

High-resolution Simulations of Strongly Coupled Coulomb Systems with a Parallel Tree Code

Mathias Winkel



IAS Series Volume 20

Forschungszentrum Jülich GmbH Institute for Advanced Simulation (IAS) Jülich Supercomputing Centre (JSC)

High-resolution Simulations of Strongly Coupled Coulomb Systems with a Parallel Tree Code

Mathias Winkel

Schriften des Forschungszentrums Jülich

IAS Series Volume 20

Bibliographic information published by the Deutsche Nationalbibliothek. The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Publisher and Forschungszentrum Jülich GmbH

Distributor: Zentralbibliothek

52425 Jülich

Phone +49 (0) 24 61 61-53 68 · Fax +49 (0) 24 61 61-61 03

e-mail: zb-publikation@fz-juelich.de Internet: http://www.fz-juelich.de/zb

Cover Design: Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH

Printer: Grafische Medien, Forschungszentrum Jülich GmbH

Copyright: Forschungszentrum Jülich 2013

Schriften des Forschungszentrums Jülich IAS Series Volume 20

D 82 (Diss. RWTH Aachen University, 2013)

ISSN 1868-8489 ISBN 978-3-89336-901-0

Persistent Identifier: urn:nbn:de:0001-2013091802

Resolving URL: http://www.persistent-identifier.de/?link=610

Neither this book nor any part of it may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage and retrieval system, without permission in writing from the publisher.

To those who are always with us - in our minds and in our hearts.

High-resolution Simulations of Strongly Coupled Coulomb Systems with a Parallel Tree Code

Despite intense research, the properties of strongly coupled Coulomb systems have not yet been completely understood. However, with the advent of Free Electron Lasers with wavelengths reaching down to tenths of nanometers and intensities beyond $10^{16}\,\mathrm{W/cm^{-2}}$ during the last years, it has become possible to experimentally probe the warm dense matter regime up to solid densities.

Now, systems that can be studied are reaching from hot, low-density plasmas of fusion research to cold dense solids that are dominated by quantum-mechanical effects and strong correlations. Their consistent theoretical description requires a multitude of effects to be considered. In particular, strong correlations pose significant difficulties here. Computer simulations provide a tool for bridging between experiments and theory as they do not suffer from these complications.

The experimentally accessible optical and transport properties in plasmas are primarily featured by the electronic subsystem, such as its collective behavior and interaction with the ionic background, i. e. Coulomb collisions. In this work the collisional behavior of warm dense bulk matter and collective effects in nano plasmas are investigated by means of molecular dynamics simulations. To this end, simulation experiments performed earlier on electronic resonances in metallic nano clusters are extended to significantly larger systems. The observed complex resonance structure is analyzed using a newly introduced spatially resolved spectral diagnostic. As a second field of study, the bulk collision frequency as the key parameter for optical and transport properties in warm dense matter is evaluated in a generalized Drude approach for a hydrogen-like plasma. Here, the combined high-field and strong coupling regime that is only scarcely covered by theoretical models is of primary interest.

To solve the underlying N-body problem for both applications, a highly parallel Barnes-Hut tree code is utilized and considerably extended with respect to functionality, versatility, and scalability. With its new excellent scalability to hundred thousands of processors and simulation setups consisting of up to billions of particles and its support for periodic boundary conditions with an efficient and precise real-space approach it delivers highly resolved results and is prepared for further studies on the warm dense matter regime. Here, its unique predictive capabilities can finally be used for connecting to real-world experiments.

Hoch aufgelöste Simulationen stark gekoppelter Coulomb-Systeme mit einem parallelen Tree-Code

Die Eigenschaften stark gekoppelter Coulombsysteme sind trotz intensiver Forschung nicht vollständig verstanden. Mit dem Aufkommen von Freie-Elektronen-Lasern, die Strahlungsintensitäten jenseits $10^{16}\,\mathrm{W/cm^{-2}}$ bei Wellenlängen deutlich im sub-Nanometer Bereich liefern, stehen jedoch erstmals die experimentellen Möglichkeiten zur Verfügung, mit optischen Pump-Probe Experimenten bis zu Festkörperdichten vorzudringen.

Systeme, die untersucht werden können reichen damit von den heißen, dünnen Plasmen der Fusionsforschung bis hin zu kalten, von quantenmechanischen Effekten und starken Korrelationen dominierten Festkörpern. Für ihre konsistente theoretische Beschreibung muss entsprechend eine Vielzahl physikalischer Effekte einbezogen werden. Insbesondere starke Korrelationen führen dabei zu großen Schwierigkeiten. Mit Computersimulationen, die dieser Einschränkung nicht unterliegen, stehen jedoch Werkzeuge zur Verfügung, um die Lücke zwischen experimentellen und theoretischen Möglichkeiten zu schließen.

Die experimentell erreichbaren optischen und Transporteigenschaften in Plasmen werden primär durch die Eigenschaften des elektronischen Subsystems, etwa kollektive Effekte und Wechselwirkungen mit dem Ionenhintergrund – also Coulomb-Stöße – dominiert. In dieser Arbeit werden das Stoßverhalten warmer dichter Bulk-Materie sowie kollektive Effekte in Nanoplasmen mittels Molekulardynamiksimulationen untersucht. Hierzu werden frühere Simulationsexperimente zu elektronischen Resonanzen in metallischen Nanoclustern zu deutlich größeren Systemen ausgedehnt und die gefundenen Resonanzen mit einer neuen, räumlich aufgelösten spektralen Diagnostik analysiert. Als zweites Anwendungsgebiet wird die Bulk-Stoßfrequenz als Schlüsselparameter für optische und Transporteigenschaften in warmer dichter Materie in einem verallgemeinerten Drude-Ansatz für ein wasserstoffartiges Plasma ausgewertet. Von primärem Interesse sind hier Parameterbereiche mit starkem externem Feld sowie starker Kopplung, die mit theoretischen Modellen nur unzureichend abgedeckt werden können.

Um das beiden Anwendungen zugrundeliegende N-Teilchen-Problem zu lösen wird ein paralleler Barnes-Hut Treecode genutzt und in Bezug auf Funktionalität, Vielseitigkeit sowie Skalierbarkeit deutlich erweitert. Mit seiner nun exzellenten Skalierbarkeit zu hunderttausenden von Prozessoren und simulierten Systemen aus Milliarden von Teilchen sowie der Unterstützung für periodische Randbedingungen einem effizienten und präzisen Echtraumansatz liefert er hoch aufgelöste Ergebnisse und ist auf weitere Anwendungen zum Warm-Dense-Matter Regime vorbereitet. Hier können seine einzigartigen Fähigkeiten genutzt werden um die Verbindung zu Experimenten in der realen Welt zu schlagen.

Contents

Lis	t of	Figure	S		xiii
Lis	t of	Tables			χv
Lis	t of	Algorit	thms		xvii
1		oductio			1
	1.1		amental plasma parameters		
	1.2		WELL's equations and the dielectric function		
	1.3		NTZ plasma and DRUDE model		
	1.4		lectron-ion collision frequency		
	1.5	Scope	of this work	٠	 19
2	The	N-bo	dy problem		21
	2.1	Molec	cular dynamics from basic principles		 21
	2.2	Nume	erical considerations – mathematics		 23
		2.2.1	Time discretization and general molecular dynamics		 23
		2.2.2	Leap-frog integration		 24
		2.2.3	Simulation timestep		 25
	2.3		ical constraints – computers		
	2.4	The N	V-body problem		 28
		2.4.1	Direct summation		
		2.4.2	Grid-based approaches		
		2.4.3	Multipole-based approaches		 30
3	The	e paralle	el Barnes-Hut tree code PEPC		37
	3.1		Barnes-Hut algorithm		 38
		3.1.1	Tree construction		
		3.1.2	Computation of multipole properties		
		3.1.3	Tree traversal and multipole acceptance criterion		
	3.2	The p	parallel tree algorithm		
		3.2.1	Allocation of globally unique keys		
		3.2.2	Hash table		
		3 2 3	Domain decomposition		40

	3.3 3.4 3.5	3.2.4 Parallel tree construction503.2.5 Parallel tree traversal50Performance analysis for the pure-MPI tree code PEPC50Other implementations and algorithm variations50Summary60	2 4 8
4	Peri	odic boundary conditions 6:	1
•	4.1	Periodic movement constraint	
	4.2	Periodic forces and potential	
		4.2.1 Next-neighbor periodicity	
		4.2.2 Nearest-image periodicity	3
		4.2.3 EWALD summation	3
		4.2.4 FMM-approach for periodic boundary conditions 6	5
		4.2.5 Renormalization approach for the lattice coefficients	0
		4.2.6 Near-field	6
		4.2.7 Dipole correction	6
	4.3	Verification of correctness and computational overhead	8
	4.4	Summary	1
5	Mul	ti-level parallelism 83	3
	5.1	Hybrid parallelization	3
		5.1.1 An MPI + Pthreads tree traversal	4
		5.1.2 Performance results for the new approach 90	0
		5.1.3 Intra- and inter-node load balancing	4
	5.2	Branch nodes	
		5.2.1 A-priori branch node estimation	6
		5.2.2 Hierarchical branch node clustering	
	5.3	Modularity, applications, and further prospects	
		5.3.1 Modularity and further applications	
		5.3.2 Projected technical optimizations	
		5.3.3 Further hybrid parallelism and task-based approaches 100	
		5.3.4 Additional directions of parallelization	
	5.4	Summary	2
6	Coll	ective electronic properties in nano clusters 109	
	6.1	Introduction	
	6.2	Numerical simulation setup	
	6.3	Total momentum autocorrelation function	
		6.3.1 Dynamical conductivity and its connection to simulation 109	
		6.3.2 Simulation results	
		6.3.3 Analysis and theory for the resonance shift	
		6.3.4 Conclusions	2

	6.4	Spatially resolved studies of the electronic resonances					
		6.4.1 Spatially resolved field spectra					 . 119
		6.4.2 Analysis of the spatially resolved spectra					 . 128
	6.5	Summary					 . 129
7	Coll	sion frequency in bulk plasmas					131
	7.1	The plasma heating rate					. 131
	7.2	Simulations on the heating rate in hydrogen plasmas					. 134
		7.2.1 Simulation setup					
		7.2.2 Kelbg potential					
		7.2.3 Simulation workflow					
	7.3	Method of measurement					
		7.3.1 Constant-temperature simulations					 . 141
		The Berendsen thermostat					
		The Nosé-Hoover thermostat					
		7.3.2 Instantaneous heating rate measurement					
	7.4	Simulation results					
		7.4.1 Scaling with the nonideality parameter					
		7.4.2 Dynamic behavior of the collision frequency					
		7.4.3 Dependence on field amplitude					
	7.5	Summary					
		~ annuary	•	·	•	•	. 101
8	Wra	oping things up: a conclusion					165
	8.1	Summary					 . 165
	8.2	Outlook					
Bi	bliog	aphy					173

List of Figures

1.1 1.2	Wavelength-intensity plane	
$\frac{2.1}{2.2}$	Expansion of a charge cloud's potential for a remote observer Particle-based N -body methods	33 34
	·	
3.1	Quad-tree construction for a two-dimensional example	39
3.2	Subcluster contributions to the parent cluster's multipole expansion	40
3.3	BARNES-HUT multipole acceptance criterion (MAC)	42
3.4	Precision comparison for different MAC parameters ϑ	45 48
3.5	Connection between keys of parent and children boxes	
3.6	MORTON Z- and HILBERT-spacefilling curve	50
3.7	Parallel tree construction	51 55
3.8	Detailed timings for the 3×10^7 -particle set of Figure 3.8	56
3.9 3.10	Execution trace of pure-MPI PEPC	57
3.10	Execution trace of pure-ivit 1 1 Ef C	51
4.1	Nearest-image periodicity	64
4.2	Notation for evaluating the far-field contribution	66
4.3	Near-field definition due to the criterion for absolute convergence	68
4.4	Regions in the derivation of the lattice coefficient algorithm	72
4.5	Precision comparison for a periodic system	79
4.6	Verification of energy conservation with periodic boundary conditions $$.	80
4.7	Runtime comparison for periodic and non-periodic setup $\ \ldots \ \ldots \ \ldots$	81
5.1	Hybrid parallelization uses available shared memory parallelism	84
5.2	Flowchart of the hybrid parallelized tree traversal in PEPC	86
5.3	Schematic flow of the hybrid parallelized PEPC	88
5.4	Execution trace of the hybrid parallelized PEPC	89
5.5	LLView screenshot: PEPC running on 294,912 processors	91
5.6	Strong scaling of the tree traversal and force computation	92
5.7	Total runtime for the homogeneous setup	93
5.8	Weak scaling of the full code with a homogeneous setup	93
5.9	Number of computed interactions for (un)balanced workload	95
5.10	Average number of branch nodes per MPI rank	97

5.11	Modularity of PEPC
	Combination of PEPC with PFASST
6.1	Icosahedral structure of the studied clusters
6.2	Evolution of the simulation for a cluster with $N_{\rm ion}=N_{\rm el}=2,869$ 108
6.3	Momentum autocorrelation spectra $K(\omega)$ for smaller nanoclusters 113
6.4	Momentum autocorrelation spectra $K(\omega)$ for larger nanoclusters 114
6.6	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=55.$ 121
6.7	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=147.$ 122
6.8	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=309.$ 123
6.9	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=923.$ 124
6.10	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=2,\!869.$ 125
6.11	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=28{,}741.$ 126
6.12	Spatially resolved observation for a cluster with $N_{\rm el}=N_{\rm ion}=293{,}789.$ 127
7.1	Comparison of Kelbg and Coulomb potential
7.2	Comparison of energy conservation for Kelbg and Coulomb potential . 138
7.3	Simulation scenario as used in this chapter
7.4	System energy per particle with Berendsen thermostat
7.5	System energy per particle with Nosé-Hoover thermostat $\dots \dots 150$
7.6	Demonstration of the instantaneous heating rate measurement 152
7.7	Collision frequency $\nu_{\rm ei}$ as function of the coupling parameter Γ 156
7.8	Comparison of instantaneous heating rate measurement to thermost at method
7.9	Frequency dependence fo ν_{ei} for $\Gamma = 2.5$, $\Theta = 1.18$
7.10	Frequency dependence to $\nu_{\rm ei}$ for $\Gamma=2.5,\Theta=1.18$
	Dependence of $\nu_{\rm ei}$ on quiver velocity and amplitude
1.11	Dependence of $\nu_{\rm ei}$ on quiver velocity and amplitude

List of Tables

3.1	Algorithmic scaling of major routines in PEPC		54
3.2	Overview on shared- and distributed-memory parallel tree codes $\ . \ .$.		59
6.1	Fit parameters for Figures 6.3 and 6.4		115
7.1	Densities and temperatures used in the simulations		135

List of Algorithms

2.1	Generic molecular dynamics scheme	24
	Serial tree traversal	
	Hybrid parallelized traversal: Worker-threads	

1 Introduction

" Wie verhalten sich die verschiedenen Körper gegenüber schnellen Schwingungen? [...] Von diesem Gesichtspunkt aus ergibt sich also ein enger Zusammenhang [der] Arbeiten über elektrische Schwingungen mit der optischen Untersuchung der Körper. [...] Der Gang [...], den hierbei die Optik gegangen ist, ist meines Erachtens typisch für alle Disziplinen, die [...] durch Untersuchung der inneren, mehr versteckten Eigenschaften der Körper den Zusammenhang mit Nachbardisziplinen und dadurch die Vereinheitlichung der Wissenschaft suchen. Und dabei ist eine wechselseitig Ergänzung von Experiment und Theorie unerläβlich. [1]

** How do objects behave under the influence of fast oscillations? [...] So, from this perspective a connection of the work on electrical oscillations to the optical investigation of physical objects is resulting. [...] In my opinion, the path [...] pursued by optics is typical for all disciplines that [...] are – by analyzing the inner, more hidden properties of objects – searching for connections to neighboring fields and thus a unification of science. In the course of this, mutual supplementation of experiment and theory is de rigueur. [1], translated from the German original.

Though not being explicitly labeled as such, the field of plasma physics and neighboring disciplines have captured a significant amount of interest of physicists for almost one hundred years now. Already at the beginning of the 20th century, the foundations for today's conception of science were laid. Around and after Albert Einstein's annus mirabilis, fundamental concepts for the structure and mathematical description of world's smallest scales (quantum mechanics) as well as the architecture of space and time itself (general relativity) were developed. However, besides these dominating advances uncountable other scientific contributions were made.

Since this exciting time, tremendous progress in experimental possibilities for studying the extremest states of matter in our world has been made. For example we are now able to penetrate optically opaque solid metal blocks with high-frequency and high-intensity radiation from novel light sources such as the free electron laser to study their inner structure [2]. Exciting a multitude of internal degrees of freedom, we generate hard x-rays and fast bunches of particles, discover surprising resonance phenomena and unexpected magnitudes of absorption and reflectivity. Thus, such experiments pose as many new questions as they give answers to existing problems.

Paul Drude was quite correct in stating that progress in experiment and theory are indispensably geared to each other. Progress in either field strongly depends on findings in the other. However, what he could not see at the beginning of the 20th century is a third tool for investigation of matter's inner structure, a tool that is widely used today. With the development of digital computers and their incredible progress during the last 60 years, it has become possible to *simulate* the behavior of physical objects at any scale. As perfectly controllable experiments, computer simulations have been rapidly established as the bridge between models developed in theory and data obtained in real-world experiments.

In computer experiments, virtually all parameters can be adjusted, models can be simplified or made more complex, and experiments can be repeated at any time. In that sense, they prove or falsify model images of the real world by applying a multitude of theoretical concepts at once. Naturally, computer simulations can neither replace pen-and-paper theory nor physical measurements of real objects supplement both and are de rigueur in science today.

