more basic process first. The expansion of a heated solid-density plasma into vacuum is the prerequisite to the most robust ion acceleration process, TNSA, and is therefore a suitable phenomenon to characterize. In particular, the temporal evolution of the plasma scale length L_p is of interest, as it allows direct validation of particle-in-cell simulations which are widely adopted to model laser-ion acceleration processes. In two pilot experiments, Kluge et al.[G16] showed the feasibility of Small Angle X-ray Scattering (SAXS, see 4.3.2) on plasmas for validating physics models and simulations on the aforementioned scales. In the first experiment, aluminum wires were irradiated with a circularly polarized Ti:Sa laser pulse ($a_0 = 2.45$, $\tau = 60\,\mathrm{fs}$) and the sharpness as well as the degree of indentation of the wire's surface were then assessed by probing the target perpendicularly with the LCLS X-ray beam. In the second experiment, a shock was induced in a multi-layer silicon slab via a nanosecond laser pulse. It was possible to measure the change in the crystalline structure from the target compression that way. Both experiments already established a spatial resolution on the order of nanometers. To achieve a significant sensitivity to the femtosecond timescales of nanometer changes in plasma expansion, silicon targets were specifically designed to exploit the increased scattering signal and characteristic scattering pattern of a pre-inscribed surface grating.

In the following, preparatory 2D PIC-simulations are presented that tested the feasibility of the SAXS experiment performed at LCLS which was published in [G8]. Resulting from these simulations, the expansion profiles, and –velocity were extracted and the necessary detector sensitivity, q–range and time resolution could be derived. As a main result next to the general surface-structured expansion dynamics, priorly unexpected plasma jets were found that have a major influence on the applicability of simple analytic models used for forward calculation and

fits of the scattering patterns. With this knowledge, several experimental shots were later identified that strongly hint at the presence of these plasma jets. In a second step, post-experiment simulations were used to interpret the experimental results and explain the spatiotemporal expansion dynamics that were observed.

X-ray Scattering from Plasmas

X-rays have long been an important tool for studying the structure of matter which is mainly owed to their ability to penetrate materials of solid density due to short wavelengths on the order of Ångströms (1 Å $= 10^{-10}$ m). As such, significant scientific discoveries were and are continued to be made regarding the lattice structure of crystals, the architecture of our DNA and more recently, complex protein structures and even functional building blocks of living organisms (for a review, see [324]). But unlike the direct photographic X-ray images that first revealed the interior of the human body and its bones[325], all of the above benefit from the property of X-rays to resolve structures on the nanometer to micrometer scale via X-ray scattering and diffraction. Here, the initial propagation direction of a particle or wave, respectively, is changed due to the interaction with an obstacle in its path and images in reciprocal space are created on a detector that measures the scattered intensity.

Small-Angle Scattering with hard X-rays is a well-known technique[326, 327] in the retrieval of structural information between 1 to 1000 nm in the fields mentioned above, but was only recently proposed for application to laser-driven solid-density plasmas[50, 51]. First experiments already demonstrated that the technique qualifies to study the ultra-fast (fs), ultra-small (nm) scales of transient plasma processes[G8, G16].

Thomson-Scattering from Free and Bound Electrons The basis for SAXS from solid-density plasmas is the elastic scattering of X-rays from single charged particles. The Compton effect[328] describes the energy that is transmitted from the scattered photon to the charged particle, $\Delta\mathcal{E}_{\mathrm{ph}} = \Delta p_{\mathrm{ph}} \cdot c = \left(p_{\mathrm{ph,i}}^2 \cdot c^2/m_{\mathrm{S}}\right)\left(1 - \cos\left(1 - 2\theta\right)\right)$ in the rest frame of the charged particle with mass m_{S}. The main contribution to the scattering signal from such plasmas comes from free and bound electrons. Under small scattering angles ($\sin\left(2\theta\right) \approx 2\theta$) and with X-ray photon energies $\mathcal{E}_{\mathrm{ph}}$ of a few kiloelectronvolts the scattering can be treated as fully elastic. For extended information on the following derivation see [G16] or [287, 327]. The interaction is now independent of the particular value of the photon energy and can be described by the *Thomson scattering* cross section.

$$\left(\frac{\partial\sigma}{\partial\Omega}\right)_{\mathrm{Th}} = r_{\mathrm{e}}^2 \frac{1}{2}\left(1 + \cos^2\left(2\theta\right)\right) \approx r_{\mathrm{e}}^2 \tag{4.8}$$

The *classical electron radius* $r_{\mathrm{e}} = e^2/4\pi\varepsilon_0 m_{\mathrm{e}} c^2 \approx 2.82$ fm sets the *scattering length* scale and the rest of the right side of Eq. 4.8 is the *polarization factor* that scales the amplitude of the scattered wave ϕ_{out}. In first Born approximation, the phase difference of two scattered plane waves whose scatterers are separated by $\vec{r}$ can be expressed via $\varphi = -\vec{r}(\vec{k} - \vec{k}_0) \equiv -\vec{r}\vec{q}$.

$$\phi_{\mathrm{out}}(\vec{q}, t) = \frac{r_{\mathrm{e}}}{d} \int \phi_{in}(\vec{r}, t'(z))\, n(\vec{r}, t'(z))\, \mathrm{e}^{-i\vec{q}\vec{r}}\, \mathrm{d}^3 r \tag{4.9}$$

$$n = n_{\mathrm{e,free}} + \sum_{\Gamma} \Gamma n_{\mathrm{ion},\Gamma} + \sum_{\gamma} n_{\mathrm{ion},\gamma} F_{\mathrm{ion},\gamma}^{\mathrm{res}}\left(\mathcal{E}_{\mathrm{ph}}\right) \tag{4.10}$$

If each photon is scattered only once and the distance to the detector is much larger than the distance between individual scatterers, the scattered waves are given by Eq. 4.9. The

density of scatterers $n(\vec{r}, t'(z))$ is defined with respect to an arbitrary origin at $t' = t + \frac{z}{c}$. In conventional SAXS, the variation of n with time is negligible. While it is still possible to get structural information of a plasma sample before it is destroyed, the notion is valid that due to the femtosecond timescales of interest, this change should possibly be taken into account. Section 4.3.2 presents Monte-Carlo photon scattering simulations where this was realized in a first approach.

Equation 4.10 illustrates that in general the density of scatterers is composed of both free electrons and a number of $\Gamma = Z - Q$ electrons that are still bound to the ion, the latter of which might be susceptible to resonant *bound-bound* or *bound-free* transitions in the ionic potential, which depend strongly on the X-ray photon energy $\mathcal{E}_{\mathrm{ph}}$ and the number of shielding electrons, y. The scattering intensity I_{out} is the square of the scattering amplitude and the product of *structure*– and *form factor* of the scatterer distribution.

$$I_{\mathrm{out}}(\vec{q}) = \left(\frac{r_{\mathrm{e}}\phi}{d}\right)^2 \left(\mathcal{S}_{\mathrm{free}}(\vec{q})\mathcal{F}_{\mathrm{free}} + \mathcal{S}_{\mathrm{bound}}(\vec{q})\mathcal{F}_{\mathrm{bound}}\right)^2 \tag{4.11}$$

$$\mathcal{S}_{\mathrm{bound}} = \int n_{\mathrm{i}}(\vec{r}\exp\left(-i\vec{q}\vec{r}\right))\,\mathrm{d}^3 r \tag{4.12}$$

$$\mathcal{F}_{\mathrm{bound}} = \int_{\mathrm{ion}} n_{\mathrm{e,bound}}(\vec{r}\exp\left(-i\vec{q}\vec{r}\right))\,\mathrm{d}^3 r \tag{4.13}$$

It is most important to note that the scattering signal is not sensitive to absolute scattering length densities[8] but *modulations* of it. The structure factor $\mathcal{S}(\vec{q})$ describes the degree of spatial correlations between these modulations and the form factor $\mathcal{F}(\vec{q})$ holds information about the averaged geometrical shape of a single feature. SAXS on solid density plasmas relies heavily on the coherence of the source and has therefore a lot of similarity with the conventional *coherent X-ray diffraction* (CXD). Apart from any resonant transitions, free and bound electrons contribute similarly to the scattering signal which means that it mainly depends on the total electron density fluctuations $\Delta n_{\mathrm{e,total}}$ in the plasma. The resonant interaction depends on the complex form factor $\mathcal{F}^{\mathrm{res}}_{\mathrm{ion},y} = \mathcal{F}'_{\mathrm{ion},y} + i\mathcal{F}''_{\mathrm{ion},y}$. The imaginary part is proportional to the absorption and so to the opacity which can be calculated with 0D atomic physics codes like SCFLY[184, 185] based on LTE plasma conditions. The real part is then usually inferred from the imaginary part via the Kramers-Kronig relations if the energy dependence of the absorption cross-section is known (see Ch. 8.3 in [327]). For an experiment that aims at studying just the spatial structure or plasma modulations, the X-ray photon energy is tuned to be far off from any resonances. But resonant scattering (or RCXD) also promises to deliver spatial information about plasma temperatures and ionization states[51] which could be a direct benchmark to the collisional and atomic physics modules in PIC codes.

Reconstruction of Scattering Images and the Importance of Forward Calculations

With the assumptions of X-ray source coherence, small scattering angles and only a single scattering event per photon, one finds that the scattering amplitude corresponds to the Fourier transform of the scattering length densities.

$$\phi_{\mathrm{out}}(\vec{q}) \sim r_{\mathrm{e}} \cdot \mathrm{FT}\left[n(\vec{r})\right] \tag{4.14}$$

$$I(\vec{q}) = \phi_{\mathrm{out}}\phi^*_{\mathrm{out}} \tag{4.15}$$

$$= I_0 \cdot r_{\mathrm{e}}^2 \left| \int n(\vec{r}) \mathrm{e}^{-i\vec{q}\vec{r}}\,\mathrm{d}^3 r \right|^2 \tag{4.16}$$

[8] i.e. the product of the number density of scatterers and the individual scattering length

As mentioned above, the scattering signal is the absolute square of the scattering amplitude which results in the *loss of the phase information* $e^{-i\vec{q}\vec{r}}$ once the intensity is measured on a detector. This problem is well-known in the X-ray scattering and diffraction community. Several techniques exist to retrieve this information and reconstruct the full 3D structure of the sample. Already established optical probing techniques overcome the phase problem by either probing the same target multiple times, like it is done with ptychography[9], or by added initial information from a regular substructure, like in crystal diffraction. Alternatively, holographic methods use the interference of a probe and a reference beam and allow to either record or reconstruct the 3D phase information. X-ray Fourier-transform holography (FTH) is already a robust single-shot measurement technique but the known reconstruction algorithms suffer greatly from missing information such as it would occur for high-brilliance XFELs where the unscattered beam is typically cut out by a beam block. X-ray FTH is in principle promising for the study of transient phenomena and even though several challenges remain, continous progress is made that hopefully allow its use for study of laser-solid plasmas in the future[329].

Due to the extremely transient nature of plasmas that are driven by ultra-high laser intensities, they neither exhibit an inherently regular structure, nor do they form structures that exist long enough to be readily probed multiple times. Therefore, it is important that not only the experimental methodology but also simulation tools are developed for these new experiments, to predict both the target dynamics as well as the scattering signals. For many cases, forward calculations using Fourier transforms and simple analytic models provide a valuable tool for first studies of plasma expansion into vacuum that exploit the enormous spatiotemporal resolution that SAXS experiments are able to provide. For these experiments, targets have to be specifically designed to increase sensitivity for the expansion process while lowering the signal-to-noise ratio. While previous studies have shown that front-side structures can enhance the laser absorption and X-ray emission[330], higher-harmonics generation[331] and creation of resonant plasma waves[332], their main purpose for this study lies elsewhere. The regular *ridge-valley* structuring of the surface acts like a simple line grating when the total electron density is projected in XFEL-probing direction.

A schematic view of a total electron density grating profile and its defining parameters are shown in Fig. 4.41. The ridge edges, softened via plasma expansion ensuing the interaction with the high power laser, can be analytically described with sigmoid functions.

$$n_{\mathrm{r}_\perp} = \frac{n_{\mathrm{e},0}}{2}\left(\mathrm{erf}(x/(\sqrt{2}\sigma)) + 1\right) \tag{4.17}$$

or approximately

$$n_{\mathrm{r}_\perp} = n_{\mathrm{e},0}(1 + \exp(-x/\sigma))^{-1} \tag{4.18}$$

The simple forms of Eqs. 4.17 and 4.18 agree very well with the numerically obtained expansion profiles seen later in this section and therefore allow to describe the grating fully with only 3 parameters. These are the *ridge width b*, the *edge softening scale length σ* and the ridge periodicity, or *pitch g*.

When such a target is irradiated with an ultrashort X-ray pulse, it produces a well-known and analytically defined scattering pattern on the detector. In reality the idealized solution from an infinite line grating is additionally convolved with the finite illumination of the target, resulting in a finite grating pattern $G(q)$, the scattering pattern of a single slit $S(q)$ due to the finite width of the ridges and, vitally, the effect of the edge sharpness $E(q)$ of the ridges. Equation 4.19

[9] post-measurement, a numeric algorithm extracts the phase information from a superposition of all measurements

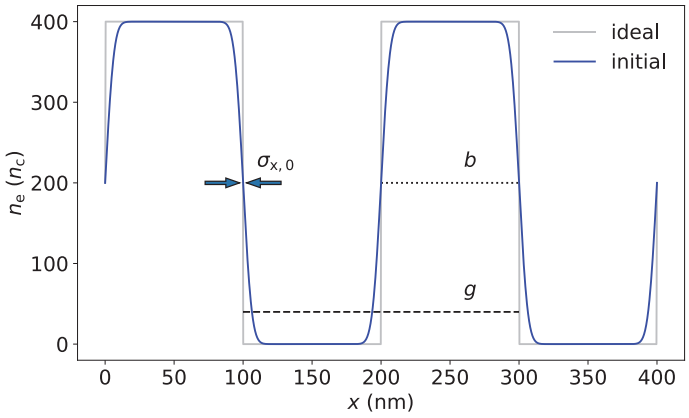

Figure 4.41.: Schematic depiction of the total electron density grating profile used in 2D PIC simulations and its characteristic parameters that define the analytic model used in forward calculations and fits for the experimental data obtained in the experiment at LCLS. Plasma expansion following high-power laser irradiation softens the ridge edges of the grating and is modeled via the sigmoid function $1/(1+\exp(-x/\sigma))$ as an approximation to the error function profile $\frac{1}{2}\left(\mathrm{erf}(x/(\sqrt{2}\sigma))+1\right)$. To account for realistic surface roughness, an initial scale length of $\sigma_{x,0} = 5\,\mathrm{nm}$ was configured.

shows the one-dimensional function that fully characterizes the expected signal at unit length distance. In Fourier space, the convolution of the aforementioned influences can be expressed as the product of the normalized Fourier transform squares of each structure, that all depend on the scattering vector q.

$$I(q) = \frac{I_0}{q^2} \cdot G(q) \cdot S(q) \cdot E(q) \tag{4.19}$$

with (4.20)

$$I_0 = I_X r_e^2 \left(2N\frac{N_b}{b}\right)^2 \tag{4.21}$$

$$G(q) = \left[\frac{\sin(Nqg/2)}{N\sin(qg/2)}\right]^2 \tag{4.22}$$

$$S(q) = \sin^2\left(\frac{qb}{2}\right) \tag{4.23}$$

$$E(q) = e^{-q^2\sigma^2} \tag{4.24}$$

With a finely controlled delay between pump-pulse and X-ray probe, the temporal evolution of the decrease in the edge-sharpness of the ridges, i.e. the plasma expansion, can be directly inferred from the measurement. The edge scale length σ can be retrieved from fits of the envelope $E(q) = e^{-q^2\sigma^2}$ of the overall scattering pattern that suppresses the grating peaks at higher q-values. With this, predictions for plasma expansion by particle-in-cell simulations can be directly benchmarked and, hopefully, better modeling will allow for quantitative predictions of the averaged electron energy and ion sound speed that enable accurate quantitative predictions for the ion energies in TNSA.

PIC Simulation of Grating Targets Irradiated with Near-Relativistic Laser Pulses

In preparation for the experimental realization of a pump-probe SAXS experiment on a laser-driven grating target, the author performed a simulation study on the *Hypnos* cluster[10] at HZDR. The simulations confirmed the feasibility of the experiment[G8, G9] that was then performed at the MEC end station of LCLS at SLAC[333]. Figure 4.42 displays the general setup of the experiment, a scanning electron microscopy image of the target surface, and shows a timeline of typical scattering patterns obtained in the campaign.

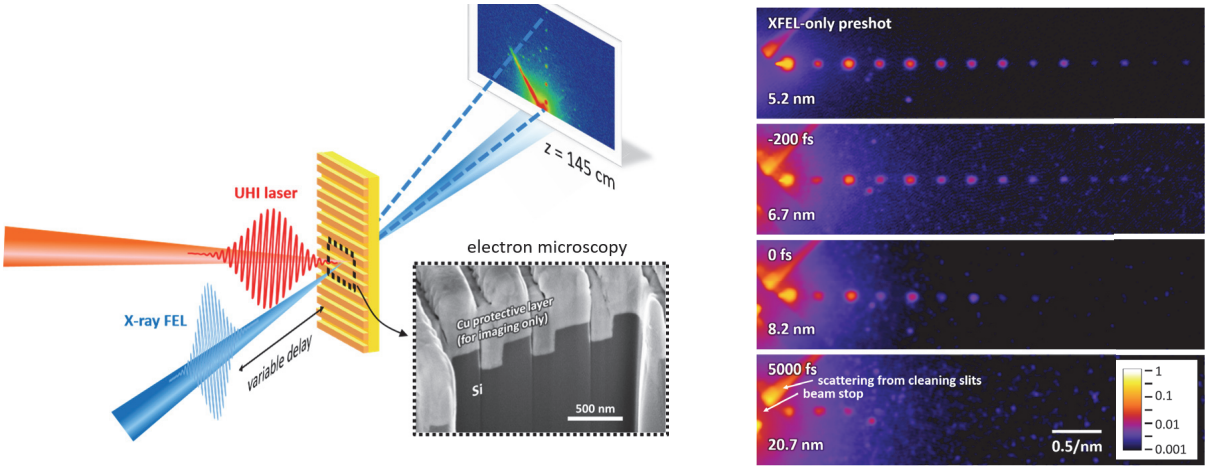

Figure 4.42.: **Schematic of the experimental SAXS setup. left** The near-infrared drive laser pulse is focused on a silicon grating target where it creates a plasma that is subsequently heated. With a controlled delay ($\Delta t_{\text{delay}} = \pm 119$ fs), the X-ray pulses probe the plasma dynamics. The pump laser is incident under 45° and the probe laser is offset by 90° in the same plane that is parallel to the grating ridges. Both lasers were polarized perpendicularly to the grating ridges and the pump-probe plane. A scanning electron microscopy image of a representative target sample with grating period 500 nm is shown in the inset (targets were covered with Cu only for imaging). **right** Examples of scattering patterns of a grating with period $g \cong 300$ nm and ridge width $b \cong 150$ nm for different XFEL delays (rotated by 90°, normalized to the respective intensity in the second scattering peak - the first one was sometimes over-saturated). Originally in [G8].

The questions that were answered in detail concern the performance of the planned grating targets, different target materials and laser intensities, varying pulse durations and features in the laser contrast. For these, the plasma expansion evolution was observed, the expansion speed was measured and the necessary temporal resolution in the experiment could be determined. Of special importance is the finding of an unexpected expansion feature[11] that is stable against the change of parameters and provides unique opportunities to further diagnose the expansion process even though it complicates the analysis of the scattering signals because it ends the applicability of the single-grating model past a certain point in the expansion history. The 2D particle-in-cell code PICLS[54, 182], including field and impact ionization [94], was used for the simulations. [12]

Pre-Experiment Results Three different target compositions were tested: *T1* - pure silicon targets, *T2* - silicon gratings buried under a layer of copper, *T3* - same as T2 but flipped with a flat silicon front side and a grating on the back that is covered with copper. However, all considerations described in this chapter are based on the pure silicon gratings T1 for which

[10] decommissioned in 2020, since 2018 gradually superseded by the supercomputer *Hemera*
[11] see dedicated discussion of the electron jet feature at the end of this section
[12] See the appendix B for access to a copy of the code and the input parameter sets.

the studies, both experimentally and in simulation, were most extensive. Table 4.3 shows the physical parameters that were varied for this study.

a_0	τ_L (fs)	n_e (n_c)	$\sigma_{x,0}$ (nm)	$L_x \times L_y$ (µm × µm)	$b \mid g$ (nm)	t_{sim} (ps)
$3.0 \cdot 10^{-2}$ (pre*)						
0.216					25 | 50	
0.257	45		**5**		50 | 100	**2**
0.305	**82**	**403**	10	**1.6 × 2.4**	**100 | 200**	5
0.38	30 (pre*)		50	6.4 × 12.8	200 | 400	20
3						

Table 4.3.: Physical properties of the preparatory 2D PICLS parameter scan of surface-structured silicon foils performed on Hypnos at HZDR. The bold-faced values are the ones for the simulation that is most representative for the experiment. Experimental and simulation results are reported in [G8, G9]. The values marked with *(pre*)* signify the pre-pulse study that was based on the initially specified laser parameters which later were corrected for the lower values from the experiment.

All simulations were performed in a *quasi-1D* manner, where the transverse simulation size was reduced in order to increase performance while keeping the resolution high. Since the target surface structure is periodic, the box size was chosen to be an integer multiple of the grating pitch g and the transverse boundary conditions were set to periodic. The assumptions for this manner of simulation are that the simulated volume is in the very center of the laser spot, while the laser spot size is also large compared to the transverse box size. Therefore, the laser pulse was simplified to a plane wave with only an explicit temporal pulse shape. These assumptions were also valid for later simulations that followed the experiment since the actual spot sizes at MEC were 16 µm × 30 µm for the near-infrared drive laser and 5 to 10 µm for the LCLS XFEL probe.

The first batch of simulations started out with a laser a_0 of 3 and pulse duration $\tau_L = 45$ fs, since this was the initial expectation of the maximum laser amplitude during the experiment at MEC. During the simulation campaign and throughout the analysis phase after the experiment, the exact laser main pulse parameters and the contrast curve were subject to several changes. The campaign at MEC was the first experiment that required a tight IR laser focus from the laser that was still in its commissioning phase and the Sequoia contrast measurement tool had broken down shortly before the beamtime. The early parameter scan from table 4.3 was for the envisioned laser specifications which had to be corrected for lower intensity values later. Thus, the electron temperatures, expansion speeds and density gradient evolution observed in the first simulations were used as a first estimate of the upper energy limit to the dynamics that were likely to occur in the experiment, until the final laser measurements were made later.

Fig. 4.43 shows two example cases of the temporal evolution of average electron kinetic energy T_e and *free* electron density n_e^{free} at different grating feature sizes b of 100 nm (4.43a) and 200 nm (4.43b), where $b = h = g/2$ applies for both cases. The depicted quantities represent lineouts through the grating ridges parallel to the target plane at the height of $h/2$. Already these lineouts show the emergence of a transient grating structure (marked by pink arrows) between the original grating ridge positions, but at different times with respect to the arrival of the laser maximum. These structures also become visible in the projection of the total electron density in probe direction. Otherwise, the general evolution and shapes of the density distributions were observed to be similar between different grating feature sizes.

Finally, as reported in the methods section of [G8], a spatially plane wave laser ($\lambda_L = 800$ nm) with Gaussian temporal shape, $a_0 = 0.38$, $\tau = 80$ fs was used. The grating pitch g was varied

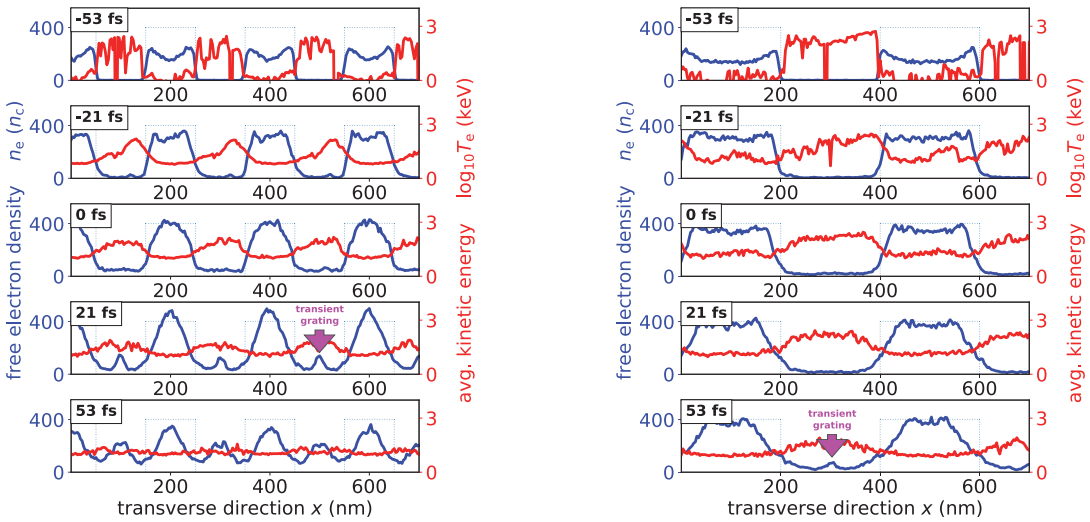

(a) $(b\,|\,g) = (100\,|\,200)$ nm, $h = 100$ nm.

(b) $(b\,|\,g) = (200\,|\,400)$ nm, $h = 200$ nm.

Figure 4.43.: Example lineouts through $(b\,|\,g) = (100\,|\,200)$ nm (*left*) and $(200\,|\,400)$ nm (*right*) gratings at $h/2$ of the grating height h from 2D PICLS simulation results at different points in time t_{rel} relative to the laser maximum on target. **blue:** Free electron density n_e in critical densities $n_c = 1.74 \cdot 10^{21}$ cm^{-3}. Perfect grating total electron density is shown in gray dotted lines. **red:** Free electron average kinetic energy (logarithmic) $\log_{10} T_e$. Laser parameters: $a_0 = 3, \lambda_L = 800$ nm, $\tau = 45$ fs. Single example features of emerging transient grating structures are marked with pink arrows, respectively.

between 50, 100, 200 and 500 nm, the ridge width b and height h were chosen to be $g/2$. Initially, it was also reported that the laser contrast would have a pre-pulse at $t_{rel} = -10$ ps of $I_{pp}/I_0 = 10^{-4}$ intensity. In order to resolve the plasma oscillations well, resolutions of $\lambda_L/768$

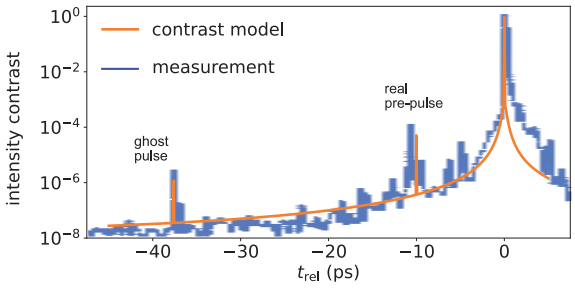

Figure 4.44.: Early contrast model built from measurements of the high-intensity Ti:Sa laser at MEC. The early pre-pulse at –37 ps was identified as a *ghost pulse*, a mirrored image of a post-pulse with the square root of the actual intensity. The later pre-pulse at –10 ps was likely real and its influence was studied in the earlier campaign where $a_0^{main} = 3$.

and $\lambda_L/512$ per cell were chosen. For the silicon targets, the maximum electron density at full ionization is $n_e = 403\, n_c$. Using the relation $\lambda_{pe} = \sqrt{n_c/n_e}\lambda_L$, these resolutions correspond to 38.4 and 25.6 cells per electron plasma wavelength λ_{pe}, respectively. Since in PICLS the length of the time step is directly coupled to the cell width in a 1 : 1 ratio, the temporal sampling of the plasma period is identical. Due to the accumulation of small numerical errors and impending instabilities of the algorithms over millions of iterations, the 10 ps history prior to the main pulse is usually not simulated with PIC. To estimate the influence of the ∼ 30 fs pre-pulse at –10 ps on the main interaction, dedicated simulations with just the latter were performed as well.

At four orders of magnitude less intensity than the main pulse, the pre-pulse does still have $10^{15}\,\mathrm{W\,cm^{-2}}$ (in case $a_0 = 3$) and be therefore above the ionization threshold at around $10^{14}\,\mathrm{W\,cm^{-2}}$. In dedicated simulations with PICLS, a sigmoid shape of the expanding plasma density of the form $n(x) = n_0\,(1 + \exp\{-x/\sigma\})^{-1}$ was observed. Therefore, the longitudinal pre-expansion with the scale length σ_z was varied as shown in Fig. 4.45. Later, the whole grating was convolved in all directions with this s-shape to model the initial surface roughness of $\sigma_{x,0} = 5.4\,\mathrm{nm}$ that was measured in the XFEL pre-shots in the experiment.