In this thesis, we will make use of molecular dynamics simulations on highly parallel computers for examining collective phenomena in plasmas which are connected to experimentally accessible observables. The systems to be studied cover both nano clusters from a few dozen to hundred thousands of particles in vacuum as well as bulk matter with its surfaces at infinity. Particle densities and temperatures in our experiments will range from gaseous to solid state. In the latter, the particles are close enough to form long-range correlated systems that pose significant difficulties when approached with theoretical methods. Further details on the phenomena to be examined are given in Section 1.5 and the subsequent chapters. Hitherto, we will describe some fundamental concepts and parameters necessary for dealing with particle-based plasma physics and electrodynamics and establish the connection between microscopic theoretical models and the experimentally accessible world. We will see in the following, that the envisaged simulations on microscopic properties of many-particle systems pose a computationally challenging N-body problem. Its solution requires today's most powerful computers and extremely efficient algorithms. Thus, covering topics from physics and computer science, this work aims at bridging between theory and experiment, weak and strong coupling, nano scale and bulk matter by means of creating a new workhorse tool and applying it in two different regimes of warm dense matter.

1.1 Fundamental parameters for characterization of warm dense matter

The greek word πλάσμα ("structure") has been introduced by LANGMUIR in the description of his experiments with gas discharge tubes:

** [...] the ionized gas contains ions and electrons in about equal numbers so that the resultant space charge is very small. We shall use the name plasma to describe this region containing balanced charges of ions and electrons. **[3]

His definition of a state of matter involving free charges has essentially been kept until today. However, density and temperature ranges to which this description applies have been significantly broadened. From very thin interstellar matter to electrons in bulk metals and even the quark-gluon matter in nuclei contain free charges. They are evident inside ultra-cold superconductors as well as in the hot solar fusion zone. Research on the ubiquitous plasma state is performed in a multitude of physical disciplines from astronomic scales down to the vicinity of the smallest particles we currently can imagine.

Plasmas – or many-particle systems subject to COULOMB interaction in general – are essentially characterized by only a few statistical observables. First of all, the composition is of interest. Usually, electrons are the dominant charge carriers due to their small mass and hence large mobility in comparison to the ions, which are at least 1,800 times heavier. For the ions, every charge state of each chemical element that is present in the plasma can be considered as an additional particle species. In this chemical picture [4], conversion from one same-element species to another is possible via ionization or electron capture processes and can be described using rate equations and coupled mass-action laws, see e. g. [5–8]. Each species c, including electrons, has its own physical features, such as $mass\ m_c$ and $charge\ q_c$ as well as statistical observables. Among them are the $number\ density$

$$n_c = \frac{N_c}{V} \tag{1.1}$$

with the total number of particles N_c of species c and the system volume V; the temperature T_c that is, via the kinetic or thermal energy

$$E_{\text{therm},c} = \frac{3}{2} N_c k_{\text{B}} T_c = \sum_{i=1}^{N_c} \frac{m_c}{2} |\vec{v}_{c,i}|^2$$
(1.2)

with the velocity $\vec{v}_{c,i}$ of the *i*th particle of species c connected to the thermal velocity [9]

$$v_{\text{therm},c} = \sqrt{\frac{3k_{\text{B}}T_c}{m_c}} \tag{1.3}$$

where $k_{\rm B}$ is Boltzman's constant. The sum in (1.2) runs over all particles of species c. Considering only macroscopically neutral systems, we have

$$q_{\rm el}n_{\rm el} = \sum_{c \neq \rm el} q_c n_c , \qquad (1.4)$$

where the sum runs over all ion species present and the subscript $_{\rm el}$ denotes electronic observables. The species charge number Z_c is defined via

$$q_c = -Z_c \cdot q_{\rm el} \,, \tag{1.5}$$

an average ionic charge state can be given as

$$Z_{\text{avg}} = \frac{n_{\text{el}}}{\sum_{c \neq el} n_c} \,. \tag{1.6}$$

Neglecting ionization processes, this – in general fractional – average charge can be used to apply two-species treatments to situations where multiple ionic charge states are present. Besides these macroscopic and statistically motivated observables, phenomenological considerations stimulate the definition of a number of additional parameters. These provide information on the prevailing interaction regime – quantum or classical, weak or strong coupling, high or low frequency – that is actually dominant in the system.

The widely used degeneracy parameter

$$\Theta = \frac{k_{\rm B}T_{\rm el}}{E^{({\rm Fermi})}}$$
 with $E^{({\rm Fermi})} = \frac{\hbar^2}{2m_{\rm el}}(3\pi^2 n_{\rm el})^{2/3}$ (1.7)

is the ratio of electron thermal to Fermi energy and accordingly indicates the importance of degeneracy effects. If the thermal energy is much larger than $E^{(\text{Fermi})}$, i.e. $\Theta \gg 1$, a classical treatment is sufficient, whereas $\Theta \lesssim 1$ requires inclusion of quantum effects such as Pauli blocking and Heisenberg uncertainty.

An analogous definition holds for the coupling or nonideality parameter

$$\Gamma_c = \frac{E_c^{\text{(COULOMB)}}}{k_{\text{B}}T_c}$$
 with $E_c^{\text{(COULOMB)}} = \frac{q_c^2}{4\pi\varepsilon_0}\frac{1}{d_c}$, (1.8)

where the average interparticle distance

$$d_c = \left(\frac{4}{3}\pi n_c\right)^{-1/3} \tag{1.9}$$

is called Wigner-Seitz radius. Being the ratio of Coulomb to thermal energy, the value of Γ indicates the importance of interparticle correlations. For $\Gamma \ll 1$, species c

can be treated as an ideal gas. In the case of $\Gamma \gtrsim 1$, the species is approaching fluid and finally solid state. Then, the *structure factor* [10]

$$S_c(\vec{q}) = \left\langle \frac{1}{N_c} \sum_{j,k=1}^{N_c} e^{i\vec{q} \cdot (\vec{R}_j - \vec{R}_k)} \right\rangle$$
 (1.10)

with the time or (via ergodicity equivalent) ensemble average $\langle \cdot \rangle$ and the particle positions $\vec{R}_{1...N_c}$ cannot be set to unity but has to be used explicitly in theoretical derivations. With the exception of static regular lattices, $S_c(\vec{q})$ cannot be given analytically. This makes the extension of theories into the region of $\Gamma \gtrsim 1$ at best cumbersome and normally requires numerical computation of the relevant observables.

When examining the plasma systems we are studying here on a macroscopic scale, they appear to be neutral, i.e. not charged at all. On the microscopic scale, though, electrically charged particles – electrons and ions – are present. This apparent contradiction results from the balancing of charges when going beyond the scale of individual particles and considering larger volumes. Such a compensation effect is already present on a mesoscopic scale, where the electrons are moving freely in front of the ionic background. When looking onto an individual ion or any other system-internal charge from a sufficient distance, its charge is effectively screened by the surrounding electrons. The corresponding effective potential can be characterized with an additional exponential factor to yield the DEBYE or Yukawa potential [4, 11–13]

$$\Phi^{(\text{Debye})}(r) = \frac{q_c}{4\pi\varepsilon_0} \frac{1}{r} \cdot e^{-r/\lambda_D} . \qquad (1.11)$$

Formally, this is the solution to the linearized Poisson-Boltzmann equation which is derived by replacing the charge density in the Poisson equation (1.25) with the Boltzman-distributed particle density for finite temperature and Coulomb interaction [9]. The Debye-Hückel screening length

$$\lambda_{\rm D} = \left(\sum_{c} \frac{n_c q_c^2}{\varepsilon_0 k_{\rm B} T_c}\right)^{-1/2} \tag{1.12}$$

represents the distance where the effective (screened) potential has decreased by a factor of 1 /e with respect to the usual Coulomb interaction (1.27). Despite its analogous structure, the Yukawa potential is used for modeling interaction in particle physics with massive gauge bosons which is a completely different physical effect. There, $\lambda = \frac{2\pi\hbar}{mc}$ is the Compton wavelength of the virtual charge carrier. For photons that mediate the electromagnetic interaction, m=0 and the Yukawa potential reduces to the Coulomb interaction.

When driving a free electron gas to collective motion, the *plasma frequency* or LANG-MUIR frequency [3, 11, 14]

$$\omega_{\rm pl} = \sqrt{\frac{n_{\rm el}q_{\rm el}^2}{\varepsilon_0 m_{\rm el}}} \tag{1.13}$$

is its characteristic frequency if thermal motion is neglected. Since

$$\lambda_{\rm D} \cdot \omega_{\rm pl} = \sqrt{\frac{k_{\rm B} T_{\rm el}}{m_{\rm el}}} = v_{\rm therm} ,$$
 (1.14)

a thermal electron moves by one DEBYE length $\lambda_{\rm D}$ during one plasma period. Maximum energy transfer from an external wave with frequency $\omega_{\rm laser}$ to the electronic system results for resonant excitation, i. e. $\omega_{\rm laser} = \omega_{\rm pl}$. As this is only valid for free electrons, a multitude of effects that influence plasma oscillation can be studied. For example, due to interaction of the electrons with plasma ions, this resonant coupling suffers from damping and resonance shifts.

Furthermore, if the thermal movement of the electrons is included, the electron pressure contributes an additional restoring force which modifies the frequency of collective oscillation (1.13) to read

$$\omega_{\text{Bohm-Gross}}^2(k) = \omega_{\text{pl}}^2 + 3 k^2 v_{\text{therm,el}}^2$$
 (1.15)

This is the BOHM-GROSS dispersion relation [15]. On spatial scales larger than $\lambda_{\rm D}$, the electron pressure only plays a secondary role. However, for small wavelength scales the thermal velocity can be comparable to the phase velocity $v_{\rm ph} := \omega/k$. Electrons moving with $v_{\rm ph}$ move in phase with the plasma wave. Hence, they experience a stationary field that does not vanish in the cycle average and energy transfer between wave and electrons is possible. This effect is called Landau damping. In contrast to collisional damping, due to its reversibility this collisionless effect is not related to real dissipation and increase of entropy [11].

For nonrelativistic intensities, propagation of electromagnetic waves in a plasma is in general only possible for $\omega_{\text{laser}} > \omega_{\text{pl}}$. Going below the plasma frequency, the electrons are fast enough to follow the field oscillation. Then, the wave energy is converted to oscillatory electronic motion that itself leads to emission of Bremsstrahlung and finally total reflection of the incident wave. As $\omega_{\text{pl}} \propto \sqrt{n_{\text{el}}}$, there is a *critical electron density*

$$n_{\rm crit} = \frac{\varepsilon_0 m_{\rm el}}{q_{\rm el}^2} \cdot \omega_{\rm laser}^2 \,.$$
 (1.16)

For subcritical (underdense) plasmas $n_{\rm el} < n_{\rm crit}$, wave propagation is possible. For supercritical (overdense) plasmas with $n_{\rm el} > n_{\rm crit}$, it is prohibited. In density gradients, classical total reflection in the image of geometrical optics occurs at the $n_{\rm el} = n_{\rm crit}$

surface. Several further details on propagation and effects in under- and overdense systems can be found in [16].

Finally, though not being a plasma parameter, the strength of the external driver field has to be characterized. This is by convention done via the *quiver velocity*

$$v_{\rm osc} = \frac{q_{\rm el} E_{\rm L}}{m_{\rm el} \omega_{\rm laser}} \,, \tag{1.17}$$

the quiver amplitude

$$x_{\rm osc} = \frac{q_{\rm el} E_{\rm L}}{m_{\rm el} \omega_{\rm locor}^2} \,, \tag{1.18}$$

and the ponderomotive potential

$$\Phi_{\text{pond}} = \frac{q_{\text{el}} E_{\text{L}}^2}{4m_{\text{el}} \omega_{\text{locor}}^2} \,. \tag{1.19}$$

These are the velocity, amplitude and quiver energy acquired by a resonantly moving electron in an electromagnetic wave with amplitude (maximum field strength) $E_{\rm L}$ [16]. The corresponding intensity is

$$I_{\rm L} = \frac{\varepsilon_0 c E_{\rm L}^2}{2} \ . \tag{1.20}$$

The ratio between quiver and thermal energy separates the high- and low-intensity regimes. For $v_{\rm osc} \ll v_{\rm therm}$, the random thermal motion dominates; for $v_{\rm osc} \gg v_{\rm therm}$, the directed quiver movement is prevailing. In addition, as soon as the normalized quiver velocity

$$\tilde{v}_{\rm osc} = \frac{v_{\rm osc}}{c} \tag{1.21}$$

approaches unity, relativistic motion of the electrons has to be considered consistently. In contrast to $v_{\rm osc}$ and $\tilde{v}_{\rm osc}$ that specify the fundamental relevance of the external field, the ratio between quiver amplitude and average particle distance $d_{\rm ion}$, see (1.9), indicates the relevance of correlated collisions in microscopic considerations. For $x_{\rm osc} \gg d_{\rm ion}$, every electron will be scattered at several ions during a laser cycle which results in a randomization of their trajectories. In contrast, for $x_{\rm osc} \ll d_{\rm ion}$ the electrons effectively oscillate in front of an individual ion and thus experience multiple scattering events at the same center. This effect, which is referred to as *correlated collisions* can lead to enhanced particle deflection [17].

Figures 1.1 and 1.2 give an overview of the different regimes that are specified by these plasma parameters. The region of warm dense matter that lies in the primary focus of contemporary plasma physics, covers densities of $n_{\rm el} = 1 \times 10^{18} \dots 1 \times 10^{27} \, {\rm cm}^{-3}$, i.e.

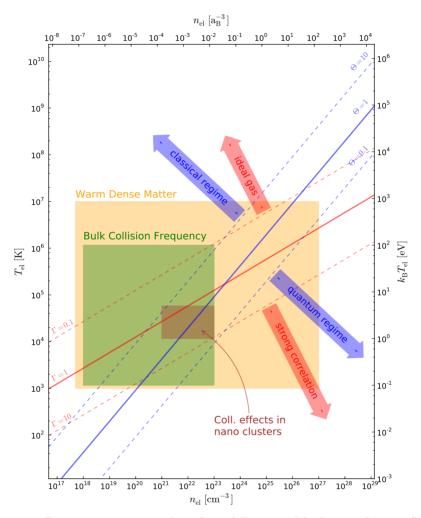


Figure 1.1: Density-temperature plane for a fully ionized hydrogen plasma. Several isolines for the degeneracy parameter Θ (1.7) and nonideality parameter Γ (1.8) as well as the regimes covered in this thesis are highlighted.

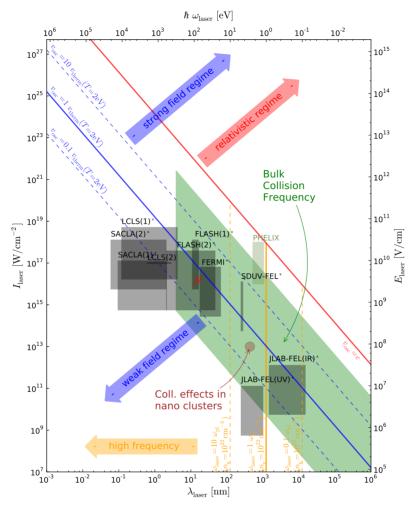


Figure 1.2: Wavelength-intensity plane with isolines $v_{\rm osc} = v_{\rm therm}$ for $T_{\rm el} = 2\,{\rm eV}$, the relativistic intensity $\tilde{v}_{\rm osc} = 1$, and $\omega = \omega_{\rm pl}$ for $n_{\rm el} = 1\times 10^{22}\,{\rm cm}^{-3}$. The intensities used for the applications in this work lie sufficiently below the relativistic threshold and cover weak and strong field regimes as well as the high- and low-frequency region. In addition, some prominent light sources are included to show today's experimental possibilities. The meaning of the abbreviations and a list of the depicted systems is given in the text. The region highlighted in red corresponds to the aluminum experiment by NAGLER et al. [2].

up to and beyond solid density, and temperatures of $k_{\rm B}T_{\rm el}=0.1\dots1,000\,{\rm eV}$. This embraces ideal and weakly coupled systems as well as strongly correlated quantum matter and hence applications for example in fields such as solar physics and the ICF (inertial confinement fusion). The phenomena to be studied in Chapters 6 and 7 are highlighted in the figures. With a degeneracy parameter $\Theta=0.1\dots10,000$, they are non-degenerate or weakly degenerate – a classical or pseudo-potential treatment is possible to include quantum effects in lower order. Furthermore, they are subject to nonideality parameters $\Gamma=0.01\dots50$ and thus cover regimes from ideal to strongly correlated matter. As discussed in Section 1.4, this broad range is an ideal application for molecular dynamics simulations as performed in this thesis. As visible in Figure 1.2 our applications comprise the non-relativistic regime $v_{\rm osc}\ll c$. Accordingly, relativistic effects can be neglected in the following.

The introductory example on absorption saturation in aluminum [2] is included in Figure 1.2 to highlight today's experimental possibilities on warm dense matter. In addition, Figure 1.2 shows regions of wavelengths and intensities that can currently be covered by Free Electron Laser (FEL) installations as representatives of the most powerful light sources available today. The abbreviations in the figure represent

- Free Electron Laser at DESY, Hamburg, Germany: FLASH(1) [18], FLASH(2) [19],
- FEL at Thomas Jefferson Lab, Newport News, USA, Infrared and Ultraviolet Beamline: JLAB-FEL(IR), JLAB-FEL(UV) [20],
- Linac Coherent Light Source at SLAC National Accelerator Laboratory, Menlo Park, CA, USA: LCLS(1) [21], LCLS(2) [22],
- SPring-8 Angstrom Compact Free Electron Laser, Riken/Harima, Japan: SACLA(1) [23], SACLA(2) [24],
- Free Electron Laser for Multidisciplinary Investigations, Trieste, Italy: FERMI [25–27],
- Shanghai Deep-Ultraviolet FEL, Shanghai, China: SDUV-FEL [28, 29].