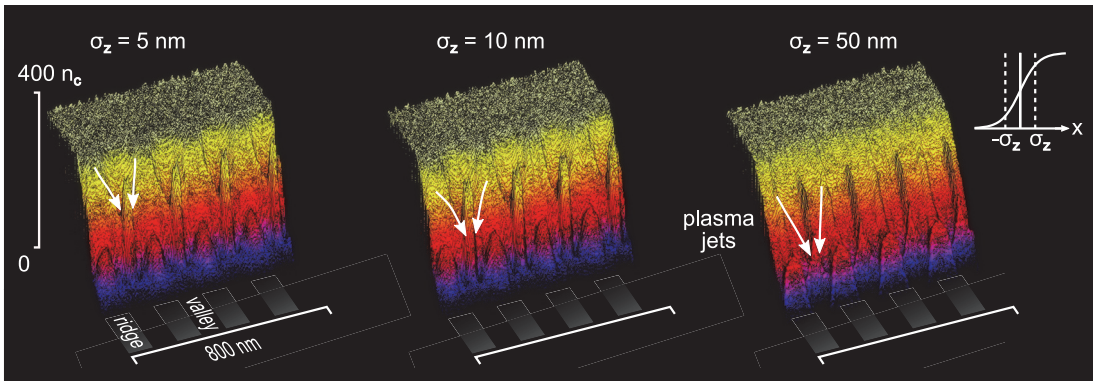

Figure 4.45.: Free electron density at $t_{\mathrm{rel}} = 192\,\mathrm{fs}$ after the laser pulse ($\tau = 45\,\mathrm{fs}$, $a_0 = 3$) would hit the unexpanded, sharp grating surface. Pure silicon grating targets are pre-expanded in target-normal direction with $n(x) = n_0\,(1 + \exp\{-x/\sigma_{\mathrm{L}}\})^{-1}$, and (*l. to r.*) $\sigma_{\mathrm{L}} = 5\,\mathrm{nm}$, 10 nm and 50 nm. Simulations performed with PICLS[54].

Given the microscopic geometry of the target, two profiles with characteristic expansion scale lengths L_{p} can be defined. One is in target-normal direction and the other is on the side walls of the grating ridges, in x-direction parallel to the foil plane. Analytically, they can be be expressed by sigmoid functions such as the error function. However, the scattering signal depends on the projection of the electron density variation in XFEL-probing direction, here called z.

$$n(z) = \frac{A}{2}\left[1 + \mathrm{erf}\left(\frac{z}{\sqrt{2}\sigma_z}\right)\right] \tag{4.25}$$

$$n(x) = \frac{A}{2}\left[1 + \mathrm{erf}\left(\frac{x}{\sqrt{2}\sigma_x}\right)\right] \tag{4.26}$$

The scaling parameters σ_z and σ_x in Eqs. 4.25 and 4.26 need to be transformed to the plasma scale lengths $L_{\mathrm{p,x/z}}$ after extraction from the fits with respect to the probing angle.

Throughout all simulations the following dynamics were observed. When the grating is irradiated by the laser pulse its surface is ionized and the plasma expands from all edges of the grating. While the free electron density increases at these edges (as seen in Fig. 4.43), the total electron density which is relevant for the Thomson-scattering signal keeps its shape. Since the drive laser is polarized perpendicularly to the grating ridges, the electric field can easily pull out electrons which quickly heat the material in the ridges upon reentering. In the valley regions, the average electron energy quickly reaches several 10 to 100 keV due to the missing bulk plasma which would otherwise thermalize these electrons. Reaching the main pulse maximum, the grating starts to expand more rapidly in both target-normal and parallel direction with approximately $c_{\mathrm{s}} \approx 0.76\,\mathrm{nm\,fs^{-1}}$. When the laser a_0 is relativistic, the plasma compression caused by the $\vec{v} \times \vec{B}$-force causes an increase of the maximum electron density to about $600\,n_{\mathrm{c}}$ or $3/2\,n_{\mathrm{e}}^{\mathrm{total}}$ in the density profile along the target normal. As the grating quickly disintegrates over the next 100 to 200 fs, the formation of a secondary, regular electron density structure

can be observed. The ablated and expanding plasma of the original grating overlap in the valley regions which can be seen in the lower two panels of the lineouts in Fig. 4.43.

Over the next ∼100 fs, plasma jets, emanating from the corners of the original grating, create a transient grating whose ridge density profiles look very similar to the softened original ridges. The structure is longitudinally stable enough that also in the projection these plasma jets produce a significant density modulation. For a very brief time interval (∼ 10 fs), the grating ridges of the newly formed grating as well as the ablated original one appear to have the same height, while the edge shape is also very similar. This leads to an effect where the whole projected structure can be described with a grating pitch of $g/2$. This was an unexpected feature that occurred very stably over all the simulations that were performed. Later, the projected plasma jet density exceeds the magnitude of the remainder of the expanding ridges. After the first overlap and crossing of the jets they overlap with the jet from the ridge after the next and so forth, causing an oscillating density pattern. Once the rest of the original grating has ablated, the jets and secondary grating mostly dissipate over the course of the following half picosecond. The rest of the interaction is characterized by a plasma compression wave that has no grating features, anymore, and keeps ablating the target for the next several ∼ 10 ps.

Post-Experiment Simulations In contrast to the earlier simulations with the relativistic laser pulse, the pre-pulse of the drive laser at the final parameters did not cause significant target pre-expansion. Therefore, the initial grating sharpness was set to $\sigma = 5.4\,\mathrm{nm}$. That is the experimentally determined average value that was obtained in the two XFEL *pre-shots* that were performed before each *main shot* event. Pre-shots are the two lower-intensity XFEL-only shots that 1.) measure the cold target structure and 2.) ensure that the first shot did not degrade the target.

Due to the lower – now non-relativistic – intensity, the speed of target ionization and plasma compression due to the ponderomotive force were lower. This allowed for a reduction of the resolution to $\lambda_L/256$, thus reducing the calculation time for 1 ps to 27 h. The later simulations were complemented by low-resolution ($\lambda_L/32$) runs that each took about one and a half weeks for 600 000 steps. Irrespective of the possible algorithmic instabilities leading to inaccurate detailed dynamics, they allowed for an observation of the more general plasma evolution over 50 ps.

With a decrease of the original grating period g, the oscillations of the density modulation become faster which can be explained with the shorter distance until the plasma jets overlap. At the same time, the kinetic energy equilibration between ridge and valley regions also takes less time. When the longitudinal pre-expansion was varied in the earlier simulations, a larger plasma scale length σ_z effected a less pronounced transient grating structure later on (compare Fig. 4.45). The decrease in the laser intensity from a relativistic $a_0 = 3$ to the sub-relativistic $a_0 = 0.38$ did not result in a qualitative change of the observed plasma dynamics apart from the lack of initial plasma compression during the main pulse interaction phase. This compression would be invisible in the scattering experiment since it is uniform along the target surface. The smaller laser intensity did, however, increase the resiliency of the original grating structure. Transient grating structures appeared about 30 fs later than in the preparatory simulations.

Figure 4.46 shows fits of the analytic edge and grating models to the projected total electron density from the PICLS simulations. The dissertation of Melanie Rödel [287] extensively describes the fitting method and also the analysis of the experimental data. Up until the appearance of the plasma jets, the single grating model describes the projected density very well. During the early time of expansion, forward calculations produce an unambiguous signal. Once the expansion, i.e. the grating softness σ_x surpasses $b_0/8$, the edges of the original grat-

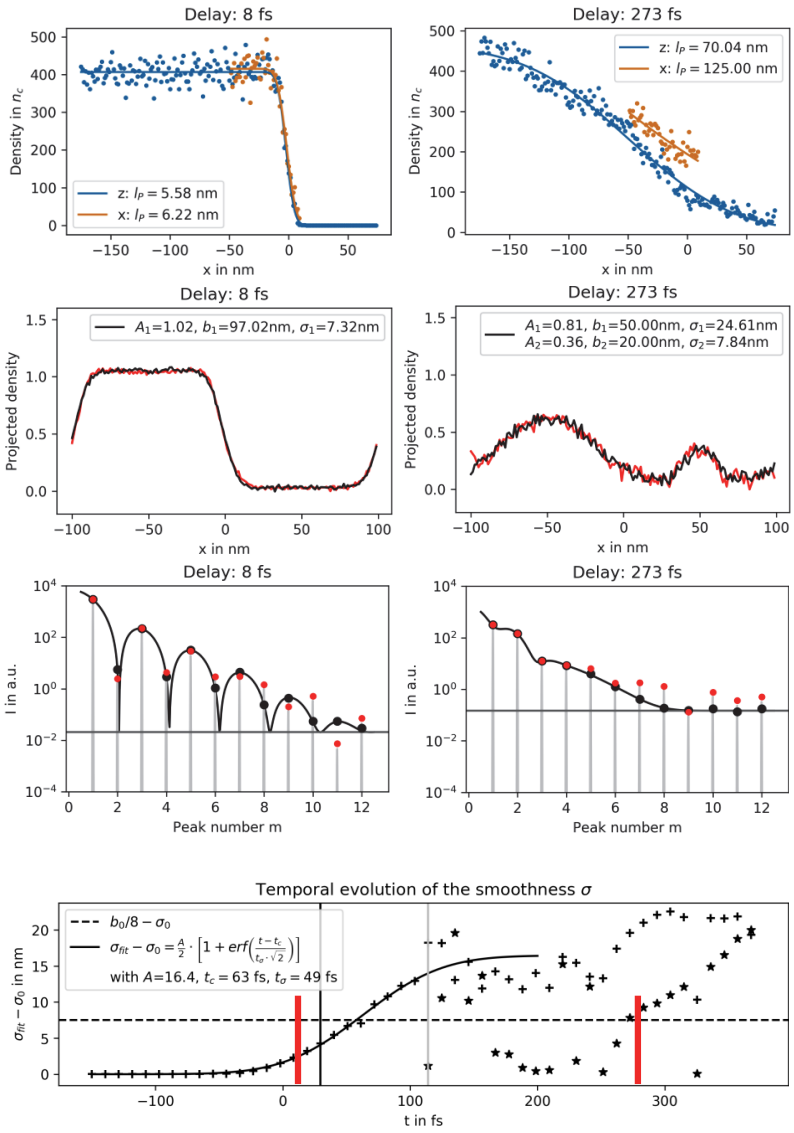

Figure 4.46.: *1st row:* Plasma scale profiles from lineouts of the PIC simulations and scale length (L_p) fits in z'– (blue) and x-direction (orange). *2nd row:* Error function fits (black) from the projected density (red) as defined in Eqs. 4.25 and 4.26. *3rd row:* Forward calculations of the analytical model using the error function fits (black), the projected density from PICLS as a scaled lineout (grey) and the integrated scattering peak signals from this lineout as a scatter plot (red). The horizontal line represents the Fourier transform of noise from the real space projection that was added to the analytical model. *bottom panel:* Temporal evolution of the error function fits (black line and plus markers). The red vertical bars mark the times shown in the three rows above. Once the edge softness reaches $\sigma = b_0/8$ (dashed black line) the edges merge and the plateau of the ridges vanishes. Together with the appearance of plasma jets at $t = 30$ fs (vertical black line) the single grating model becomes unreliable to still describe the projected grating. Then, a two-grating model is used. However, the number of visible scattering maxima becomes too small to identify a unique solution with the 6 parameters of the model. Analysis performed by M. Rödel [287], based on initial PICLS simulation campaign of the author.

ing merge. The flat ridge-top vanishes and now it becomes possible to describe the remainder of the original grating with a Dirac comb of Gaussians. The original model can be corrected for the now different height $\Delta_x \int n_e^{\text{total}} \, dz$ between valley-bottom and ridge-top but as soon as the transient grating from the plasma jets arrives at a significant magnitude, at $t = 30$ fs, the single-grating model is not reliable, anymore. Unfortunately, as the expansion continues, the forward-calculated scattering signal is strongly suppressed for higher q-values. This results in ambiguous predictions for the still visible intensity ranges due to the 6 parameters of the two-grating model. In other words, measurements and reconstructions of the expansion scale lengths do not yield a unique solution within the confines of the model, anymore.

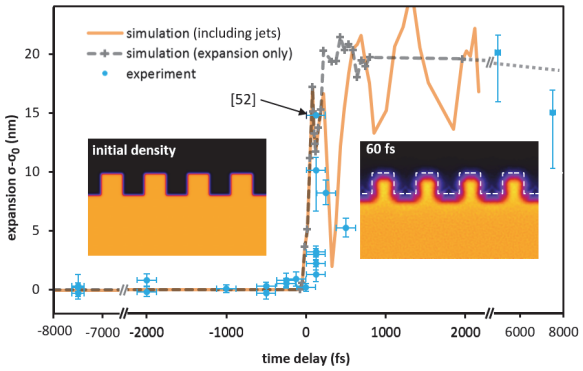

(a) Plasma expansion is shown as the decrease of the edge sharpness $\sigma - \sigma_0$, accounting for the initial ridge sharpness σ_0 versus the delay of the XFEL probe to the IR drive laser. **blue circles** Experimental data obtained from fits of the envelope through the peaks of the scattering maxima. **gray crosses** Data from PIC simulation, omitting influence of transient grating due to plasma jets. The gray dotted line extrapolates σ beyond the maximum simulated time of 2 ps. **orange solid lines** Simulation data including the signal of the plasma jets where it dominates the remaining original grating.

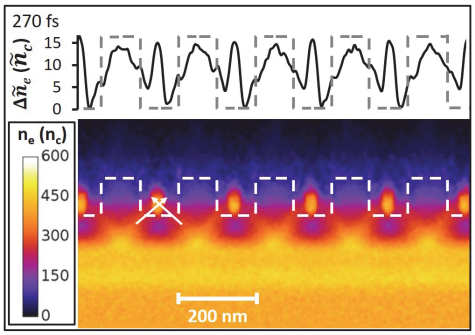

(b) Total electron density from PIC simulation in units of $n_c = 1.74 \cdot 10^{21}$ cm^{-3}. A transient grating formed by plasma jets emerging from the valley corners of the original grating (white dashed lines) as seen at $t_{\text{rel}} = 270$ fs.

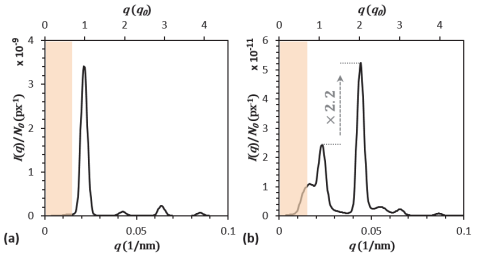

(c) Lineouts through experimental scattering pattern for **(a)** a XFEL-only preshot and **(b)** a main shot 6 ps after the 400 mJ HI laser pulse hit the target. The second peak in the main shot shows a much larger signal than the first peak which is forbidden given there is only a single grating with the scattering vector q_0.

Figure 4.47.: Temporal evolution of plasma expansion extracted from PIC simulations and compared to experimental X-ray scattering results. Figures originally in Kluge et al.[G8].

Comparison to scattering experiment

In the 2016 beamtime at LCLS, silicon grating targets that were specifically fabricated for the purpose of diagnosing the plasma expansion were shot with an XFEL pulse of $\tau_{\text{XFEL}} = 40$ fs and a spot size w_{XFEL} of 5 to 10 µm. Since this is much smaller than the drive laser spot w_{IR} of 16 to 30 µm, it is valid to assume a homogeneous expansion of the grating over the probed area and only consider the XFEL illumination function in the model of the scattering signal. Figure 4.47 shows the major results of the experimental campaign complemented by the numerical results the author obtained in the PICLS simulation study. In order to compare both studies, the scattering peaks obtained in the experiment were first located semi-automatically by an

algorithm (see [G9, 287] for details) and the random underground was removed. The peaks were then fitted to an analytical model (Eq. 4.19) containing a regular grating with softened ridges. The analysis of the edge sharpness parameter σ was found to be very robust with respect to the exact shape of the density gradient (see [334] for further reading). From the simulations, σ could be directly measured after the total electron density was integrated along the probing direction and the plasma expansion was defined as $\sigma - \sigma_0$, where σ_0 was the averaged value of the initial grating sharpness.

Figure 4.47a shows very good agreement of the grating expansion versus the delay between pump and probe laser pulse. The simulation shows that it takes only 30 fs after the main pulse for the ridges to soften enough to intersect. After another 100 fs, the original ridges are mostly ablated and just remain as density modulations inscribed on the plasma layer behind which is then persistent for at least a picosecond, leaving σ virtually unchanged from during this time. These observations from the simulation fit very well to the time scales of expansion seen in the experiment. A major result is the experimental evidence found for the existence of the very transient secondary grating that is formed by the plasma jets mentioned before. With the timing jitter between XFEL and HI laser being $\Delta t = 119$ fs, and the oscillations of the transient grating on the same order but confined to mostly happening within a picosecond, the phenomenon was less likely to be captured in one of the shots. After a careful search of the data, a scattering image could be identified where the second scattering peak signal was significantly larger than the first peak. As such a pattern is forbidden in the model of a single grating and with the stable prediction of the feature throughout all the simulations, there is a strong indication that the projected transient grating created via the expanding plasma jets had the same ridge height and shape as the disintegrating original one which would have caused the observed double-grating-frequency signal. Owing again to the timing jitter, only a lower limit for the plasma expansion speed of 0.04 nm fs^{-1} could be identified. The simulations yielded an estimate of 0.19 nm fs^{-1}.

Conclusions The combined simulation- and experimental study demonstrated clearly that probing the fs-nm scales of solid density plasmas is possible with Small-Angle X-ray Scattering at the new and powerful XFEL machines. For the first time, numerical simulations could be validated directly by the experiment at simultaneously high temporal and spatial resolution. Additionally, a transient density modulation spawning from non-equilibrium plasma dynamics and the specific target geometry was predicted and convincing evidence for its existence was then found in the experimental scattering pattern that cannot be explained in the more simple model of a single grating. The transient structures and ongoing plasma expansion also caused the failing applicability of the simple models for larger interaction times.

While these results are promising, they show that even with *simple* models and specifically designed targets the expectation of clear signals might still not be met. The geometric origin of the plasma jets and their transient nature make the grating targets a very suitable candidate for diagnosing the laser intensities and interaction delays if the experiments are repeated with better time resolution. Recent developments[G13] in the analysis of the connected experiment and its successor in 2017 reveal that there were scattering images in which the signal intensities mirrored at the $q = 0$ direction were not always symmetric. This could only be caused by a non-zero imaginary part of the scattering form factor $\mathcal{F}_{\text{ion}}$ which is an indication for absorption processes in the plasma. As such, more complex processes can make the interpretation of these novel pump-probe experiments much more challenging. The PICLS study has also shown that many conditions need to be improved that currently limit the number, the scale and the reproducibility of supporting simulations to achieve predictive capabilities within a reasonable time prior to an experiment.

Previous other studies[155] have shown that quantitative predictions for e.g. acceleration time scales, field strengths and maximum particle energies require a full-scale modeling in 3D. The study of this section was performed in 2D, though, since this version of the code contained advanced solid density plasma physics. Furthermore, the resolution necessary to resolve the plasma oscillations would result in an increase in computational effort by easily three orders of magnitude were the simulations performed in 3D. That would quickly lead beyond the scope of a common university-scale computing cluster and fit only to the current Top500[203] petascale HPC systems. Since the code also was only parallelized in laser transverse direction, the simulation of a more extended target in propagation direction would have increased the time to solution linearly. With calculation times of days to weeks, the possible reaction speed to new measurements of experimental parameters and the whole extent of the parameter scan is severely limited. At HZDR, a very modern, scalable and fast particle-in-cell code PIConGPU[52][G7] is constantly being developed and improved. It has been proven to scale very well[155] on the largest supercomputers to date that lead the Top500 lists of the past 5 years. It is also developed in an open-source, open-science, collaborative fashion which makes it less prone to errors not being traceable and therefore especially promotes reproducibility of simulations using the versioned releases of the code. Following this initial study, it is therefore very desirable that advanced methods for treating solid density plasmas and extracting the relevant physics also be implemented in this code. Lastly, to complement these developments, new and scalable methods for predicting the XFEL probe signal beyond Fast Fourier Transforms for novel SAXS experiments are needed as a next step that goes beyond simple analytic models and allows for explorative start-to-end simulations.

Further Discussion of Plasma Jets While the occurrence of previously unexpected plasma jets caused the breakdown of the simple description of the grating projection, the scattering signal from the transient grating provides the unique opportunity of an extraordinarily clear signal that could lead to increased precision in the expansion measurement. Normally, the even scattering peaks at $n \cdot 2\,q_0$ are strongly suppressed. They only begin to appear when the expansion causes the exact relation of $b = g/2$ to fail. However, the transient grating is an oscillating structure that only for a very brief moment in time (< 10 fs) exhibits the same magnitude as the ablating original grating. During that instance, the $2\,q_0$ signal is much stronger than the peak at $1\,q_0$. The appearance of the jets depends strongly on the grating geometry and feature sizes as well as on the laser intensity and plasma parameters. With more in-depth study of this phenomenon it is likely possible that simulation and experiments can be compared even more closely, yielding a spatio-temporal benchmark of previously unknown quality.

On a more speculative note, it could even be possible to use this phenomenon as an ultra-fast switch for the XFEL pulse. Since the divergence of the scattered beam is still low and the signal from a line grating structure are well-defined scattering peaks, the transient grating could potentially be used as an optical element that shifts the direction of the XFEL pulse for a duration on the order of 10 fs or below. The oscillation frequency of the grating is faster, the smaller the grating features are. As such, a new class of ultra-fast experiments could be opened up if this phenomenon could be made use of.

Forward Calculations for X-ray Scattering off Grating Targets with X-ray Free Electron Lasers

One of the most significant challenges of X-ray probing is the loss of depth information from the target due to the well-known *phase problem*. The previous section already stressed the importance of hand-in-hand development of the methods for predicting and analyzing both

the highly transient plasma dynamics as well as the scattering signals.

A relativstic ($a_0 \geq 1$) laser pulse creates bunches of *prompt electrons* at the surface of the target that are highly directional. The laser pressure and counter-flow of "hot" and "cold" electrons is known to be the seed for plasma instabilities that rapidly evolve on the same time- and length scales intended to study with SAXS experiments. The increasingly complex density structure is not only rapidly changing while it is being probed, X-ray photons are also more likely to be scattered multiple times or absorbed in the highly overdense plasma when its extent is comparable or larger than the photon mean free path (typically micron-scale). Traditional forward calculations of the probing signal like Fourier transforms cannot account for these effects. Hence, proof-of-concept simulations with the PIC-like photon tracing code ParaTAXIS[56, G17] are presented. With this code, photons move as particles through a simulation volume, offering the ability to take these additional physics into account. As a first step prior to a fully integrated pump-probe simulation, time-resolved density output from the particle-in-cell simulations mentioned before serve as an input to the photon scattering code. These simulations are embedded in the larger framework of complete start-to-end simulations for future experiments at European Advanced Laser Light Sources (*simex_platform*[G18] of the SIMEX[G19, G20] workpackage of the EUCALL project). In this context, the work presented here is also a preparation for upcoming plasma pump-probe experiments of the Helmholtz International Beamline for Extreme Fields (HIBEF) at the recently commissioned HED end-station of the European XFEL[335].

Predicting Scattering Signals Beyond Employing Fourier Transforms Predicting the signal of an X-ray scattering experiment is often done via Fourier transform of the projected scatterer density, as was mentioned before in Sec. 4.3.2. . However, the targets that are most often used for laser-ion acceleration are micron-sized metal or plastic foils. The Fourier transform approach for predicting the signal of such a process, if it were probed, has several shortcomings. The first and most striking issue is the loss of depth information about the target. This stems from *phase problem* that is well-known from 2D image reconstruction (see Eq. 4.16). Since the signal on the detector panel is an intensity measurement and $I \propto |\exp(i\varphi)|^2$, the depth information, that is also encoded in the phase φ of the scattered wave, is lost. Second, the X-ray pulses that modern XFEL machines can produce are ultrashort but with $\sim 10\,\mathrm{fs}$ duration they are on the same order of magnitude as the transient processes in the plasmas that are driven by equally short laser pulses. Furthermore, for an X-ray pulse to traverse a micron-sized target, a duration on the order of 1 to 10 femtoseconds is necessary, as well. This has two implications: the target changes while it is being probed and different parts of the probe pulse can see different states of the target. Additionally to these complications, the realistic X-ray pulse is likely to not have a top-hat spatial and temporal profile which results in non-equally weighted scattering signal from the illuminated area and temporal plasma dynamics. In plasmas of high density and energy density it becomes increasingly likely that an X-ray photon is scattered multiple times. A plasma with this property is called *optically thick*. Multiple-scattering, however, would violate the initial assumption of far-field propagation that allows for the forward calculation of the scattering signal by Fourier transform of the assumed scatterer density distribution. A different method of forward calculation is necessary to adequately support novel experiments designed to probe the transient fs-nm dynamics of HED plasmas. A promising new approach is the PARAllel Tracer for Arbitrary X-ray Interaction and Scattering-ParaTAXIS[56, G17]. It discretizes the wavefronts of the probing X-ray pulse to PIC photons, effectively simulating the scattering process on a grid where these photons interact with the scatterer densities on a numerical cell level. Just like in a particle-in-cell code, this approach avoids the unfeasibly massive computational effort of direct particle-particle interaction while

allowing to address the challenges mentioned above.

ParaTAXIS ParaTAXIS is an open-source code that allows for multiple scattering of X-ray photons on arbitrarily complex 3D density distributions from PIC simulations. It is based on the same particle-mesh methods as PIConGPU and directly enables reading density data of the latter into ParaTAXIS via the `openPMD` standard[255]. The reason for why such an approach is viable can be seen from the radiation transport equation which describes the flow of radiation through a medium[336].

$$\frac{1}{c}\frac{\partial I_\nu}{\partial t} + \vec{\Omega} \cdot \nabla I_\nu = -\alpha_\nu I_\nu + \sqrt{\eta_\nu} \tag{4.27}$$

It can also be interpreted as a Boltzmann equation for photons of frequency ν with α_ν the absorption coefficient and η_ν the emissivity.

$$\frac{1}{c}\frac{\partial f_\gamma}{\partial t} + \vec{\Omega} \cdot \vec{\nabla} f_\gamma = \left(\frac{\partial f_\gamma}{\partial t}\right)_{\text{coll}} \quad , I_\nu = \frac{2h\nu^3}{c^2} f_\gamma(\vec{r}, \nu, \vec{\Omega}, t) \tag{4.28}$$

Equation 4.28 is a direct analogy to the governing equations for describing charged particles in a plasma through PIC. But here, the left side describes the flow of radiation in phase space and the right hand side describes absorption and emission which depend on atomic physics. The macroscopic specific intensity I_ν is here represented via the microscopic photon phase space distribution function f_γ. In this framework, Thomson scattering is modeled self-consistently through the local interaction cross-section. Elastic scattering is the simplest process for it conserves the number of photons. There are, however, other caveats that such a Monte-Carlo sampling technique introduces. As in PIC, the ParaTAXIS photons are not physical particles but rather samples of the distribution that carry photon properties like energy $\mathcal{E}$, direction $\vec{k}$, phase φ and weighting. This implies that they take only one *classical path* to the detector. As a *path* classifies also: being absorbed or scattering off a charged particle. While real photons can take many paths at the same time, this sampling technique actually requires *more photons than in reality* to satisfy the same statistics. Macro-photons in the simulation are propagated towards a detector once they leave the simulation volume. They then hit a distinct virtual pixel on the detector plane. Only photons that hit the same pixel can interfere, adding their contribution to the signal in that pixel and accounting for their phase.

In order to arrive at a physical scattering pattern, coherent interference of each of the photons is necessary. That requires a massive summation of photon contributions while correctly considering each of their phases. As such, phase summation errors could be a show-stopper for this technique. Therefore, in ParaTAXIS, once the photons leave the simulation volume around the direct vicinity of the target, a ray-tracing approach is employed to overcome this problem[56]. This approach avoids that after photons traverse the highly-resolved solid-density plasma region (physically, about 1 to 10 µm) in around 10^4 simulation steps, the same cell-wise propagation has to be computed for the whole, substantially larger, vacuum distance to the detector. With the latter usually being centimeters to meters away in an experiment, the number of necessary steps would exceed 10^7 to 10^9. The numerical precision of the phase and the summation errors would be significantly changing the result, rendering it useless for interpretation. To overcome the described challenges and enable a kinetic scattering treatment, GPUs turned out to be a very suitable hardware for ParaTAXIS. Their large memory bandwidth facilitates the massively parallel summation of phases, the simultaneously high floating point performance keeps cross section calculations fast and a high numerical precision suppresses numerical phase errors.