In cases where only pulse powers instead of intensities are published, a focal spot diameter of $1\dots 15\,\mu\mathrm{m}$ are assumed for the data. These cases are marked with a star * in the figure. Up-to-date reviews on FEL technology and the current status of installations worldwide can be found in [30, 31]. The *Petawatt Hoch-Energie Laser für SchwerIoneneXperimente* (PHELIX, Darmstadt, Germany) [32] is included in the figure as one representative for conventional, optical high-power laser systems. Clearly, with the radiation from these installations it is now possible to probe the warm dense matter regime in laboratory experiments, providing fresh impetus to theory and modeling of this complex state.

1.2 Maxwell's equations and the dielectric function – connection to optical and transport properties

Classical electrodynamics, the theoretical framework for treating plasmas of electrically charged particles is described by MAXWELL's equations [33–35]

$$\operatorname{div} \vec{E} = \frac{\varrho}{\varepsilon_0} , \qquad \operatorname{div} \vec{B} = 0 , \qquad (1.22)$$

$$\operatorname{curl} \vec{E} = -\frac{\partial \vec{B}}{\partial t} , \qquad \operatorname{curl} \vec{B} = \mu_0 \vec{j} + \mu_0 \varepsilon_0 \frac{\partial \vec{E}}{\partial t}$$
 (1.23)

with the electric and magnetic fields $\vec{E} \equiv \vec{E}(\vec{r},t)$ and $\vec{B} \equiv \vec{B}(\vec{r},t)$, the charge and current densities $\varrho \equiv \varrho(\vec{r},t)$ and $\vec{j} \equiv \vec{j}(\vec{r},t)$, and the permittivity ε_0 and permeability μ_0 , respectively.

In the simplified electrostatic picture, all differentiations with respect to time – including the current density (particle velocity) – vanish and the induction law in (1.23) reduces to curl $\vec{E} \equiv 0$. This is equivalent to writing \vec{E} as the gradient of a scalar potential $\Phi \equiv \Phi(\vec{r},t)$

$$\vec{E} = -\operatorname{grad}\Phi. \tag{1.24}$$

With the Gauss law (1.22), this yields Poisson's equation

$$\Delta \Phi = -\frac{\varrho}{\varepsilon_0} \ . \tag{1.25}$$

The solution to (1.25) finally reads

$$\Phi(\vec{r}) = \frac{1}{4\pi\varepsilon_0} \int \frac{\varrho(\vec{r}')}{|\vec{r} - \vec{r}'|} d^3 \vec{r}' , \qquad (1.26)$$

which for a single isolated charge q at the origin with $\varrho(\vec{r}) = q \cdot \delta(\vec{r})$ yields [34, 35]

$$\Phi^{\text{(Coulomb)}}(\vec{r}) = \frac{1}{4\pi\varepsilon_0} \frac{q}{|\vec{r}|} . \tag{1.27}$$

This is the COULOMB potential.

Naturally, the assumption of electrostatics neglects the generation of magnetic fields via the induction law 1.23 and is becoming invalid for rapidly changing electric fields or fast moving particles where the time derivatives and currents cannot be assumed to be small.

Being exposed to radiation fields, the matter under examination is not only passively influenced, but modifies the electromagnetic wave itself. For example, the dispersion relation inside any medium strongly depends on its microscopic properties. Experimentally accessible observables are for example the frequency and wave number dependent

- absorption coefficient $\alpha(\vec{k},\omega)$ and emission coefficient $j(\vec{k},\omega)$,
- index of refraction $n(\vec{k},\omega)$,
- reflectivity R(ω),
- conductivity $\sigma(\vec{k},\omega)$.

These observables are all closely connected to the dielectric function $\varepsilon(\vec{k},\omega)$ that is introduced when transferring MAXWELL's equation into a medium. Assuming harmonic time dependence for the magnetic field $\vec{H} \equiv \vec{H}(\vec{k},\omega)$ and the electric field $\vec{E} \equiv \vec{E}(\vec{k},\omega)$, they read in wave number/frequency space [34, 35]

$$i\omega\mu \vec{H} = c \operatorname{curl} \vec{E}$$
 (1.28)

$$i\omega\varepsilon\vec{E} = -c\operatorname{curl}\vec{H}$$
, (1.29)

where the dielectric function $\varepsilon \equiv \varepsilon(\vec{k},\omega)$ and the magnetic susceptibility $\mu \equiv \mu(\vec{k},\omega)$ are in general also functions of wave vector \vec{k} and frequency ω . Neglecting the magnetic response, i.e. setting $\mu \equiv 1$ and after eliminating the magnetic field, the wave equation for \vec{E} reads

$$\left(\Delta + \varepsilon(\vec{k},\omega)\frac{\omega^2}{c^2}\right)\vec{E}(\vec{k},\omega) = 0, \qquad (1.30)$$

that for harmonic plane waves with amplitude $E_{\rm L}$

$$\vec{E}(\vec{k},\omega) = E_{\rm L} \cdot e^{i(\vec{k}\cdot\vec{r}-\omega t)}$$
(1.31)

yields the dispersion relation

$$k^2 = \varepsilon(\vec{k}, \omega) \frac{\omega^2}{c^2} \,. \tag{1.32}$$

Thus, the wave vector \vec{k} is in general a complex quantity $\vec{k} = \text{Re}\{\vec{k}\} + i \,\text{Im}\{\vec{k}\}$ and it is expedient to split the space-dependent part of the phase factor

$$e^{i\vec{k}\cdot\vec{r}} = e^{i\operatorname{Re}\{\vec{k}\}\cdot\vec{r}} + e^{-\operatorname{Im}\{\vec{k}\}\cdot\vec{r}}$$
 (1.33)

Clearly, the imaginary part of \vec{k} describes damping of the wave due to presence of the medium. Utilizing a convention used in classical electrodynamics, the wave number can be written as [34, 35]

$$k = \left(n(k,\omega) + i\frac{c}{2\omega}\alpha(k,\omega)\right)\frac{\omega}{c} = \sqrt{\varepsilon(k,\omega)}\frac{\omega}{c}$$
 (1.34)

with the index of refraction n and absorption coefficient α . Expressed in terms of $\text{Re}\{\varepsilon\}$ and $\text{Im}\{\varepsilon\}$ they read

$$\alpha = -\frac{\omega}{c \, n} \text{Im}\{\varepsilon\} \tag{1.35}$$

$$n = \frac{1}{\sqrt{2}} \sqrt{\text{Re}\{\varepsilon\} + |\varepsilon|} \ . \tag{1.36}$$

Governing wave propagation in media, the dielectric function is also related to reflectivity. For example, when studying a step-like transition from vacuum ($\varepsilon \equiv 1$) into a medium with nontrivial dielectric function $\varepsilon(\omega)$, the ratio of reflected to incident intensity is given by Fresnel's formula [36]

$$\mathcal{R}(\omega) = \left| \frac{\sqrt{\varepsilon(\omega)} - 1}{\sqrt{\varepsilon(\omega)} + 1} \right|^2 \tag{1.37}$$

for normal incidence, i.e. \vec{k} perpendicular to the surface.

Finally, the dynamic (ac-)conductivity $\sigma(\omega)$ is also connected to the dielectric function. In the long-wavelength limit $\varepsilon(\omega) = \lim_{k\to 0} \varepsilon(k,\omega)$ this relation reads [36]

$$\sigma(\omega) = i\omega \varepsilon_0 (1 - \varepsilon(\omega)) . \tag{1.38}$$

The static conductivity $\sigma^{(dc)} = \lim_{\omega \to 0} \sigma(\omega)$ is also an important observable in experimental studies.

As we have seen, knowledge about the dielectric function $\varepsilon(k,\omega)$ is essential to describe the optical observables. Accordingly, a first-principles theory for the dielectric function is necessary to describe interaction of light with matter. The fundamental concepts presented here will be supplemented by such a microscopic theory for the dielectric function ε and the electric conductivity $\sigma(k,\omega)$ in the next Section 1.3.

1.3 Lorentz plasma and Drude model

Already directly after publication of the plum-pudding model of the atom by Thomson [37] in 1904 and only few years after the actual discovery of the electron in 1897 [38], LORENTZ published his first-principles model on microscopic properties of metals and dielectrics [39–41] during Einstein's annus mirabilis 1905. It is based on the idea of electrons that are bound to their atom by a linear restoring force. They are allowed to perform damped harmonic oscillations around their position of rest. The equation of motion for the electron's displacement \vec{r} reads

$$m\frac{\mathrm{d}^{2}}{\mathrm{d}t^{2}}\vec{r} = \underbrace{q_{\mathrm{el}} \cdot \vec{E}(t)}_{\text{Force from external field}} - \underbrace{\frac{m}{\tau} \frac{\mathrm{d}}{\mathrm{d}t}\vec{r}}_{\text{Relaxation due to friction}} - \underbrace{m\omega_{0}^{2}\vec{r}}_{\text{Harmonic restoring force}}. \tag{1.39}$$

Despite its simplicity, the model already allows for a fundamental microscopic understanding of optical features of metals and dielectrics, e. g. dispersion, absorption, optical

line width. An instructive analysis is given in [42]. Furthermore, the model gives intuitive explanations for results from rigorous quantum mechanical analyses that will not be discussed in further detail here.

Concentrating on plasma properties that result from the presence of free charge carriers, we omit the restoring force by setting $m\omega_0^2\vec{r}\equiv 0$. With the relaxation time τ , this leads to the DRUDE model [43–45]. For the electron momentum $\vec{p}=m\cdot\frac{\mathrm{d}}{\mathrm{d}t}\vec{r}$ it reads

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{p} = q_{\mathrm{el}} \cdot \vec{E}(t) - \frac{1}{\tau} \cdot \vec{p} \,. \tag{1.40}$$

Assuming a harmonic external field $\vec{E}(t) = \vec{E} \cdot e^{-i\omega t}$ and a harmonic isotropic momentum response $\vec{p}(t) = \vec{p} \cdot e^{-i\omega t}$, the equation of motion (1.40) yields

$$\vec{p} = \frac{1}{-i\omega + \nu} \vec{E} \,, \tag{1.41}$$

where $\nu := 1/\tau$ has been defined. Being only a shorthand notation for now, the relevance and physical interpretation of ν will be discussed in Section 1.4. With the current density

$$\vec{j} = q_{\rm el} \cdot n \frac{p}{m} \tag{1.42}$$

and plasma frequency (1.13), equation (1.41) reads

$$\vec{j} = \frac{\varepsilon_0 \omega_{\rm pl}^2}{-i\omega + \nu} \vec{E} \ . \tag{1.43}$$

Thus, the electric conductivity σ , that – in the sense of a susceptibility – connects internal currents to the external fields via a simple OHM's law

$$\vec{j} = \sigma \cdot \vec{E} \tag{1.44}$$

is defined by

$$\sigma(\omega) = \frac{\varepsilon_0 \omega_{\rm pl}^2}{-i\omega + \nu} \tag{1.45}$$

in the DRUDE model. With this approach, good explanations for ac- and dc- conductivity σ in metals, the electron contribution to the thermal conductivity κ and the Hall effect could already be given at the beginning of the 20th century. Also, the Wiedemann-Franz law that postulates proportionality of the ratio of thermal conductivity κ to electric conductivity σ with temperature $\kappa/\sigma \propto T$ [46] could be explained.

Accordingly, this theoretical foundation and derived theories are relevant for a much broader field than just optical properties. With appropriate extensions, they can also

be related to plasma transport features, such as the thermal conductivity, questions of relaxation of non-equilibrium systems, dynamics of magnetized plasmas, the equation of state, etc. However, concentrating on collective electronic phenomena in this thesis, we do not branch too deeply, into these fields.

Derived from the assumption that all free electrons contribute to the charge transport and neglecting quantum mechanical effects such as PAULI blocking and HEISENBERG uncertainty as well as bound states, the DRUDE model cannot describe semiconductors and dielectrics correctly and overestimates the dc-conductivity. Only with later enhancements by SOMMERFELD and BETHE [47] and the band model [10], has this become possible. However, since this thesis primarily covers classical plasmas, i.e. the non-degenerate electron gas in front of the ionic background, we will not discuss these extensions here.

The only free parameter in (1.45) is the relaxation time $\tau = 1/\nu$ which is the damping time parameter in (1.41). In the following section, we study collisions of the charged particles as one of the main causes for their damped motion in plasmas.

1.4 The electron-ion collision frequency

Collisions between the charges will be the primary contribution to a loss of momentum and hence friction in a plasma. Consequently, it seems natural to identify $\nu=1/\tau$ as a collision frequency. Using Kirchhoff's rule, different physical effects that lead to damping or relaxation, e.g. electron-electron and electron-ion collisions, can be collected additively here: $\nu=\nu_{\rm ei}+\nu_{\rm ee}$. In normal gases which are composed primarily of neutral particles, close encounters between them are dominant. In contrast, the slowly decreasing nature of Coulomb interaction already leads to a considerable number of distant collisions for charged particles and results in strong deflection. They outweigh the importance of close encounters by far [48]. Normally, every particle in a plasma experiences a large number of such far collisions at every instant in time. If all these led to deflection into the same direction, these would have an enormous effect on all trajectories. In reality, due to the random character of the collisions their directions will cancel out. However, they lead to a transfer of momentum and energy between the collision partners that significantly influences the macroscopic transport and optical properties.

The problem of determining ν in a broad parameter range of density and temperature with inclusion of different physical effects has kept scientists busy already since publication of the model. In the collision term of the Boltzmann equation, L. D. Landau applies the Rutherford scattering cross section and an effective collision frequency

that depends on particle velocity v [11, 49]

$$\operatorname{St} f = -\nu_{ei}(v)\delta f \tag{1.46}$$

with

$$\nu_{\rm ei}(v) = n_{\rm ion} \cdot v \cdot \sigma_{\rm ei}^{\rm (RUTHERFORD)} = \frac{4\pi Z e^4 n_{\rm el}}{m_{\rm el}^2 v^3} \mathcal{L}$$
 (1.47)

and the so-called COULOMB logarithm

$$\mathcal{L} = \int_{\chi_{\min}}^{\chi_{\max}} \frac{\mathrm{d}\chi}{\chi} = \ln \frac{\chi_{\max}}{\chi_{\min}}$$
 (1.48)

which is an integral from minimal to maximal deflection angle χ . The integration limits can be defined intuitively via the maximum and minimum collision parameter. Here, Heisenberg's uncertainty and Pauli blocking have to be mentioned for near and Debye screening for distant encounters. Using similar arguments, Spitzer gives the dc-conductivity as [48, 50]

$$\sigma^{(dc)} = \frac{4\sqrt{\pi}}{\pi^{3/2}} \frac{(4\pi\varepsilon_0)^2 (k_B T)^{3/2}}{\sqrt{m_{el}} e^2} \frac{n_{el}}{n_{ion}} \cdot \mathcal{L}^{(SPITZER)^{-1}}, \qquad (1.49)$$

$$\mathcal{L}^{(\text{SPITZER})} = \frac{1}{2} \ln \frac{3}{2} \Gamma^{-3} \ .$$
 (1.50)

In this model,

$$\tau = \tau_{90^{\circ}} = \frac{m^2 v^3}{8\pi n_{\text{ion}} Z^2 e^4 \mathcal{L}^{\text{(Spitzer)}}}$$
(1.51)

is defined as the cumulative 90° deflection time [48, 51]. When including screening effects, the collision integral can also be given with better-than-logarithmic precision [11, 52, 53]. This dynamically screened approach, that avoids divergences for large impact parameters, is called Lenard-Balescu approximation. Including higher-order self-energy diagrams, Gould and DeWitt give a collision integral that also avoids the divergence for small impact parameters and treats close collisions in a consistent fashion [54].

PEREL and ELIASHBERG apply a diagram technique to derive an expression for the collision frequency in the high-frequency and low-intensity limit [55]. A consistent treatment of ac-conductivity for frequencies embracing $\omega_{\rm pl}$ is given by DAWSON and OBERMAN [56]. They find an enhanced collision frequency right above $\omega = \omega_{\rm pl}$, which is attributed to excitation of longitudinal plasma modes. Finally, RAND [57] and SILIN [58] extend the kinetic approach under the assumption of a MAXWELLian distribution to arbitrary – in particular also strong – high-frequency fields and find the asymptotic

behavior $\nu_{\rm ei}(v_{\rm osc} \to \infty) \propto v_{\rm osc}^{-3}$. A derivation of these results based on a Lorentz collision operator is given by Catto and Speziale [59]. Pert [60], Schlessinger and Wright [61], and Silin and Uryupin [62] give extensions to the strong-field regime with inclusion of quantum corrections. It is later shown by Decker et al., that Silin's low-intensity limit $v_{\rm osc} \ll v_{\rm therm}$ corresponds to the results of Dawson and Oberman [17]. Furthermore, using particle-in-cell simulations, Decker et al. indicate an — in contrast to theory — enhanced collision frequency for $v_{\rm osc} \gg v_{\rm therm}$ and $x_{\rm osc} < \lambda_{\rm D}$. This is interpreted as the result of correlated collision from multiple deflections at the same scattering center. We will return to this picture in Chapter 7 when analyzing our simulation results on bulk collision frequency. A similar increase of the collision frequency due to correlated collisions is found by Brantov et al. using test-particle simulations [63] for beam-like particle configurations and anisotropic velocity distributions.