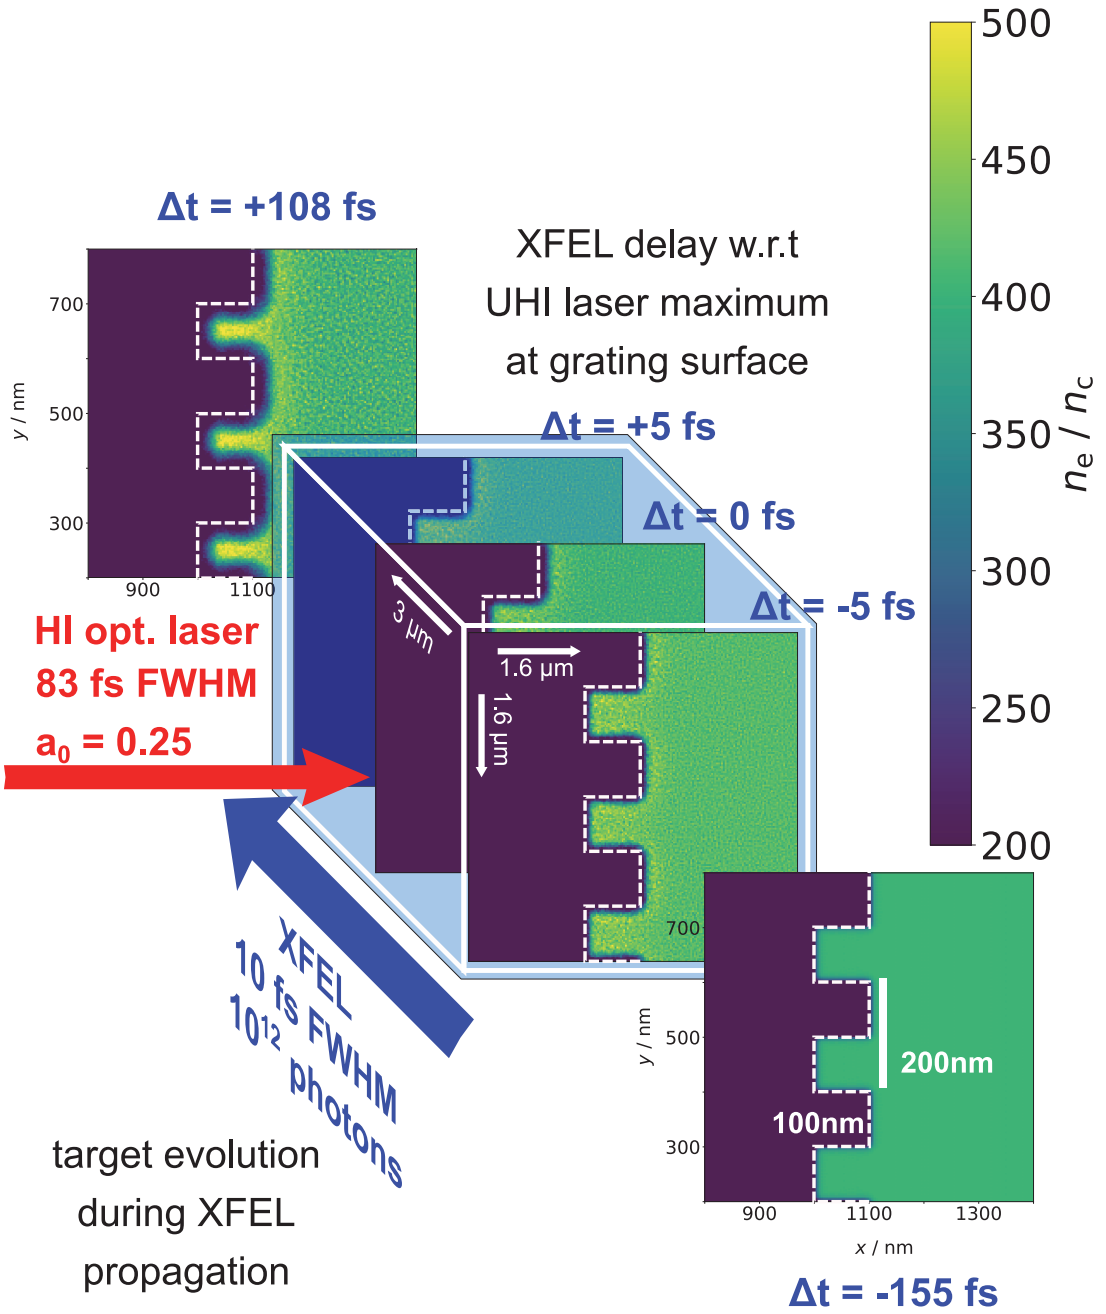

Figure 4.48.: ParaTAXIS scattering geometry for a silicon plasma grating as described in 4.3.2 that is driven by a HI optical laser and probed with an XFEL pulse in perpendicular direction. As the XFEL pulse traverses the target, the total electron density n_e changes due to the HI laser interacting with the grating which degrades over time. Delay times Δt are given between HI pump and XFEL probe. The area illuminated by the XFEL pulse is $2\lambda_{HI} \times 2\lambda_{HI}$, with the corresponding SAXS image, assuming 3 µm target depth, seen in Fig. 4.50. Adapted from image published in [G10].

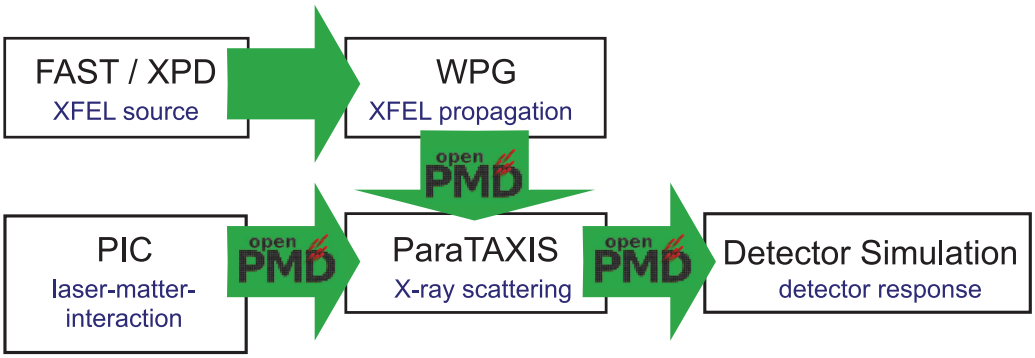

Figure 4.49.: Workflow demonstrating the very first successful connection of simulation codes for X-ray scattering from laser-driven plasmas within `simex_platform`. The simulation data was flavored in the openPMD standard[255], thus greatly facilitating the exchange between different codes. Figure published in [G10].

Loosely Coupled Start-to-End Simulation of SAXS on Silicon Gratings

In first proof-of-concept simulations, ParaTAXIS was employed during this thesis work to demonstrate the viability of its Monte-Carlo scattering treatment to model a realistic SAXS experiment performed at LCLS (see 4.3.2). The simulations were done in the greater scope of `simex_platform`[G18, G20], the modular software environment for start-to-end simulations of EUCALL. The international EUCALL[258] collaboration[13] brought together the major photon science laboratories in Europe to develop common strategies for experiments and simulations at existing and novel light-source facilities. The code project `simex_platform`[14] is a collection of photon science codes from various optical laser, free-electron laser or synchrotron light sources. It covers beam creation and propagation, light-matter interaction simulations, transport of the interaction products (e.g. charged particles or scattered radiation) to the detectors, and ultimately detector simulations. The main incentive is the complexity of real experiments and the need to mirror this complexity in virtual experiments to be able to account for all the different influences that can affect the final scattering images. These virtual experiments are complemented by a python-based framework with common interfaces that connect these codes, aiming at a unification of input and output data formats. This unification is vital for the interoperability of the simulation codes and to ensure that parallel development of the codes does not break the whole chain.

The workflow employed for the X-ray scattering from hot, dense plasmas is depicted in figure 4.49. In particular, the open particle-mesh data standard `openPMD`[255] was adopted by many of the participating codes, also PIConGPU and ParaTAXIS.

ParaTAXIS was used by the author to model a SAXS experiment with silicon gratings aimed at measuring plasma expansion with simultaneous femtosecond-nanometer resolution. At first, the Thomson cross section was implemented into the code so that realistic electron density input could be used. The input data was obtained from the PIC simulations described in 4.3.2. To account for the temporal evolution of the target during the probing process, 959 subsequent PIC outputs ($\Delta t_{\mathrm{PIC}} = 1.042 \cdot 10^{-17}$ s) were stacked in propagation direction of the XFEL pulse. The density slices were combined into a 3D simulation volume that measured 512×512 cubic cells ($\Delta x_{\mathrm{PT}} = \Delta x_{\mathrm{PIC}} = 3.125$ nm) in transverse direction. Due to the original PIC simulations being 2D, and ParaTAXIS so far being able to probe only directly along one axis,

[13] European Cluster of Advanced Laser Light Sources – was active from 2015 to 2018 and was continued in 2020 as PaNOSC[259], the Photon and Neutron Open Science Cloud

[14] a result of EUCALL WP4 SIMEX – Simulation of Experiments at Advanced Light Sources, now adopted into PaNOSC `https://panosc-vinyl.github.io/SimEx/`

the XFEL propagation direction is exactly aligned with the grating ridge-lines and not tilted by 45° as in the experiment (compare Fig.4.42). The scattering pattern was, therefore, expected to differ and show structure in the direction parallel to the target normal while the original only showed a pattern along one direction. As probing time the interval from –5 fs to 5 fs around the arrival of the drive laser main pulse maximum on target was chosen, which corresponds to $\Delta t_{\mathrm{delay}} = 0$. The thickness of the target in probing direction was 3 µm owing to the cubic cells and the number of density slices such that they would account for the spatial length of the 10 fs probe pulse. The virtual detector was put at $d_{\mathrm{det}} = 1.4$ m distance with 2048 × 2048 pixels of size $a_{\mathrm{D}} = 13.5$ µm. The scattering angle within ParaTAXIS is randomly determined, but follows a uniform distribution over the whole solid angle that is user-limited by θ_{max}. The maximum scattering angle for these simulations was chosen to be $\theta_{\mathrm{max}} = 0.01$ rad. Such a limit is crucial for the *time to convergence* of the code (see also Ch. 4 of [56]) as it increases sampling of the relevant detector area. All input and output data are openPMD-compliant and were published at [G21].

Two conceptually different cases were simulated, i.e. the *optically thin* and the *optically thick* case. A plasma is optically thin (thick) when its dimensions are smaller (larger) than the radiation mean free path $\lambda_{\mathrm{mfp}} = 1/n_x\sigma_x(\nu)$, see also [183]. Here, n_x is the density of scatterers and σ_x is the interaction cross section that depends on the frequency ν of the radiation. To model the optically thin case and actually suppress multiple scattering while also drastically decreasing computation time, the density slices were all integrated into a single slice, reducing the simulation box to 12 cells in probing direction with only one non-zero cell. For the optically thick case, the simulation box measured 1024 cells in probing direction and contained all the 959 density layers surrounded by vacuum. Furthermore, the scattering cross-section was increased by a factor of 10^3, as would be reasonable in a resonant scattering (RCXD) scenario where a bound-bound or bound-free scenario is targeted by the specifically chosen X-ray wavelength. While these two scenarios were still relatively idealized, they serve well to show the two extremes in first proof-of-principle simulations. The interoperability of workflows of the loosely coupled simulation codes was published in [G10–G12, 288].

Comparison of the Scattering Signals from ParaTAXIS to Experimental Data

Figure 4.50 compares the virtual detector signals of two simulations, one in the optically thin and one in the optically thick setup. The grating target in real space is structured in y-direction while the target normal is aligned with the x-direction. According to the expectation, both the optically thick and thin cases show line-shaped scattering maxima at $q_y = q_0 = 2\pi/g$, with g the grating period of 200 nm. In the experiment from [G8], the probing direction was under 45° with respect to the target plane and the grating was laser-heated throughout the whole field-of-view of the X-ray pulse, thus integrating the signal along the grating ridge direction, resulting in a line of bright scattering maxima. In the simulation, due to the different projection direction, only the central x-region of the actual grating is similar to the experiment, while the y-direction can be described by a step function. The Fourier transform of a step is a *Sinc-function* which explains the in q_x-direction elongated scattering peaks on the $q_x = 0$ position all over q_y seen in Fig. 4.50. At the stripe that marks the first scattering maximum of the grating ($q_0 = 2\pi/g$), a minimum in q_x-direction appears at $q_x = 2\pi/h$ which corresponds to the grating height of $h = 100$ nm. To be able to compare the simulation results to the experimental result and the analytical solution in Eq. 4.19, the 2D scattering image was integrated over q_x. The comparison will mostly be qualitative since the absolute peak heights and the signal-to-noise ratio depend also on the thickness of the target and its density perturbations in the direction of the probe beam. Since the projections are different, the scattering signal lineouts need to be scaled to

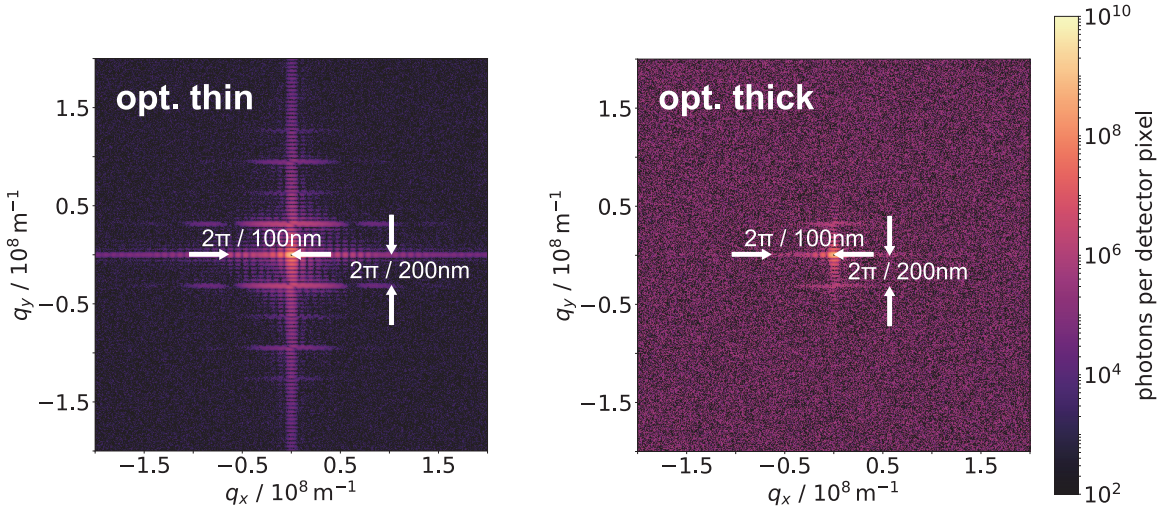

Figure 4.50.: **Left (right)** Simulated SAXS image for the optically thin (thick) target at 1.4 m detection distance using ParaTAXIS. Detector pixel size $a_D = 13.5\,\mu m$, X-ray wavelength $\lambda_{XFEL} = 1.47\,Å$ and 10^{12} photons in the illuminated area. The vertical separation of scattering lines corresponds to the grating period g of 200 nm, the horizontal to the grating depth h of 100 nm. For the optically thick target, the scattering cross section σ_{Th} was increased by a factor of 1000 to account for resonant scattering at the ion density. All other parameters remain the same. Published in [G10].

allow for a comparison.

Figure 4.51 compares the scattering pattern in the experimental shot (violet) with the analytic solution (dark blue), the q_x-integrated Fourier transform of the projected density in the ParaTAXIS simulation (orange) and the simulation results in the optically thin (annotated; dark green) and two simulations of the optically thick case (annotated; turquoise and light green). All the 1-dimensional data is normalized to the height of the first scattering maximum of the experimental data which is given in units of 10^{-10} scattered photons per incoming photon and pixel. The background-corrected example data[15] chosen from the experiment was measured at a pump-probe delay of $\Delta t = 0$ fs and from the fit, the expansion parameter was determined to be $\sigma = 7.64$ nm which served as input for the analytical single-grating model. The peaks that might seam like maxima in the experimental data that occur at smaller q than the first major scattering maxima are results of the light scattering around the beam block that is usually inserted to avoid the central maximum from mostly unscattered X-ray laser light to destroy the detector or over-saturate the image.

The scattering peak locations of the scaled data from the simulations (annotated) are in very good agreement with the analytical solution (dark blue) at nq_0. The heights of the even scattering peaks of the density Fourier transform (orange) also agree very well with the experimental results (violet) and the analytical solution up until the third scattering order. It can be seen, however, that the scattering peaks at larger absolute q-values have not developed fully in the simulation and the level of noise between the visible scattering peaks is almost as large as the peaks. Compared to the analytical solution where the even numbered peaks are completely suppressed, the simulations and the experimental data show peaks there, too. In the simulation, the grating is already slightly expanded and therefore the initially perfect ratio of $b/g = 0.5$ is now smaller which causes the even-numbered peaks to appear. In the experiment, this cannot be determined with certainty, but it is very likely that the reason for even-numbered peaks to appear is also the b/g ratio.

A most striking difference between the optically thick and optically thin simulations is that

[15] Shot 150:3 of LN04

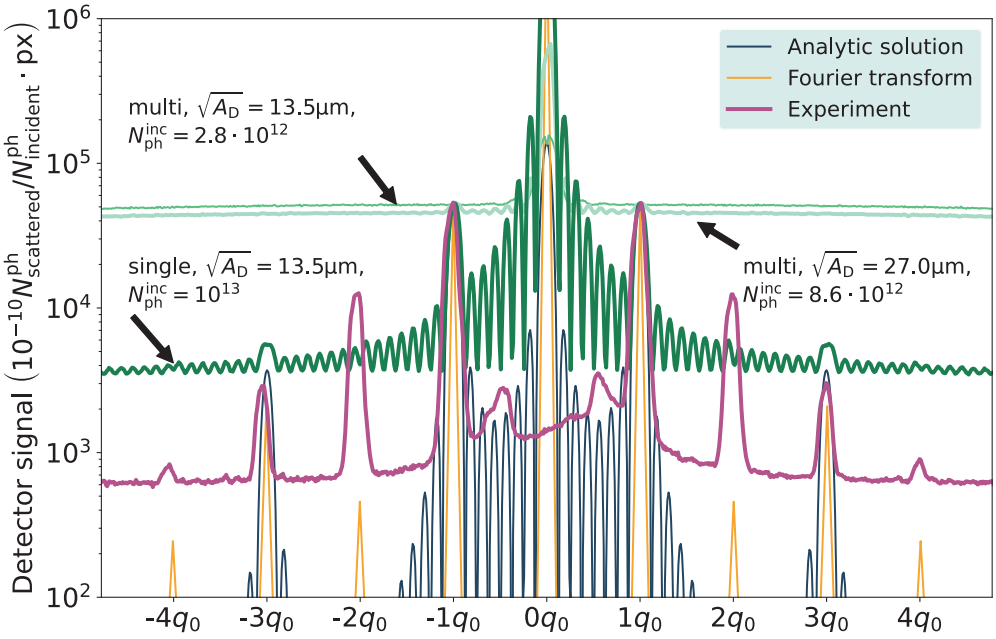

Figure 4.51.: Comparison of ParaTAXIS-simulated scattering patterns (annotated with arrows) against the Fourier transform (*orange*) of the projected electron density in these simulations. One simulation enforced single scattering from a quasi-2D density while the other two featured fully 3D multiple scattering. In *dark blue* is shown the analytic solution to the finite grating (see Eq. 4.19) and an experimental shot (*violet*), both reported in [G8]. All signals are taken at $\Delta t_{delay} = 0$ and have been normalized to the height of the first scattering peak from the experiment $I(q_0)^{exp}$. The signal in between for the ParaTAXIS simulations is dominated by the noise owing to the finite number of quasi-photons in the simulation. For the experimental example shot at, an imperfect peak-to-valley aspect ratio $b/g \neq 0.5$ of the front side grating may be responsible for the emergence of even-numbered scattering peaks.

the signal-to-noise ratio is much lower in the former than in the latter. This can be explained by two reasons: the overall scattering cross-section was increased by a factor of 10^3 to simulate a resonant scattering case. Furthermore, the optically thin case enforces single-scattering by just one non-vacuum layer of integrated, projected density whereas multiple scattering is possible in the optically thick case with the 3D electron density. The latter better represents the experiment where the projection of the target is really a density grating, offset by a constant density, creating additional incoherent background. However, since in reality the scattering pattern is created by the non-local collapse of the photon wave function on the detector, the scattering pattern is created much more quickly than in the simulation where a signal is only counted coherently if two macro-photons reach the same detector cell. A faster emergence of the scattering signal was found in a simulation where the detector pixel size was four times as large as in the other optically thick case ($\sqrt{A_D} = 27\,\mu m$ instead of $13.5\,\mu m$) which comes at a trade-off in q-space resolution.

This confirms the prediction that actually many more simulated photons are necessary than the real number of photons in an X-ray pulse to approximate the scattering pattern to a level of signal-to-noise ratio that is acceptable. This means that a predictive simulation would need to run longer if the structure size or correlation length of interest is small and the expected signal would appear at larger q values.

Conclusion The results shown above constitute to the first-ever proof-of-concept loosely coupled start-to-end simulations that model the XFEL probing of a laser-driven solid-density plasma. Small-angle X-ray scattering patterns from kinetic simulations with ParaTAXIS prove to resolve nanometer-scale density modulations well and were shown to agree qualitatively to a Fourier transform of the 3D electron density projection as well as the analytical scattering pattern solution built with experimental shot data. However, a quantitative agreement was not achieved, yet, since the proof-of-concept simulations varied in the scattering geometry. At the European XFEL, generated pulses contain of about 10^{12} photons[49]. In the simulations shown here, roughly 10^{13} macro-photons were used and it can be seen that a much larger number is still necessary to increase the signal-to-noise ratio to see the scattering features, depending on the location in q-space. As such it can be foreseen that a 3D pump-probe simulation of a laboratory laser-solid density plasma experiment could easily be an exascale effort, exceeding the computing capabilities available today.

Outlook While the target was assumed to be invariant over time here, the method itself promises to be capable of covering the temporal evolution of it as well as the probe pulse as it traverses the plasma. For better feasibility, the convergence of the scattering signal needs to be improved, possibly with a treatment of the non-local collapse of the photon wave function. Integrating ParaTAXIS into PIConGPU is desirable for a truly "online" pump-probe simulation. While ParaTAXIS is already based on `libpmacc`, the particle-mesh library that PIConGPU uses, the densities are still static and *offline* from the plasma dynamics simulation. For a full modeling of an X-ray pump-probe experiment, including the reaction of the plasma to the probe pulse, the scattering framework needs to be incorporated into PIConGPU. However, to achieve this goal, additional challenges have to be overcome. One of these is to propagate a massive number of photons through the particle-in-cell volume which would very likely require sub-stepping of the PIC-cycle as both the plasma particles as well as photons are held in the same memory. Currently, the radiation transport implementation into PIConGPU is an ongoing effort within the Master's thesis of Ordyna [337] and beyond. Meanwhile, two projects have already emerged to calculate X-ray scattering alongside a running PIConGPU simulation. The `xrayScattering` plugin calculates the complex scattering amplitude of available particle species from their respective density field data which can post-run be summed up coherently at a virtual detector[338] to produce a scattering signal. This approach supports also the implementation of absorption processes. On the other hand, PIConGPU streaming output was tightly coupled to the GPU-accelerated code `GAPD`[339] that calculates an X-ray scattering signal via an Ewald-sphere approach from each particle and its position.

Both of these projects are immensely important steps toward a full pump-probe start-to-end simulation of a laser-driven ion acceleration experiment and will be able to cover many scenarios while offering significant performance and reasonable accuracy. ParaTAXIS and the ensuing implementation in PIConGPU, however, still remain the first true prospective solutions to include multiple-scattering. The detailed example for a typical PIC-setup given in section 3.2.1 demonstrates how easily today's (petascale) top 10 supercomputers[203] can still be brought to their limits with laser-solid simulations. Adding an X-ray probing module with $\geq 10^{12}$ photons to such simulations would easily make them an exascale[16] effort. On the other hand, local particle splitting and merging[17] are promising techniques that could be used to mitigate the computational demand somewhat by putting emphasis on interactions and phase space regions that are otherwise not sampled well.

[16] 1 exaFLOPS = 10^{18} *floating point operations per second*, (peta– = 10^{15})

[17] These methods are also developed in and for use with PIConGPU: particle splitting [337], probabilistic particle merging[340] and particle reduction library[341].

5. Summary

This thesis investigated the role of realistic intensity ramps for the interaction between ultra-short, ultra-high intensity laser pulses and ultra-thin foil targets with respect to their effect on the laser-ion acceleration performance in the (extended) TNSA-regime. To achieve this, 3D3V particle-in-cell simulations of massive scale and very high resolution were performed on the Swiss supercomputer Piz Daint and in the wake of the project, around 4 Petabytes of data were produced and analyzed. It was found that a realistic laser contrast during the last few hundred femtoseconds alters the intra-pulse dynamics of the ion acceleration process significantly when compared to a perfectly Gaussian laser pulse. By including the intensity ramps, maximum proton energies could be optimized to the highest overall values in excess of 160 MeV, whereas an ideally shaped Gaussian pulse resulted in generally lower maximum energies (110 MeV at the optimum thickness). However, the spectral shape in the realistic laser contrast cases exhibited a TNSA-cutoff at reduced energy, typically followed by a higher energy tail. Premature expansion of the rear-side ions leads to this reduced TNSA cutoff-energy in both the Gaussian and the ramp cases. But the latter also create an ion population that is placed at the origin of the sheath when it is strongest. There, the $2\omega_L$–electron bunches accelerated from the target front side are still very dense and have the largest effect on the accelerating field the ions experience. This effect is not covered by any of the common analytical or empirical models that are used to predict ion energies. For pre-expanded protons further away from the sheath origin, electron bunches have already dispersed more and analytical models are able to reproduce the remaining ion acceleration well if the pre-expansion is accounted for, as the adapted model in this thesis shows. Whereas the behavior described above is clearly visible in copper targets, the simulated plastic targets did not have the same observed high-energy tail. While a phase space analysis of the plastic targets shows that the highest energies also occur in protons briefly accelerated in the front and then injected into the sheath, there is no clear spectral separation since originally protons occur everywhere throughout the target. Only with the initial spatial separation with metal targets having their proton source layer in their organic contaminants, this becomes visible.

Within this thesis project, the highly-parallel ionization methods in the open-source particle-in-cell code PIConGPU were improved and extended to statistically model collisional ionization, thus enabling the aforementioned simulation campaign. In support of an optical-pump– X-ray-probe experiment that studied the melting process of colloidal crystals at LCLS, the Thomas-Fermi collisional ionization implementation was directly used by another simulation campaign, that the author contributed to. From these and accompanying MHD simulations, a three-stage model was built and plasma expansion speeds could be obtained that matched the experimen-

tal observations well, thus also benchmarking the PIConGPU code.

The enhanced field ionization framework, on the other hand, served a simulation campaign that provided detailed insight into the injection process of a laser wakefield accelerator operating in the self-truncated ionization injection regime, which had previously been demonstrated in experiments at HZDR.

Furthermore, the large-scale 3D PIC simulations on Piz Daint featured self-generated bremsstrahlung which was recorded by synthetic radiation diagnostics that produced a high-quality dataset that is the first of its kind with the inclusion of realistic pulse contrasts at simultaneously high simulation resolution. In particular, it was found that the high-energy signatures ($\mathcal{E}_{ph} > 20\,\mathrm{MeV}$) encode the target state and plasma conditions during the brief period of maximum laser intensity on target. Whereas electrons are deflected by the sheath fields that are accelerating ions, bremsstrahlung photons sample the plasma and then leave the interaction region virtually unperturbed. Yet, bremsstrahlung at such high energies is currently still very challenging to measure and the technical development of on-shot detectors is ongoing. Such passive methods of probing laser-driven solid-density plasmas on extreme scales are gradually being complemented by active probing using large-scale X-ray Free Electron Lasers that push the resolution to the femtosecond-nanometer level.

However, experiments are still rare and extremely elaborate while the results are obtained in the form of scattering images that encrypt spatio-temporal information of highly transient processes in Fourier space. As such they are hard to interpret and often allow for ambiguous real-space representation. Therefore, a separate a-priori PIC simulation campaign performed during this thesis enabled and guided the first experimental observation of an expanding, optical laser-driven, solid-density plasma distribution at simultaneous femtosecond and nanometer resolution. The accompanying experiment with structured grating targets was performed at LCLS. A previously unexpected electron jet signature was first observed in the simulations which, in XFEL projection, leads to the formation of a transient, staggered grating structure. In the experimental data, convincing evidence was found that such a structure and thereby possibly the transient plasma jets were observed. At the same time, the existence of this feature limits the valid time frame during which the initially conceived analytical model can be applied to reconstruct density distributions from experimental scattering images. As such, a refinement of the model was necessary that accounts for the newly found feature.

Analytic modeling and Fourier transforms of the target density projection have further limits in their predictive power of X-ray scattering experiments, however. For these techniques to apply reasonably well, single-scattering of the X-ray photons is assumed. Furthermore, the probe pulse duration as well as the time to traverse the target need to be much shorter than the time scale of the probed plasma dynamics. But for micrometer solid-density plasmas driven by ultra-intense optical lasers these assumptions can be violated. Spatial inhomogeneity of the X-ray laser profile, a limited coherence time and the interaction of X-ray photons with the target as it is probed further increase the challenges to predict the scattering image. In the final part of this thesis, the author used the PIC simulation data of the structured grating targets to perform the first start-to-end simulations of Small-Angle X-ray Scattering from a laser-driven solid density plasma with the Monte-Carlo photon scattering code ParaTAXIS. As part of the EUCALL collaboration, these simulations proved the interoperability of several simulation codes from the creation of the X-ray pulse via the optical laser-matter interaction and the probing of the laser-driven target, dynamically creating a virtual scattering image on a synthetic detector. As a proof-of-concept they showed that analytical and Fourier transform results could be reproduced well by the code in single-scattering mode. With activated multi-scattering, the amount of virtual photons during the campaign was just large enough to reproduce the main scattering features but the signal-to-noise ratio was still too low to resolve peaks at large scattering vectors. Many more virtual photons are necessary to suppress the level of statis-

tical noise far enough that the relative peak heights of the scattering maxima of the grating become comparable. Nevertheless, with the connection of a multitude of different codes into `simex_platform`, the Python library providing control over the codes and communicating their outputs between each other based on the openPMD standard, an essential step has been taken towards complete start-to-end simulations of laser-plasma pump-probe experiments that can be performed by a single researcher and allow for scattering image predictions of an evolving target beyond the Born approximation.

6. Conclusions

The particle-in-cell simulation campaign run by the author on Piz Daint remains to this date one of the largest ever performed on the matter of laser-ion acceleration from a solid-density target where realistic laser pulse shapes are included. It produced a wealth of high-quality data that still provides opportunity for further analysis pursuing questions like the Z-dependence of the acceleration performance under the influence of last-picosecond intensity ramps. The campaign was originally motivated to explain the significant energy gains seen in the experiments performed at the DRACO laser system at the HZDR DRACO laser that were recently published by Ziegler et al. [10]. SPIDER measurements of the few hundred femtoseconds around the peak intensity indicated that a pulse that is initially shallower than the shortest possible pulse leads to higher proton cutoff energies and that especially the last picosecond of laser contrast is decisive when working at a target's optimum thickness. Indeed, the results obtained in this thesis confirm that the last picosecond intensity contrast has a substantial effect on the proton acceleration performance. However, whereas the same manipulations of the temporal pulse shape lead to higher experimental proton energies for both metal and plastic targets in experiments, there was no direct evidence for an energy increase from hydro-carbon-based targets in the simulations. The laser absorption was consistently lower in cases that had added energy in front of the Gaussian main pulse. On the other hand, metal targets showed an improved proton maximum energy from pulses preceded by an intensity ramp which is caused by better placement of protons in the accelerating sheath field. A direct measurement of these increased proton energies could currently still be lying below the detection limit of commonly used Thomson-parabola spectrometers (TPS) or radiochromic film stacks (RCF). On the other hand, experimentally measured proton beam spots, especially from ultrathin targets, can be angularly shifted when the target becomes transparent during the interaction[5, 342]. They are often warped or structured by surface– or counter-streaming instabilities[23, 320, 343] or field structures in the residual gas in the vacuum chamber which gets ionized by light leakage past mass-limited or wire-targets[G22]. All of these effects make it especially challenging to reliably align the high-energetic parts of such ion beams towards diagnostics through, e.g. pinholes for TPS measurements. However, a characteristic spectral feature with a distinct cutoff and the high-energy tail, such as was found in this thesis, was also observed in other previous works utilizing metal targets[23, 141], but at micrometer target thicknesses. With the newly obtained results it could be conceived that such proton spectra may result from protons of differently accelerated populations from the target front and rear sides. Specifically designed multi-layer targets could be promising in increasing the proton number when the effect found in this thesis work is specifically exploited and a variation in target thickness would result in

a delay between the differently accelerated proton populations, thus possibly probing undisturbed or structured sheath fields.

Moreover, the methods developed in this thesis improve the capability of doing wide-ranging large-scale 3D parameter scans of laser-solid interactions with PIConGPU including the laser intensity ramp phase and realistic ionization dynamics. Since the results obtained in this thesis work show how sensitive the laser-ion acceleration performance is under realistic pulse conditions, more simulation work at full scale and dimensionality is necessary to reach predictive capabilities. The workflows for large-scale data transport an analysis that were established between the HZDR, ZIH and CSCS continue to be valuable for the increasing demand of large-scale simulations.

Furthermore, this thesis contributed to the development of simulations and workflows for passive and active probing of transient plasma dynamics at femtosecond-nanometer scales. With X-ray free electron lasers and especially the newly commissioned ReLaX laser at the HIBEF end station[344] of the European XFEL, the plasma dynamics at such scales have finally become accessible. But as was initially motivated, the interpretation of X-ray scattering images is not straightforward and therefore simulations have been performed that guided the first experiment to ever dynamically measure plasma expansion of a known grating structure at the fs-nm level and showed how previously unexpected dynamics require newly designed experiments that can deliver clear signatures as a basis for plasma models. The study of more complex laser-target interaction dynamics, such as instabilities, requires dedicated forward simulation tools that allow to model pump-probe processes, including, e.g., the earlier stages of pre-plasma formation due to a realistic laser temporal intensity distribution of the drive laser, as well as X-ray photon scattering beyond the Born approximation. With this thesis, a step toward start-to-end simulations operated by single researchers was taken in a proof-of-concept manner. It showed that a connection between different simulation codes from the sources of laser radiation via beamlines and target interaction towards detection is possible within the framework of EUCALL's `simex_platform`. With the end of the project, this simulation suite was continued under the umbrella of PanOSC[G18]. The incorporation of Monte-Carlo scattering of photons during a running particle-in-cell simulation is possible but requires new solutions for the modeling of the photons' quantum character to arrive at good signal-to-noise ratios with reasonable photon numbers under realistic scattering conditions. Next to the active probing techniques, radiation signatures from self-generated bremsstrahlung and synchrotron radiation promise to be interesting plasma diagnostics once the higher energy parts become accessible for measurements in experiment, especially on a shot-to-shot basis. But to access the full potential of deriving a target's state during the FWHM of laser main pulse interaction, again 3D simulations are necessary to correctly model the influence of laser polarization and energy dissipation.

7. Outlook

The results obtained from the work of this thesis show that the study of the intricate laser-plasma dynamics connected to the temporal laser contrast, specifically during the last picosecond, promises to remain of great interest. A strong influence of the last picosecond intensity ramp on the laser-ion acceleration performance was shown and even higher energies than with a perfectly Gaussian shaped laser pulse could be achieved under the right circumstances. In a most recent study beyond the scope of this thesis, new quasi-1D[1] simulations performed by the author using the PIC code Smilei[224] have shown very promising results that mirror the trend of proton cutoff energies with temporal laser shape manipulation as reported in [10]. In contrast to the strategy employed before, where the measured temporal pulse contrast was translated into characteristic features like ramps and pre-pulses, the laser contrast curve was constructed via Fourier transform from the spectral phase and spectral amplitude (see 2.2.2). This approach allows to create more realistically complex temporal pulse shapes that can trigger interactions that are more challenging to understand but they allow for the direct change of ΔGVD and ΔTOD to manipulate the pulse shape, as was done in [10]. While the parameter changes are not directly translatable between the low-dimensional, still idealized simulations and the experimental values, a similar trend was observed as can be seen in Fig. 7.2. Simulations were first performed with pulses from an idealized 4th order super-Gaussian spectral amplitude $A(\omega)$ for which an increase in ΔTOD did not result in overall higher proton cutoff energies. They were then repeated with a modified spectral intensity modeled after measurements at DRACO, featuring a more narrow 10th order super-Gaussian with a parabolic dip which was blue-shifted by 12.5 nm to be exactly at the central wavelength of $\lambda = 800$ nm. Examples for pulses from the idealized (more realistic) laser spectral amplitude input and $\Delta GVD = 0$ are shown in Fig. 7.1.

The simulation results in Fig. 7.2a of an unexpanded 400 nm Formvar target, irradiated by a laser of maximum $a_0 = 50$ (decreases by up to 40 % for $\Delta TOD \neq 0$ and much more when $\Delta GVD \neq 0$), show that the maximum proton energies increase with increasing values of ΔTOD. The same behavior was observed in the experiment (see Fig. 7.2b). Early analyses suggest that the non-ideal spectrum reduces the quality of the temporal compression of the laser pulse. Especially after the plasma mirror, the spectrum is narrower than before which causes a longer pulse. Under these conditions the leading and falling intensity ramps are less steep than previously expected.

These new developments give rise to more simulation scans that systematically study how the interaction with these more complex laser pulse shapes can again be replaced by analytic

[1] quasi-1D means here two-dimensional simulations with narrow transverse dimensions of $L_{\text{transverse}} < \lambda_L$.

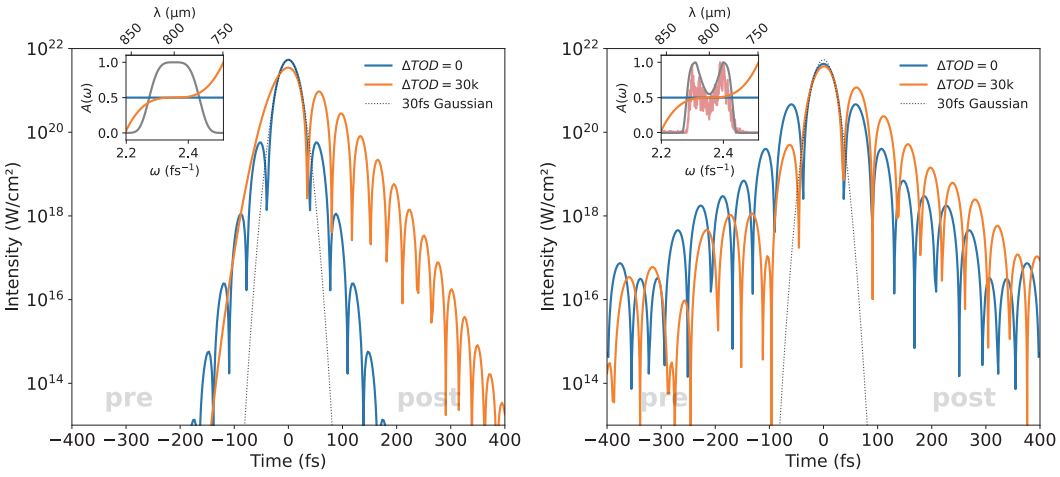

Figure 7.1.: Artificially created laser intensity envelopes from spectral phase and spectral amplitude input. *Left:* Pulses created with an idealized spectral amplitude of 4th order super-Gaussian shape. *Right:* Pulses created with $A(\omega)$ as a 10th order super-Gaussian with a parabolic dip, modeled after an example shot at DRACO at a position behind the plasma mirror (red line in the inset). The gray lines in the insets show synthetic spectral amplitudes $A(\omega)$. The spectral phase $\varphi(\omega)$ in the ideally flat (blue, $\Delta TOD = 0\,\mathrm{fs}^3$) and modified (orange, $\Delta TOD = 30\,000\,\mathrm{fs}^3$) cases are schematically shown in arbitrary units.

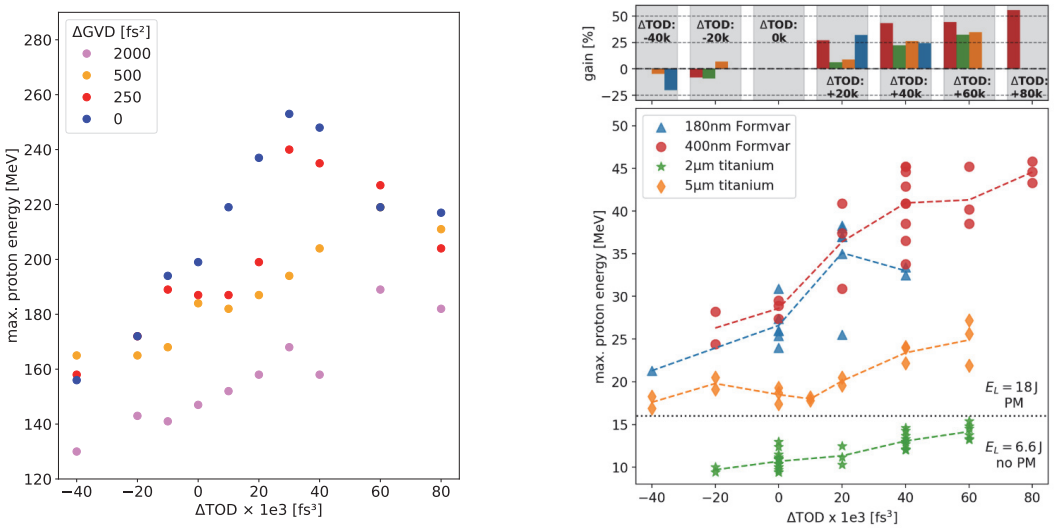

(a) Maximum proton energies from quasi-1D PIC simulations with the code Smilei[224]. The target was a 400 nm unexpanded Formvar foil. Laser pulses were created via Fourier transform from a spectral amplitude, modeled after a measurement behind the plasma mirror in the experiments of Fig. 7.2b, and a spectral phase with manipulations of ΔTOD and ΔGVD. Examples for $\Delta GVD = 0$ are shown in the right panel of Fig. 7.1.

(b) Maximum proton energies measured in experiments at DRACO with varying target types and laser energies (separated by color and marker). The laser temporal profile was manipulated with changes of ΔTOD. However, the values cannot be directly compared to Fig. 7.2a where the pulse starts from an ideally flat spectral phase. The upper panel shows relative energy gain compared to the settings at $\Delta TOD = 0$. Originally in [10].

Figure 7.2.: Comparison of maximum proton energies from quasi-1D PIC simulations (left) and laser-proton acceleration experiments at DRACO with varying target materials.

features such as ramps and pre-pulses or initial density and temperature distributions. These studies are in progress and while the remaining work is beyond the scope of this thesis, the results are very promising with respect to explaining the substantial energy gain that was observed in the experimental work of Ziegler et al. [10]. However, since the new parameter scans require a much longer time excursion of possibly several picoseconds around the main pulse maximum, mainly quasi-1D parameter simulations are performed first due to their low computational cost. Additionally, magneto-hydrodynamic simulations (MHD) can deliver the initial conditions that arise from the interaction with the nanosecond ASE and picosecond pedestal of the laser pulse. From the resulting parameter maps, chosen scans have to be repeated in higher dimensions for which PIConGPU is well-suited with its exceptionally fast computing speed compared to other PIC codes. Here, modern machine learning methods like transfer learning[345, 346] can guide the search for trends by transferring them from 1D via 2D to large-scale 3D simulations that will be necessary to create surrogate models or analytical descriptions of the acceleration process with predictive capability. While still presenting a challenging endeavor for the next decades in general, start-to-end representations of at least selected laser-ion acceleration experimental setups by digital twins seem to be in reach and promise to finally provide the necessary understanding to give control over the beam parameters of laser-accelerated protons for applications. The increasing fidelity in laboratory experiment diagnostics and their parallel operation on a shot-to-shot basis is a most vital part to achieve this goal. It is foreseeable that both in experiment as well as simulation increased output rates and output sizes are demanding the continued development of in-situ data analysis and reduction methods, especially where the computation will require next generation supercomputers of the exascale.

Appendices

A. Equation conversion

This is an example calculation for how the equation for laser energy is converted to a unit system where $e = m_e = c = \omega_L = k_L = 1$. The laser propagation direction is the y-direction like in PIConGPU. Note here that $w_r^{\mathrm{FWHM}} = \sqrt{2\ln 2} w_0 = 2\sqrt{2\ln 2}\sigma_r$. The parameter w_0 is called *beam waist* (in the focus) or also *spot size* which is not to be confused with the FWHM spot size w_r^{FWHM}.

$$E_L = \sqrt{2\pi}^3 I_0 \sigma_x \sigma_z \sigma_t \qquad \text{(integration over 3 Gaussians)} \tag{A.1}$$

$$\overset{\sigma_x=\sigma_z=\sigma_r}{=} \sqrt{2\pi}^3 I_0 \sigma_r^2 \sigma_t \tag{A.2}$$

$$= \left(\frac{\sqrt{2\pi}}{2\sqrt{2\ln 2}} \right)^3 I_0 (\sqrt{2\ln 2} w_0)^2 t_p \tag{A.3}$$

$$= I_0 \pi w_0^2 t_p \frac{\sqrt{2\pi}}{4\sqrt{2\ln 2}} \tag{A.4}$$

Calculating the Laser Energy

$$I_0 = \frac{1}{2} c \varepsilon_0 E_0^2, \qquad a_0 = \frac{e E_0}{m_e c \omega_L} \tag{A.5}$$

$$= \frac{1}{2} c \varepsilon_0 \left(\frac{a_0 m_e c \omega_L}{e} \right)^2 \tag{A.6}$$

$$E_L = I_0 \cdot A \cdot t_p \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}} \qquad A = \pi w_0^2 \tag{A.7}$$

$$= \frac{1}{2} c \varepsilon_0 a_0^2 \frac{m_e c^2 \omega_L^2}{e^2} \cdot \pi w_0^2 \cdot t_p \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}}, \qquad \omega_L = \frac{2\pi}{T_L} \tag{A.8}$$

$$\underbrace{\frac{E_L}{m_e c^2}}_{\tilde{E}_L} = \frac{a_0^2}{2} \pi w_0^2 \underbrace{\left(t_p \frac{2\pi}{T_L} \right)}_{\tilde{t}_p} \varepsilon_0 \cdot \frac{m_e c \omega_L}{e^2} \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}} \qquad \left| \cdot \frac{\omega_L c^2}{\omega_L c^2} = \frac{2\pi c^2}{\lambda_L \omega_L c} \right. \tag{A.9}$$

$$\tilde{E}_L = \frac{a_0^2}{2} \pi \left(w_0 \frac{2\pi}{\lambda_L} \right)^2 \tilde{t}_p \cdot \varepsilon_0 \cdot \underbrace{\frac{m_e c^3 T_L}{e^2 2\pi}}_{\tilde{\varepsilon}_0} \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}} \tag{A.10}$$

$$= \frac{a_0^2}{2} \pi \tilde{w}_0^2 \tilde{t}_p \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}} \cdot 8.854 \cdot 10^{-12}\,\mathrm{As/Vm} \cdot 4.063 \cdot 10^{17}\,\mathrm{Vm/As} \tag{A.11}$$

$$= \underline{\frac{a_0^2}{2} \pi \tilde{w}_0^2 \tilde{t}_p \cdot 3.6 \cdot 10^6 \cdot \frac{\sqrt{2\pi}}{4\sqrt{2 \ln 2}}} \tag{A.12}$$

B. PICLS LN04 Campaign Input Files and Code

The version of the code PICLS that was used for the simulations of the LN04 experimental beamtime as described in 4.3.2 was uploaded to the private Github repository `https://github.com/ComputationalRadiationPhysics/projects-ions/` under the branch name *prod-2016-LN04-picls*. The input data sets are furthermore available upon reasonable request from the HZDR archive with the ID 3273. The density profiles that were used can be found in the file `picls/code/src/density_profile.f`. Different density profiles are chosen via an ID that is configured in the main input file `input.in` New IDs corresponding to the cases from this campaign carry the date of their first creation.

- *T1 (pure Si)* : `20160616`
- *T2/T3 (buried layer)* : `20160617`
- *T4 (pure Si - grating on back)* : `20160618`
- *T1 - longitudinal pre-expansion* : `20161025`
- *T1 - pre-expanded in all directions* : `20161123`

C. PRACE Piz Daint Campaign Simulation Archives

ID	Set	Mat.	d (nm)	Contrast	E (J)	Notes	Size (TB)
3791	003	Cu	30	PC	3	res.: $dx = 1/12\,\lambda_{pe}$, I/O not optimal yet	3
3792	005	Cu	30	PC	3	res.: $dx = 1/12\,\lambda_{pe}$	11
3472	006	Cu	10	PC	3		13
3801	006	Cu	10	PC	3	w/o full I/O dumps	2
3477	006	Cu	30	PC	3		13
3806	006	Cu	30	PC	3	w/o full I/O dumps	3
3478	006	Cu	300	PC	3		30
3800	006	Cu	300	PC	3	only few full I/O dumps	11
3479	007	Cu	10	R3	3		14
3488	007	Cu	10	R4	3		17
3507	007	Cu	10	R5	3		9
3805	007	Cu	30	R2	3		2
3798	007	Cu	30	R3	3		16
3808	007	Cu	30	R4	3		6
3809	007	Cu	30	R5	3		6
3804	007	Cu	300	R2	3		3
3508	007	Cu	300	R3	3		32
4029	007	Cu	300	R4	3		10
4042	007	Cu	300	R5	3		24
3844	008	Cu	6	PC	30		43
3827	008	Cu	10	PC	30		4
3816	008	Cu	10	R3	30		9
3817	008	Cu	10	R4	30		4
3814	008	Cu	10	R5	30		7
3821	008	Cu	20	PC	30		44
4082	008	Cu	30	PC	30		6
3813	008	Cu	20	R3	30		53
4080	008	Cu	30	R3	30		36
3830	008	Cu	30	R4	30		7
3829	008	Cu	300	PC	30		7
4012	008	Cu	300	R3	30		11
3828	008	Cu	300	R4	30		21

Table C.1.: Archive IDs, description and size for simulation data obtained on Piz Daint during the large-scale campaign detailed in Sec. 4.2 (Part 1).

ID	Set	Mat.	d (nm)	Contrast	E (J)	Notes	Size (TB)
4135	009	Cu	9	PC	30	with bremsstrahlung	12
4136	009	Cu	9	PC	30	with bremsstrahlung, circularly polarized laser	11
3846	009	Cu	30	PC	30	with bremsstrahlung	8
3845	009	Cu	100	PC	30	with bremsstrahlung	8
4041	010	Cu	30	R4	30	with bremsstrahlung	1
3786	011	LCT	10	PC	30	with bremsstrahlung	13
3788	011	LCT	63	PC	30	with bremsstrahlung	14
3787	011	LCT	221	PC	30	with bremsstrahlung	18
4077	015	Cu	6	PC	30	contaminant layers distinguishable (front/back)	12
4062	015	Cu	17	R4	30	contaminant layers distinguishable (front/back)	19
4063	015	Cu	17	PC	30	contaminant layers distinguishable (front/back)	13
3715	016	FV	37	PC	10		16
3700	016	FV	37	R4	10		19
3623	016	FV	37	R5	10		18
3763	016	FV	62	PC	10		17
3780	016	FV	62	R5	10		8
3782	016	FV	62	R4	10		13
3764	016	FV	309	PC	10		14
3622	016	FV	105	R4	10		21
3767	016	FV	105	R5	10		14
3773	016	FV	185	R4	10		17
3775	016	FV	185	PC	10		16
3779	016	FV	309	R4	10		22
3757	016	FV	617	PC	10	early crash (out of memory)	7
3765	016	FV	617	R5	10	early crash (out of memory)	3
5070	Compilation Archive (all run directories w/ inputs, synthetic diagnostic output, analysis scripts and analysis results, jupyter notebooks for analysis from Piz Daint (CSCS), Taurus (ZIH), Hypnos & Hemera (HZDR)data reduction scripts, log files)						26

Table C.2.: Archive IDs, description and size for simulation data obtained on Piz Daint during the large-scale campaign detailed in Sec. 4.2 (Part 2)

References

[1] RA Snavely, MH Key, SP Hatchett, TE Cowan, M Roth, TW Phillips, MA Stoyer, EA Henry, TC Sangster, MS Singh, SC Wilks, A MacKinnon, A Offenberger, DM Pennington, K Yasuike, AB Langdon, BF Lasinski, J Johnson, MD Perry, and EM Campbell. "Intense high-energy proton beams from petawatt-laser irradiation of solids". In: *Physical Review Letters* 85.14 (2000), pp. 2945–2948. ISSN: 00319007. DOI: 10.1103/PhysRevLett.85.2945 (cit. on pp. 3, 18).

[2] A Maksimchuk, S Gu, K Flippo, D Umstadter, and VY Bychenkov. "Forward Ion Acceleration in Thin Films Driven by a High-Intensity Laser". In: *Physical Review Letters* 84.18 (2000), pp. 4108–4111. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.84.4108 (cit. on pp. 3, 18).

[3] SP Hatchett, CG Brown, TE Cowan, EA Henry, JS Johnson, MH Key, JA Koch, AB Langdon, BF Lasinski, RW Lee, AJ Mackinnon, DM Pennington, MD Perry, TW Phillips, M Roth, TC Sangster, MS Singh, RA Snavely, MA Stoyer, SC Wilks, and K Yasuike. "Electron, photon, and ion beams from the relativistic interaction of Petawatt laser pulses with solid targets". In: *Physics of Plasmas* 7.5 (2000), pp. 2076–2082. ISSN: 1070-664X. DOI: 10.1063/1.874030 (cit. on pp. 3, 18).

[4] F Wagner, O Deppert, C Brabetz, P Fiala, A Kleinschmidt, P Poth, VA Schanz, A Tebartz, B Zielbauer, M Roth, T Stöhlker, and V Bagnoud. "Maximum Proton Energy above 85 MeV from the Relativistic Interaction of Laser Pulses with Micrometer Thick CH2 Targets". In: *Physical Review Letters* 116.20 (2016), pp. 1–5. ISSN: 10797114. DOI: 10.1103/PhysRevLett.116.205002 (cit. on p. 3).

[5] A Higginson, RJ Gray, M King, RJ Dance, SDR Williamson, NMH Butler, R Wilson, R Capdessus, C Armstrong, JS Green, SJ Hawkes, P Martin, WQ Wei, SR Mirfayzi, XH Yuan, S Kar, M Borghesi, RJ Clarke, D Neely, and P McKenna. "Near-100 MeV protons via a laser-driven transparency-enhanced hybrid acceleration scheme". In: *Nature Communications* 9.1 (2018), p. 724. ISSN: 2041-1723. DOI: 10.1038/s41467-018-03063-9 (cit. on pp. 3, 145).

[6] C Danson, D Hillier, N Hopps, and D Neely. "Petawatt class lasers worldwide". In: *High Power Laser Science and Engineering* 3 (2015), e3. ISSN: 2095-4719. DOI: 10.1017/hpl.2014.52 (cit. on p. 3).

[7] S Gales, KA Tanaka, DL Balabanski, F Negoita, D Stutman, O Tesileanu, CA Ur, D Ursescu, I Andrei, S Ataman, MO Cernaianu, L D'Alessi, I Dancus, B Diaconescu, N Djourelov, D Filipescu, P Ghenuche, DG Ghita, C Matei, K Seto, M Zeng, and NV Zamfir. "The extreme light infrastructure-nuclear physics (ELI-NP) facility: new horizons in physics with 10 PW ultra-intense lasers and 20 MeV brilliant gamma beams". In: *Reports on Progress in Physics* 81.9 (2018), p. 094301. ISSN: 0034-4885. DOI: 10.1088/1361-6633/aacfe8 (cit. on p. 3).

[8] L Obst et al. "Efficient laser-driven proton acceleration from cylindrical and planar cryogenic hydrogen jets". In: *Scientific Reports* 7.1 (2017), p. 10248. ISSN: 2045-2322. DOI: 10.1038/s41598-017-10589-3 (cit. on pp. 3, 19).

[9] WJ Ma, IJ Kim, JQ Yu, IW Choi, PK Singh, HW Lee, JH Sung, SK Lee, C Lin, Q Liao, JG Zhu, HY Lu, B Liu, HY Wang, RF Xu, XT He, JE Chen, M Zepf, J Schreiber, XQ Yan, and CH Nam. "Laser Acceleration of Highly Energetic Carbon Ions Using a Double-Layer Target Composed of Slightly Underdense Plasma and Ultrathin Foil". In: *Physical Review Letters* 122.1 (2019), p. 014803. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.122.014803 (cit. on p. 3).

[10] T Ziegler, D Albach, C Bernert, S Bock, FE Brack, TE Cowan, NP Dover, M Garten, L Gaus, R Gebhardt, I Goethel, U Helbig, A Irman, H Kiriyama, T Kluge, A Kon, S Kraft, F Kroll, M Loeser, J Metzkes-Ng, M Nishiuchi, L Obst-Huebl, T Püschel, M Rehwald, HP Schlenvoigt, U Schramm, and K Zeil. "Proton beam quality enhancement by spectral phase control of a PW-class laser system". In: *Scientific Reports* 11.1 (2021), p. 7338. ISSN: 20452322. DOI: 10.1038/s41598-021-86547-x. arXiv: 2007.11499 (cit. on pp. 3, 5, 78, 81, 102, 106, 107, 145, 147–149).

[11] M Borghesi, J Fuchs, SV Bulanov, AJ MacKinnon, PK Patel, and M Roth. "Fast Ion Generation by High-Intensity Laser Irradiation of Solid Targets and Applications". In: *Fusion Science and Technology* 49.3 (2006), pp. 412–439. ISSN: 1536-1055. DOI: 10.13182/FST06-A1159 (cit. on p. 3).

[12] SD Kraft, C Richter, K Zeil, M Baumann, E Beyreuther, S Bock, M Bussmann, TE Cowan, Y Dammene, W Enghardt, U Helbig, L Karsch, T Kluge, L Laschinsky, E Lessmann, J Metzkes, D Naumburger, R Sauerbrey, M Schürer, M Sobiella, J Woithe, U Schramm, and J Pawelke. "Dose-dependent biological damage of tumour cells by laser-accelerated proton beams". In: *New Journal of Physics* 12.8 (2010), p. 085003. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/8/085003 (cit. on p. 3).

[13] H Daido, M Nishiuchi, and AS Pirozhkov. "Review of laser-driven ion sources and their applications". In: *Reports on Progress in Physics* 75.5 (2012), p. 056401. ISSN: 0034-4885. DOI: 10.1088/0034-4885/75/5/056401 (cit. on pp. 3, 24, 27).

[14] F Albert, ME Couprie, A Debus, MC Downer, J Faure, A Flacco, LA Gizzi, T Grismayer, A Huebl, C Joshi, M Labat, WP Leemans, AR Maier, SP Mangles, P Mason, F Mathieu, P Muggli, M Nishiuchi, J Osterhoff, PP Rajeev, U Schramm, J Schreiber, AG Thomas, JL Vay, M Vranic, and K Zeil. "2020 Roadmap on Plasma Accelerators". In: *New Journal of Physics* (2020), pp. 0–31. DOI: 10.1088/1367-2630/abcc62/pdf (cit. on p. 4).