A reduced Inverse Bremsstrahlung absorption is found by Langdon [64] as the result of a non-Maxwellian electron velocity distribution function already for $Zv_{\rm osc}^2 \gtrsim v_{\rm therm}^2$. For small electron temperatures, electron-ion collisions are preferred to those with high electron velocities. This prevents the velocity distribution function from relaxing into its equilibrium shape and finally suppresses collisions. The effect is already present before nonlinearities as studied by Silin become evident for $v_{\rm osc} \gg v_{\rm therm}$. The results of Langdon are successfully verified by Jones and Lee using Monte Carlo simulations [65]. Using the Langdon distribution function, Pfaline [66] finds – in contrast to Silin and Catto – an increase of the collision frequency below $v_{\rm osc}/v_{\rm therm} \approx 0.1$. For higher intensities, $v_{\rm ei}$ decreases for $v_{\rm osc} \gg v_{\rm therm}$ as in the other theories.

Using a ballistic model, MULSER and SAEMANN study the electron-ion collision frequency in laser-plasma interaction and stress that for $\nu_{\rm ei} \approx \omega_{\rm laser}$ the widely used laser cycle average is not applicable and a time-dependent theory of collisional absorption has to replace the time-averaging [67]. Comparisons of different approaches to the collision frequency are given in [68].

Numerical results for moderate and strongly coupled systems $\Gamma \gtrsim 1$ were first derived by Cauble and Rozmus [69]. They include spatial distribution functions from hypernetted chain (HNC) calculations [70, 71] to model the spatial structure of strongly coupled systems into the theory.

Founded on a rigorous generalized linear response theory for the non-equilibrium many-particle system, RÖPKE, WIERLING, REINHOLZ, et al. [72–77] propose a quantum statistical approach for the electron-ion collision frequency. Electron-electron collisions can be accounted for by including higher moments in terms of a renormalization factor [77]. Interpolation expressions for the static conductivity are also derived [78–80]. The quantum statistical approach can be shown to lead to a generalized DRUDE expression where the collision frequency itself becomes a frequency-dependent complex quantity $\nu_{\rm ei} \equiv \nu_{\rm ei}(\omega) \in \mathbb{C}$, see [81] for details, [82–84] for an extensive derivation, and

[85, 86] for the algebraic connection and numerical comparison to the originally used kinetic description. With this modification the – now frequency-dependent – collision frequency can be shown to fulfill the usual sum rules for the dielectric function. In this framework it is calculated with inclusion of dynamic screening and strong collisions via a statically screened T-matrix summation using a diagram technique [4, 87, 88] through a GOULD-DEWITT approach [89, 90]. Comparisons of the quantum statistical approach for different pseudopotentials [83, 91, 92] and with molecular dynamics simulations for the collision frequency and the structure factor [93–95] are available. An approach of including dynamical screening into the full T-Matrix by BORNATH et al. shows algebraic difficulties and is restricted to first order in the screened potential [96]. MOLL et al. include a second-order BORN approximation into the theory and find a plateau in the Inverse Bremsstrahlung heating rate when considering its dependence on the DEBYE screening length $\lambda_{\rm D}$ for weak fields [97]. This is absent in the case of strong external fields.

Theories for partially ionized plasmas with inclusion of bound-free transitions are also available [7, 98, 99].

A quantum kinetic approach that is also valid for strong fields is given by KREMP, BORNATH, BONITZ, SCHLANGES, and HILSE [100, 101] and shown to be in compliance with molecular dynamics simulations [102, 103], e. g. by PFALZNER and GIBBON [104]. GRINENKO and GERICKE split contributions of weak and strong collisions [105]. Exploiting similarities of the strong-collision term to stopping-power calculations, they find agreement with analytical and simulation results of HILSE et al. [103]. Additionally including an ion-ion structure factor from HNC calculations, SCHLANGES et al. show an enhanced electron-ion collision frequency for strong coupling $\Gamma \gtrsim 3$ [106, 107].

Based on a quantum-mechanical dielectric model in the framework of a quantum-VLASOV theory, KULL and PLAGNE [108] provide an expression for the strong-field scaling of the COULOMB logarithm and find agreement with the findings of SCHLANGES et al. as well as a connection to the classically obtained results of DAWSON and DECKER. Furthermore, applying wave-packet scattering via solution of the time-dependent SCHRÖDINGER equation with inclusion of multi-photon effects, KULL et al. also find enhanced scattering rates due to correlated collisions in the strong-field regime [109–111]. The obtained energy spectra are found to be comparable to those of above-threshold ionization.

Experimental data for validation of the many different models and approaches is still patchy. For example, Kim et al. perform measurements of dc-conductivity via THz pump-probe reflectivity experiment on aluminum [112]. A similar technique is applied by Widmann et al. who study the ac-conductivity of gold by means of reflectivity measurements [113]. Measurements on reflectivity of shock-compressed xenon could be fitted to theory by Reinholz, Zaporozhets, et al. through modifications to the anticipated shock front density profile [84, 114–122]. The resistivity of a titanium/aluminum

vapor discharge is evaluated by Renaudin et al. [123]. Mostovych et al. report data on Inverse Bremsstrahlung absorption and the Coulomb logarithm for strongly coupled, optically thin aluminum plasmas [124]. This overview gives an only very short and incomplete excerpt from the current state of experiments. More detailed reviews and extensive data tables were for example compiled by Benage [125] and Clérouin et al. [126].

Most of the theories listed above are limited to regimes of either weak coupling ($\Gamma \lesssim 1$) or weak external fields ($v_{\rm osc} \ll v_{\rm therm}$) for hydrogen-like matter. The experimental data, though, concentrates on higher-Z elements, such as aluminum, titanium, xenon, gold – mostly in bulk state, i.e. for high Γ . Despite first attempts, such as [127], these still cannot be described by closed theories. This is also due to the fact that experimental conditions in the high- Γ regime are difficult to control and cannot be easily matched by theory.

These restrictions can be overcome to a certain extent using molecular dynamics simulations, which are in principle not limited to a specific range of the coupling parameter or field strength. However, published results on molecular dynamics simulations for the collision frequency in correlated media are scarce. Often, the data obtained by PFALZNER and Gibbon [104] via classical molecular dynamics with a soft core Coulomb potential is used as reference. In addition, also Hilse et al. [103], Reinholz et al. [94] and Mo-ROZOV et al. [95] report on own simulation results for comparison with their theories. They include quantum effects into their simulations by using pseudo-potentials instead of the pure Coulomb interaction. As an alternative, Filinov, Bonitz, et al. propose path integral Monte Carlo methods to consistently treat uncertainty effects [128–130]. Another approach for inclusion of quantum effects into computer simulations is the wave packet molecular dynamics method as for example used by ZWICKNAGEL and Pschiwul |131|. Classical pseudo-potential molecular dynamics simulations were also performed by RAITZA et al. for studying collective effects in finite nano-plasmas [132– 140. In particular, they observe plasmon resonance shifts and a complex mode structure inside metallic nanoclusters.

1.5 Scope of this work

This work aims at gaining new insight into the collisional behavior of warm dense bulk matter and on collective effects in nano-plasmas. The density, temperature, laser wavelength and intensity regions relevant to this work are indicated in Figures 1.1 and 1.2 as green and brown areas.

In particular, simulation experiments given by RAITZA et al. in [134–136] on electronic resonances in metallic nano-clusters were limited to systems with up to $N_{\rm el}=N_{\rm ion}=$

1,000 electrons and ions. In Chapter 6, we extend these studies – with respect to particle number – to 300 times larger systems. While confirming the previous results for small systems, we find a much more complex resonance structure for the larger clusters. This is analyzed by means of a new spatially resolved diagnostic. Being able to step towards these large systems, we start closing the gap between nano-scale and bulk matter.

Following the work of Pfalzner [104] and Hilse [103], the bulk collision frequency is evaluated using classical molecular dynamics simulations for a hydrogen-like plasma in Chapter 7. Quantum effects are accounted for via inclusion of the Kelbg pseudopotential. Primarily, we are interested in the combined high-field and strong coupling, regime – a scenario which is well within reach of present day FEL facilities, but which has only scarcely been covered theoretically.

To solve the underlying N-body problem for both applications, we utilize the highly-parallel Barnes-Hut tree code that has been previously developed at Jülich Super-computing Centre, but which was significantly extended with respect to functionality, versatility, and scalability in the course of preparing this thesis – Chapters 4 and 5. With the new developments presented here, the simulation code now offers

- excellent scaling on hundred thousands of processors,
- the possibility to simulate systems with $> 10^9$ particles,
- the option to virtually extend the simulation region to infinity by means of accurate periodic boundary conditions.

These features are prerequisites for the envisaged applications. To align these developments and their implications with large-scale particle simulations, Chapters 2 and 3 summarize the N-body problem, several algorithms for dealing with it and finally concentrate on multipole approaches as the methods of choice in this thesis.

2 The N-body problem

Besides the real-world experiment, the most straightforward approach for studying the dynamical behavior of physical systems is their simulation, e.g. on computer systems. Here, using a direct one-to-one correspondence of physical to simulation objects, their evolution in time and space can be studied by immediate inspection of bodies or phenomena in the computer-generated data.

While this idea of mapping the real world and all its details to simulation objects would be the most realistic and hence most favorable approach, it is in general not feasible due to the actual complexity of nature itself. Instead, a mathematical model has to be constructed, that reduces reality to manageable systems. Details that are not necessary for describing the effects to be studied are omitted. Others are simplified by casting them into mathematical expressions that can finally be examined and evaluated in a closed form. In the first part of this chapter, we will describe the fundamental concepts, principles, and limits that are foundations for every particle-based computer simulation experiment. Subsequently, different approaches for solving the N-body problem arising in such experiments are discussed. Concentrating on multipole algorithms, we finally lay the foundations for the simulations performed within this work.

2.1 Molecular dynamics from basic principles

Within this work, the dynamic properties of COULOMB plasmas are examined. These consist of differently charged particles in gaseous, fluid or bulk state. They can be treated in computer simulations by means of either a macroscopic continuum-mechanical approach using hydrodynamics codes or with a microscopic picture on the level of individual (pseudo-) particles.

Since it is intended to study the effects of individual COULOMB collisions on optical and transport properties in strongly correlated media here, the latter method has been chosen. It deals with particle collisions in a direct, *ab initio* fashion without the need for any *ad hoc* approximations or reliance on theory for computing transport properties. Furthermore, the individual particles that can also be interpreted as LAGRANGian discretization nodes of the real particle density. Thus, they are able to reproduce arbitrary density contrasts. Accordingly, simulating a set of colliding particles surrounded

by vacuum is much easier with particle methods than with the continuum-mechanical hydrodynamics method.

As stated in Section 1.1, we will concentrate on the non-relativistic case, where the general Hamilton (total energy) function for a system of N particles can be written as [141]

$$\mathcal{H}(\{\vec{p}_{1}(t)\dots\vec{p}_{N}(t)\},\{\vec{q}_{1}(t)\dots\vec{q}_{N}(t)\},t) = T(\{\vec{p}_{1}(t)\dots\vec{p}_{N}(t)\}) + U(\{\vec{q}_{1}(t)\dots\vec{q}_{N}(t)\},t) = \sum_{i=1}^{N} \frac{|\vec{p}_{i}(t)|^{2}}{2m_{i}(t)} + \frac{1}{2}\sum_{i=1}^{N} Q_{i} \cdot \Phi(\vec{q}_{i},\{\vec{q}_{1}(t)\dots\vec{q}_{N}(t)\},t) . \quad (2.1)$$

With the canonical particle coordinates $\vec{q}_i(t)$ and momenta $\vec{p}_i(t)$, their charge Q_i , the kinetic energy T and the system's potential energy U, the dynamics of a physical system can be deduced by evaluating the Hamilton equations. They connect momenta and positions via

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{q}_i(t) = \frac{\partial \mathcal{H}}{\partial \vec{p}_i(t)} = \frac{\vec{p}_i(t)}{m_i(t)}$$
(2.2)

and

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{p}_{i}(t) = -\frac{\partial \mathcal{H}}{\partial \vec{q}_{i}(t)} = -Q_{i}\nabla_{\vec{q}}\Phi\left(\vec{q}, \{\vec{q}_{1}(t)\dots\vec{q}_{N}(t)\}, t\right)\Big|_{\vec{q}=\vec{q}_{i}(t)}$$

$$\equiv \vec{F}\left(\vec{q}_{i}(t), \{\vec{q}_{1}(t)\dots\vec{q}_{N}(t)\}, t\right) . \tag{2.3}$$

For readability the explicit time dependence of particle coordinates and momenta will be omitted in the following. Furthermore, the particle masses are assumed to be constant over time, i.e. $\frac{d}{dt}m_i(t) \equiv 0$. In addition, the canonical coordinates $\vec{q_i}$ are identified as the usual spatial positions $\vec{r_i}$, and charges Q_i are printed as q_i . Then, (2.2) and (2.3) can be recast into Newton's law

$$m_i \cdot \vec{a}_i = \vec{F}(\vec{r}_i, \{\vec{r}_1 \dots \vec{r}_N\}, t) = -q_i \cdot \nabla_{\vec{r}} \Phi(\{\vec{r}_1 \dots \vec{r}_N\}, t) \Big|_{\vec{r} = \vec{r}_i(t)}$$
 (2.4)

with

$$\vec{a}_i = \frac{\mathrm{d}^2}{\mathrm{d}t^2} \vec{r}_i \,, \tag{2.5}$$

which is the basis for any classical molecular dynamics (MD) simulation. Provided, the forces $\vec{F}(\vec{r_i}, \{\vec{r_1} \dots \vec{r_N}\}, t)$ on each particle can be determined, (2.5) can be used with a simple time discretization scheme to compute the individual particle trajectories as outlined in Section 2.2. However, the set of N coupled second-order partial differential

equations (2.4) and (2.5) that governs the dynamics of the complete system already points to several shortcomings of all microscopic, particle-based simulation techniques.

First of all for macroscopic objects with $N \propto 10^{23}$ particles, the amount of necessary storage for particle coordinates, velocities, and other properties vastly exceeds the capabilities of any available computer system. This will be discussed in Section 2.3.

Furthermore, since the forces $\vec{F}_i \equiv \vec{F} (\vec{r}_i, \{\vec{r}_1 \dots \vec{r}_N\}, t)$ in general depend on all coordinates, their determination is a global problem as soon as the interaction law has a long-range character. Then, for every particle all other particles have to be considered during the force computation and the amount of work required grows quadratically with the size of the simulated system. Without significant technical advances, though, computing power can at most grow linearly with the amount of resources invested. Hence, there is a point where going to even larger computers would not allow us to noticeably increase system size. Then, building more powerful computers would not make any sense from the economical point of view. This very fundamental problem and some solution approaches will be topic of Section 2.4.

2.2 Numerical considerations – mathematics

2.2.1 Time discretization and general molecular dynamics

As already mentioned, solving the set of N coupled second-order differential equations (2.5) does not pose a fundamental problem as soon as the accelerations, i.e. the forces, are known. Usually, the particle velocity \vec{v}_i is introduced as an auxiliary variable to yield

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{r}_{i} = \vec{v}_{i} ,$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{v}_{i} = \frac{\vec{F}\left(\vec{r}_{i}, \{\vec{r}_{1} \dots \vec{r}_{N}\}\right)}{m_{i}} ,$$
(2.6)

where the explicit time dependence has been dropped. This set of 2N coupled differential equations of first order can be solved by discretization of the time variable. Therefore, differentiation with respect to time is replaced by a difference quotient

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) \approx \frac{x(t) - x(t - \Delta t)}{\Delta t} \tag{2.7}$$

and after some reordering (2.6) reads

$$\vec{r}_i(t) = \vec{r}_i(t - \Delta t) + \vec{v}_i(t) \cdot \Delta t , \qquad (2.8)$$

$$\vec{v}_i(t) = \vec{v}_i(t - \Delta t) + \frac{\vec{F}_i(t)}{m_i} \cdot \Delta t , \qquad (2.9)$$

Algorithm 2.1 Generic molecular dynamics scheme

```
t \leftarrow 0
\mathbf{call} \ \ \mathrm{Initialize\_Particles}(\vec{r}_{1...N}, \vec{v}_{1...N})
\mathbf{while} \ t < t_{\mathrm{end}} \ \mathbf{do}
t \leftarrow t + \Delta t
\mathbf{call} \ \ \mathrm{Compute\_Forces}(\vec{F}_{1...N})
\mathbf{call} \ \ \mathrm{Update\_Velocities}(\vec{v}_{1...N})
\mathbf{call} \ \ \mathrm{Diagnostics}()
\mathbf{call} \ \ \mathrm{Update\_Positions}(\vec{r}_{1...N})
\mathbf{end} \ \ \mathbf{while}
```

which is the EULER integration with the discrete timestep Δt . Actually, when applying (2.7), all positions, velocities and the time variable would have to be renamed since we are stepping from a continuous exact solution of (2.6) to a discrete and only approximatively correct parameter and variable space. However, as the distinction between discrete and continuous variables is usually obvious from the context – and in most cases even would not lead to formal differences – we do not use different notations here.

Again, it has to be emphasized, that (2.8) and (2.9) can be evaluated conveniently in an iterative scheme as soon as the forces $\vec{F_i}(t)$ are known or can be computed rapidly. Chapter 2.3 will deal with approaches for this in general and finally, Chapter 3 will present the method that was employed in this work. It must be noted that in general the EULER scheme (2.8) and (2.9) is not stable. Thus, other methods are used in practical applications such as the leap-frog scheme that will be discussed in Section 2.2.2.

Algorithm 2.1 gives a general overview on a very basic molecular dynamics method. The routine Initialize_Particles() sets up the system configuration of particle positions and velocities (temperatures). After evaluating the force law (2.4) for every particle inside Compute_Forces(), the velocity integration (2.9) is used in Update_Velocities(). All physical Diagnostics() are performed before finally computing the new positions with (2.8) in Update_Positions(). This is done to ensure, that all relevant properties (forces, velocities, positions) are evaluated at the same instant in time, e.g. when computing kinetic, potential, and total energy in Diagnostics(). This iterative scheme is continued until the simulation end time $t_{\rm end}$ is reached.

2.2.2 Leap-frog integration

The very simple Eulerian time discretization from the previous section is only of first-order accuracy in time and in general not stable. Naturally, better and higher-order approaches are available, e.g. Runge-Kutta integrators [142]. These make

use of the particle trajectory history or – equivalently – evaluation of the force at intermediate instants in time within a single timestep. This is undesirable since memory and computing capabilities are limited resources – see Section 2.3.

Although memory efficient Runge-Kutta schemes are available, see e.g. [143], we restrict ourselves to the leap-frog method that requires only minor modifications to (2.8) and (2.9) and delivers second-order integration in time [144, 145]. Its pictorial name results from the fact that positions and velocities are evaluated at different times – usually, velocity is computed at $t \pm \Delta t/2$. Thus, both variables leap-frog over each other:

$$\vec{r}_i(t) = \vec{r}_i(t - \Delta t) + \vec{v}_i(t - \Delta t/2) \cdot \Delta t , \qquad (2.10)$$

$$\vec{v}_i(t + \Delta t/2) = \vec{v}_i(t - \Delta t/2) + \frac{\vec{F}_i(t)}{m_i} \cdot \Delta t$$
 (2.11)

Being similar to the velocity-Verlet scheme [146] and thus Verlet integration [147], this method is time-reversible and symplectic. Especially the latter feature is essential for molecular dynamics simulations, since it is equivalent to energy conservation. Furthermore, provided $\Delta t \leq 2/\omega$ and constant, it is stable for oscillatory motion with angular frequency ω [144]. It is worth noting that with staggered-grid approaches [148] very similar half-step methods are used for spatial discretization in grid-based continuum calculations.

Naturally, the particular choice of time parameters in (2.11) requires special care in the Diagnostics() routine in Algorithm 2.1. For obtaining positions and velocities at the same instant in time, for the velocities one half step has to be performed backwards. The respective expression can be derived as the average of previous and current value:

$$\vec{v}_i(t) = \frac{1}{2} \left(\vec{v}_i(t - \Delta t/2) + \vec{v}_i(t + \Delta t/2) \right)$$
 (2.12)

$$= \vec{v}_i(t + \Delta t/2) - \frac{\vec{F}_i(t)}{m_i} \cdot \frac{\Delta t}{2} . \qquad (2.13)$$

When calling Diagnostics() in Algorithm 2.1, $\vec{v_i}(t+\Delta t/2)$ as well as $\vec{F_i}(t)$ are available. Thus, no necessary information is missing for computing dynamic properties at every full timestep t.

2.2.3 Simulation timestep

As already indicated in the previous section, the choice of timestep is not totally arbitrary. For example, for preserving stability in harmonic problems with the leap-frog integrator, a constant timestep $\Delta t \leq 2/\omega$ has to be chosen. Clearly, with the characteristic system frequencies $\omega_{\rm pl}$ and $\omega_{\rm laser}$ and particle velocities, several additional constraints for Δt arise.

Due to the mass ratio $m_{\rm p} \approx 1836 m_{\rm el}$ between protons and electrons, ion dynamics is significantly slower than electron movement. Consequently, it is sufficient to consider the electronic system to find an upper limit for Δt .

For resolving resonant movement in the external field, $\Delta t \ll 2/\omega_{\text{laser}}$ and for the collective plasmon resonance, $\Delta t \ll 2/\omega_{\text{pl}}$ are required. Based upon the discretization (2.7) in (2.6), further constraints can be derived:

$$\Delta t \le \frac{\Delta \vec{r}}{\vec{v}} \,, \tag{2.14}$$

$$\Delta t \le m \frac{\Delta \vec{v}}{\vec{F}} \tag{2.15}$$

with appropriate upper limits for $\Delta \vec{r}$ and $\Delta \vec{v}$ as well as estimates for the particle's velocities and forces. The division and inequalities have to be interpreted component-wise or using an appropriate vector norm. All of them have to be fulfilled simultaneously by the timestep Δt that is finally chosen.

It seems natural to estimate the particle velocity in (2.14) to $v \approx v_{\rm therm}$. The maximum velocity update in (2.15) can then be constrained with the average (thermal) velocity, e. g. $\Delta v \leq v_{\rm therm}/10$. A reliable choice for the maximum position update can be based on considerations about preventing too close encounters. Therefore, we do not allow the particles to pass the average interparticle separation a_{ee} in a single timestep by setting $\Delta r \leq \frac{a_{ee}}{10}$. An upper limit for the force can be guaranteed by using a Plummer/soft core interaction potential

$$\Phi^{(\text{PLUMMER})}(r) = \frac{1}{4\pi\varepsilon_0} \cdot \frac{q}{\sqrt{r^2 + \alpha^2}}$$
 (2.16)

instead of the Coulomb interaction (1.27). Originally used to describe globular clusters in theoretical astrophysics [149], this form of regularization has already proved its applicability in molecular dynamics simulations, for example in in calculations on the electron-ion collision frequency in warm dense matter [104]. The cutoff parameter α can be used as closest interaction distance between two particles. To treat classical COULOMB collisions correctly, it should be chosen smaller thank the classical distance of minimum approach

$$\alpha < b_{\min} = \frac{q_{\text{el}}q_{\text{ion}}}{\mu_{\text{ei}} \cdot v_{\text{rel}}^2} \qquad \text{with} \qquad \mu_{\text{ei}} = \frac{m_{\text{el}}m_{\text{ion}}}{m_{\text{el}} + m_{\text{ion}}}$$
(2.17)

and the particles' relative velocity $v_{\rm rel}$.

In an external oscillating field, besides the thermal velocity $v_{\rm therm}$, the electron quiver velocity $v_{\rm osc}$ also plays an an important role. It requires for another set of constraints to be constructed to the given upper limits for Δt in full analogy to those where $v_{\rm therm}$ is involved.

Of course, these constraints on Δt are more or less empirical. Their validity has to be checked continually during each simulation for example by observing conserved quantities such as linear and angular momentum and total energy.

With the numerical foundation from this section, a framework for classical molecular dynamics simulations can be built. However, there are further, technical constraints to be considered. These primarily affect the size of any system to be simulated due to limited memory and computing time resources. The according consequences for molecular dynamics simulations will be drawn in the next section.

2.3 Technical constraints – computers

Macroscopic systems consist of $\mathcal{O}(10^{23})$ particles. The amount of storage needed to represent them can be estimated as

$$8 \,^{\text{byte}/\text{variable}} \cdot \left(3 \,^{\text{coordinates}/\text{particle}} + 3 \,^{\text{momenta}/\text{particle}}\right) \cdot 1 \times 10^{23} \,^{\text{particles}}$$

$$= 48 \times 10^{23} \,^{\text{byte}}$$

$$= 4.8 \times 10^{12} \,^{\text{Tbyte}} \,^{\text{theorem}}$$
(2.18)

vastly exceeds the capacities of any conventional computer system in present and (near) future. Recently, the global data storage that was available in the year 2007 has been estimated to 2.95×10^{20} byte [150]. Even though this number already includes all non-volatile storage – even ancient analog material – this is far less than the estimate in (2.18). Questions of rapid access to such storage as needed for a dynamic simulation are not even covered.

Due to technical and physical constraints, this situation will not change dramatically with future technologies. Even if it was possible to store one variable inside a hypothetical memory unit by making use of some atomic property, the pure physical dimensions of the memory without the necessary wiring, connection, stabilizing structure, etc. would have to be six times larger than the simulated system itself as for every particle at least the coordinates and momenta have to be stored.

Clearly, a direct one-to-one particle-based simulation of macroscopic systems is and will be impossible. Consequently, the number of particles has to be reduced either by introducing pseudo-particles with effective physical properties or by down-scaling the system via decreased particle number N and simulation volume V while keeping densities N/V and other relevant properties constant. Obviously, both approaches further reduce the simulated physics towards a simpler model. Pseudo-particles mask fine-grained structure and prevent the particle-particle interactions from being treated correctly. However, the reduction of system size also creates artifacts. First, with decreasing V, its

ratio to the simulation region's surface area $A \propto V^{2/3}$ grows:

$$\frac{A}{V} \propto \frac{1}{V^{1/3}} \longrightarrow \infty \quad \text{for} \quad V \to 0 \; . \eqno(2.19)$$

Accordingly, surface effects start to dominate in finite-sized systems. Second, very small particle numbers N render a study collective or resonance phenomena of the electronic subsystem in a plasma difficult. Consequently, the number of particles in a simulation should be as large as possible.

2.4 The N-body problem

Unfortunately, besides the memory limit already discussed, there is another constraint on the particle number N. In (2.4), the force $\vec{F}(\vec{r_i}, \{\vec{r_1} \dots \vec{r_N}\}, t)$ onto each individual particle is in general a function of the full set of 3N particle coordinates $\{\vec{r_1} \dots \vec{r_N}\}$ of all other particles, which makes its computation extremely costly. For example, for particles interacting via the COULOMB potential (1.27) and experiencing some general external potential

$$\Phi(\vec{r}, \{\vec{r}_1 \dots \vec{r}_N\}) = \sum_{i=1}^{N} \frac{q_i}{4\pi\varepsilon_0} \frac{1}{|\vec{r} - \vec{r}_i|} + \Phi_{\text{ext}}(\vec{r}) , \qquad (2.20)$$

a total of N-1 distances has to be evaluated per particle. Accordingly, for each particle the force

$$\vec{F}(\vec{r}_{i}, \{\vec{r}_{1} \dots \vec{r}_{N}\}, t) = -q_{i} \cdot \nabla_{\vec{r}} \Phi(\{\vec{r}_{1} \dots \vec{r}_{N}\}, t) \bigg|_{\vec{r} = \vec{r}_{i}(t)}$$
(2.21)

$$= \frac{q_i}{4\pi\varepsilon_0} \sum_{j\neq i} q_j \frac{\vec{r_i} - \vec{r_j}}{|\vec{r_i} - \vec{r_j}|^3} + \vec{F}_{\text{ext}} \left(\vec{r_i}, t\right)$$
 (2.22)

requires the evaluation of N-1 distances per particle. Hence, in total $\mathcal{O}(N^2)$ interactions have to be evaluated.

Furthermore, in addition to (2.4) and (2.5), initial conditions

$$\vec{r}_i(t=t_0) = \vec{r}_i^{(0)},$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{r}_i(t=t_0) = \dot{\vec{r}}_i^{(0)}$$
(2.23)

have to be given for uniqueness of the solution. They are usually given by the definition of the system that is to be simulated. The question of finding a global solution to the 6N coupled partial differential equations (2.4) and (2.5) with the 6N initial conditions (2.23) is considered as the N-body problem.

2.4.1 Direct summation

The most straightforward way of numerically solving the N-body problem (2.4, 2.5, 2.23) involves the direct computation of all N-1 terms in (2.22) for each of the N particles. As shown above, this requires $\mathcal{O}(N^2)$ operations, rendering the direct method impractical on today's computers for moderate particle numbers N > 10,000, where already $> 10^8$ pairwise interactions are involved.

However, apart from round-off errors due to limited machine precision, the direct method gives exact results for all forces in the considered system. Furthermore, a data-parallel implementation is trivial by distributing equal chunks of the particle list across all available processors and passing them to the next neighbor in a round-robin scheme. These features make this method a handy tool for generating "exact" data for test runs with small particle numbers N, or when evaluating (2.22) for a subset of M among the N particles, leading to $\mathcal{O}(M \cdot N)$ scaling of the necessary computational effort, which is feasible for few test cases. However, since the trajectories of all N system particles are needed for meaningful physical studies, this approach is not followed here.

2.4.2 Grid-based approaches

One popular way to get around the enormous computational effort of the direct evaluation of (2.22) already dates back to the 1960s and is known as Particle-In-Cell (PIC) approach. Tutorial introductions can be found in [16, 51, 144, 151]. Here, we will only give a very short overview on the method. It is based on the idea of resolving the computational domain using a sufficiently fine grid and projecting charges and currents onto grid nodes at position \vec{R}_k via

$$\varrho\left(\vec{R}_{k}\right) = \sum_{i} q_{i} \cdot S\left(\vec{r}_{i} - \vec{R}_{k}\right) \tag{2.24}$$

$$\vec{j}\left(\vec{R}_k\right) = \sum_{i}^{i} q_i \vec{v}_i \cdot S\left(\vec{r}_i - \vec{R}_k\right) \tag{2.25}$$

with the weighting function $S(\vec{r})$ and the particle velocity $\vec{v}_i = \frac{\vec{p}_i}{m_i}$. Then, MAXWELL's equations (1.22) and (1.23) are solved on the grid, which can efficiently be done in parallel using finite difference schemes, (parallel) FOURIER transform methods (PFFT) or a multigrid approach. Detailed overviews on these and related methods can for example be found in [142, 145, 152–156] and references therein. The resulting fields and potentials can be interpolated back to the actual particle positions using the same weighting function $S(\vec{r})$ where they are used for the individual velocity and position updates.

Normally, such PIC methods treat the simulated particles as discrete elements of the single-particle distribution function $f(\vec{r}, \vec{v})$ rather than real physical entities. Thus, they

essentially solve the VLASOV equation

$$\frac{\partial f}{\partial t} + \vec{v} \cdot \frac{\partial f}{\partial \vec{r}} + \vec{F} (\vec{r}, \vec{v}) \cdot \frac{\partial f}{\partial \vec{v}} = 0$$
 (2.26)

and the weighting function $S(\vec{r})$ represents the shape of the considered particle clouds.

Despite the PIC approach's significant advantage of treating not only particle evolution but also propagation of the electromagnetic fields, i.e. including radiation from an external source and the self-consistent internal electrodynamics, by means of a full solution of Maxwell's equations, it still has some drawbacks.

Firstly, grid resolution is a difficult question. To sufficiently resolve rapidly evolving systems, e.g. laser-driven particle acceleration, a wide dynamic range in particle densities from vacuum to bulk matter with steep density gradients has to be covered. This usually needs adaptive grid methods that make the numerics very complex - mathematically as well as from the algorithmic and implementation side. Furthermore, with grid-based techniques it is difficult to handle systems with open boundaries corresponding to DIRICHLET boundary conditions at infinity. For finiteness, usually

$$\Phi\left(\left|\vec{r}\right| \to \infty\right) = 0\tag{2.27}$$

is required which corresponds to conventions from electrodynamics [34, 35]. While in bulk simulations on laser-plasma interactions this does not impose a limitation, finite-sized systems can be studied more conveniently using real physical particles directly instead of using a grid-based method.

Furthermore, the use of pseudo-particles (clouds and/or grid nodes) for the interaction of nearby partners precludes the exact treatment of particle collisions. Since this is one of the main objectives of this work, the models here will be constructed around a grid-free multipole-based technique of evaluating (2.22) for large particle ensembles. Two closely related representatives of this class of algorithm will be discussed in the following section. Concepts from both methods will be applied in the following chapters.

2.4.3 Multipole-based approaches

As already mentioned, the use of pseudo-particles for reducing the effective number of interactions to be computed prevents individual close encounters for being sufficiently resolved. However, it can be observed that interactions of particles with distant particle groups do not have to be resolved down to the level of individual interactions. Instead, the particle group can be considered as a single, spatially expanded charge distribution. Due to its large distance, the particle only experiences the far-field expansion of this particle group's electric field. This approach is frequently used in standard electrodynamics, see e. g. [34, 35], and has found its way into computer simulations in the early

1980's [157–159]. A textbook overview on multipole methods in physics is available in [160]. In this section, we will briefly summarize Cartesian and spherical multipole expansions and describe their particular value for evaluating forces in large sets of particles with long-range interaction. The mathematical foundation for the expressions given here can be found in any textbook on advanced calculus or mathematical methods for physicists, such as [161, 162]. The TAYLOR series

$$f(x) = \sum_{n=0}^{\infty} \frac{(x-a)^n}{n!} \cdot \frac{\mathrm{d}^n}{\mathrm{d}x^n} f(x) \bigg|_{x=a}$$
(2.28)

for a function $f: \mathbb{R} \to \mathbb{R}$ which is infinitely differentiable in the neighborhood of a can be generalized to functions $f: \mathbb{R}^d \to \mathbb{R}$ with $\vec{r}, \vec{a} \in \mathbb{R}^d$ and superscripts denoting vector components as

$$f(\vec{r}) = \sum_{n_1=0}^{\infty} \cdots \sum_{n_d=0}^{\infty} \frac{(r^{(1)} - a^{(1)})^{n_1} \cdots (r^{(d)} - a^{(d)})^{n_d}}{n_1! \cdots n_d!} \cdot \frac{\partial^{n_1+\dots+n_d}}{(\partial x^{(1)})^{n_1} \cdots (\partial x^{(d)})^{n_d}} f(x) \bigg|_{\vec{x}=\vec{a}}.$$
(2.29)

Considering for now just the first three terms and d=3, i.e. real physical space, this can be cast into the more compact but identical expression

$$f(\vec{r}) = f(\vec{a}) + (\vec{r} - \vec{a})^{\mathrm{T}} \nabla f(\vec{a}) + \frac{1}{2} (\vec{r} - \vec{a})^{\mathrm{T}} \bar{H}_f(\vec{a}) (\vec{r} - \vec{a}) + \dots$$
 (2.30)

with the gradient $\nabla f(\vec{a})$ and the HESSian $\bar{H}_f(\vec{a})$ of f in \vec{a} . Using (2.30), an expansion for the inverse distance can then be written as

$$\frac{1}{|\vec{r} - \vec{a}|} = \frac{1}{|\vec{a}|} + \frac{\vec{a}^{\mathrm{T}}}{|\vec{a}|^{3}} \vec{r} + \frac{1}{2} \vec{r}^{\mathrm{T}} \left(3 \frac{\vec{a} \otimes \vec{a}}{|\vec{a}|^{5}} - \frac{\bar{I}}{|\vec{a}|^{3}} \right) \vec{r} + \dots$$
(2.31)

with the identity matrix \bar{I} and the dyadic (outer) product $\vec{a} \otimes \vec{a} = \vec{a}\vec{a}^{\mathrm{T}}$, see Figure 2.1.