[15] A Huebl. "PIConGPU : Predictive Simulations of Laser-Particle Accelerators with Manycore Hardware". PhD thesis. 2019, p. 171. DOI: 10.5281/zenodo.3266820 (cit. on pp. 4, 15, 41, 42).

[16] D Strickland and G Mourou. "Compression of amplified chirped optical pulses". In: *Optics Communications* 55.6 (1985), pp. 447–449. ISSN: 00304018. DOI: 10.1016/0030-4018(85)90151-8 (cit. on pp. 3, 27).

[17] T Esirkepov, M Borghesi, SV Bulanov, G Mourou, and T Tajima. "Highly Efficient Relativistic-Ion Generation in the Laser-Piston Regime". In: *Physical Review Letters* 92.17 (2004), p. 175003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.92.175003 (cit. on pp. 3, 23).

[18] APL Robinson, M Zepf, S Kar, RG Evans, and C Bellei. "Radiation pressure acceleration of thin foils with circularly polarized laser pulses". In: *New Journal of Physics* 10.1 (2008), p. 013021. ISSN: 1367-2630. DOI: 10.1088/1367-2630/10/1/013021. arXiv: 0708.2040 (cit. on p. 3).

[19] A Henig, D Kiefer, M Geissler, SG Rykovanov, R Ramis, R Hörlein, J Osterhoff, Z Major, L Veisz, S Karsch, F Krausz, D Habs, and J Schreiber. "Laser-Driven Shock Acceleration of Ion Beams from Spherical Mass-Limited Targets". In: *Physical Review Letters* 102.9 (2009), p. 095002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.102.095002 (cit. on pp. 3, 24).

[20] A Macchi, S Veghini, and F Pegoraro. ""Light Sail" Acceleration Reexamined". In: *Physical Review Letters* 103.8 (2009), p. 085003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.103.085003 (cit. on pp. 3, 23, 24).

[21] F Pegoraro and SV Bulanov. "Photon Bubbles and Ion Acceleration in a Plasma Dominated by the Radiation Pressure of an Electromagnetic Pulse". In: *Physical Review Letters* 99.6 (2007), p. 065002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.99.065002 (cit. on pp. 3, 23, 107).

[22] CAJ Palmer, J Schreiber, SR Nagel, NP Dover, C Bellei, FN Beg, S Bott, RJ Clarke, AE Dangor, SM Hassan, P Hilz, D Jung, S Kneip, SPD Mangles, KL Lancaster, A Rehman, APL Robinson, C Spindloe, J Szerypo, M Tatarakis, M Yeung, M Zepf, and Z Najmudin. "Rayleigh-Taylor Instability of an Ultrathin Foil Accelerated by the Radiation Pressure of an Intense Laser". In: *Physical Review Letters* 108.22 (2012), p. 225002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.108.225002 (cit. on pp. 3, 23).

[23] J Metzkes, T Kluge, K Zeil, M Bussmann, SD Kraft, TE Cowan, and U Schramm. "Experimental observation of transverse modulations in laser-driven proton beams". In: *New Journal of Physics* 16 (2014). ISSN: 13672630. DOI: 10.1088/1367-2630/16/2/023008 (cit. on pp. 3, 107, 145).

[24] T Tajima and JM Dawson. "Laser Electron Accelerator". In: *Physical Review Letters* 43.4 (1979), pp. 267–270. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.43.267 (cit. on pp. 4, 12).

[25] CGR Geddes, C Toth, J van Tilborg, E Esarey, CB Schroeder, D Bruhwiler, C Nieter, J Cary, and WP Leemans. "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding". In: *Nature* 431.7008 (2004), pp. 538–541. ISSN: 0028-0836. DOI: 10.1038/nature02900 (cit. on pp. 4, 12).

[26] SPD Mangles, CD Murphy, Z Najmudin, AGR Thomas, JL Collier, AE Dangor, EJ Divall, PS Foster, JG Gallacher, CJ Hooker, DA Jaroszynski, AJ Langley, WB Mori, PA Norreys, FS Tsung, R Viskup, BR Walton, and K Krushelnick. "Monoenergetic beams of relativistic electrons from intense laser-plasma interactions". In: *Nature* 431.7008 (2004), pp. 535–538. ISSN: 0028-0836. DOI: 10.1038/nature02939 (cit. on pp. 4, 12).

[27] J Faure, Y Glinec, A Pukhov, S Kiselev, S Gordienko, E Lefebvre, JP Rousseau, F Burgy, and V Malka. "A laser-plasma accelerator producing monoenergetic electron beams". In: *Nature* 431.7008 (2004), pp. 541–544. ISSN: 0028-0836. DOI: 10.1038/nature02963 (cit. on pp. 4, 12).

[28] T Kurz, T Heinemann, MF Gilljohann, YY Chang, JP Couperus Cabadağ, A Debus, O Kononenko, R Pausch, S Schöbel, RW Assmann, M Bussmann, H Ding, J Götzfried, A Köhler, G Raj, S Schindler, K Steiniger, O Zarini, S Corde, A Döpp, B Hidding, S Karsch, U Schramm, A Martinez de la Ossa, and A Irman. "Demonstration of a compact plasma accelerator powered by laser-accelerated electron beams". In: *Nature Communications* 12.1 (2021), p. 2895. ISSN: 2041-1723. DOI: 10.1038/s41467-021-23000-7. arXiv: 1909.06676 (cit. on p. 4).

[29] F Grüner, S Becker, U Schramm, T Eichner, M Fuchs, R Weingartner, D Habs, J Meyer-ter-Vehn, M Geissler, M Ferrario, L Serafini, B van der Geer, H Backe, W Lauth, and S Reiche. "Design considerations for table-top, laser-based VUV and X-ray free electron lasers". In: *Applied Physics B* 86.3 (2007), pp. 431–435. ISSN: 0946-2171. DOI: 10.1007/s00340-006-2565-7. arXiv: 0612125 [physics] (cit. on p. 4).

[30] W Wang, K Feng, L Ke, C Yu, Y Xu, R Qi, Y Chen, Z Qin, Z Zhang, M Fang, J Liu, K Jiang, H Wang, C Wang, X Yang, F Wu, Y Leng, J Liu, R Li, and Z Xu. "Free-electron lasing at 27 nanometres based on a laser wakefield accelerator". In: *Nature* 595.7868 (2021), pp. 516–520. ISSN: 0028-0836. DOI: 10.1038/s41586-021-03678-x (cit. on pp. 4, 13).

[31] ET Gumbrell, RA Smith, T Ditmire, A Djaoui, SJ Rose, and MHR Hutchinson. "Picosecond optical probing of ultrafast energy transport in short pulse laser solid target interaction experiments". In: *Physics of Plasmas* 5.10 (1998), pp. 3714–3721. ISSN: 1070-664X. DOI: 10.1063/1.872981 (cit. on p. 4).

[32] A Sävert, SPD Mangles, M Schnell, E Siminos, JM Cole, M Leier, M Reuter, MB Schwab, M Möller, K Poder, O Jäckel, GG Paulus, C Spielmann, S Skupin, Z Najmudin, and MC Kaluza. "Direct Observation of the Injection Dynamics of a Laser Wakefield Accelerator Using Few-Femtosecond Shadowgraphy". In: *Physical Review Letters* 115.5 (2015), p. 055002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.115.055002. arXiv: 1402.3052 (cit. on p. 4).

[33] P Antici, J Fuchs, M Borghesi, L Gremillet, T Grismayer, Y Sentoku, E D'Humières, CA Cecchetti, A Mančić, AC Pipahl, T Toncian, O Willi, P Mora, and P Audebert. "Hot and Cold Electron Dynamics Following High-Intensity Laser Matter Interaction". In: *Physical Review Letters* 101.10 (2008), p. 105004. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.101.105004 (cit. on p. 4).

[34] MC Kaluza, MIK Santala, J Schreiber, GD Tsakiris, and KJ Witte. "Time-sequence imaging of relativistic laser-plasma interactions using a novel two-color probe pulse". In: *Applied Physics B* 92.4 (2008), pp. 475–479. ISSN: 0946-2171. DOI: 10.1007/s00340-008-3150-z (cit. on p. 4).

[35] O Jäckel, J Polz, SM Pfotenhauer, HP Schlenvoigt, H Schwoerer, and MC Kaluza. "All-optical measurement of the hot electron sheath driving laser ion acceleration from thin foils". In: *New Journal of Physics* 12.10 (2010), p. 103027. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/10/103027 (cit. on p. 4).

[36] M Bocoum, F Böhle, A Vernier, A Jullien, J Faure, and R Lopez-Martens. "Spatial-domain interferometer for measuring plasma mirror expansion". In: *Optics Letters* 40.13 (2015), p. 3009. ISSN: 0146-9592. DOI: 10.1364/OL.40.003009 (cit. on p. 4).

[37] V Bagnoud, J Hornung, T Schlegel, B Zielbauer, C Brabetz, M Roth, P Hilz, M Haug, J Schreiber, and F Wagner. "Studying the Dynamics of Relativistic Laser-Plasma Interaction on Thin Foils by Means of Fourier-Transform Spectral Interferometry". In: *Physical Review Letters* 118.25 (2017), p. 255003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.118.255003 (cit. on p. 4).

[38] T Ziegler, M Rehwald, L Obst, C Bernert, FE Brack, CB Curry, M Gauthier, SH Glenzer, S Göde, L Kazak, SD Kraft, M Kuntzsch, M Loeser, J Metzkes-Ng, C Rödel, HP Schlenvoigt, U Schramm, M Siebold, J Tiggesbäumker, S Wolter, and K Zeil. "Optical probing of high intensity laser interaction with micron-sized cryogenic hydrogen jets". In: *Plasma Physics and Controlled Fusion* 60.7 (2018), p. 074003. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aabf4f (cit. on p. 4).

[39] J Hornung. "Study of preplasma properties using time-resolved reflection spectroscopy". English. Dissertation. Friedrich-Schiller-Universität Jena, 2021, p. 129. DOI: 10.22032/dbt.50024 (cit. on p. 4).

[40] S Souri, R Amrollahi, and R Sadighi-Bonabi. "Laser-driven proton acceleration enhancement by the optimized intense short laser pulse shape". In: *Physics of Plasmas* 24.5 (2017), p. 053108. ISSN: 1070-664X. DOI: 10.1063/1.4982611 (cit. on pp. 4, 28, 29, 76).

[41] S Souri, R Amrollahi, and R Sadighi-Bonabi. "Improvement of laser-driven proton beam quality by optimized intense chirped laser pulses". In: *Physics of Plasmas* 25.1 (2018), p. 013117. ISSN: 1070-664X. DOI: 10.1063/1.5009015 (cit. on pp. 4, 29).

[42] S Kumar, K Gopal, and DN Gupta. "Proton acceleration from overdense plasma target interacting with shaped laser pulses in the presence of preplasmas". In: *Plasma Physics and Controlled Fusion* 61.8 (2019), p. 085001. ISSN: 0741-3335. DOI: 10.1088/1361-6587/ab216e (cit. on pp. 4, 28, 29, 76).

[43] P Mora. "Plasma Expansion into a Vacuum". In: *Physical Review Letters* 90.18 (2003), p. 185002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.90.185002 (cit. on pp. 5, 19, 21).

[44] J Schreiber, F Bell, F Grüner, U Schramm, M Geissler, M Schnürer, S Ter-Avetisyan, BM Hegelich, J Cobble, E Brambrink, J Fuchs, P Audebert, and D Habs. "Analytical model for ion acceleration by high-intensity laser pulses". In: *Physical Review Letters* 97.4 (2006), pp. 1–4. ISSN: 00319007. DOI: 10.1103/PhysRevLett.97.045005 (cit. on pp. 5, 7, 20–22, 39, 64, 75, 93, 97–101, 105).

[45] A Andreev, A Lévy, T Ceccotti, C Thaury, K Platonov, RA Loch, and P Martin. "Fast-Ion Energy-Flux Enhancement from Ultrathin Foils Irradiated by Intense and High-Contrast Short Laser Pulses". In: *Physical Review Letters* 101.15 (2008), p. 155002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.101.155002 (cit. on pp. 5, 22).

[46] P Emma et al. "First lasing and operation of an ångstrom-wavelength free-electron laser". In: *Nature Photonics* 4.9 (2010), pp. 641–647. ISSN: 1749-4885. DOI: 10.1038/nphoton.2010.176 (cit. on p. 5).

[47] T Ishikawa et al. "A compact X-ray free-electron laser emitting in the sub-ångström region". In: *Nature Photonics* 6.8 (2012), pp. 540–544. ISSN: 1749-4885. DOI: 10.1038/nphoton.2012.141 (cit. on p. 5).

[48] W Decking et al. "A MHz-repetition-rate hard X-ray free-electron laser driven by a superconducting linear accelerator". In: *Nature Photonics* 14.6 (2020), pp. 391–397. ISSN: 17494893. DOI: 10.1038/s41566-020-0607-z (cit. on pp. 5, 6).

[49] R Abela et al. *XFEL: The European X-Ray Free-Electron Laser - Technical Design Report*. Hamburg: DESY, 2006, pp. 1–646. ISBN: 978-3-935702-17-1. DOI: 10.3204/DESY{_}06-097 (cit. on pp. 6, 140).

[50] T Kluge, C Gutt, LG Huang, J Metzkes, U Schramm, M Bussmann, and TE Cowan. "Using X-ray free-electron lasers for probing of complex interaction dynamics of ultra-intense lasers with solid matter". In: *Physics of Plasmas* 21.3 (2014). ISSN: 10897674. DOI: 10.1063/1.4869331. arXiv: 1306.0420 (cit. on pp. 6, 64, 120).

[51] T Kluge, M Bussmann, HK Chung, C Gutt, LG Huang, M Zacharias, U Schramm, and TE Cowan. "Nanoscale femtosecond imaging of transient hot solid density plasmas with elemental and charge state sensitivity using resonant coherent diffraction". In: *Physics of Plasmas* 23.3 (2016), p. 033103. ISSN: 1070-664X. DOI: 10.1063/1.4942786. arXiv: 1508.03988 (cit. on pp. 6, 120, 121).

[52] M Bussmann, F Schmitt, U Schramm, J Schuchart, R Widera, H Burau, TE Cowan, A Debus, A Huebl, G Juckeland, T Kluge, WE Nagel, and R Pausch. "Radiative signatures of the relativistic Kelvin-Helmholtz instability". In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '13* January 2013 (2013), pp. 1–12. DOI: 10.1145/2503210.2504564 (cit. on pp. 6, 15, 36, 41, 132).

[53] Y Sentoku, K Mima, Y Kishimoto, and M Honda. "Effects of Relativistic Binary Collisions on PIC Simulation of Laser Plasmas". In: *Journal of the Physical Society of Japan* 67.12 (1998), pp. 4084–4088. ISSN: 0031-9015. DOI: 10.1143/JPSJ.67.4084 (cit. on pp. 6, 18).

[54] Y Sentoku and A Kemp. "Numerical methods for particle simulations at extreme densities and temperatures: Weighted particles, relativistic collisions and reduced currents". In: *Journal of Computational Physics* 227.14 (2008), pp. 6846–6861. ISSN: 00219991. DOI: 10.1016/j.jcp.2008.03.043 (cit. on pp. 6, 18, 36, 124, 127).

[55] A Irman, JP Couperus, A Debus, A Köhler, JM Krämer, R Pausch, O Zarini, and U Schramm. "Improved performance of laser wakefield acceleration by tailored self-truncated ionization injection". In: *Plasma Physics and Controlled Fusion* 60.4 (2018), p. 044015. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aaaef1 (cit. on pp. 7, 66).

[56] A Grund. "A highly-parallel Monte-Carlo-Simulation of X-Ray-Scattering using a Particle-Mesh-Code on GPUs". Diploma thesis. TU Dresden, 2016. DOI: 10.5281/zenodo.1164654 (cit. on pp. 8, 133, 134, 137).

[57] LD Landau and EM Lifshitz. *The Classical Theory of Fields*. 3rd ed. Vol. 2. Pergamon Press, 1971, p. 387. ISBN: 978-0-08-016019-1 (cit. on p. 10).

[58] F Chen and P Publishing. "The Ponderomotive Force". In: *Introduction to Plasma Physics and Controlled Fusion*. 2nd ed. Vol. 1. Introduction to Plasma Physics and Controlled Fusion. Springer, 1984. Chap. 8.4, pp. 305–309. ISBN: 9780306413322 (cit. on p. 10).

[59] F Chen and P Publishing. "Plasma Oscillations". In: *Introduction to Plasma Physics and Controlled Fusion*. 2nd ed. Vol. 1. Introduction to Plasma Physics and Controlled Fusion. Springer, 1984. Chap. 4.3, p. 85. ISBN: 9780306413322 (cit. on p. 11).

[60] D Levy, C Bernert, M Rehwald, IA Andriyash, S Assenbaum, T Kluge, E Kroupp, L Obst-Huebl, R Pausch, A Schultze-Makuch, K Zeil, U Schramm, and V Malka. "Laser-plasma proton acceleration with a combined gas-foil target". In: (2020), pp. 1–8. arXiv: 2007.14089 (cit. on p. 11).

[61] AJ Gonsalves, K Nakamura, J Daniels, C Benedetti, C Pieronek, TCH de Raadt, S Steinke, JH Bin, SS Bulanov, J van Tilborg, CGR Geddes, CB Schroeder, C Tóth, E Esarey, K Swanson, L Fan-Chiang, G Bagdasarov, N Bobrova, V Gasilov, G Korn, P Sasorov, and WP Leemans. "Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide". In: *Physical Review Letters* 122.8 (2019), p. 084801. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.122.084801 (cit. on p. 12).

[62] R Pausch. “Synthetic Radiation Diagnostics as a Pathway for Studying Plasma Dynamics from Advanced Accelerators to Astrophysical Observations”. Dissertation. TU Dresden, 2018, p. 177. DOI: 10.5281/zenodo.3616045 (cit. on pp. 12, 64, 66, 67, 69, 70).

[63] K Nakajima, T Kawakubo, H Nakanishi, A Ogata, Y Kato, Y Kitagawa, R Kodama, K Mima, H Shiraga, K Suzuki, T Zhang, Y Sakawa, T Shoji, Y Nishida, N Yugami, M Downer, D Fisher, B Newberger, and T Tajima. “A proof-of-principle experiment of laser wakefield acceleration”. In: *Physica Scripta* T52.T52 (1994), pp. 61–64. ISSN: 0031-8949. DOI: 10.1088/0031-8949/1994/T52/009 (cit. on p. 12).

[64] A Pukhov and J Meyer-ter-Vehn. “Laser wake field acceleration: The highly non-linear broken-wave regime”. In: *Applied Physics B: Lasers and Optics* 74.4-5 (2002), pp. 355–361. ISSN: 09462171. DOI: 10.1007/s003400200795 (cit. on p. 12).

[65] MC Downer, R Zgadzaj, A Debus, U Schramm, and MC Kaluza. “Diagnostics for plasma-based electron accelerators”. In: *Reviews of Modern Physics* 90.3 (2018), p. 35002. ISSN: 15390756. DOI: 10.1103/RevModPhys.90.035002. arXiv: arXiv:1011.1669v3 (cit. on p. 13).

[66] O Klimo. “Simulations of Ultrashort-Pulse Laser Solid-Target Interactions”. In: *PhD Thesis at CTU in Prague* April (2007), p. 190 (cit. on p. 13).

[67] F Brunel. “Not-so-resonant, resonant absorption”. In: *Physical Review Letters* 59.1 (1987), pp. 52–55. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.59.52 (cit. on p. 13).

[68] WL Kruer and K Estabrook. “JxB heating by very intense laser light”. In: *Physics of Fluids* 28.1 (1985), pp. 430–432. ISSN: 0031-9171. DOI: 10.1063/1.865171 (cit. on p. 13).

[69] P Mulser, SM Weng, and T Liseykina. “Analysis of the Brunel model and resulting hot electron spectra”. In: *Physics of Plasmas* 19.4 (2012), p. 043301. ISSN: 1070-664X. DOI: 10.1063/1.3696034 (cit. on p. 13).

[70] VL(L Ginzburg. *The propagation of electromagnetic waves in plasmas*. Oxford ;New York: Pergamon Press, 1970, p. 615. ISBN: 9780080155692 (cit. on p. 13).

[71] DW Forslund, JM Kindel, K Lee, EL Lindman, and RL Morse. “Theory and simulation of resonant absorption in a hot plasma”. In: *Physical Review A* 11.2 (1975), pp. 679–683. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.11.679 (cit. on p. 13).

[72] KG Estabrook, EJ Valeo, and WL Kruer. “Two-dimensional relativistic simulations of resonance absorption”. In: *Physics of Fluids* 18.9 (1975), p. 1151. ISSN: 00319171. DOI: 10.1063/1.861276 (cit. on pp. 13, 29).

[73] T Kluge, M Bussmann, U Schramm, and TE Cowan. “Simple scaling equations for electron spectra, currents, and bulk heating in ultra-intense short-pulse laser-solid interaction”. In: *Physics of Plasmas* 25.7 (2018), p. 073106. ISSN: 1070-664X. DOI: 10.1063/1.5037753. arXiv: 1511.00497 (cit. on pp. 14, 111).

[74] T Kluge, TE Cowan, A Debus, U Schramm, K Zeil, and M Bussmann. “Electron Temperature Scaling in Laser Interaction with Solids”. In: *Physical Review Letters* 107.20 (2011), p. 205003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.107.205003 (cit. on pp. 14, 80, 111).

[75] SC Wilks, AB Langdon, TE Cowan, M Roth, M Singh, S Hatchett, MH Key, D Pennington, A MacKinnon, and RA Snavely. “Energetic proton generation in ultra-intense laser-solid interactions”. In: *Physics of Plasmas* 8.2 (2001), pp. 542–549. ISSN: 1070-664X. DOI: 10.1063/1.1333697 (cit. on p. 14).

[76] EG Gamaly, AV Rode, B Luther-Davies, and VT Tikhonchuk. “Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics”. In: *Physics of Plasmas* 9.3 (2002), pp. 949–957. ISSN: 1070-664X. DOI: 10.1063/1.1447555 (cit. on p. 14).

[77] P Mulser and D Bauer. *High Power Laser-Matter Interaction*. Vol. 238. Springer Tracts in Modern Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 5–89. ISBN: 978-3-540-50669-0. DOI: 10.1007/978-3-540-46065-7 (cit. on pp. 15, 17).

[78] LV Keldysh. “Ionization in the field of a strong electromagnetic wave”. In: *Journal of Experimental and Theoretical Physics* 20.5 (1965), pp. 1307–1314. ISSN: 0038-5646 (cit. on pp. 15, 16, 67).

[79] HR Reiss. “Unsuitability of the Keldysh parameter for laser fields”. In: *Physical Review A* 82.2 (2010), p. 023418. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.82.023418 (cit. on pp. 15, 16).

[80] HR Reiss. “The tunneling model of laser-induced ionization and its failure at low frequencies”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 47.20 (2014), pp. 2–3. ISSN: 0953-4075. DOI: 10.1088/0953-4075/47/20/204006. arXiv: 1403.0568 (cit. on pp. 15, 16).

[81] D Bauer and P Mulser. “Exact field ionization rates in the barrier suppression-regime from numerical TDSE calculations”. In: (1998), p. 22. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.59.569. arXiv: 9802042 [physics] (cit. on pp. 16, 36).

[82] FHM Faisal. “Multiple absorption of laser photons by atoms”. In: *Journal of Physics B: Atomic and Molecular Physics* 6.4 (1973), pp. L89–L92. ISSN: 0022-3700. DOI: 10.1088/0022-3700/6/4/011 (cit. on p. 16).

[83] HR Reiss. “Effect of an intense electromagnetic field on a weakly bound system”. In: *Physical Review A* 22.5 (1980), pp. 1786–1813. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.22.1786 (cit. on p. 16).

[84] HR Reiss. “Relativistic strong-field photoionization”. In: *Journal of the Optical Society of America B* 7.4 (1990), p. 574. ISSN: 0740-3224. DOI: 10.1364/JOSAB.7.000574 (cit. on p. 16).

[85] MV Ammosov, NB Delone, and VP Krainov. “Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field”. In: *Sov. Phys. JETP* 64.December 1986 (1986), pp. 1191–1194. ISSN: 0044-4510. DOI: 10.1117/12.938695 (cit. on p. 16).

[86] HR Reiss. “Theoretical methods in quantum optics: S-matrix and Keldysh techniques for strong-field processes”. In: *Progress in Quantum Electronics* 16.1 (1992), pp. 1–71. ISSN: 00796727. DOI: 10.1016/0079-6727(92)90008-J (cit. on p. 16).

[87] HR Reiss. “Limits on Tunneling Theories of Strong-Field Ionization”. In: *Physical Review Letters* 101.4 (2008), p. 043002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.101.043002 (cit. on p. 16).

[88] S Augst, DD Meyerhofer, D Strickland, and SL Chint. “Laser ionization of noble gases by Coulomb-barrier suppression”. In: *Journal of the Optical Society of America B* 8.4 (1991), p. 858. ISSN: 0740-3224. DOI: 10.1364/JOSAB.8.000858 (cit. on p. 17).

[89] A Thomas, M Tzoufras, A Robinson, R Kingham, C Ridgers, M Sherlock, and A Bell. “A review of Vlasov-Fokker-Planck numerical modeling of inertial confinement fusion plasma”. In: *Journal of Computational Physics* 231.3 (2012), pp. 1051–1079. ISSN: 00219991. DOI: 10.1016/j.jcp.2011.09.028 (cit. on pp. 17, 32).

[90] R More. "Pressure Ionization, Resonances, and the Continuity of Bound and Free States". In: *Advances in Atomic and Molecular Physics*. Vol. 21. C. 1985, pp. 305–356. ISBN: 9780120038213. DOI: 10.1016/S0065-2199(08)60145-1 (cit. on pp. 18, 37, 72, 84).

[91] K Nanbu. "Theory of cumulative small-angle collisions in plasmas". In: *Physical Review E* 55.4 (1997), pp. 4642–4652. ISSN: 1063-651X. DOI: 10.1103/PhysRevE.55.4642 (cit. on p. 18).

[92] K Nanbu and S Yonemura. "Weighted Particles in Coulomb Collision Simulations Based on the Theory of a Cumulative Scattering Angle". In: *Journal of Computational Physics* 145.2 (1998), pp. 639–654. ISSN: 00219991. DOI: 10.1006/jcph.1998.6049 (cit. on p. 18).

[93] F Pérez, L Gremillet, A Decoster, M Drouin, and E Lefebvre. "Improved modeling of relativistic collisions and collisional ionization in particle-in-cell codes". In: *Physics of Plasmas* 19.8 (2012). ISSN: 1070664X. DOI: 10.1063/1.4742167 (cit. on p. 18).

[94] R Mishra, P Leblanc, Y Sentoku, MS Wei, and FN Beg. "Collisional particle-in-cell modeling for energy transport accompanied by atomic processes in dense plasmas". In: *Physics of Plasmas* (2013). ISSN: 1070664X. DOI: 10.1063/1.4812701 (cit. on pp. 18, 37, 124).

[95] H Wiedemann. *Particle Accelerator Physics I*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999. ISBN: 978-3-662-03829-1. DOI: 10.1007/978-3-662-03827-7 (cit. on p. 19).

[96] A Grudiev, S Calatroni, and W Wuensch. "New local field quantity describing the high gradient limit of accelerating structures". In: *Physical Review Special Topics - Accelerators and Beams* 12.10 (2009), p. 102001. ISSN: 1098-4402. DOI: 10.1103/PhysRevSTAB.12.102001 (cit. on p. 19).

[97] TE Cowan, J Fuchs, H Ruhl, A Kemp, P Audebert, M Roth, R Stephens, I Barton, A Blazevic, E Brambrink, J Cobble, J Fernández, JC Gauthier, M Geissel, M Hegelich, J Kaae, S Karsch, GP Le Sage, S Letzring, M Manclossi, S Meyroneinc, A Newkirk, H Pépin, and N Renard-LeGalloudec. "Ultralow Emittance, Multi-MeV Proton Beams from a Laser Virtual-Cathode Plasma Accelerator". In: *Physical Review Letters* 92.20 (2004), p. 204801. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.92.204801 (cit. on p. 19).

[98] S Garcia, D Chatain, and JP Perin. "Continuous production of a thin ribbon of solid hydrogen". In: *Laser and Particle Beams* 32.4 (2014), pp. 569–575. ISSN: 0263-0346. DOI: 10.1017/S0263034614000524 (cit. on p. 19).

[99] S Astbury, S Bedacht, P Brummitt, D Carroll, R Clarke, S Crisp, C Hernandez-Gomez, P Holligan, S Hook, JS Merchan, D Neely, A Ortner, D Rathbone, P Rice, G Schaumann, G Scott, C Spindloe, S Spurdle, A Tebartz, S Tomlinson, F Wagner, M Borghesi, M Roth, and MK Tolley. "In-situ formation of solidified hydrogen thin-membrane targets using a pulse tube cryocooler". In: *Journal of Physics: Conference Series* 713.1 (2016), p. 012006. ISSN: 1742-6588. DOI: 10.1088/1742-6596/713/1/012006 (cit. on p. 19).