To compute the electrical field of a particle cloud onto a remote observer at \vec{a} , we consider a subset of M particles at positions $\vec{r}_{i=1...M}$ around the origin and with $\Phi_{\rm ext} \equiv 0$. Insertion of (2.31) into the potential sum (2.20) leads to

$$\Phi(\vec{a}) = \frac{1}{4\pi\varepsilon_0} \left[\frac{1}{|\vec{a}|} \cdot \underbrace{\left(\sum_{i=1}^M q_i\right)}_{Q: \text{monopole-}} + \frac{1}{|\vec{a}|^3} \cdot \vec{a}^{\text{T}} \underbrace{\left(\sum_{i=1}^M q_i \cdot \vec{r}_i\right)}_{\vec{Q}: \text{dipole-}} + \frac{1}{2} \frac{1}{|\vec{a}|^5} \cdot \vec{a}^{\text{T}} \underbrace{\left(\sum_{i=1}^M q_i \cdot \left[\vec{r}_i \otimes \vec{r}_i - |\vec{r}|^2\right]\right)}_{\bar{Q}: \text{quadrupole-moment}} \vec{a} + \dots \right] \qquad (2.32)$$

$$= \frac{1}{4\pi\varepsilon_0} \left[\frac{1}{|\vec{a}|} \cdot Q + \frac{1}{|\vec{a}|^3} \cdot \vec{a}^{\text{T}} \vec{Q} + \frac{1}{2} \frac{1}{|\vec{a}|^5} \cdot \vec{a}^{\text{T}} \bar{Q} \vec{a} + \dots \right], \qquad (2.33)$$

which is a series in $\frac{|\vec{r_i}|}{|\vec{a}|}$. Via (2.33), the multipole moments are defined by

$$Q = \sum_{i=1}^{M} q_i \tag{2.34}$$

$$\vec{Q} = \sum_{i=1}^{M} q_i \cdot \vec{r_i} \tag{2.35}$$

$$\bar{Q} = \sum_{i=1}^{M} q_i \cdot \left[\vec{r}_i \otimes \vec{r}_i - |\vec{r}|^2 \right] . \tag{2.36}$$

For the above transformation, the property $\vec{r}^{\rm T}(\vec{a}\otimes\vec{a})=(\vec{r}^{\rm T}\vec{a})\vec{a}^{\rm T}$ of the dyadic product and commutativity of the scalar product $\vec{r}^{\rm T}\vec{a}=\vec{a}^{\rm T}\vec{r}$ has been used to rewrite the quadrupole contribution. The multipole series converges if

$$|\vec{r_i}| < |\vec{a}| \quad \text{for all } i = 1, \dots, M$$
 (2.37)

Hence, the potential of a cloud of charges around the origin can be approximated for large distance from the origin by truncating (2.33) after the first few terms. For functions $f: \mathbb{R} \to \mathbb{R}^m$, (2.29) can be applied component-wise which makes it immediately applicable for derivation of an expansion for the electric field:

$$-\vec{E}(\vec{a}) = \nabla \Phi(\vec{a}) = \frac{\vec{a} \cdot Q}{|\vec{a}|^3} + \frac{3\vec{a}(\vec{a} \cdot \vec{Q}) - |\vec{a}|^2 \vec{Q}}{|\vec{a}|^5} + \frac{\bar{Q}\vec{a}}{|\vec{a}|^5} - \frac{5}{2} \frac{\vec{a}^{\mathrm{T}} \bar{Q}\vec{a}}{|\vec{a}|^7} \vec{a} + \dots$$
 (2.38)

Again, the expression shows convergence for $\vec{a} \to \infty$.

A more rigorous approach that is formally equivalent to (2.33) can be given by expanding the inverse distance in spherical harmonics [35]

$$\frac{1}{|\vec{r} - \vec{a}|} = \sum_{l=0}^{\infty} \frac{|\vec{a}|^l}{|\vec{r}|^{l+1}} P_l(\cos \chi)$$
 (2.39)

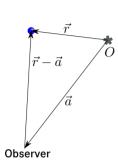
where $P_l(x)$ are Legendre polynomials. An extensive overview on their definition and properties used in the following can be found in any textbook on special functions as well as in the standard mathematical tables, such as [163]. The parameter χ is the angle enclosed by \vec{a} and \vec{r} . Here, since $-1 \leq P_l(x) \leq 1$, convergence is possible for $|\vec{a}| < |\vec{r}|$, which restricts the validity of the expansion to a region outside the particle cloud. Using the standard addition theorem for Legendre polynomials

$$P_{l}(\cos \chi) = \sum_{m=-l}^{l} \frac{(l-|m|)!}{(l+|m|)!} P_{l}^{|m|}(\cos \theta) P_{l}^{|m|}(\cos \alpha) e^{im(\phi-\beta)}$$
(2.40)

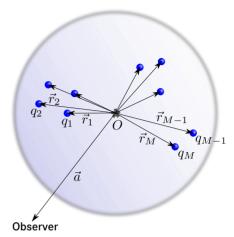
and the definition of spherical harmonics

$$Y_{lm}(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\phi}, \quad m \ge 0$$
 (2.41)

$$Y_{l,-|m|}(\theta,\phi) = (-1)^{l-m} Y_{l|m|}^*(\theta,\phi)$$
(2.42)



(a) The COULOMB potential that results from a charge at position \vec{r} as seen by an observer at \vec{a} is determined by $\frac{1}{|\vec{r}-\vec{a}|}$.



(b) By performing a multipole expansion for a cloud of charges q_i at positions \vec{r}_i , the potential at \vec{a} outside the cloud can be approximated.

Figure 2.1: Using the expansion (2.31), an approximation for the inverse distance $\frac{1}{|\vec{r}-\vec{a}|}$ is found. This can be used to perform an expansion for the potential $\Phi(\vec{a})$, experienced by a remote observer at \vec{a} due to a particle cloud around the origin O.

with the associated LEGENDRE polynomials $P_l^m(x)$, (2.39) takes the form

$$\frac{1}{|\vec{r} - \vec{a}|} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \frac{4\pi}{2l+1} \frac{a^{l}}{r^{l+1}} Y_{lm}^{*}(\alpha, \beta) Y_{lm}(\theta, \phi)$$
(2.43)

with the spherical coordinates $\vec{r} = (r, \theta, \phi)$ and $\vec{a} = (a, \alpha, \beta)$. A short but complete proof can be found in [164]. Insertion of (2.43) into (2.20) now yields an expression for the potential expanded in spherical harmonics $Y_{lm}(\theta, \phi)$:

$$\Phi(\vec{r} = (r, \theta, \phi), \{\vec{r}_{i=1...N} = (a_i, \alpha_i, \beta_i)\})
= \frac{1}{4\pi\varepsilon_0} \sum_{i=1}^{N} q_i \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \frac{4\pi}{2l+1} \frac{a^l}{r^{l+1}} Y_{lm}^*(\alpha, \beta) Y_{lm}(\theta, \phi) .$$
(2.44)

After summing over i, this reads

$$\Phi(\vec{r}) = \frac{1}{4\pi\varepsilon_0} \sum_{l=0}^{\infty} \sum_{m=l}^{l} \frac{4\pi}{2l+1} q_{lm} \frac{Y_{lm}(\theta,\phi)}{r^{l+1}}$$
(2.45)

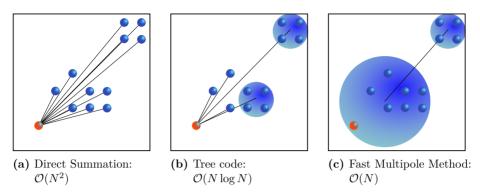


Figure 2.2: Particle-based N-body methods: When determining the force on the lower left red particle, different numbers of interactions have to be evaluated depending on the approximation method.

with appropriate coefficients q_{lm} , which take the role of the multipole coefficients here. They can be found in detail in standard literature on electrodynamics, e.g. [34, 35]. Naturally, the corresponding expression for the electrical field and thus for the force from a particle cloud on a distant particle can be derived in full analogy to (2.38). Being formally a rather closed expression even for higher orders, the spherical expansion is much more compact in comparison to the analogous Cartesian expressions. On the other hand, it contains transcendental functions that demand numerically sophisticated algorithms for their stable evaluation. Furthermore, the shifting rules for multipole expansions needed later – which can for now be seen as algorithms for combining expansions around different centers – are significantly simpler in the Cartesian case. The relevance and consequences of these properties are studied in more detail in Chapter 3, which describes a tree algorithm for the N-body problem that is based on the Cartesian expansion. In this work, the spherical expansion plays an important role in the algorithm's extension for periodic boundary conditions that demands only few evaluations, but high orders of multipoles, see Chapter 4.

Clearly, the multipole moments contain information on the internal charge distribution of the particle group under consideration. By including them, e.g. via (2.38), into the force sum, a large number of individual interactions with far particles can be performed by just one interaction with a multipole expansion. Now the idea of multipole-based summation techniques is to efficiently group particles into clusters and to decide whether to interact either with them or with their constituents using a simple distance-based criterion. In this respect, the multipole-approach is also a pseudo-particle method. However, in contrast to PIC or grid-based techniques, pseudo-particles or particle groups are only employed for far-field interactions, while near-neighbors are still treated directly.

Figure 2.2 shows – in addition to the direct $\mathcal{O}(N^2)$ sum – the interaction scheme of two multipole-based approaches for evaluating (2.22). In the BARNES-HUT tree method [158], only far particles are incorporated into clusters as interaction partners. This algorithm can be shown to be able to scale as $\mathcal{O}(N\log N)$ with particle number N, see [160]. In the Fast Multipole Method (FMM) [159], close particles form an additional local cluster, allowing cluster-cluster interactions between objects with different multipole expansions to be performed. Then, the result for the local cluster is shifted down to its constituent particles. This approach can scale as $\mathcal{O}(N)$, compare [156, 165, 166].

While being limited in multipole order when constructing an implementation, as already mentioned, the Cartesian expansion that is generally used in the Barnes-Hut algorithm leads to simpler multipole shifting rules. These are necessary tools in the tree algorithm and will be described in Section 3.1.2. In addition, the algorithm can be easily adapted to implement additional interaction laws such as regularized kernels or even a nearest-neighbor search needed for Smooth Particle Hydrodynamics (SPH). In Section 5.3 we will describe some of the applications that have evolved as side projects during the development of the code that has been used and extended in this work and propose possible directions of development in future. The flexibility to add both algorithm and physics modules straightforwardly makes the parallel Barnes-Hut tree code the tool of choice for the envisaged simulations – outweighing the more favorable $\mathcal{O}(N)$ scaling of the Fast Multipole Method.

To get a more detailed insight into the algorithm, its advantages and latest developments, the next chapter will summarize the main features of the Barnes-Hut method, the Warren & Salmon hashed-oct-tree approach that is used for parallelization, and the implementation PEPC that is developed at Jülich Supercomputing Centre. Alternative implementations and recently proposed modifications to the algorithm will also be covered. Recent developments and improvements to the algorithm that were developed as part of this work will be described in the subsequent Chapters 4 and 5.

3 The parallel Barnes-Hut tree code PEPC

The inventors of the tree algorithm, Josh Barnes and Piet Hut, proposed their idea of a hierarchical tree code in the mid-1980s for efficiently simulating gravitational N-body systems [158]. Together with the Fast Multipole Method, published independently slightly later in [159], it has since revolutionized long-range N-body simulations for scientists across a wide range of disciplines, ranging from gravitational problems and plasma physics to fluid dynamics.

Both methods reduce the number of direct particle-particle interactions through a systematic use of multipole expansions as described in Section 2.4.3. Especially when simulating dynamical systems, there is no use in computing the potential and force to higher accuracy than the error which the time discretization scheme for integrating the equations of motion imposes anyway. In these cases, multipole methods are able to significantly speed up particle simulations with many millions of physical objects.

Since the first publication of the Barnes-Hut tree method in 1986, a large number of implementations for many different applications have evolved. Among them is the Pretty Efficient Parallel Coulomb Solver (PEPC) [167], which has been developed at Jülich Supercomputing Centre over the past decade [168–175]. PEPC is the MPI-parallelized successor to treemd [104, 160, 176], a non-parallel implementation of the original serial algorithm with several additions for bulk plasma simulations and vector computers. PEPC represents a whole family of codes for a multitude of applications under the umbrella of a common tree code framework and has been used as the basis for all further developments and improvements that were implemented within this thesis.

This chapter will describe the basic concepts of the Barnes-Hut algorithm and its parallelization via the Warren and Salmon hashed-oct-tree method at the example of the original PEPC implementation. An overview of alternative implementations and algorithm variations is given in Section 3.4 before Chapters 4 and 5 describe the new developments on periodic boundary conditions and overall code scalability that have been conducted within this work.

3.1 The Barnes-Hut algorithm

The fundamental idea behind any tree code in this context is to map spatial position information onto a data structure that can be easily analyzed and exploited to classify proximity between points in space. While not being limited in the actual dimension of space, we restrict ourselves to the physically most natural 3D-version in the following. This leads to the construction of an oct-tree. However, for clarity in all figures here, the 2D-equivalent – a quad-tree – will be shown.

The algorithm itself basically consists of three steps:

- 1. Tree construction
- 2. Computation of multipole properties for tree nodes
- 3. Tree traversal with
 - Identification of interaction partners
 - Evaluation of the force law for each identified interaction

These will be described in the following.

3.1.1 Tree construction

The tree construction is depicted for an illustrative set of particles in Figure 3.1, where the tree is drawn in a bottom-up fashion. Starting from the root node, the simulation region is recursively subdivided into boxes of decreasing size until each particle resides inside its own box. These boxes are assigned to nodes within the tree data structure: Particles correspond to leaf nodes, while particle groups, i. e. larger boxes that contain several sub-boxes are referred to as twig nodes. The root node contains all particles within the simulation region. For now, the leaves only contain individual particles although this is not required as we will see later. The tree level is counted from root (1 evel = 0) towards the leaves (1 evel > 0).

Of course, particles and particle groups that are situated close to each other physically are also closely related in the tree through sharing common ancestor nodes. Hence, they are part of the same particle cluster on a lower tree level.

3.1.2 Computation of multipole properties

To compute interactions between particles and particle clusters by means of their multipole expansion via (2.33) and (2.38), the multipole moments $Q, \bar{Q}, \bar{Q}, \ldots$ have to be calculated for all tree nodes that are potential interaction partners. For sake of sim-

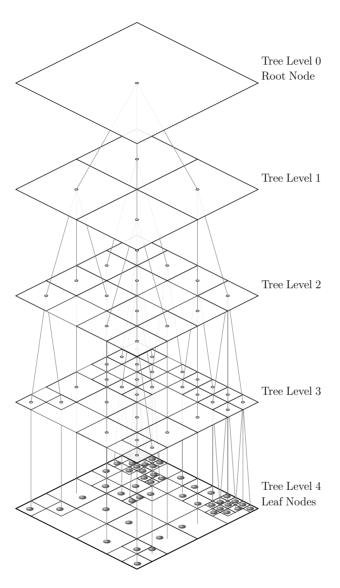


Figure 3.1: For a two-dimensional example, a quad-tree is constructed by recursively subdividing space into boxes with decreasing size from the root node towards the tree leaves containing individual particles.

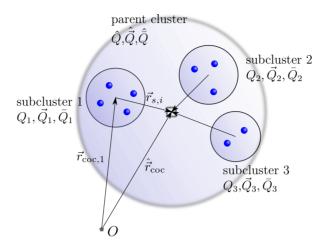


Figure 3.2: Collection of subcluster multipole moments for computing the parent cluster's properties.

plicity we restrict ourselves to at most quadrupole order, thus truncating the expansion after the term containing \bar{Q} . The consequences of this choice will be discussed later.