[100] JB Kim, S Göde, and SH Glenzer. "Development of a cryogenic hydrogen microjet for high-intensity, high-repetition rate experiments". In: *Review of Scientific Instruments* 87.11 (2016), 11E328. ISSN: 0034-6748. DOI: 10.1063/1.4961089 (cit. on p. 19).

[101] A Tebartz, S Bedacht, M Hesse, S Astbury, R Clarke, A Ortner, G Schaumann, F Wagner, D Neely, and M Roth. "Creation and characterization of free-standing cryogenic targets for laser-driven ion acceleration". In: *Review of Scientific Instruments* 88.9 (2017), p. 093512. ISSN: 0034-6748. DOI: 10.1063/1.5001487 (cit. on p. 19).

[102] A Sharma, Z Tibai, J Hebling, and JA Fülöp. "Quasi-monoenergetic proton acceleration from cryogenic hydrogen microjet by ultrashort ultraintense laser pulses". In: *Physics of Plasmas* 25.3 (2018), p. 033111. ISSN: 1070-664X. DOI: 10.1063/1.5003353. arXiv: 1709.06860 (cit. on p. 19).

[103] AV Gurevich, LV Pariiskaya, and LP Pitaevskii. "Self-Similar Motion of Rarefied Plasma". In: *J. Exptl. and Theoret. Phys. (U.S.S.R.)* 22.2 (1966), pp. 647–654 (cit. on p. 19).

[104] JE Crow, PL Auer, and JE Allen. "The expansion of a plasma into a vacuum". In: *Journal of Plasma Physics* 14.1 (1975), pp. 65–76. ISSN: 0022-3778. DOI: 10.1017/S0022377800025538 (cit. on p. 19).

[105] L Obst-Hübl. "Achieving optimal laser-proton acceleration through multi-parameter interaction control". Doctoral Thesis. TU Dresden, 2019. DOI: 10.5281/zenodo.3252952 (cit. on p. 20).

[106] J Fuchs, P Antici, E D'Humières, E Lefebvre, M Borghesi, E Brambrink, CA Cecchetti, M Kaluza, V Malka, M Manclossi, S Meyroneinc, P Mora, J Schreiber, T Toncian, H Pépin, and P Audebert. "Laser-driven proton scaling laws and new paths towards energy increase". In: *Nature Physics* 2.1 (2006), pp. 48–54. ISSN: 17452481. DOI: 10.1038/nphys199 (cit. on pp. 20, 27, 47).

[107] P Mora. "Thin-foil expansion into a vacuum". In: *Physical Review E* 72.5 (2005), p. 056401. ISSN: 1539-3755. DOI: 10.1103/PhysRevE.72.056401 (cit. on p. 20).

[108] MA True. "Fast ion production by suprathermal electrons in laser fusion plasmas". In: *Physics of Fluids* 24.10 (1981), p. 1885. ISSN: 00319171. DOI: 10.1063/1.863270 (cit. on p. 20).

[109] DS Dorozhkina and VE Semenov. "Exact Solution of Vlasov Equations for Quasineutral Expansion of Plasma Bunch into Vacuum". In: *Physical Review Letters* 81.13 (1998), pp. 2691–2694. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.81.2691 (cit. on p. 20).

[110] AV BAITIN and KM KUZANYAN. "A self-similar solution for expansion into a vacuum of a collisionless plasma bunch". In: *Journal of Plasma Physics* 59.1 (1998), pp. 83–90. ISSN: 0022-3778. DOI: 10.1017/S0022377897005916 (cit. on p. 20).

[111] VF Kovalev, VY Bychenkov, and VT Tikhonchuk. "Particle dynamics during adiabatic expansion of a plasma bunch". In: *Journal of Experimental and Theoretical Physics* 95.2 (2002), pp. 226–241. ISSN: 1063-7761. DOI: 10.1134/1.1506430 (cit. on p. 20).

[112] VF Kovalev and VY Bychenkov. "Analytic Solutions to the Vlasov Equations for Expanding Plasmas". In: *Physical Review Letters* 90.18 (2003), p. 4. ISSN: 10797114. DOI: 10.1103/PhysRevLett.90.185004 (cit. on p. 20).

[113] G Manfredi, S Mola, and MR Feix. "Rescaling methods and plasma expansions into vacuum". In: *Physics of Fluids B: Plasma Physics* 5.2 (1993), pp. 388–401. ISSN: 0899-8221. DOI: 10.1063/1.860524 (cit. on p. 20).

[114] P Mora. "Laser driven ion acceleration". In: *AIP Conference Proceedings*. Vol. 920. July. AIP, 2007, pp. 98–117. ISBN: 9780735404304. DOI: 10.1063/1.2756774 (cit. on p. 20).

[115] P Mora and T Grismayer. "Rarefaction Acceleration and Kinetic Effects in Thin-Foil Expansion into a Vacuum". In: *Physical Review Letters* 102.14 (2009), p. 145001. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.102.145001 (cit. on p. 20).

[116] J Schreiber, F Bell, and Z Najmudin. "Optimization of relativistic laser-ion acceleration". In: *High Power Laser Science and Engineering* 2 (2014), e41. ISSN: 2095-4719. DOI: 10.1017/hpl.2014.46 (cit. on pp. 21, 99).

[117] A Macchi, M Borghesi, and M Passoni. "Ion acceleration by superintense laser-plasma interaction". In: *Reviews of Modern Physics* 85.2 (2013), pp. 751–793. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.85.751. arXiv: 1302.1775 (cit. on pp. 22, 23, 27).

[118] AA Andreev, S Steinke, T Sokollik, M Schnürer, ST Avetsiyan, KY Platonov, and PV Nickles. "Optimal ion acceleration from ultrathin foils irradiated by a profiled laser pulse of relativistic intensity". In: *Physics of Plasmas* 16.1 (2009), p. 013103. ISSN: 1070-664X. DOI: 10.1063/1.3054528 (cit. on p. 22).

[119] AJ Mackinnon, M Borghesi, S Hatchett, MH Key, PK Patel, H Campbell, A Schiavi, R Snavely, SC Wilks, and O Willi. "Effect of Plasma Scale Length on Multi-MeV Proton Production by Intense Laser Pulses". In: *Physical Review Letters* 86.9 (2001), pp. 1769–1772. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.86.1769 (cit. on p. 22).

[120] T Grismayer and P Mora. "Influence of a finite initial ion density gradient on plasma expansion into a vacuum". In: *Physics of Plasmas* 13.3 (2006). ISSN: 1070664X. DOI: 10.1063/1.2178653 (cit. on pp. 22, 26, 105).

[121] T Esirkepov, M Yamagiwa, and T Tajima. "Laser Ion-Acceleration Scaling Laws Seen in Multiparametric Particle-in-Cell Simulations". In: *Physical Review Letters* 96.10 (2006), p. 105001. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.96.105001 (cit. on p. 22).

[122] R Mishra, F Fiuza, and S Glenzer. "Enhanced ion acceleration in transition from opaque to transparent plasmas". In: *New Journal of Physics* 20.4 (2018). ISSN: 13672630. DOI: 10.1088/1367-2630/aab8db (cit. on p. 22).

[123] A Henig, D Kiefer, K Markey, DC Gautier, KA Flippo, S Letzring, RP Johnson, T Shimada, L Yin, BJ Albright, KJ Bowers, JC Fernández, SG Rykovanov, HC Wu, M Zepf, D Jung, VK Liechtenstein, J Schreiber, D Habs, and BM Hegelich. "Enhanced Laser-Driven Ion Acceleration in the Relativistic Transparency Regime". In: *Physical Review Letters* 103.4 (2009), p. 045002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.103.045002 (cit. on p. 22).

[124] PL Poole, L Obst, GE Cochran, J Metzkes, HP Schlenvoigt, I Prencipe, T Kluge, TE Cowan, U Schramm, DW Schumacher, and K Zeil. "Laser-driven ion acceleration via target normal sheath acceleration in the relativistic transparency regime". In: *New Journal of Physics* 20.1 (2018), p. 013019. ISSN: 1367-2630. DOI: 10.1088/1367-2630/aa9d47 (cit. on pp. 22, 79, 80).

[125] SC Wilks, WL Kruer, M Tabak, and AB Langdon. "Absorption of ultra-intense laser pulses". In: *Physical Review Letters* 69.9 (1992), pp. 1383–1386. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.69.1383 (cit. on p. 23).

[126] R Sauerbrey. "Acceleration in femtosecond laser-produced plasmas". In: *Physics of Plasmas* 3.12 (1996), pp. 4712–4716. ISSN: 1070-664X. DOI: 10.1063/1.872038 (cit. on p. 23).

[127] J Hornung, Y Zobus, S Roeder, A Kleinschmidt, D Bertini, M Zepf, and V Bagnoud. "Time-resolved study of holeboring in realistic experimental conditions". In: *Nature Communications* 12.1 (2021), p. 6999. ISSN: 2041-1723. DOI: 10.1038/s41467-021-27363-9 (cit. on p. 23).

[128] A Sgattoni, S Sinigardi, L Fedeli, F Pegoraro, and A Macchi. "Laser-driven Rayleigh-Taylor instability: Plasmonic effects and three-dimensional structures". In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 91.1 (2015), pp. 1–6. ISSN: 15502376. DOI: 10.1103/PhysRevE.91.013106 (cit. on p. 23).

[129] VK Tripathi, CS Liu, X Shao, B Eliasson, and RZ Sagdeev. "Laser acceleration of monoenergetic protons in a self-organized double layer from thin foil". In: *Plasma Physics and Controlled Fusion* 51.2 (2009), p. 024014. ISSN: 0741-3335. DOI: 10.1088/0741-3335/51/2/024014 (cit. on p. 24).

[130] APL Robinson, P Gibbon, M Zepf, S Kar, RG Evans, and C Bellei. "Relativistically correct hole-boring and ion acceleration by circularly polarized laser pulses". In: *Plasma Physics and Controlled Fusion* 51.2 (2009). ISSN: 07413335. DOI: 10.1088/0741-3335/51/2/024004 (cit. on p. 24).

[131] DA Tidman and NA Krall. *Shock waves in collisionless plasmas*. 1971. Chap. 6 (cit. on p. 24).

[132] K Matsukado, T Esirkepov, K Kinoshita, H Daido, T Utsumi, Z Li, A Fukumi, Y Hayashi, S Orimo, M Nishiuchi, SV Bulanov, T Tajima, A Noda, Y Iwashita, T Shirai, T Takeuchi, S Nakamura, A Yamazaki, M Ikegami, T Mihara, A Morita, M Uesaka, K Yoshii, T Watanabe, T Hosokai, A Zhidkov, A Ogata, Y Wada, and T Kubota. "Energetic Protons from a Few-Micron Metallic Foil Evaporated by an Intense Laser Pulse". In: *Physical Review Letters* 91.21 (2003), p. 215001. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.91.215001 (cit. on pp. 24, 26).

[133] L Yin, BJ Albright, BM Hegelich, and JC Fernández. "GeV laser ion acceleration from ultrathin targets: The laser break-out afterburner". In: *Laser and Particle Beams* 24.2 (2006), pp. 291–298. ISSN: 02630346. DOI: 10.1017/S0263034606060459 (cit. on p. 25).

[134] BJ Albright, L Yin, KJ Bowers, BM Hegelich, KA Flippo, TJT Kwan, and JC Fernández. "Relativistic Buneman instability in the laser breakout afterburner". In: *Physics of Plasmas* 14.9 (2007), p. 094502. ISSN: 1070-664X. DOI: 10.1063/1.2768933 (cit. on p. 25).

[135] L Yin, BJ Albright, BM Hegelich, KJ Bowers, KA Flippo, TJT Kwan, and JC Fernández. "Monoenergetic and GeV ion acceleration from the laser breakout afterburner using ultrathin targets". In: *Physics of Plasmas* 14.5 (2007), p. 056706. ISSN: 1070-664X. DOI: 10.1063/1.2436857 (cit. on p. 25).

[136] L Yin, BJ Albright, KJ Bowers, D Jung, JC Fernández, and BM Hegelich. "Three-dimensional dynamics of breakout afterburner ion acceleration using high-contrast short-pulse laser and nanoscale targets". In: *Physical Review Letters* 107.4 (2011), pp. 1–4. ISSN: 00319007. DOI: 10.1103/PhysRevLett.107.045003 (cit. on p. 25).

[137] O Buneman. "Dissipation of currents in ionized media". In: *Physical Review* 115.3 (1959), pp. 503–517. ISSN: 0031899X. DOI: 10.1103/PhysRev.115.503 (cit. on p. 25).

[138] L Obst, J Metzkes-Ng, S Bock, GE Cochran, TE Cowan, T Oksenhendler, PL Poole, I Prencipe, M Rehwald, C Rödel, HP Schlenvoigt, U Schramm, DW Schumacher, T Ziegler, and K Zeil. "On-shot characterization of single plasma mirror temporal contrast improvement". In: *Plasma Physics and Controlled Fusion* 60.5 (2018), p. 054007. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aab3bb (cit. on pp. 25, 76).

[139] KB Wharton, CD Boley, AM Komashko, AM Rubenchik, J Zweiback, J Crane, G Hays, TE Cowan, and T Ditmire. "Effects of nonionizing prepulses in high-intensity laser-solid interactions". In: *Physical Review E* 64.2 (2001), p. 025401. ISSN: 1063-651X. DOI: 10.1103/PhysRevE.64.025401 (cit. on p. 26).

[140] A Yogo, H Daido, SV Bulanov, K Nemoto, Y Oishi, T Nayuki, T Fujii, K Ogura, S Orimo, A Sagisaka, JL Ma, TZ Esirkepov, M Mori, M Nishiuchi, AS Pirozhkov, S Nakamura, A Noda, H Nagatomo, T Kimura, and T Tajima. "Laser ion acceleration via control of the near-critical density target". In: *Physical Review E* 77.1 (2008), p. 016401. ISSN: 1539-3755. DOI: 10.1103/PhysRevE.77.016401 (cit. on p. 26).

[141] M Kaluza, J Schreiber, MIK Santala, GD Tsakiris, K Eidmann, J Meyer-ter-Vehn, and KJ Witte. "Influence of the Laser Prepulse on Proton Acceleration in Thin-Foil Experiments". In: *Physical Review Letters* 93.4 (2004), p. 045003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.93.045003 (cit. on pp. 26, 84, 145).

[142] TZ Esirkepov, JK Koga, A Sunahara, T Morita, M Nishikino, K Kageyama, H Nagatomo, K Nishihara, A Sagisaka, H Kotaki, T Nakamura, Y Fukuda, H Okada, AS Pirozhkov, A Yogo, M Nishiuchi, H Kiriyama, K Kondo, M Kando, and SV Bulanov. "Prepulse and amplified spontaneous emission effects on the interaction of a petawatt class laser with thin solid targets". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 745 (2014), pp. 150–163. ISSN: 01689002. DOI: 10.1016/j.nima.2014.01.056. arXiv: 1310.0568 (cit. on p. 26).

[143] A Flacco, T Ceccotti, H George, P Monot, P Martin, F Réau, O Tcherbakoff, P D'Oliveira, F Sylla, M Veltcheva, F Burgy, A Tafzi, V Malka, and D Batani. "Comparative study of laser ion acceleration with different contrast enhancement techniques". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 620.1 (2010), pp. 18–22. ISSN: 01689002. DOI: 10.1016/j.nima.2010.01.053 (cit. on pp. 26–28).

[144] P Hadjisolomou, IP Tsygvintsev, P Sasorov, V Gasilov, G Korn, and SV Bulanov. "Preplasma effects on laser ion generation from thin foil targets". In: *Physics of Plasmas* 27.1 (2020), p. 013107. ISSN: 1070-664X. DOI: 10.1063/1.5124457 (cit. on p. 26).

[145] R Nuter, L Gremillet, P Combis, M Drouin, E Lefebvre, A Flacco, and V Malka. "Influence of a preplasma on electron heating and proton acceleration in ultraintense laser-foil interaction". In: *Journal of Applied Physics* 104.10 (2008), p. 103307. ISSN: 0021-8979. DOI: 10.1063/1.3028274 (cit. on p. 27).

[146] M Carrié, E Lefebvre, A Flacco, and V Malka. "Influence of subpicosecond laser pulse duration on proton acceleration". In: *Physics of Plasmas* 16.5 (2009), p. 053105. ISSN: 1070-664X. DOI: 10.1063/1.3138742 (cit. on p. 27).

[147] D Neely, P Foster, A Robinson, F Lindau, O Lundh, A Persson, CG Wahlström, and P McKenna. "Enhanced proton beams from ultrathin targets driven by high contrast laser pulses". In: *Applied Physics Letters* 89.2 (2006), pp. 0–4. ISSN: 00036951. DOI: 10.1063/1.2220011 (cit. on p. 27).

[148] T Ceccotti, A Lévy, H Popescu, F Réau, P D'Oliveira, P Monot, JP Geindre, E Lefebvre, and P Martin. "Proton acceleration with high-intensity ultrahigh-contrast laser pulses". In: *Physical Review Letters* 99.18 (2007), pp. 1–4. ISSN: 00319007. DOI: 10.1103/PhysRevLett.99.185002 (cit. on p. 27).

[149] H Kiriyama, M Mori, Y Nakai, T Shimomura, H Sasao, M Tanoue, S Kanazawa, D Wakai, F Sasao, H Okada, I Daito, M Suzuki, S Kondo, K Kondo, A Sugiyama, PR Bolton, A Yokoyama, H Daido, S Kawanishi, T Kimura, and T Tajima. "High temporal and spatial quality petawatt-class Ti:sapphire chirped-pulse amplification laser system". In: *Optics Letters* 35.10 (2010), p. 1497. ISSN: 0146-9592. DOI: 10.1364/OL.35.001497 (cit. on p. 27).

[150] H Kiriyama, AS Pirozhkov, M Nishiuchi, Y Fukuda, K Ogura, A Sagisaka, Y Miyasaka, M Mori, H Sakaki, NP Dover, K Kondo, JK Koga, TZ Esirkepov, M Kando, and K Kondo. "High-contrast high-intensity repetitive petawatt laser". In: *Optics Letters* 43.11 (2018), pp. 12–15 (cit. on p. 27).

[151] J Fuchs, Y Sentoku, E D'Humières, TE Cowan, J Cobble, P Audebert, A Kemp, A Nikroo, P Antici, E Brambrink, A Blazevic, EM Campbell, JC Fernández, JC Gauthier, M Geissel, M Hegelich, S Karsch, H Popescu, N Renard-LeGalloudec, M Roth, J Schreiber, R Stephens, and H Pépin. "Comparative spectra and efficiencies of ions laser-accelerated forward from the front and rear surfaces of thin solid foils". In: *Physics of Plasmas* 14.5 (2007), p. 053105. ISSN: 1070-664X. DOI: 10.1063/1.2720373 (cit. on pp. 27, 28).

[152] K Zeil, J Metzkes, T Kluge, M Bussmann, T Cowan, S Kraft, R Sauerbrey, and U Schramm. "Direct observation of prompt pre-thermal laser ion sheath acceleration". In: *Nature Communications* 3.1 (2012), p. 874. ISSN: 2041-1723. DOI: 10.1038/ncomms1883 (cit. on p. 28).

[153] NV Didenko, AV Konyashchenko, AP Lutsenko, and SY Tenyakov. "Contrast degradation in a chirped-pulse amplifier due to generation of prepulses by postpulses". In: *Optics Express* 16.5 (2008), p. 3178. ISSN: 1094-4087. DOI: 10.1364/OE.16.003178 (cit. on p. 28).

[154] A Stockem Novo, MC Kaluza, RA Fonseca, and LO Silva. "Optimizing laser-driven proton acceleration from overdense targets". In: *Scientific Reports* 6.1 (2016), p. 29402. ISSN: 2045-2322. DOI: 10.1038/srep29402 (cit. on p. 28).

[155] 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, T. Rösch, F Wagner, B Zielbauer, and J Schreiber. "Isolated proton bunch acceleration by a petawatt laser pulse". In: *Nature Communications* 9.1 (2018), p. 423. ISSN: 2041-1723. DOI: 10.1038/s41467-017-02663-1 (cit. on pp. 28, 50, 132).

[156] P McKenna, D Carroll, O Lundh, F Nürnberg, K Markey, S Bandyopadhyay, D Batani, R Evans, R Jafer, S Kar, D Neely, D Pepler, M Quinn, R Redaelli, M Roth, CG Wahlström, X Yuan, and M Zepf. "Effects of front surface plasma expansion on proton acceleration in ultraintense laser irradiation of foil targets". In: *Laser and Particle Beams* 26.4 (2008), pp. 591–596. ISSN: 0263-0346. DOI: 10.1017/S0263034608000657 (cit. on p. 28).

[157] Y Glinec, G Genoud, O Lundh, A Persson, and CG Wahlström. "Evolution of energy spectrum from laser-accelerated protons with a 100 fs intense prepulse". In: *Applied Physics B: Lasers and Optics* 93.2-3 (2008), pp. 317–321. ISSN: 09462171. DOI: 10.1007/s00340-008-3172-6 (cit. on p. 28).

[158] APL Robinson, D Neely, P McKenna, and RG Evans. "Spectral control in proton acceleration with multiple laser pulses". In: *Plasma Physics and Controlled Fusion* 49.4 (2007), pp. 373–384. ISSN: 0741-3335. DOI: 10.1088/0741-3335/49/4/002 (cit. on p. 28).

[159] CM Brenner, APL Robinson, K Markey, RHH Scott, RJ Gray, M Rosinski, O Deppert, J Badziak, D Batani, JR Davies, SM Hassan, KL Lancaster, K Li, IO Musgrave, PA Norreys, J Pasley, M Roth, HP Schlenvoigt, C Spindloe, M Tatarakis, T Winstone, J Wolowski, D Wyatt, P McKenna, and D Neely. "High energy conversion efficiency in laser-proton acceleration by controlling laser-energy deposition onto thin foil targets". In: *Applied Physics Letters* 104.8 (2014), p. 081123. ISSN: 0003-6951. DOI: 10.1063/1.4865812 (cit. on p. 28).

[160] K Markey, P McKenna, CM Brenner, DC Carroll, MM Günther, K Harres, S Kar, K Lancaster, F Nürnberg, MN Quinn, APL Robinson, M Roth, M Zepf, and D Neely. "Spectral Enhancement in the Double Pulse Regime of Laser Proton Acceleration". In: *Physical Review Letters* 105.19 (2010), p. 195008. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.105.195008 (cit. on p. 28).

[161] GG Scott, JS Green, V Bagnoud, C Brabetz, CM Brenner, DC Carroll, DA MacLellan, APL Robinson, M Roth, C Spindloe, F Wagner, B Zielbauer, P McKenna, and D Neely. "Multi-pulse enhanced laser ion acceleration using plasma half cavity targets". In: *Applied Physics Letters* 101.2 (2012), p. 024101. ISSN: 0003-6951. DOI: 10.1063/1.4734397 (cit. on p. 28).

[162] JR Smith, C Orban, GK Ngirmang, JT Morrison, KM George, EA Chowdhury, and WM Roquemore. "Particle-in-cell simulations of density peak formation and ion heating from short pulse laser-driven ponderomotive steepening". In: *Physics of Plasmas* 26.12 (2019), p. 123103. ISSN: 1070-664X. DOI: 10.1063/1.5108811 (cit. on p. 29).

[163] M Tayyab, S Bagchi, JA Chakera, RA Khan, and PA Naik. "Effect of temporally modified ultra-short laser pulses on ion acceleration from thin foil targets". In: *Physics of Plasmas* 25.8 (2018), p. 083113. ISSN: 1070-664X. DOI: 10.1063/1.5037260 (cit. on p. 29).

[164] M Schollmeier, AB Sefkow, M Geissel, AV Arefiev, KA Flippo, SA Gaillard, RP Johnson, MW Kimmel, DT Offermann, PK Rambo, J Schwarz, and T Shimada. "Laser-to-hot-electron conversion limitations in relativistic laser matter interactions due to multi-picosecond dynamics". In: *Physics of Plasmas* 22.4 (2015), p. 043116. ISSN: 1070-664X. DOI: 10.1063/1.4918332 (cit. on p. 31).

[165] L Buikstra. *Relativistic fluid simulation of plasma-vacuum expansion with LTE closure*. 2013. URL: http://essay.utwente.nl/63283/ (cit. on p. 32).

[166] J Nikl, I Göthel, M Kuchařík, S Weber, and M Bussmann. "Implicit reduced Vlasov-Fokker-Planck-Maxwell model based on high-order mixed elements". In: *Journal of Computational Physics* 434 (2021), p. 110214. ISSN: 00219991. DOI: 10.1016/j.jcp.2021.110214 (cit. on p. 32).

[167] RW Hockney and JW Eastwood. *Computer simulation using particles*. CRC Press, 1988, p. 540. ISBN: 9780852743928 (cit. on pp. 34, 37).

[168] C Birdsall and A Langdon. *Plasma Physics via Computer Simulation*. Vol. 42. 1. CRC Press, 1991, p. 504. ISBN: 9781315275048. DOI: 10.1201/9781315275048 (cit. on pp. 34, 50).

[169] KS Yee. "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media". In: *IEEE Transactions on Antennas and Propagation* 14.3 (1966), pp. 302–307. ISSN: 0018-926X. DOI: 10.1109/TAP.1966.1138693 (cit. on p. 34).

[170] A Taflove and M Brodwin. "Numerical Solution of Steady-State Electromagnetic Scattering Problems Using the Time-Dependent Maxwell's Equations". In: *IEEE Transactions on Microwave Theory and Techniques* 23.8 (1975), pp. 623–630. ISSN: 0018-9480. DOI: 10.1109/TMTT.1975.1128640 (cit. on p. 35).

[171] D Tskhakaya, K Matyash, R Schneider, and F Taccogna. "The particle-in-cell method". In: *Contributions to Plasma Physics* 47.8-9 (2007), pp. 563–594. ISSN: 08631042. DOI: 10.1002/ctpp.200710072 (cit. on p. 35).

[172] H Fehske, R Schneider, and A Weiße. *Computational Many-Particle Physics*. Ed. by H Fehske, R Schneider, and A Weiße. Vol. 739. Lecture Notes in Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, p. 409. ISBN: 978-3-540-74685-0. DOI: 10.1007/978-3-540-74686-7 (cit. on p. 35).

[173] TH Maiman. "Stimulated Optical Radiation in Ruby". In: *Nature 1960 187:4736* 187.4736 (1960), pp. 493–494. ISSN: 1476-4687. DOI: 10.1038/187493a0 (cit. on p. 36).

[174] SFJ Larochelle, A Talebpour, and SL Chin. "Coulomb effect in multiphoton ionization of rare-gas atoms". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 31.6 (1998), pp. 1215–1224. ISSN: 0953-4075. DOI: 10.1088/0953-4075/31/6/009 (cit. on p. 36).

[175] S Hankin, D Villeneuve, P Corkum, and D Rayner. “Intense-field laser ionization rates in atoms and molecules”. In: *Physical Review A* 64.1 (2001), pp. 1–12. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.64.013405 (cit. on p. 36).

[176] E Gubbini, U Eichmann, M Kalashnikov, and W Sandner. “Core relaxation in atomic ultrastrong laser field ionization”. In: *Physical Review Letters* 94.5 (2005), pp. 11–14. ISSN: 00319007. DOI: 10.1103/PhysRevLett.94.053602 (cit. on p. 36).

[177] P Wustelt, M Möller, T Rathje, AM Sayler, T Stöhlker, and GG Paulus. “Momentum-resolved study of the saturation intensity in multiple ionization”. In: *Physical Review A* 91.3 (2015), p. 031401. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.91.031401 (cit. on p. 36).

[178] MS Schöffler, X Xie, P Wustelt, M Möller, S Roither, D Kartashov, AM Sayler, A Baltuska, GG Paulus, and M Kitzler. “Laser-subcycle control of sequential double-ionization dynamics of helium”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 93.6 (2016), pp. 1–10. ISSN: 10941622. DOI: 10.1103/PhysRevA.93.063421 (cit. on p. 36).

[179] CM Maharjan, AS Alnaser, XM Tong, B Ulrich, P Ranitovic, S Ghimire, Z Chang, IV Litvinyuk, and CL Cocke. “Momentum imaging of doubly charged ions of Ne and Ar in the sequential ionization region”. In: *Physical Review A* 72.4 (2005), p. 041403. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.72.041403 (cit. on p. 36).

[180] P Eckle, M Smolarski, P Schlup, J Biegert, A Staudte, M Schöffler, HG Muller, R Dörner, and U Keller. “Attosecond angular streaking”. In: *Nature Physics* 4.7 (2008), pp. 565–570. ISSN: 1745-2473. DOI: 10.1038/nphys982 (cit. on p. 36).

[181] A Fleischer, HJ Wörner, L Arissian, LR Liu, M Meckel, A Rippert, R Dörner, DM Villeneuve, PB Corkum, and A Staudte. “Probing Angular Correlations in Sequential Double Ionization”. In: *Physical Review Letters* 107.11 (2011), p. 113003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.107.113003 (cit. on p. 36).

[182] Y Sentoku, AJ Kemp, R Presura, MS Bakeman, and TE Cowan. “Isochoric heating in heterogeneous solid targets with ultrashort laser pulses”. In: *Physics of Plasmas* 14.12 (2007). ISSN: 1070664X. DOI: 10.1063/1.2816439 (cit. on pp. 36, 37, 124).