An efficient way of computing all relevant multipole moments is to perform a node-wise expansion around the respective particle cluster's center of all N_c subclusters

$$\hat{\vec{r}}_{\text{coc}} = \frac{\sum_{i=1}^{N_c} Q_i \cdot \vec{r}_{\text{coc},i}}{\sum_{i=1}^{N_c} Q_i} \ . \tag{3.1}$$

Subscript indices denote particle numbers inside the cluster, variables with circumflex accent, such as $\hat{\vec{r}}_{\rm coc}$ are parent node properties (here its center-of-charge), those without belong to subclusters as shown in Figure 3.2. Defining the shift vector

$$\vec{r}_{\mathrm{s},i} = \hat{\vec{r}}_{\mathrm{coc}} - \vec{r}_{\mathrm{coc},i} \tag{3.2}$$

between the center-of-charge $\vec{r}_{\text{coc},i}$ of a sub-cluster and that of its parent node $\hat{\vec{r}}_{\text{coc}}$, the multipole moments of the parent node can be computed via

$$\hat{Q} = \sum_{i=1}^{N_c} Q_i \,, \tag{3.3}$$

$$\hat{\vec{Q}} = \sum_{i=1}^{N_c} \left[\vec{Q}_i - Q_i \cdot \vec{r}_{s,i} \right] , \qquad (3.4)$$

$$\hat{\bar{Q}} = \sum_{i=1}^{N_c} \left[\bar{Q}_i - \vec{r}_{s,i} \otimes \vec{Q}_i - (\vec{r}_{s,i} \otimes \vec{Q}_i)^{\mathrm{T}} + (\vec{r}_{s,i} \otimes \vec{r}_{s,i}) Q_i \right] . \tag{3.5}$$

Algorithm 3.1 Serial tree traversal

```
for all particle in particle list do
   todo stack.clear()
   node \leftarrow root
                                        > traversal for the particle starts at root node
   repeat
       if (MAC OK(particle, node)) and not (particle \in node) then \triangleright MAC eval.

    interaction allowed

          call Interact(particle, node)
                                                   ▷ due to interacting with this node
                                           > its children do not have to be considered
       else
                                  ▶ the MAC requires the node to be further resolved
          todo stack.push(node.children)
                                                            ▷ proceed vertically in tree
       end if
   until IS INVALID(node \leftarrow todo stack.pop())
                   ▷ todo_stack is empty for this particle – its traversal is complete
end for
```

A detailed derivation is available in [160]. It makes use of the additivity of multipole expansions when performed with respect to same expansion center. All multipole moments are shifted to the common center-of-charge $\hat{r}_{\rm coc}$ and summed up to obtain the parent properties. Individual particles can be treated as clusters with vanishing dipole and quadrupole moments and center-of-charge localized at their own position, so the multipole expansion of all nodes on any tree level can be computed via a bottom-up sweep through the tree. Starting from the leaf nodes with their trivial multipole expansion, the nodes are visited level-by-level while combining properties of the previous (higher) level clusters until finally the root node is reached. Instead of gathering all particles for every cluster individually and having to perform $\mathcal{O}(\mathbf{n_treelevels} \cdot N)$ operations, this technique reduces the amount of work to $\mathcal{O}(\mathbf{n_treelevels}) \propto \mathcal{O}(N)$.

3.1.3 Tree traversal and multipole acceptance criterion

Finally, after constructing the tree and dressing all tree nodes with their multipole coefficients, interaction partners have to be identified and the actual interactions have to be computed. This is performed during the next step, where – starting from the root node – for every particle at position \vec{R} , a depth-first traversal through the tree is performed. A pseudo-code overview is given in Algorithm 3.1. During this traversal, a series of decisions are made whether the particle may interact with the node that is currently under consideration and the traversal into the particular tree branch can be terminated. Otherwise, it is further resolved into its constituent nodes. A natural

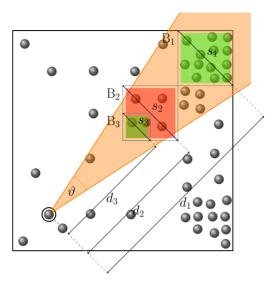


Figure 3.3: Barnes-Hut multipole acceptance criterion (MAC)

basis for deciding this is to compare the cluster distance to its size. Far clusters may be larger, while nearer particle groups have to be smaller to yield a consistently accurate evaluation of (2.38). Obviously, the particle itself may not be part of the considered cloud, since this would violate (2.37) and would include unphysical self-interaction of particles.

The decision for accepting interaction with a tree node at a certain tree level is called $Multipole\ Acceptance\ Criterion\ (MAC)$. Among the many flavors of different MACs, the classical Barnes-Hut-MAC is the simplest and most well-established [158, 160, 177, 178]. In this MAC, for every tree node j touched during the traversal its center-of-charge distance from the particle

$$d_j = |\vec{R} - \vec{r}_{\text{coc},j}| \tag{3.6}$$

is determined. If the ratio of its box size s_j to d_j is larger than a fixed parameter $\vartheta > 0$, interaction with the cluster is not allowed and it has to be further resolved. If

$$\frac{s_j}{d_j} \le \vartheta \;, \tag{3.7}$$

interaction is allowed and the cluster's substructure can be ignored. Figure 3.3 depicts the multipole acceptance criterion for a randomly chosen particle from the previous 2D example in Figure 3.1. While box B_1 is far enough and does not have to be resolved,

 B_2 is of the same size $s_2 = s_1$, but lies closer to the particle. It does not fulfill (3.7) and has to be decomposed. One of its components, B_3 is significantly smaller $s_3 = \frac{s_2}{2}$ and is an acceptable partner for interaction. Thus, this MAC can be interpreted as a cone with opening angle ϑ . Only nodes that fit into this cone are acceptable interaction partners.

Setting $\vartheta = 0$ corresponds to resolving all nodes up to leaf level, which results in performing the unfavorable $\mathcal{O}(N^2)$ sum again in a roundabout fashion. It can be shown that values of $\vartheta > 0$ lead to $\mathcal{O}(\log N)$ interactions per particle and hence $\mathcal{O}(N \log N)$ interactions in total, see [160, 177] for a detailed argumentation.

Besides the ratio of cluster distance and size being large enough, one also has to ensure that the particle itself is not contained in the box. This can arise if the MAC parameter is smaller than the diagonal of a unit box, i. e. $\vartheta < 1/\sqrt{2} \approx 0.707$ in 2D, or $\vartheta < 1/\sqrt{3} \approx 0.577$ in 3D. These are worst-case estimates for the center-of-charge and the particle lying in opposite corners of the box. However, to prevent ϑ from – for prevention of self-interactions – having to be smaller than necessary, it is favorable to avoid using the MAC for particle exclusion. Instead, forcing to resolve the particle's parent boxes independent of the MAC is expedient.

In addition to the very simple yet effective Barnes-Hut MAC, several other variants are possible [172, 178]

- the minimum distance MAC, where d_j is computed as nearest distance to the cluster instead as distance to its center-of-charge,
- the b_{max} MAC, that replaces the box size s_j with the maximum distance of the cluster's center-of-charge from the box edges,
- the \vec{E} -field MAC, that estimates and limits the relative error that results from inclusion of a cluster on a lower level.

to name but a few. The $b_{\rm max}$ and \vec{E} -field MAC include an additional data dependency to limit the actually introduced error from each individual interaction. This is favorable to derive usable expressions for the precision of the multipole approximation. In [172] these multipole acceptance criteria have been compared with respect to performance and achieved precision. Further details can also be found in [178], which is the standard work on evaluation of different multipole acceptance criteria.

During the tree traversal, the MAC is evaluated for every cluster, visited for the particle. Hence, it has to be significantly cheaper from the computational side than the actual interaction to prevent it from dominating the total traversal time. Therefore, we adhere to the very simple BARNES-HUT MAC (3.7), which can efficiently be evaluated by just comparing $d_j^2 \cdot \vartheta^2$ with s_j^2 . It does not require any square roots, divisions or computationally more demanding expressions. Furthermore, s_j^2 can be pre-tabulated as a function of the tree level and the distance d_j^2 is re-used during the interaction computation if the tree node was accepted.

A rigorous first-principles error estimation for the BARNES-HUT tree code with generalized algebraic kernels of the form

$$\Phi(r) = \frac{1}{(r^2 + \sigma^2)^{\tau}}, \quad \tau > 0 \tag{3.8}$$

has been derived by SPECK [179]. It is particularly useful for the b_{max} multipole acceptance criterion

$$\frac{b}{d} < \vartheta , \tag{3.9}$$

where b is the maximum distance between a cluster's center-of-charge $\vec{r}_{\rm coc}$ and all its members and d the distance between the point of observation and $\vec{r}_{\rm coc}$. In [179] it is shown that the error $\varepsilon_{(p+1)}^{\tau}$ due to truncating a multipole expansion of (3.8) after p terms is bounded by

$$|\varepsilon_{(p+1)}^{\tau}| < \mathcal{D}_{2\tau+1}(p+1) \cdot \frac{\bar{M}_{(0)}}{(d-b)^{2\tau}} \cdot \left(\frac{b}{d-b}\right)^{p+1}$$
 (3.10)

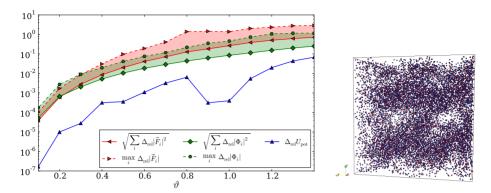
with a polynomial $\mathcal{D}_{2\tau+1}(p+1)$ of order $2\tau+1$ in p+1. The coefficient $\bar{M}_{(0)}=\sum_{i\in \text{cluster}}|q_i|$ is the total absolute charge of the cluster. Clearly, the truncation error is smaller for higher orders p of terms included in the multipole expansion. In addition, convergence is only guaranteed for b< d-b, i. e. if the evaluation position is not situated inside a ball that tightly encloses the particle group. Further details, interpretation, and connection to previous work by Salmon can be found in [179].

As (3.10) is more of academic than practical use and we will utilize the standard Barnes-Hut multipole acceptance criterion (3.7), Figure 3.4 gives an empirical error estimate. Here, the root-mean-square and maximum relative errors

$$\Delta_{\text{rel}} A = \left| \frac{A_{\text{pepc}} - A_{\text{direct}}}{A_{\text{direct}}} \right| \tag{3.11}$$

are given for the field, potential and total potential energy for different MAC parameters ϑ . The example system considered consists of 12,960 positively and negatively charged ions in a melting NaCl crystal with open boundaries. It is taken from the ScaFaCoS testsuite [156, 180]. Values found by PEPC are compared the results of the direct N^2 sum.

We will not track individual trajectories or rare events in our applications in Chapters 6 and 7 and are primarily interested in statistical measures such as heating rate and time-averaged collective movement of particles. Accordingly, a maximum relative force and potential error of $\mathcal{O}(1\,\%)$ is expected to be tolerable. As seen in the figure, this fixes the value of the MAC parameter to $\vartheta=0.3$.



(a) Maximum and root-mean-square relative errors for poten- (b) Example system tial, field and total potential energy compared between PEPC and the direct N² sum.

Figure 3.4: Precision comparison for different multipole acceptance parameters ϑ . The example system considered for this comparisons is a 12,960 particles open-boundary extract of a melting NaCl crystal that includes some density variations. Reference values were produced using the direct N^2 sum for the N-body problem. For $\vartheta=0.0$ the relative errors reach 10^{-14} which is expected as the numerical limit for double precision computations and thus not included in the figure.

3.2 The parallel tree algorithm

Clearly, the hierarchical ordering of particle clusters with nodes inside a quad- or oct-tree is an algorithmically attractive technique for tackling the problem of finding interaction partners of appropriate size and distance. However, an efficient data-parallel implementation of algorithms that incorporate hierarchical data structures is difficult without an extensive redesign.

One approach of parallelization – the *Hashed Oct-Tree* (HOT) with the *Latency Hiding Tree Traversal* was proposed by WARREN and SALMON [181]. It keeps the original structure intact and requires only few changes and additions. In its original implementation, it requires the following modifications to the data structure and algorithm:

- Allocation of a globally unique key to all nodes in the tree
- Storage of the tree node properties, such as multipole properties and administrative data in a hash table structure with $\mathcal{O}(1)$ access complexity
- Domain decomposition using a space-filling curve
- Tree construction:

- Construction of *local trees*
- Identification and exchange of branch nodes as entry points for remote processors
- Construction of the global tree
- Latency-hiding tree traversal

The Warren and Salmon HOT is the most well-known distributed memory parallelism extension for the Barnes-Hut algorithm. However, today several variations and other algorithmic modifications do exist. An overview will be given in Section 3.4 after describing the HOT parallelization steps as implemented in PEPC. These are the basis for the code's further evolution.

3.2.1 Allocation of globally unique keys

Firstly, every particle is assigned a unique key. To do this, the physical simulation region $B_{\text{sim}} = [x_{\text{min}} \dots x_{\text{max}}] \times [y_{\text{min}} \dots y_{\text{max}}] \times [z_{\text{min}} \dots z_{\text{max}}]$ is mapped to a $B_{\text{int}} = [0 \dots 2^{21} - 1] \times [0 \dots 2^{21} - 1] \times [0 \dots 2^{21} - 1]$ cube of integer coordinates via the operation

$$B_{\text{int}} \ni \bar{\vec{r}} := \begin{bmatrix} \frac{2^{21} - 1}{\left(\frac{x_{\text{max}} - x_{\text{min}}}{y_{\text{max}} - y_{\text{min}}} \right)} \vec{r} - \left(\frac{x_{\text{min}}}{y_{\text{min}}} \right) \\ z_{\text{min}} \end{bmatrix}, \qquad (3.12)$$

where the right hand side has to be evaluated component-wise and $\vec{r} \in B_{\text{sim}}$. Now, the components of $\bar{r} = (\bar{x}.\bar{y},\bar{z})$ are integer numbers that can for example be bitwise interleaved to generate a unique key

$$k(\bar{r}) = 2^{3\cdot21} + \sum_{i=0}^{20} \left(2^2 \cdot \bar{x}_i + 2^1 \cdot \bar{y}_i + 2^0 \cdot \bar{z}_i \right) \cdot 2^{3i}$$
 (3.13)

$$= (1 x_{21}y_{21}z_{21} \dots x_2y_2z_2 x_1y_1z_1 x_0y_0z_0)_2$$
 (3.14)

$$= (1 k_{21} \dots k_2 k_1 k_0)_8 \tag{3.15}$$

where subscript indices i denote bits counted from the LSB (least significant bit, i=0) to the MSB (most significant bit, i=20). Subscripts i and i indicate the used number system – binary or octal, respectively. The leading bit at position 63 is called the placeholder bit and ensures that each key has a defined length even if the most significant bits of the coordinates are zero. This will be necessary when also assigning keys to tree nodes later. Conveniently, k is a 64-bit number which nicely fits into an 8-byte integer. It is unique as long as the fractional part that has to be discarded in (3.12) is smaller than the nearest inter-particle distance. This restricts the approach to a spatial resolution

of $\frac{1}{2^{21}} \approx 0.47 \times 10^{-7}$ in units of the box lengths in every spatial direction. Particles that are closer to each other are assigned the same key during tree construction. To avoid this, one of them is usually slightly displaced in case of a key collision. By using longer keys, an even higher density contrast can be modeled. However, this increases the necessary storage and often prevents use of standard routines for key manipulation. It is also possible to simply limit the maximum tree level by reinterpreting tree leaves as particle containers instead of single particles. It must be stressed that this restriction on a minimum particle distance only applies to the tree construction. The maximum density contrast among physical particles is not limited at all and as indicated, the issues that arise during the tree construction phase can be mitigated straightforwardly.

Obviously, in (3.14) every group of three bits represents an additional level of refinement with respect to spatial resolution. Therefore, it is natural to use an octal notation as in (3.15). Furthermore, keys that differ only in the rightmost octal digit are all contained in a common coarser box and after truncating the rightmost digit, a unique key for this parent box is immediately found. Here, the placeholder bit comes into play: If it did not exist, a key on a higher level starting with many zeros could not be distinguished from a lower-level key. The placeholder bit guarantees all keys to start with a set bit. The extraction of the level for a given key is simply determined by the position of the leading non-zero digit:

$$level(k) = \left| \frac{\log_8(k)}{3} \right| = \frac{bit_size(k) - leadz(k) - 1}{3}. \tag{3.16}$$

In the right version, the functions $\mathtt{bit_size}(k)$ and $\mathtt{leadz}(k)$ return the total number of bits and the number of leading zeroes in k. Since both functions are part of the Fortran standard and can usually be evaluated rapidly with underlying machine instructions, this version is significantly faster than the cumbersome floating point evaluation of the logarithm.

This technique of addressing nodes inside the tree is shown in Figure 3.5. In this 2D quad-tree example, every tree level corresponds to two bits and the placeholder bit has been omitted. For finding all children of a box, its key has to be augmented by two additional bits identifying the child number. To get the parent key for some child, the last two bits of its key have to be removed. A sibling can be found by increasing/decreasing the last two bits by one without carry or borrow to/from the next level's bits. Thus, a set of simple rules for an up-/downwards and horizontal traversal is available and – although not explicitly constructed until now – the keys already define the topology of the tree. Since the keys are unique to particles and tree nodes, they present a convenient way of addressing and accessing information about them across processor boundaries, which will be used in the domain decomposition, parallel tree construction, and traversal.

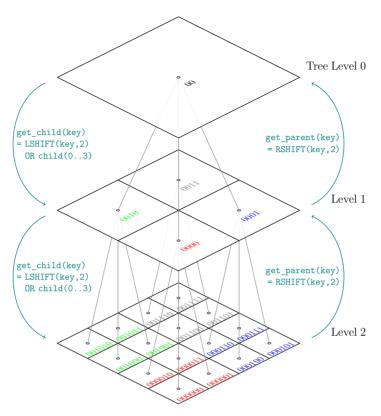


Figure 3.5: Connection between binary keys of parent and children boxes. In the 2D example, two bits correspond to each tree level; the placeholder bit has been omitted.

3.2.2 Hash table

Directly using the 63-bit key as a memory address would require $2^{64}-1$ storage locations which were mostly empty, since $B_{\rm int}$ is in general sparsely occupied. This is why, a mapping function from key- to address space has to be used. To still ensure $\mathcal{O}(1)$ access time to node data for a given key k, the very simple hashing function

$$address(k) = k & (00000000000000077777)_8, (3.17)$$

where the ampersand represents a bitwise AND() operation, is used [181]. It simply extracts the trailing part of the key. By construction, this part is mostly different for nearby particles, since they only share common leading parts of their key. Hence,

(3.17) allows for more or less uniform utilization of the available memory. Naturally, this also depends on the actual particle configuration. Especially very regular setups, i.e. particles on a regular grid will have identical trailing bits in their address. Such a configuration can stimulate collisions and larger gaps in the hash table. Another cause for duplicate storage addresses is the insertion of remote tree nodes into the local hash table during the later parallel traversal. These collisions can be resolved by introducing simple linked lists or other collision resolution mechanisms [182].