[183] D Salzmann. *Atomic Physics in Hot Plasmas*. 1998, p. 109. ISBN: 0195109309 (cit. on pp. 37, 137).

[184] HK Chung, M Chen, W Morgan, Y Ralchenko, and R Lee. “FLYCHK: Generalized population kinetics and spectral model for rapid spectroscopic analysis for all elements”. In: *High Energy Density Physics* 1.1 (2005), pp. 3–12. ISSN: 15741818. DOI: 10.1016/j.hedp.2005.07.001 (cit. on pp. 37, 72, 83, 84, 121).

[185] HK Chung, M Chen, and R Lee. “Extension of atomic configuration sets of the Non-LTE model in the application to the Kα diagnostics of hot dense matter”. In: *High Energy Density Physics* 3.1-2 (2007), pp. 57–64. ISSN: 15741818. DOI: 10.1016/j.hedp.2007.02.001 (cit. on pp. 37, 72, 121).

[186] IH Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 2002. ISBN: 9780521803892. DOI: 10.1017/CBO9780511613630 (cit. on p. 37).

[187] AJ Kemp, REW Pfund, and J Meyer-ter-Vehn. “Modeling ultrafast laser-driven ionization dynamics with Monte Carlo collisional particle-in-cell simulations”. In: *Physics of Plasmas* 11.12 (2004), pp. 5648–5657. ISSN: 1070-664X. DOI: 10.1063/1.1814367 (cit. on p. 37).

[188] BE Marré. “Coupling of atomic states to particle in cell simulations”. English. Master thesis. Technische Universität Dresden, 2020, p. 78. DOI: 10.14278/rodare.1203 (cit. on p. 37).

[189] BA Shadwick and CB Schroeder. “Physical fidelity in particle-in-cell modeling of small debye-length plasmas”. In: *AIP Conference Proceedings* 1086 (2009), pp. 321–327. ISSN: 0094243X. DOI: 10.1063/1.3080926 (cit. on p. 37).

[190] J Vyskočil, O Klimo, and S Weber. “Simulations of bremsstrahlung emission in ultra-intense laser interactions with foil targets”. In: *Plasma Physics and Controlled Fusion* 60.5 (2018), p. 054013. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aab4c3 (cit. on pp. 38, 108, 114).

[191] TS Daykin, H Sawada, Y Sentoku, FN Beg, H Chen, HS McLean, AJ Link, PK Patel, and Y Ping. “Characterization of fast electron divergence and energy spectrum from modeling of angularly resolved bremsstrahlung measurements”. In: *Physics of Plasmas* 25.12 (2018), p. 123103. ISSN: 1070-664X. DOI: 10.1063/1.5055257 (cit. on p. 38).

[192] S Singh, CD Armstrong, N Kang, L Ren, H Liu, N Hua, DR Rusby, O Klimo, R Versaci, Y Zhang, M Sun, B Zhu, A Lei, X Ouyang, L Lancia, A Laso Garcia, A Wagner, TE Cowan, J Zhu, T Schlegel, S Weber, P McKenna, D Neely, V Tikhonchuk, and D Kumar. “Bremsstrahlung emission and plasma characterization driven by moderately relativistic laser-plasma interactions”. In: *Plasma Physics and Controlled Fusion* 63.3 (2021), p. 035004. ISSN: 0741-3335. DOI: 10.1088/1361-6587/abcf7e. arXiv: 2009.12307 (cit. on pp. 38, 118).

[193] JD Jackson. *Classical electrodynamics*. 3rd ed. New York, NY: Wiley, 1999. ISBN: 9780471309321 (cit. on p. 38).

[194] H Burau. “Entwicklung und Überprüfung eines Photonenmodells für die Abstrahlung durch hochenergetische Elektronen”. Diploma Thesis. Technische Universität Dresden, 2016, p. 105. DOI: 10.5281/zenodo.192116 (cit. on pp. 38, 108).

[195] WH Furry. “On bound states and scattering in positron theory”. In: *Physical Review* 81.1 (1951), pp. 115–124. ISSN: 0031899X. DOI: 10.1103/PhysRev.81.115 (cit. on p. 38).

[196] A Gonoskov, S Bastrakov, E Efimenko, A Ilderton, M Marklund, I Meyerov, A Muraviev, A Sergeev, I Surmin, and E Wallin. “Extended particle-in-cell schemes for physics in ultrastrong laser fields: Review and developments”. In: *Physical Review E* 92.2 (2015), p. 023305. ISSN: 1539-3755. DOI: 10.1103/PhysRevE.92.023305 (cit. on p. 38).

[197] G Wentzel. “Zwei Bemerkungen über die Zerstreuung korpuskularer Strahlen als Beugungserscheinung”. In: *Zeitschrift für Physik* 40.8 (1926), pp. 590–593. ISSN: 1434-6001. DOI: 10.1007/BF01390457 (cit. on p. 38).

[198] J Babaei, LA Gizzi, P Londrillo, S Mirzanejad, T Rovelli, S Sinigardi, and G Turchetti. “Rise time of proton cut-off energy in 2D and 3D PIC simulations”. In: *Physics of Plasmas* 24.4 (2017). ISSN: 10897674. DOI: 10.1063/1.4979901. arXiv: 1611.08381 (cit. on pp. 39, 85).

[199] C Benedetti, A Sgattoni, G Turchetti, and P Londrillo. “ALaDyn: A High-Accuracy PIC Code for the Maxwell–Vlasov Equations”. In: *IEEE Transactions on Plasma Science* 36.4 (2008), pp. 1790–1798. ISSN: 0093-3813. DOI: 10.1109/TPS.2008.927143 (cit. on p. 39).

[200] *PRACE 15th Call for Proposals for Project Access - PRACE*. URL: https://prace-ri.eu/call/call15-for-proposals-for-project-access/ (visited on July 6, 2020) (cit. on p. 41).

[201] insideHPC. *PRACE Awards 1.7 Thousand Million Core Hours for Research Projects in Europe*. 2017. URL: https://insidehpc.com/2017/10/prace-awards-1-7-million-core-hours-research-projects-europe/ (cit. on p. 41).

[202] *CSCS Swiss National Supercomputing Centre*. URL: https://www.cscs.ch/ (visited on July 6, 2020) (cit. on p. 41).

[203] H Meuer, E Strohmaier, and J Dongarra. *TOP500 Supercomputer Sites*. 2019. URL: https://www.top500.org/list/2019/11/ (cit. on pp. 41, 132, 140).

[204] *PIConGPU*. 2020. URL: https://github.com/ComputationalRadiationPhysics/picongpu/ (cit. on p. 41).

[205] A Huebl. "Injection Control for Electrons in Laser-Driven Plasma Wakes on the Femtosecond Time Scale". Diploma thesis. Technische Universität Dresden, 2014, p. 117. DOI: 10.5281/zenodo.15924 (cit. on pp. 42, 44).

[206] *PIConGPU Plugin Documentation*. 2021. URL: https://picongpu.readthedocs.io/en/0.5.0/usage/plugins.html (cit. on pp. 42, 101).

[207] TL Veldhuizen and ME Jernigan. "Will C++ be faster than Fortran?" In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 1343. Springer Verlag, 1997, pp. 50–56. ISBN: 354063827X. DOI: 10.1007/3-540-63827-X_43 (cit. on p. 42).

[208] A Alexandrescu. *Modern C++ Design: Generic Programming and Design Pattern Applied*. 2001 (cit. on p. 42).

[209] P Gottschling. *Discovering Modern C++: An Intensive Course for Scientists, Engineers, and Programmers*. 1st. Addison-Wesley Professional, 2015. ISBN: 0134383583 (cit. on p. 42).

[210] 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: 2017. arXiv: 1706.10086 (cit. on p. 42).

[211] E Zenker, B Worpitz, R Widera, A Huebl, G Juckeland, A Knüpfer, WE Nagel, and M Bussmann. "Alpaka - An Abstraction Library for Parallel Kernel Acceleration". In: IEEE Computer Society, 2016. arXiv: 1602.08477 (cit. on pp. 42, 49).

[212] B Worpitz. *Investigating performance portability of a highly scalable particle-in-cell simulation code on various multi-core architectures*. Master Thesis. 2015. DOI: 10.5281/zenodo.49768 (cit. on p. 42).

[213] *ONRL Frontier - Landing Page*. 2021. URL: https://www.olcf.ornl.gov/frontier/ (cit. on p. 42).

[214] CSCS. *"Piz Daint", One of the Most Powerful Supercomputers in the World*. 2017. URL: https://www.cscs.ch/publications/news/2017/piz-daint-one-of-the-most-powerful-supercomputers-in-the-world/ (cit. on pp. 43, 54).

[215] *Frontier Center for Accelerated Application Readiness*. 2021. URL: https://www.olcf.ornl.gov/caar/frontier-caar/ (cit. on p. 43).

[216] *HZDR Press Release: Aufbruch in die Exaflops-Welt*. 2019. URL: https://www.hzdr.de/db/Cms?pOid=59532 (cit. on p. 43).

[217] *HZDR Press Release: Digital frontier runners*. 2021. URL: https://www.hzdr.de/db/Cms?pOid=59532 (cit. on p. 43).

[218] *PIConGPU Community Map*. 2021. URL: https://computationalradiationphysics.github.io/picongpu-communitymap/ (cit. on p. 43).

[219] *HZDR Press Release: HZDR und Weizmann-Institut gründen Laser-Labor in Israel*. 2017. URL: https://www.hzdr.de/db/Cms?pOid=50425 (cit. on p. 43).

[220] C Eckert. *Enhancements of the massively parallel memory allocator ScatterAlloc and its adaption to the general interface mallocMC*. Tech. rep. 2014. DOI: 10.5281/zenodo.34461 (cit. on p. 45).

[221] M Steinberger, M Kenzel, B Kainz, and D Schmalstieg. "ScatterAlloc: Massively parallel dynamic memory allocation for the GPU". In: *2012 Innovative Parallel Computing, InPar 2012* (2012). DOI: 10.1109/InPar.2012.6339604 (cit. on p. 45).

[222] K Germaschewski, W Fox, S Abbott, N Ahmadi, K Maynard, L Wang, H Ruhl, and A Bhattacharjee. "The Plasma Simulation Code: A modern particle-in-cell code with patch-based load-balancing". In: *Journal of Computational Physics* 318 (2016), pp. 305–326. DOI: 10.1016/j.jcp.2016.05.013 (cit. on p. 46).

[223] M Sauget and G Latu. "Dynamic Load Balancing for PIC code using Eulerian/Lagrangian partitioning". In: December (2017), pp. 1–15. arXiv: 1706.08362 (cit. on p. 46).

[224] J Derouillat, A Beck, F Pérez, T Vinci, M Chiaramello, A Grassi, M Flé, G Bouchard, I Plotnikov, N Aunai, J Dargent, C Riconda, and M Grech. "SMILEI: a collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation". In: *Computer Physics Communications* 222 (2017), pp. 351–373. ISSN: 00104655. DOI: 10.1016/j.cpc.2017.09.024. arXiv: 1702.05128 (cit. on pp. 46, 56, 147, 148).

[225] P Ortwein, T Binder, S Copplestone, A Mirza, P Nizenkov, M Pfeiffer, CD Munz, and S Fasoulas. "A Load Balance Strategy for Hybrid Particle-Mesh Methods". In: (2018). arXiv: 1811.05152 (cit. on p. 46).

[226] JA Roden and SD Gedney. "Convolution PML (CPML): An efficient FDTD implementation of the CFS-PML for arbitrary media". In: *Microwave and Optical Technology Letters* 27.5 (2000), pp. 334–339. ISSN: 0895-2477. DOI: 10.1002/1098-2760(20001205)27:5<334::AID-MOP14>3.0.CO;2-A (cit. on p. 47).

[227] A Taflove and SC Hagness. *Computational electrodynamics: the finite-difference time-domain method*. Artech house, 2005 (cit. on p. 47).

[228] H Meuer, E Strohmaier, and J Dongarra. *TOP500 Supercomputer Sites*. 2017. URL: https://www.top500.org/list/2017/06/ (cit. on pp. 48, 49).

[229] JJ Dongarra, P Luszczek, and A Petitet. "The LINPACK Benchmark: past, present and future". In: *Concurrency and Computation: Practice and Experience* 15.9 (2003), pp. 803–820. ISSN: 1532-0626. DOI: 10.1002/cpe.728 (cit. on p. 48).

[230] T Plewa, TJ Linde, and VG Weirs, eds. *Adaptive Mesh Refinement, Theory and Applications*. 2005th ed. Lecture notes in computational science and engineering, 41. Springer, 2005. ISBN: 3540211470 (cit. on p. 49).

[231] W Zhang, A Almgren, V Beckner, J Bell, J Blaschke, C Chan, M Day, B Friesen, K Gott, D Graves, M Katz, A Myers, T Nguyen, A Nonaka, M Rosso, S Williams, and M Zingale. "AMReX: a framework for block-structured adaptive mesh refinement". In: *Journal of Open Source Software* 4.37 (2019), p. 1370. ISSN: 2475-9066. DOI: 10.21105/joss.01370 (cit. on p. 49).

[232] A Matthes. Private Communication. Dresden, Germany, 2020 (cit. on p. 49).

[233] ME Rowan, A Huebl, KN Gott, J Deslippe, M Thévenet, R Lehe, and JL Vay. "In-Situ Assessment of Device-Side Compute Work for Dynamic Load Balancing in a GPU-Accelerated PIC Code". In: *Proceedings of TBD* 1.1 (2021). arXiv: 2104.11385 (cit. on p. 49).

[234] JL Vay and BB Godfrey. "Modeling of relativistic plasmas with the Particle-In-Cell method". In: *Comptes Rendus Mécanique* 342.10-11 (2014), pp. 610–618. ISSN: 16310721. DOI: 10.1016/j.crme.2014.07.006 (cit. on p. 50).

[235] TD Arber, K Bennett, CS Brady, A Lawrence-Douglas, MG Ramsay, NJ Sircombe, P Gillies, RG Evans, H Schmitz, aR Bell, and CP Ridgers. "Contemporary particle-in-cell approach to laser-plasma modelling". In: *Plasma Physics and Controlled Fusion* 57.11 (2015), p. 113001. ISSN: 0741-3335. DOI: 10.1088/0741-3335/57/11/113001 (cit. on p. 50).

[236] BB Godfrey. "Numerical Cherenkov instabilities in electromagnetic particle codes". In: *Journal of Computational Physics* 15.4 (1974), pp. 504–521. ISSN: 00219991. DOI: 10.1016/0021-9991(74)90076-X (cit. on p. 50).

[237] R Nuter, M Grech, P Gonzalez de Alaiza Martinez, G Bonnaud, and E D'Humières. "Maxwell solvers for the simulations of the laser-matter interaction". In: *The European Physical Journal D* 68.6 (2014), p. 177. ISSN: 1434-6060. DOI: 10.1140/epjd/e2014-50162-y (cit. on p. 50).

[238] R Lehe, M Kirchen, IA Andriyash, BB Godfrey, and JL Vay. "A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm". In: *Computer Physics Communications* 203 (2016), pp. 66–82. ISSN: 00104655. DOI: 10.1016/j.cpc.2016.02.007. arXiv: 1507.04790 (cit. on pp. 50, 56).

[239] J Park, SS Bulanov, J Bin, Q Ji, S Steinke, JL Vay, CGR Geddes, CB Schroeder, WP Leemans, T Schenkel, and E Esarey. "Ion acceleration in laser generated megatesla magnetic vortex". In: *Physics of Plasmas* 26.10 (2019), p. 103108. ISSN: 1070-664X. DOI: 10.1063/1.5094045. arXiv: 1904.03281 (cit. on p. 50).

[240] *About the Lustre File System*. 2020. URL: http://lustre.org/about/ (cit. on p. 52).

[241] *CSCS File Systems*. 2020. URL: https://user.cscs.ch/storage/file_systems/ (cit. on p. 52).

[242] *Lustre Software Release 2.x - Operations Manual*. 2020. URL: https://doc.lustre.org/lustre_manual.pdf (cit. on p. 52).

[243] Q Liu, J Logan, Y Tian, H Abbasi, N Podhorszki, JY Choi, S Klasky, R Tchoua, J Lofstead, R Oldfield, M Parashar, N Samatova, K Schwan, A Shoshani, M Wolf, K Wu, and W Yu. "Hello ADIOS: the challenges and lessons of developing leadership class I/O frameworks". In: *Concurrency and Computation: Practice and Experience* 26.7 (2014), pp. 1453–1473. ISSN: 15320626. DOI: 10.1002/cpe.3125 (cit. on pp. 52, 56).

[244] WF Godoy, N Podhorszki, R Wang, C Atkins, G Eisenhauer, J Gu, P Davis, J Choi, K Germaschewski, K Huck, A Huebl, M Kim, J Kress, T Kurc, Q Liu, J Logan, K Mehta, G Ostrouchov, M Parashar, F Poeschel, D Pugmire, E Suchyta, K Takahashi, N Thompson, S Tsutsumi, L Wan, M Wolf, K Wu, and S Klasky. "ADIOS 2: The Adaptable Input Output System. A framework for high-performance data management". In: *SoftwareX* 12 (2020), p. 100561. ISSN: 23527110. DOI: 10.1016/j.softx.2020.100561 (cit. on pp. 52, 60).

[245] 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: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 10524 LNCS. 654220. 2017, pp. 15–29. ISBN: 9783319676296. DOI: 10.1007/978-3-319-67630-2_2. arXiv: arXiv:1706.00522v1 (cit. on pp. 53, 58, 59).

[246] F Alted. *blosc 1.12.1*. 2017. URL: https://github.com/Blosc/c-blosc (cit. on p. 53).

[247] N Liu, J Cope, P Carns, C Carothers, R Ross, G Grider, A Crume, and C Maltzahn. "On the role of burst buffers in leadership-class storage systems". In: *012 IEEE 28th Symposium on Mass Storage Systems and Technologies (MSST)*. IEEE, 2012, pp. 1–11. ISBN: 978-1-4673-1747-4. DOI: 10.1109/MSST.2012.6232369 (cit. on p. 54).

[248] Y Collet and MS Kucherawy. "Zstandard Compression and the 'application/zstd' Media Type". In: *RFC* 8878 (2021), pp. 1–45. DOI: 10.17487/RFC8878 (cit. on p. 55).

[249] T Silva, F Massimo, and J Vieira. "Summary of WG6: Theory and simulations". In: *Journal of Physics: Conference Series* 1596.1 (2020), p. 012051. ISSN: 1742-6588. DOI: 10.1088/1742-6596/1596/1/012051 (cit. on p. 56).

[250] J Schreiber and Contributors. *Interactive Graph of Proton Spectra*. 2021. URL: https://www.alpa.physik.uni-muenchen.de/protons.html (cit. on p. 56).

[251] *Horizon 2020*. 2021. URL: https://ec.europa.eu/programmes/horizon2020/en (cit. on p. 56).

[252] *Horizon Europe*. 2021. URL: https://ec.europa.eu/info/research-and-innovation/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en (cit. on p. 56).

[253] *Particle-in-Cell Modeling Interface*. 2021. URL: https://github.com/picmi-standard/picmi (cit. on p. 56).

[254] JL Vay, R Lehe, and D Grote. *Particle-In-Cell/Accelerator Modeling Interface*. 2018. DOI: 10.11578/dc.20181121.1. URL: https://doi.org/10.11578/dc.20181121.1 (cit. on p. 56).

[255] A Huebl, R Lehe, JL Vay, DP Grote, I Sbalzarini, S Kuschel, and M Bussmann. *openPMD 1.0.0: A meta data standard for particle and mesh based data.* 2015. DOI: 10.5281/zenodo.33624 (cit. on pp. 56, 134, 136).

[256] A Huebl, F Poeschel, F Koller, and J Gu. "openPMD api - C++ and Python API for Scientific I/O with openPMD". In: (2021). DOI: 10.14278/RODARE.1234 (cit. on p. 56).

[257] *Projects using openPMD*. 2020. URL: https://github.com/openPMD/openPMD-projects (cit. on pp. 56, 57).

[258] *EUCALL - European Cluster of Advanced Laser Light Sources*. 2015. URL: https://www.eucall.eu/ (cit. on pp. 56, 136).

[259] *PaNOSC - photon and neutron open science cloud*. 2019. URL: https://www.panosc.eu/ (cit. on pp. 56, 136).

[260] *FBPIC: Spectral, quasi-3D Particle-In-Cell code, for CPU and GPU*. URL: https://github.com/fbpic/fbpic (cit. on p. 56).

[261] JL Vay, A Huebl, A Almgren, LD Amorim, J Bell, L Fedeli, L Ge, K Gott, DP Grote, M Hogan, R Jambunathan, R Lehe, A Myers, C Ng, M Rowan, O Shapoval, M Thévenet, H Vincenti, E Yang, N Zaïm, W Zhang, Y Zhao, and E Zoni. "Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs". In: *Physics of Plasmas* 28.2 (2021), p. 023105. ISSN: 1070-664X. DOI: 10.1063/5.0028512 (cit. on p. 56).

[262] *WarpX Github Repository*. URL: https://github.com/ECP-WarpX/WarpX (cit. on p. 56).

[263] RA Fonseca, LO Silva, FS Tsung, VK Decyk, W Lu, C Ren, WB Mori, S Deng, S Lee, T Katsouleas, and JC Adam. "OSIRIS: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2331 LNCS.PART 3 (2002), pp. 342–351. ISSN: 16113349. DOI: 10.1007/3-540-47789-6_36 (cit. on p. 56).

[264] The HDF Group. *Hierarchical Data Format version 5 (C-API: 1.8.20)*. 2017. URL: https://www.hdfgroup.org/solutions/hdf5 (cit. on p. 56).

[265] *dcp - A tool to copy file(s) in parallel on a distributed system*. 2017. URL: https://github.com/hpc/dcp (cit. on p. 57).

[266] *BBCP - Securely and quickly copy data from source to target*. 2017. URL: https://www.slac.stanford.edu/~abh/bbcp/ (cit. on p. 57).

[267] W Allcock, J Bresnahan, R Kettimuthu, and M Link. "The Globus Striped GridFTP Framework and Server". In: *ACM/IEEE SC 2005 Conference (SC'05)*. Vol. 2005. c. IEEE, 2005, pp. 54–54. ISBN: 1-59593-061-2. DOI: 10.1109/SC.2005.72 (cit. on p. 57).

[268] *AWS Snow Family*. 2021. URL: https://aws.amazon.com/snow/ (cit. on p. 58).

[269] *Globus-Url-Copy (deprecated)*. 2021. URL: https://user.cscs.ch/storage/transfer/external/#globus-url-copy-deprecated- (cit. on p. 57).

[270] *Globus Toolkit [1998-2018]*. 2021. URL: https://toolkit.globus.org/ (cit. on p. 57).

[271] *Globus Transfer API*. 2021. URL: https://docs.globus.org/api/transfer/ (cit. on p. 57).

[272] B Corcoran, M Tan, X Xu, A Boes, J Wu, TG Nguyen, ST Chu, BE Little, R Morandotti, A Mitchell, and DJ Moss. "Ultra-dense optical data transmission over standard fibre with a single chip source". In: *Nature Communications* 11.1 (2020), p. 2568. ISSN: 2041-1723. DOI: 10.1038/s41467-020-16265-x (cit. on p. 58).

[273] L Galdino, A Edwards, W Yi, E Sillekens, Y Wakayama, T Gerard, WS Pelouch, S Barnes, T Tsuritani, RI Killey, D Lavery, and P Bayvel. "Optical Fibre Capacity Optimisation via Continuous Bandwidth Amplification and Geometric Shaping". In: *IEEE Photonics Technology Letters* 32.17 (2020), pp. 1021–1024. ISSN: 1041-1135. DOI: 10.1109/LPT.2020.3007591 (cit. on p. 58).

[274] T Kluyver, B Ragan-Kelley, F Pérez, B Granger, M Bussonnier, J Frederic, K Kelley, J Hamrick, J Grout, S Corlay, P Ivanov, D Avila, S Abdalla, and C Willing. *Jupyter Notebooks – a publishing format for reproducible computational workflows*. 2016. DOI: 10.3233/978-1-61499-649-1-87 (cit. on p. 59).

[275] A Shan. "Heterogeneous processing: a strategy for augmenting moore's law: Linux Journal: Vol 2006, No 142". In: *Linux Journal* 2006.142 (2006) (cit. on p. 59).

[276] GE Moore. "Cramming more components onto integrated circuits, Reprinted from Electronics, volume 38, number 8, April 19, 1965, pp.114 ff." In: *IEEE Solid-State Circuits Society Newsletter* 11.3 (2006), pp. 33–35. ISSN: 1098-4232. DOI: 10.1109/N-SSC.2006.4785860 (cit. on p. 59).

[277] H Sim, S Vazhkudai, and A Khan. *An Analysis of System Balance and Architectural Trends Based on Top500 Supercomputers*. Tech. rep. Oak Ridge, TN (United States): Oak Ridge National Laboratory (ORNL), 2020. DOI: 10.2172/1649132 (cit. on p. 59).

[278] S van der Walt, SC Colbert, and G Varoquaux. "The NumPy Array: A Structure for Efficient Numerical Computation". In: *Computing in Science & Engineering* 13.2 (2011), pp. 22–30. ISSN: 1521-9615. DOI: 10.1109/MCSE.2011.37. arXiv: 1102.1523 (cit. on p. 59).

[279] JD Hunter. "Matplotlib: A 2D Graphics Environment". In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. ISSN: 1521-9615. DOI: 10.1109/MCSE.2007.55 (cit. on p. 59).

[280] *Urika-XC: Advanced Analytics, Analytics Software - Cray*. URL: http://www.cray.com/products/analytics/urika-xc (visited on May 19, 2021) (cit. on p. 59).

[281] JV Vermaas, A Sedova, MB Baker, S Boehm, DM Rogers, J Larkin, J Glaser, MD Smith, O Hernandez, and JC Smith. “Supercomputing Pipelines Search for Therapeutics Against COVID-19”. In: *Computing in Science & Engineering* 23.1 (2021), pp. 7–16. ISSN: 1521-9615. DOI: 10.1109/MCSE.2020.3036540 (cit. on p. 60).

[282] *RAPIDS: GPU-Accelerated Data Analytics & Machine Learning*. URL: https://developer.nvidia.com/rapids (visited on May 19, 2021) (cit. on p. 60).

[283] *Open-Source SQL in Python*. URL: https://blazingsql.com/ (visited on May 19, 2021) (cit. on p. 60).

[284] M Sherlock, EG Hill, RG Evans, SJ Rose, and W Rozmus. “In-depth plasma-wave heating of dense plasma irradiated by short laser pulses”. In: *Physical Review Letters* 113.25 (2014), pp. 1–5. ISSN: 10797114. DOI: 10.1103/PhysRevLett.113.255001 (cit. on pp. 64, 73).

[285] AJ Kemp and Y Sentoku. “Comment on "In-depth Plasma-Wave Heating of Dense Plasma Irradiated by Short Laser Pulses"”. In: *Physical Review Letters* 116.15 (2016), p. 159501. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.116.159501 (cit. on pp. 64, 73).

[286] M Sherlock, W Rozmus, EG Hill, and SJ Rose. “Sherlock et al. Reply:” in: *Physical Review Letters* 116.15 (2016), p. 159502. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.116.159502 (cit. on pp. 64, 73).

[287] M Rödel. “Coherent X-ray Diffraction of Laser-Driven Matter”. PhD thesis. 2020 (cit. on pp. 65, 120, 128, 129, 131).

[288] C Fortmann-Grote, A Andreev, R Briggs, M Bussmann, A Huebl, T Kluge, S Pascarelli, A Sharma, and AP Mancuso. *Milestone M4.1: Delivery of individual simulation modules and common interfaces for interoperability*. 2017. DOI: 10.5281/ZENODO.896329 (cit. on pp. 65, 137).

[289] C Fortmann-Grote. *EUCALL Milestone M4.2: Demonstration of a first example simulation*. 2017. DOI: 10.5281/ZENODO.998696 (cit. on p. 65).

[290] C Fortmann-Grote. *EUCALL Milestone M4.3: Simulations interoperable*. 2017. DOI: 10.5281/ZENODO.998654 (cit. on p. 65).

[291] O Lundh, J Lim, C Rechatin, L Ammoura, A Ben-Ismaïl, X Davoine, G Gallot, JP Goddet, E Lefebvre, V Malka, and J Faure. “Few femtosecond, few kiloampere electron bunch produced by a laser-plasma accelerator”. In: *Nature Physics* 7.3 (2011), pp. 219–222. ISSN: 1745-2473. DOI: 10.1038/nphys1872 (cit. on p. 66).

[292] YF Li, DZ Li, K Huang, MZ Tao, MH Li, JR Zhao, Y Ma, X Guo, JG Wang, M Chen, N Hafz, J Zhang, and LM Chen. “Generation of 20 kA electron beam from a laser wakefield accelerator”. In: *Physics of Plasmas* 24.2 (2017), p. 023108. ISSN: 1070-664X. DOI: 10.1063/1.4975613 (cit. on p. 66).

[293] M Tzoufras, W Lu, FS Tsung, C Huang, WB Mori, T Katsouleas, J Vieira, RA Fonseca, and LO Silva. “Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration”. In: *Physical Review Letters* 101.14 (2008), p. 145002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.101.145002 (cit. on pp. 66, 69).

[294] E Guillaume, A Döpp, C Thaury, A Lifschitz, JP Goddet, A Tafzi, F Sylla, G Iaquanello, T Lefrou, P Rousseau, KT Phuoc, and V Malka. “Physics of fully-loaded laser-plasma accelerators”. In: *Physical Review Special Topics - Accelerators and Beams* 18.6 (2015), p. 061301. ISSN: 1098-4402. DOI: 10.1103/PhysRevSTAB.18.061301 (cit. on p. 66).

[295] M Zeng, M Chen, ZM Sheng, WB Mori, and J Zhang. "Self-truncated ionization injection and consequent monoenergetic electron bunches in laser wakefield acceleration". In: *Physics of Plasmas* 21.3 (2014), p. 030701. ISSN: 1070-664X. DOI: 10.1063/1.4868404. arXiv: 1312.5825 (cit. on p. 66).

[296] M Mirzaie, S Li, M Zeng, NA Hafz, M Chen, GY Li, QJ Zhu, H Liao, T Sokollik, F Liu, YY Ma, LM Chen, ZM Sheng, and J Zhang. "Demonstration of self-truncated ionization injection for GeV electron beams". In: *Scientific Reports* 5 (2015), pp. 1–9. ISSN: 20452322. DOI: 10.1038/srep14659 (cit. on p. 66).

[297] JP Couperus. "Optimal beam loading in a nanocoulomb-class laser wakefield accelerator". PhD thesis. 2018. DOI: 10.5281/zenodo.1463710 (cit. on pp. 66, 70).

[298] N Delone and V Krainov. "Tunneling and barrier-suppression ionization of atoms and ions in a laser radiation field". In: *Uspekhi Fizicheskih Nauk* 168.5 (1998), p. 531. ISSN: 0042-1294. DOI: 10.3367/UFNr.0168.199805c.0531 (cit. on p. 67).

[299] LB Fletcher, HJ Lee, T Döppner, E Galtier, B Nagler, P Heimann, C Fortmann, S LePape, T Ma, M Millot, A Pak, D Turnbull, DA Chapman, DO Gericke, J Vorberger, T White, G Gregori, M Wei, B Barbrel, RW Falcone, CC Kao, H Nuhn, J Welch, U Zastrau, P Neumayer, JB Hastings, and SH Glenzer. "Ultrabright X-ray laser scattering for dynamic warm dense matter physics". In: *Nature Photonics* 9.4 (2015), pp. 274–279. ISSN: 1749-4885. DOI: 10.1038/nphoton.2015.41 (cit. on p. 70).

[300] N Mukharamova, S Lazarev, JM Meijer, M Chollet, A Singer, R Kurta, D Dzhigaev, O Gorobtsov, G Williams, D Zhu, Y Feng, M Sikorski, S Song, A Shabalin, T Gurieva, E Sulyanova, O Yefanov, and I Vartanyants. "Probing Dynamics in Colloidal Crystals with Pump-Probe Experiments at LCLS: Methodology and Analysis". In: *Applied Sciences* 7.5 (2017), p. 519. ISSN: 2076-3417. DOI: 10.3390/app7050519 (cit. on pp. 70, 74).

[301] J MacFarlane, I Golovkin, and P Woodruff. "HELIOS-CR – A 1-D radiation-magnetohydrodynamics code with inline atomic kinetics modeling". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 99.1-3 (2006), pp. 381–397. ISSN: 00224073. DOI: 10.1016/j.jqsrt.2005.05.031 (cit. on pp. 71, 73).

[302] A Kemp and J Meyer-ter-Vehn. "An equation of state code for hot dense matter, based on the QEOS description". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 415.3 (1998), pp. 674–676. ISSN: 01689002. DOI: 10.1016/S0168-9002(98)00446-X (cit. on p. 73).

[303] B Cohen, A Kemp, and L Divol. "Simulation of laser-plasma interactions and fast-electron transport in inhomogeneous plasma". In: *Journal of Computational Physics* 229.12 (2010), pp. 4591–4612. ISSN: 00219991. DOI: 10.1016/j.jcp.2010.03.001 (cit. on p. 73).

[304] S Lyon and J Johnson. "Sesame: the Los Alamos National Laboratory equation of state database". In: *Los Alamos National Laboratory, Los Alamos, NM, LA-...* (1992), p. 43 (cit. on p. 73).

[305] G Doumy, F Quéré, O Gobert, M Perdrix, P Martin, P Audebert, JC Gauthier, JP Geindre, and T Wittmann. "Complete characterization of a plasma mirror for the production of high-contrast ultraintense laser pulses". In: *Physical Review E* 69.2 (2004), p. 026402. ISSN: 1539-3755. DOI: 10.1103/PhysRevE.69.026402 (cit. on p. 76).

[306] PL Poole, A Krygier, GE Cochran, PS Foster, GG Scott, LA Wilson, J Bailey, N Bourgeois, C Hernandez-Gomez, D Neely, PP Rajeev, RR Freeman, and DW Schumacher. "Experiment and simulation of novel liquid crystal plasma mirrors for high contrast, intense laser pulses". In: *Scientific Reports* 6.June (2016), pp. 1–8. ISSN: 20452322. DOI: 10.1038/srep32041 (cit. on pp. 76, 80).

[307] BH Shaw, S Steinke, J Van Tilborg, and WP Leemans. “Reflectance characterization of tape-based plasma mirrors”. In: *Physics of Plasmas* 23.6 (2016). ISSN: 10897674. DOI: 10.1063/1.4954242 (cit. on p. 76).

[308] T Oksenhendler, P Bizouard, O Albert, S Bock, and U Schramm. “High dynamic, high resolution and wide range single shot temporal pulse contrast measurement”. In: *Optics Express* 25.11 (2017), p. 12588. ISSN: 1094-4087. DOI: 10.1364/OE.25.012588 (cit. on p. 78).

[309] F Verluise, V Laude, Z Cheng, C Spielmann, and P Tournois. “Amplitude and phase control of ultrashort pulses by use of an acousto-optic programmable dispersive filter: pulse compression and shaping”. In: *Optics Letters* 25.8 (2000), p. 575. ISSN: 0146-9592. DOI: 10.1364/OL.25.000575 (cit. on p. 78).

[310] P Sommer, J Metzkes-Ng, FE Brack, TE Cowan, SD Kraft, L Obst, M Rehwald, HP Schlenvoigt, U Schramm, and K Zeil. “Laser-ablation-based ion source characterization and manipulation for laser-driven ion acceleration”. In: *Plasma Physics and Controlled Fusion* 60.5 (2018). ISSN: 13616587. DOI: 10.1088/1361-6587/aab21e (cit. on pp. 80, 102).

[311] M Allen, PK Patel, A Mackinnon, D Price, S Wilks, and E Morse. “Direct experimental evidence of back-surface ion acceleration from laser-irradiated gold foils”. In: *Physical Review Letters* 93.26 I (2004), pp. 1–4. ISSN: 00319007. DOI: 10.1103/PhysRevLett.93.265004 (cit. on p. 80).

[312] S Steinke, A Henig, M Schnürer, T Sokollik, P Nickles, D Jung, D Kiefer, R Hörlein, J Schreiber, T Tajima, X Yan, M Hegelich, J Meyer-ter-Vehn, W Sandner, and D Habs. “Efficient ion acceleration by collective laser-driven electron dynamics with ultra-thin foil targets”. In: *Laser and Particle Beams* 28.1 (2010), pp. 215–221. ISSN: 0263-0346. DOI: 10.1017/S0263034610000157 (cit. on p. 80).

[313] *Ultra Thin Copper Nanofoil*. 2020. URL: https://www.americanelements.com/ultra-thin-copper-nanofoil-7440-50-8 (cit. on p. 80).

[314] PL Poole, C Willis, GE Cochran, RT Hanna, CD Andereck, and DW Schumacher. “Moderate repetition rate ultra-intense laser targets and optics using variable thickness liquid crystal films”. In: *Applied Physics Letters* 109.15 (2016), p. 151109. ISSN: 0003-6951. DOI: 10.1063/1.4964841 (cit. on p. 80).

[315] A Jahn. *Entwicklung und Charakterisierung ultradünner Targets zur Beschleunigung von Protonen und Ionen am DRACO Laser unter ultrahohem zeitlichen Kontrast*. Master Thesis. 2017 (cit. on p. 81).

[316] C Bernert, S Kraft, M Löser, J Metzkes-Ng, L Obst-Huebl, M Rehwald, HP Schlenvoigt, M Siebold, K Zeil, T Ziegler, and U Schramm. *Off-harmonic optical probing of high intensity laser interaction with cryogenic hydrogen jet target*. European Advanced Accelerator Concepts Workshop. 2019. URL: https://agenda.infn.it/event/17304/contributions/99107/ (cit. on p. 83).

[317] GA Becker, S Tietze, S Keppler, J Reislöhner, JH Bin, L Bock, FE Brack, J Hein, M Hellwing, P Hilz, M Hornung, A Kessler, SD Kraft, S Kuschel, H Liebetrau, W Ma, J Polz, HP Schlenvoigt, F Schorcht, MB Schwab, A Seidel, K Zeil, U Schramm, M Zepf, J Schreiber, S Rykovanov, and MC Kaluza. “Ring-like spatial distribution of laser accelerated protons in the ultra-high-contrast TNSA-regime”. In: *Plasma Physics and Controlled Fusion* 60.5 (2018), p. 055010. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aab319 (cit. on p. 102).

[318] FE Brack, F Kroll, L Gaus, C Bernert, E Beyreuther, TE Cowan, L Karsch, S Kraft, LA Kunz-Schughart, E Lessmann, J Metzkes-Ng, L Obst-Huebl, J Pawelke, M Rehwald, HP Schlenvoigt, U Schramm, M Sobiella, ER Szabó, T Ziegler, and K Zeil. "Spectral and spatial shaping of laser-driven proton beams using a pulsed high-field magnet beamline". In: *Scientific Reports* 10.1 (2020), p. 9118. ISSN: 2045-2322. DOI: 10.1038/s41598-020-65775-7. arXiv: 1910.08430 (cit. on p. 106).

[319] S Kar, KF Kakolee, B Qiao, A Macchi, M Cerchez, D Doria, M Geissler, P McKenna, D Neely, J Osterholz, R Prasad, K Quinn, B Ramakrishna, G Sarri, O Willi, XY Yuan, M Zepf, and M Borghesi. "Ion Acceleration in Multispecies Targets Driven by Intense Laser Radiation Pressure". In: *Physical Review Letters* 109.18 (2012), p. 185006. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.109.185006 (cit. on p. 107).

[320] S Krishnamurthy, K Makur, and B Ramakrishna. "Observation of resistive Weibel instability in intense laser plasma". In: *Laser and Particle Beams* (2020), pp. 1–7. ISSN: 0263-0346. DOI: 10.1017/S0263034620000154 (cit. on pp. 107, 145).

[321] SA Gaillard, T Kluge, KA Flippo, M Bussmann, B Gall, T Lockard, M Geissel, DT Offermann, M Schollmeier, Y Sentoku, and TE Cowan. "Increased laser-accelerated proton energies via direct laser-light-pressure acceleration of electrons in microcone targets". In: *Physics of Plasmas* 18.5 (2011), p. 056710. ISSN: 1070-664X. DOI: 10.1063/1.3575624 (cit. on p. 108).

[322] T Kluge, SA Gaillard, KA Flippo, T Burris-Mog, W Enghardt, B Gall, M Geissel, A Helm, SD Kraft, T Lockard, J Metzkes, DT Offermann, M Schollmeier, U Schramm, K Zeil, M Bussmann, and TE Cowan. "High proton energies from cone targets: electron acceleration mechanisms". In: *New Journal of Physics* 14.2 (2012), p. 023038. ISSN: 1367-2630. DOI: 10.1088/1367-2630/14/2/023038 (cit. on p. 108).

[323] Juniks/HZDR. *HZDR Press Release: Extremely small and fast: laser ignites hot plasma*. 2018. URL: https://www.hzdr.de/db/Cms?pOid=56858 (cit. on p. 119).

[324] JC Brooks-Bartlett and EF Garman. "The Nobel Science: One Hundred Years of Crystallography". In: *Interdisciplinary Science Reviews* 40.3 (2015), pp. 244–264. ISSN: 0308-0188. DOI: 10.1179/0308018815Z.000000000116 (cit. on p. 120).

[325] WC Röntgen. "On a New Kind of Rays". In: *Nature* 53.1369 (1896), pp. 274–276. ISSN: 0028-0836. DOI: 10.1038/053274b0 (cit. on p. 120).

[326] O Glatter and O Kratky. *Small Angle X-ray Scattering*. Ed. by O Glatter and O Kratky. 1982, p. 515. ISBN: 9780122862809 (cit. on p. 120).

[327] J Als-Nielsen and D McMorrow. *Elements of Modern X-ray Physics*. Wiley, 2011. ISBN: 9780470973950. DOI: 10.1002/9781119998365 (cit. on pp. 120, 121).

[328] "Historical Introduction to the Elementary Particles". In: *Introduction to Elementary Particles*. 2nd ed. John Wiley & Sons, Ltd, 2008. Chap. 1, p. 17. ISBN: 9783527406012 (cit. on p. 120).

[329] AV Martin, AJ D'Alfonso, F Wang, R Bean, F Capotondi, RA Kirian, E Pedersoli, L Raimondi, F Stellato, CH Yoon, and HN Chapman. "X-ray holography with a customizable reference". In: *Nature Communications* 5.1 (2014), p. 4661. ISSN: 2041-1723. DOI: 10.1038/ncomms5661 (cit. on p. 122).

[330] S Kahaly, SK Yadav, WM Wang, S Sengupta, ZM Sheng, A Das, PK Kaw, and GR Kumar. "Near-Complete Absorption of Intense, Ultrashort Laser Light by Sub- <math display="inline"> <mi>λ</mi> </math> Gratings". In: *Physical Review Letters* 101.14 (2008), p. 145001. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.101.145001 (cit. on p. 122).

[331] M Cerchez, AL Giesecke, C Peth, M Toncian, B Albertazzi, J Fuchs, O Willi, and T Toncian. "Generation of Laser-Driven Higher Harmonics from Grating Targets". In: *Physical Review Letters* 110.6 (2013), p. 065003. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.110.065003 (cit. on p. 122).

[332] T Ceccotti, V Floquet, A Sgattoni, A Bigongiari, O Klimo, M Raynaud, C Riconda, A Heron, F Baffigi, L Labate, LA Gizzi, L Vassura, J Fuchs, M Passoni, M Květon, F Novotny, M Possolt, J Prokůpek, J Proška, J Pšikal, L Štolcová, A Velyhan, M Bougeard, P D'Oliveira, O Tcherbakoff, F Réau, P Martin, and A Macchi. "Evidence of Resonant Surface-Wave Excitation in the Relativistic Regime through Measurements of Proton Acceleration from Grating Targets". In: *Physical Review Letters* 111.18 (2013), p. 185001. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.111.185001 (cit. on p. 122).

[333] B Nagler, B Arnold, G Bouchard, RF Boyce, RM Boyce, A Callen, M Campell, R Curiel, E Galtier, J Garofoli, E Granados, J Hastings, G Hays, P Heimann, RW Lee, D Milathianaki, L Plummer, A Schropp, A Wallace, M Welch, W White, Z Xing, J Yin, J Young, U Zastrau, and HJ Lee. "The Matter in Extreme Conditions instrument at the Linac Coherent Light Source". In: *Journal of Synchrotron Radiation* 22.November 2014 (2015), pp. 520–525. ISSN: 16005775. DOI: 10.1107/S1600577515004865 (cit. on p. 124).

[334] M Zacharias. "Model-Driven Parameter Reconstructions From Small Angle X-Ray Scattering Images". Master thesis. TU Dresden, 2017. DOI: 10.5281/zenodo.1208410 (cit. on p. 131).

[335] M Nakatsutsumi, K Appel, G Priebe, I Thorpe, A Pelka, B Muller, and T Tschentscher. *Technical Design Report: Scientific Instrument High Energy Density Physics (HED)*. Tech. rep. XFEL.EU TR-2014-001. OA. 2014, pp. 1–196. DOI: 10.3204/XFEL.EU/TR-2014-001 (cit. on p. 133).

[336] HA Scott. "Chapter 4 – Collisional-Radiative Modeling for Radiation Hydrodynamics Codes". In: *Modern Methods of Radiative Modeling in Plasmas*. Ed. by Y Ralchenko. 2016, p. 86. ISBN: 9783319275123 (cit. on p. 134).

[337] P Ordyna. "X-ray Radiation Transport in GPU-Accelerated Particle-in-Cell Simulations". [working title, in preparation]. Master Thesis. Technische Universität Dresden, 2022 (cit. on p. 140).

[338] P Ordyna. *PIConGPU "xrayScattering" plugin*. [Online; accessed 20-December-2021]. 2021 (cit. on p. 140).

[339] JC E, L Wang, S Chen, YY Zhang, and SN Luo. "GAPD: A GPU-accelerated atom-based polychromatic diffraction simulation code". In: *Journal of Synchrotron Radiation* 25.2 (2018), pp. 604–611. ISSN: 16005775. DOI: 10.1107/S1600577517016733 (cit. on p. 140).

[340] K Bastrakova. *PIConGPU "particleMerger" plugin (probabilistic version)*. [Online; accessed 20-December-2021]. 2020 (cit. on p. 140).

[341] K Bastrakova. *Particle Reducer*. [Online; accessed 20-December-2021]. 2021 (cit. on p. 140).

[342] F Lindau, O Lundh, A Persson, P McKenna, K Osvay, D Batani, and CG Wahlström. "Laser-Accelerated Protons with Energy-Dependent Beam Direction". In: *Physical Review Letters* 95.17 (2005), p. 175002. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.95.175002 (cit. on p. 145).

[343] S Göde, C Rödel, K Zeil, R Mishra, M Gauthier, FE Brack, T Kluge, MJ Macdonald, J Metzkes, L Obst, M Rehwald, C Ruyer, HP Schlenvoigt, W Schumaker, P Sommer, TE Cowan, U Schramm, S Glenzer, and F Fiuza. "Relativistic Electron Streaming Instabilities Modulate Proton Beams Accelerated in Laser-Plasma Interactions". In: *Physical Review Letters* 118.19 (2017), pp. 1–5. ISSN: 10797114. DOI: 10.1103/PhysRevLett.118.194801. arXiv: 1704.04311 (cit. on p. 145).

[344] U Zastrau et al. "The High Energy Density Scientific Instrument at the European XFEL". In: *Journal of Synchrotron Radiation* 28.5 (2021), pp. 1393–1416. ISSN: 1600-5775. DOI: 10.1107/S1600577521007335 (cit. on p. 146).

[345] SJ Pan and Q Yang. "A Survey on Transfer Learning". In: *IEEE Transactions on Knowledge and Data Engineering* 22.10 (2010), pp. 1345–1359. ISSN: 1041-4347. DOI: 10.1109/TKDE.2009.191 (cit. on p. 149).

[346] C Wang. "A geometric framework for transfer learning using manifold alignment". Dissertation. University of Massachusetts Amherst, 2010 (cit. on p. 149).

Acknowledgments

Am Ende dieser Arbeit möchte ich mich nun bei all den wunderbaren Menschen herzlich bedanken, die mich auf dem Weg hierher begleitet oder diesen Weg für mich überhaupt möglich gemacht haben.

Zuallererst möchte ich Thomas Cowan und Ulrich Schramm dafür danken, durch ihre Leitung des Instituts für Strahlenphysik am HZDR und der Abteilung Laser-Teilchenbeschleunigung das exzellente Wissenschaftsumfeld geschaffen zu haben, in dem ich meine Forschung Durchführung durfte. Vielen Dank für die zahlreichen spannenden Diskussionen, insbesondere auch zur Lösung von Fermi-Problemen bspw. bezüglich der Frage nach der Anzahl Wassermoleküle, die die Körper berühmter religiöser Persönlichkeiten vor Tausenden von Jahren durchströmten und sich heute in einer Flasche Mineralwasser wiederfinden. Eure Erfahrung, Geduld und euer reichhaltiges Wissen haben mir Richtung, Ansporn und wertvolles Feedback zur Erstellung dieser Arbeit geliefert.

Meinen herzlichen Dank möchte ich Michael Bussmann aussprechen, der mich im dritten Semester meines Bachelor-Studiengangs während des Physik-Proseminars bei meinem Vortrag zu dunkler Materie betreut und im Anschluss überhaupt erst in die Gruppe am HZDR eingebracht hat. Danke Micha, auch für die Betreuung meiner Bachelor- und Master-Arbeiten und die 10 Jahre spannender Plasmaphysik an den fortschrittlichsten Supercomputern unserer Zeit, die du durch unermüdliches Einwerben neuer Projekte möglich gemacht hast. Danke für die vielen Herausforderungen, dein Mentoring und dass durch dein Wirken, die Gruppe *Computational Radiation Physics* stetig angewachsen ist und ich so die vielen tollen Menschen kennenlernen durfte, mit denen ich über die Jahre zusammengearbeitet und von denen ich so viel gelernt habe.

Meinem wissenschaftlichen Betreuer Thomas Kluge möchte ich dafür danken, geduldig sein umfangreiches Wissen zur Theorie der Laser-Ionenbeschleungung mit mir geteilt, mich bei der Durchführung der einjährigen Rechenzeit auf Piz Daint leitend unterstützt und mir vor allem immer den Rücken freigehalten zu haben, wenn die schiere Menge unserer wertvollen Daten der automatischen Löschung nah war, bevor deren sicherer Abtransport effizient geregelt werden konnte. Thomas, hab Dank für deinen stetigen Rat, dein Vertrauen in mich und die Freundschaft, die über die Arbeit hinaus entstanden ist.

Mein tiefer Dank gebührt meinen lieben Kollegen und Freunden Axel Huebl und René Widera, ohne deren Einsatz, Vorarbeit und Betreuung mein Projekt sicher weder möglich gewesen wäre noch so erfolgreich hätte abgeschlossen werden können. Danke Axel für deine Inspiration, Vorbildfunktion, deine Hingabe zu *Open Science*, deine stets positive Einstellung und deine Geduld trotz unfassbar vieler gleichzeitiger Projekte, die du betreust und bearbeitest.

René, danke dass du immer geduldig Antworten auf meine vielen Fragen bezüglich Parallelisierung, Programmierung, und Supercomputer-Architektur zu geben wusstest. Für das Projekt *PIConGPU* bist und bleibst du unersetzlich.

Meinen langjährigen Bürogenossen Klaus und Richard möchte ich vielmals dafür danken, mich herzlich in ihre Mitte aufgenommen und bei all meinen Tätigkeiten mit der Beantwortung täglicher Fragen zu Physik und Supercomputing, dem Lesen von Abstracts und Durchspinnen von Ideen unterstützt zu haben. Klaus, danke dir vor allem auch für deine Einsichten in den Werdegang des Doktoranden und die Teilhabe an deinen köstlichen Brau-Spezialitäten. Richard, vielen Dank für unermüdlichen Rat beim Verfassen meiner Bash-, Batch- und Python-Skripte und danke besonders auch für den Probe-Durchlauf am Vorabend meines Rigorosums.

Vielen Dank auch an Alexander Debus, Jan, Sergei, Ilja, João, Alexander Matthes, Brian, Franzi, und Thomas M. für die wichtigen Dikussionen zu Physik und Simulationen, für euer wertvolles Feedback und Hilfe bei der Ausarbeitung von Vorträgen. Danke Ilja, dass du meine Simulationen mit deinen Studien parallel untermauert und mich bei der Bewältigung der Datenmengen untertützt hast.

Mein besonderer Dank geht an die Mitglieder des HZDR Ionenteams unter Leitung von Karl Zeil, die mich als "Computer-Experimentator" in ihre eng zusammengeschweißte Gruppe aufgenommen haben. Danke Karl für die tollen Gespräche, deine Erfahrung und Betreuung im Bezug auf die Realität der Physik im Labor, die dafür gesorgt haben, dass meine Theorien und Ideen hoffentlich nie ganz den Bezug zum Machbaren verloren haben. Vielen Dank auch dir Tim, dass ich mit dir tagein und tagaus über meine Simulationen sprechen und von dir Einblicke in die vielen Laser-Plasma Diagnostiken, die frischesten Resultate vom Beschleuniger und die Deutung der spektralen und temporalen Eigenschaften des Lasers erhalten durfte.

Auch bei dem Elektronen-Team unter Leitung von Arie Irman möchte ich mich herzlich für die vielen fruchtbaren Diskussionen und Meetings bedanken, die mir unter anderem tieferes Verständnis und Einblick in Laser-Elektronen-Beschleuniger ermöglicht haben.

Da natürlich ohne den Laser auch keine Experimente möglich sind, möchte ich auch nicht vergessen, mich an dieser Stelle herzlich bei den Laser-Teams von DRACO und PENELOPE zu bedanken.

Ulf Markwardt und Henrik Schulz sowie der ZIH- und HZDR IT möchte ich dafür danken, alle meine Service-Anfragen klaglos entgegen genommen zu haben und dass ich hunderte Terabytes meiner Daten über ihre Leitungen jagen und auf ihren Platten sicher ablegen durfte.

For fruitful collaborations, I would like to thank Nastya, Dan, Igor, Andrei, Jan Nikl, Mohammadreza, Carsten Fortmann-Grote, Ronnie, and Sunita Chandrasekaran.

Ich möchte auch vielmals Anne, Petra und Nicki dafür danken, dass alle Konferenzbesuche und Dienstreisen immer reibungsfrei funktioniert haben, da alle administrativen Aufgaben bei ihnen in sicheren Händen lagen.

In meiner Zeit am HZDR habe ich nicht nur von der wissenschaftlichen Erfahrung meiner Kollegen profitiert, sondern viele der Kollegen sind mir auch als enge Freunde ans Herz gewachsen. Daher möchte ich hier besonders als erstes der Quiz-Gruppe *WDAWKT* für die vielen wunderbaren Mittwochabende danken, die mir so wichtig als Ausgleich von den Sorgen des Doktoranden-Alltags waren. Besonders aber in der Zeit der allgemeinen Absonderung und gleichzeitigen Schreib-Phase dieser Arbeit während der Corona-Pandemie waren Online-Quizze am Freitagabend über Zoom in so mancher Woche die nahezu einzige Gelegenheit zur Sozialisierung mit echten Menschen. Ohne euch wäre mir sicher die Decke auf den Kopf gefallen! Weiterhin möchte ich meinen Dank an die Mitglieder der *Mandala*-Boulder-Gang und der *Mysa*-Gruppe aussprechen, mit denen ich so manchen Abend an der Kletterwand, auf dem Sofa beim GoT-Schauen oder philosophischen Gesprächen verbracht habe. Vielen Dank euch, Flo B., Flo K., Patti, Viki, Malte, Melanie, Pablo, Andrea, Max, Brant, Lotti, Nick, Martin, Alex K.,

Josephine, Raphael, Rengin, Jurjen, Kushal, Anja, Katja. Jurjen, Max, and Raphael, thank you for yet one more Feierabend-Bier at the Kreuzung.

Auch den Mitgliedern der *Dangerous Twins*, Lennart, Elias, Marvin, Milenko, Stefan A., Angela, und besonders Radka, möchte ich für schöne Stunden mit Cookies, World-Beer-Championships und länderspezifischen Essensabenden danken. In particular, thank you Pablo, for your ever-positive, ever-supportive words of encouragement and faith in me. Radka, thank you for your warmth, happiness and unwavering support that helped me push this thesis past the finish line.

Ein besonderes Highlight waren mir auch die jährlichen Ausflüge mit den Kanu-Chaoten Sebbi, Kupschi, Hauer, Elchi, Daniel, Felix, Sarah, Timon, Jan und Scotch. Auf dass wir niemals untergehen! An meine Nachbarn Mirko, Karsten, und Fabian, danke dass ihr mich von Zeit zu Zeit aus meiner Höhle geholt gehabt und wir eine wunderbar funktionierende Waldi 20 Hinterhaus-Gemeinschaft waren.

Meinen beiden wundervollen Elternpaaren, meinem Bruder, meinen Großeltern und meiner lieben Familie möchte ich für die jahrelange Unterstützung, die aufmunternden und inspirierenden Worte und Taten danken, die mir diesen Weg schon seit früher Kindheit eröffnet haben, indem sie mir die Entfaltung aller meiner Neugierden erlaubten und alle Förderung ermöglichten. Danke für die liebevolle, warme Umgebung, in die ich immer wieder gern nach Hause kommen kann, um neue Kraft und Mut zu tanken.

I would like to express my gratitude to the Helmholtz-Zentrum Dresden – Rossendorf and Technical University Dresden for the excellent facilities and staff providing the environment for my studies in Dresden. The GPU Center of Excellence Dresden provided a framework for many fruitful projects and travel funding, sponsored by the Nvidia Corporation. This project has been enabled by many people in open-source and open-science communities. Great thanks to the communities and developers of: PIConGPU, Jupyter, yt, the SciPy ecosystem, ADIOS, HDF5, CMake, openPMD, Spack, … This research has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 654220. We acknowledge PRACE for awarding us access to Piz Daint at CSCS, Switzerland. 15th PRACE call, Project ID 2016163983.

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 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 23. Mai 2018 erkenne ich an.

Dresden, 20. Januar 2022

Marco Garten

HELMHOLTZ ZENTRUM
DRESDEN ROSSENDORF

Bautzner Landstr. 400
01328 Dresden, Germany
Phone +49 351 260-2618
Fax +49 351 260-12618
Email t.kluge@hzdr.de
http://www.hzdr.de