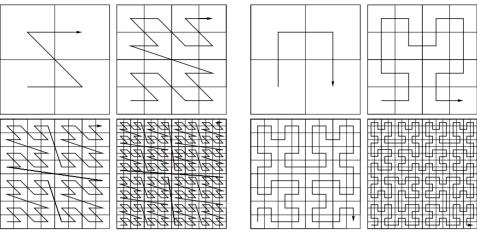
3.2.3 Domain decomposition

In addition to the global addressing feature described above, the method of assigning a unique key to each particle is a convenient starting point for the domain decomposition, i.e. a solution to the problem how the particles are distributed across the processors. This map $S: \mathbb{R}^3 \to \mathbb{N}$ defines a continuous 1D line, reaching every relevant point in space, i.e. a space-filling curve [183, 184].

Now, the curve can simply be cut into portions of equal length, that will be processed by the nodes of a distributed-memory compute cluster. Therefore, the keys have to be globally sorted and repartitioned. This can efficiently be done, for example using a parallel radix sort. An early implementation used an adaptive modification of a parallel regular sampling method [185, 186]. This proved to prevent efficient scaling beyond P=4,096 processors [187], where the sampling resolution eventually becomes insufficient without using excessive statistics, and is also handicapped by $\mathcal{O}(P^2)$ collective operations during the merge step. For this reason it has been replaced by a more sophisticated parallel sort library developed at TU Chemnitz. This uses a parallel extension of a most-significant-digit-first radix sort [182]. Details on the algorithm as well as comparisons of the new parallel sorting by partitioning to the previous approach are given in [173].

For sake of efficiency during the later parallel tree traversal, it is important, that the space-filling curve conserves locality as much as possible. That is, particles that are close to each other in real physical space should also lie nearby on S. Then, the processors get particle chunks with a large volume-to-surface ratio, which is favorable during the tree traversal since more interaction partners will already be locally available, minimizing the amount of multipole information to be communicated.

The construction rule (3.14) for $k(\vec{r})$ and hence S defines the so-called MORTON- or Z-curve, due to its characteristic form – Figure 3.6a. Naturally, there are many other ways of constructing S. For example, the HILBERT-curve [183, 184, 189], Figure 3.6b, shows much better locality since with increasing level all distances scale continuously and large jumps as in the Z-curve do not appear. In [188], both curves were implemented and analyzed. Using a domain decomposition and load balancing algorithm for the particle



- (a) The Morton- or Z-curve contains jumps of constant length even for increasing level.
- (b) For the Hilbert-curve, all internal distances scale continuously with increasing level.

Figure 3.6: Spacefilling curves can be used to map the real physical space into continuous 1D space for conveniently addressing particles and tree nodes as well as for a locality-preserving distribution of the particles across multiple processors, from [188].

keys that prefers to cut the space-filling curve at the boundaries of low-level boxes can avoid processor domains to lie across large jumps in the Z-curve to some extent. This leads to roughly equal performance using either curve.

Figure 3.7 picks up the previous example from Figure 3.1 again and uses a Z-curve for distributing the particles across three processors, which is shown by different colors here.

3.2.4 Parallel tree construction

As depicted by different colors in Figure 3.7, after performing the domain decomposition every processor possesses its own subset of particles that is as spatially compact as possible. After exchanging the boundary particles of the local domain with – with respect to the space-filling curve – both neighboring processors, all local particles are inserted into a *local tree*, that is built as far as possible towards the root until remote particles or multipole information is needed for continuation.

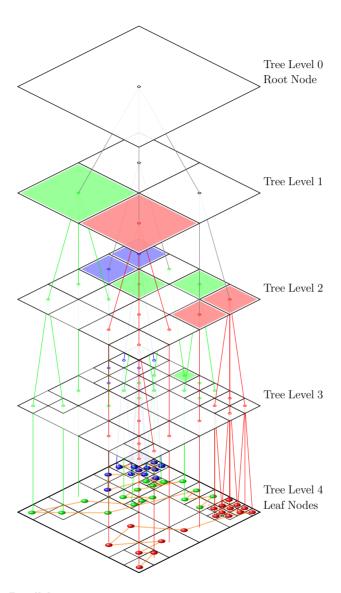


Figure 3.7: Parallel tree construction

In general, not all local particles belong to a single common branch per processor, since the local domain usually does not have the ideal shape of a low-level box. Hence, every processor can own several unconnected local trees. This is for example the case for the red processor in Figure 3.7 that owns three local trees up to their respective root nodes. The union of the local tree nodes on every processor covers the whole local domain and by adding a single lower-level node, remote domains would be touched. The roots of these local trees are considered as branch nodes and have to be globally exchanged between all processors. They are shown as colored boxes in Figure 3.7. Then, every processor can construct its global tree using fill nodes (black-colored twigs above the colored branch nodes in Figure 3.7) until reaching the common global root node of the complete simulation domain.

Now, the global tree up to the branch nodes is identical on all processors, while the local trees below the branch nodes only reside on their respective individual owner. This is the starting point for the parallel tree traversal.

3.2.5 Parallel tree traversal

The tree traversal in the parallel implementation of the BARNES-HUT tree code is in principle identical to the serial single-processor approach, see Algorithm 3.2. Again, for every particle, a depth-first search through the tree is performed to identify interaction partners by applying a multipole acceptance criterion.

However, during the traversal, it might become necessary to further resolve a branch node for which children information is locally unavailable. Here, the traversal at this particular node cannot be continued until the child data is available. In [181] it was suggested to put the unavailable node onto a request_list and to defer the traversal for the current particle. The traversal can be continued with another particle accordingly until no more traversal steps are possible for all local particles. Then, the request lists are exchanged globally and the requests are answered. Afterwards, the traversal can continue with the deferred particles. Using this technique, the exchange of tree nodes can be performed with global message exchange operations that reduce latency significantly in comparison to asynchronously sending individual requests and waiting for the answers. When interaction partners are identified during the traversal, they are put onto an interaction_list. The actual force computation is performed after the traversal to prevent the expensive floating point computations from leading to asynchronous execution of the traversal on the different ranks.

Algorithm 3.2 Parallel tree traversal: Latency Hiding Scheme

```
for all particle in particle list do
   particle.interaction list.clear()
   particle.defer list.clear()
   particle.defer list.add(root) ▷ traversals for all particles start at the root node
end for
repeat
   request list.clear()
   for all particle in particle_list do
      todo stack.clear()
      todo stack.push(particle.defer list.pop all())
      while IS INVALID(node \leftarrow todo stack.pop()) do
          if not is locally available(node) then
             request list.add(node)
                                                   ▷ node data has to be requested
             particle.defer list.add(node)

    it will be processed again later

          else
                                                                ▶ MAC evaluation
             if (MAC\_OK(particle, node)) and not (particle \in node) then
                particle.interaction list.add(node)
                                                              ▷ interaction allowed
                                     > node's children do not have to be considered
             else
                                ▶ the MAC requires the node to be further resolved
                 todo stack.push(node.children)
                                                  > proceed vertically in tree
          end if
      end while
                                            ▶ todo stack is empty for this particle
   end for
                         ▷ all particle's traversals have been done as far as possible
                    ▶ must exchange tree node data to proceed with deferred nodes
   call Send Requests (request list)
   call Answer Incoming Requests()
   call Receive Answers()
until is EMPTY(particle list[1:N].defer list))
call Process Interactions(particle list[1: N].interaction list)
```

3.3 Performance analysis for the pure-MPI tree code PEPC

In its original version, the tree code part implementation of PEPC strictly followed the ideas that were presented in Sections 3.1 and 3.2 and was thus capable of efficiently computing interactions between larger numbers of particles on distributed memory compute clusters with few thousand processors. However, despite its name, PEPC was not only a tree code for Coulomb interaction but also contained different application frontends for diverse open-boundary systems governed by long-range potentials, for example in astrophysical and fluid dynamics contexts. Thus, when analyzing the code's scalability and performance, not only the tree code kernel but also the frontend applications have to be considered. Table 3.1 lists the main parts of the complete molecular dynamics suite PEPC and their respective algorithmic scaling in terms of number of processors P and particles N. The construction of interaction lists is given there independently from the actual force evaluation, since it was implemented separately as in Algorithm 3.2. This way of implementation was expected to further reduce latency during the tree traversal by deferring the expensive floating point operations to afterwards and to exploit vectorization of the force computation.

Concerning scaling in total particle number N, the tree traversal is the dominating part. On the one hand, this is important since otherwise the administrative overhead would

Frontend	Initialization of particle properties	$\mathcal{O}(N/P)$
	$ec{r_i},ec{p_i},q_i,m_i$	
Kernel	Key construction	$\mathcal{O}(N/P)$
	$ec{r_i} ightarrow ar{ar{r}_i} = (ar{x}_i, ar{y}_i, ar{z}_i) ightarrow k_i$	
	Key sorting	$\mathcal{O}(N/P\log N)$
	k_1, k_2, k_N	
	Domain decomposition/ Key repartioning	$\mathcal{O}(N/P)$
	$k_1, k_n; k_{n+1}, k_{2n};; k_{N-n} k_N$	
	Construction of local trees and branch node exchange	$\mathcal{O}(P \log N/P)$
	Construction of global tree	$\mathcal{O}(\log P)$
	Construction of interaction lists (tree traversal)	$\mathcal{O}(N/P\log N)$
	Evaluation of forces and potentials	$\mathcal{O}(N/P\log N)$
Frontend	Particle trajectory integration	$\mathcal{O}(N/P)$

Table 3.1: Algorithmic scaling of major routines in PEPC including a generic MD frontend. The symbols N and P represent the total number of particles and processors respectively, and n = N/P. From [186].

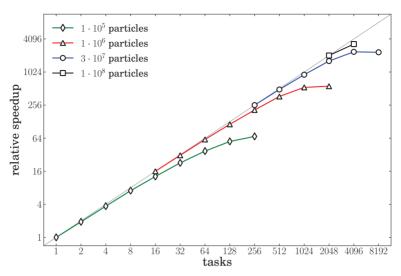


Figure 3.8: Strong scaling of pure-MPI PEPC on the IBM Blue Gene/P machine JUGENE for homogeneous setups, from [175].

dominate, which would render the method very inefficient. On the other hand, this is the part of the code where a large communication bottleneck is expected. Furthermore, it is noticeable that the branch node exchange, which scales as $\mathcal{O}(P \log N/P)$ might also play an important role for large numbers of processors – an issue, we will return to in the later Section 5.2.

Figure 3.8 shows the strong scaling of the original code for a homogeneous setup on an IBM Blue Gene/P supercomputer. Being capable to perform simulations with up to 100 million particles, the code made efficient use of up to 8,192 processors as long as N/P was large enough. However, further speedup beyond these 8,192 MPI ranks was not obtained [175].

The reduction of efficiency for the smaller setups and moderate processor numbers can be explained by the weak scaling limit, e. g. for 500,000 particles on 256 processors there are only $^{N}/P \approx 2,000$ particles per processor which is not enough work in comparison to the administrative overhead due to parallelization. For the larger configurations, scaling already saturates for $^{N}/P > 10,000$. The reason for this can be seen in Figure 3.9, which gives separate timings for the different algorithmic steps. Apparently, the tree traversal and force computation have not yet reached their weak scaling limit since both runtimes decrease linearly with the number of processors. Instead, the tree construction is beginning to dominate. Further analysis shows that due to the large number of processors and hence local trees, the average branch node level significantly

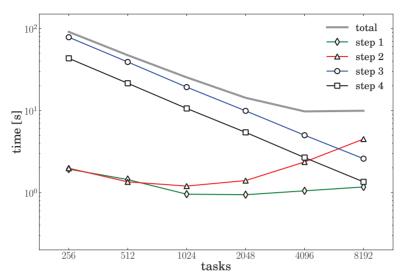


Figure 3.9: Detailed timings for the 3×10^7 -particle set of Figure 3.8. The four main steps of PEPC are also shown separately: domain decomposition (step 1), local and global tree construction (step 2), tree traversal (step 3) and force summation (step 4), from [175].

shifts towards the leaves. Consequently, their global exchange becomes more important due to the large amount of data shipped between the processors. Moreover, the global tree constructed from the branch nodes becomes significantly larger because they are lying on deeper levels. As this work has to be done by every processor and is not split across them, the non-parallel part of the algorithm increases with a larger number of processors. Several strategies were implemented to avoid or at least relocate this effect towards considerably larger processor numbers – see Section 5.2.

Apart from the branch node issue, weak scaling of the tree traversal was also unsatisfactory, which prevented the code from being used with smaller particle numbers on larger supercomputer partitions. This particular problem arises from inherent load-balancing issues between the synchronous communication stages during the traversal. As can be seen in Figure 3.10, during the alternating traversal and communication stages of a single complete traversal, load-balancing problems occur. These prevent synchronicity especially for later stages. Finally, the waiting time for other (slower) processors to start the global communication starts to dominate. For larger numbers of processors, this becomes considerably worse, since the slowest among them causes all others to wait.

In PEPC, a sophisticated load-balancing mechanism was included through distributing unequal-sized particle chunks when performing the domain decomposition. This could



Figure 3.10: Execution trace of pure-MPI PEPC on 8 processors (one per line, time is running from left to right). The excerpt only shows one complete tree traversal, that essentially consists of the traversal itself (green blocks) and synchronous data exchange (red blocks and black lines). Time is running form left to right. While the first traversal stage is finished by the processors rather synchronously, different amount of work after the first data exchange block leads to long waiting times in subsequent communications. Finally, on almost all processors communication dominates the traversal. This image was obtained using VAMPIR and VAMPIR-TRACE [190].

be achieved by using the number of interactions in the previous timestep as a heuristic measure for the chunk size instead of the number of particles. Since the particle configuration and hence the interaction partners do not change significantly between any two steps of an MD simulation, this is a good measure for equalizing the amount of work concerning the force computation. Although this is still connected to the work during the tree traversal, some randomness remains. An estimate for the necessary amount of communication and number of non-local nodes to be touched for every single particle is still missing here, which can be notably different from the number of interactions. Hence, load balancing for the tree traversal is supported but not guaranteed through this approach.

3.4 Other implementations and algorithm variations

As the Barnes-Hut tree code and its derivative algorithms have proven to be very successful in simulating N-body systems, numerous implementations, variations, and applications have been published after the original work in 1986 [158]. Historically, the Barnes-Hut method was not the very first tree algorithm of its kind. Already in 1985, Appel published a proposal for a binary tree based algorithm where interactions were performed via the monopole of particle clusters [157]. For astrophysical simulations where only positive charges (masses) are evident, this already yields good accuracy and was even shown to scale linearly in the total particle number N [191].

While for later applications, the order of multipoles has increased, the question of the optimal choice of tree topology is still controversial. Besides the original oct-tree scheme of the BARNES-HUT approach that is also used in PEPC, kd-trees were utilized and compared to the original topology [192–194]. In particular, binary trees constructed by orthogonal recursive bisection (ORB) [194–198] yield much better balanced trees than the oct-tree method but suffer from other difficulties, e.g. in general higher tree levels [199]. However, leaving the oct-tree allows for non-uniform cells and hence better adaptivity [200].

Furthermore, bringing the algorithm closer to the fast multipole method, several attempts were made to include cell-cell interactions. Already the early formulations of Warren and Salmon was prepared for these [201]. However, an implementation using the Cartesian expansion with cell-cell interaction is only reported by Dehnen [202] and later reused [198]. With the addition of an adaptive multipole acceptance criterion even sublinear scaling in the number of particles N is achieved [203]. A related approach to cell-cell interactions is the reuse of interaction lists for nearby particles [204, 205] that significantly reduces the tree traversal time by avoiding very similar traversals to be performed. The idea of grouping force evaluation points instead of particles was shown to be efficient if they outnumber the charges in the system [206].

In addition to the adaptive multipole acceptance criterion by Dehnen [203], a multitude of other criteria has been studied and error estimates were given by several groups [172, 177–179, 200, 207, 208], see also Section 3.1.3. Besides these theoretical works on tree code precision, comparative computations of PIC and tree codes for plasma applications are also reported by different authors [209, 210].

One important foundation for the tree code's versatility is its independence of the choice of interaction kernel. Accordingly, apart from the COULOMB force law it has also been used for screened COULOMB (DEBYE-HÜCKEL) potentials [211] or fluid modeling by means of the vortex particle methods [179, 212–215], for example. By construction, the tree code even allows to separate the near-field and identify close neighbor particles as necessary for the Smooth Particle Hydrodynamics (SPH) method [177, 198, 216–220].

Code	Parallelization strategy	Max. # odistrib.	of compute nodes shared mem.	$\#$ particles $\times 10^6$	
Warren	vendor spec.	512	-	17.15	[201, 224–226]
Dubinski	MPI	256		10	[196]
WDSPH-PT	MPI	128	-	101,000	[205]
GADGET	MPI	32	-	75	[217]
FLY	MPI, one-sided	1,024	_	64	[227]
Parallel Gravity	Charm++	1,024	_	5	[228]
ChaNGa	Charm++	40,960	_	730	[229]
TREE	MPI + OpenMP	255	4	10	[209, 230]
PEPC	MPI	16,384	-	256	[173, 175]
RCB-tree	OpenMP	-	24	100	[198]

Table 3.2: Overview on shared- and distributed-memory parallel Barnes-Hut tree code implementations in chronological order. Variations such as TreePM are omitted for brevity. Numbers are given as reported in the respective publications. Fields which are not applicable are marked with a hyphen "-". PEPC is included in its original version as described in this chapter.

Furthermore, the far-field can also be treated separately which lead to the development of the Tree-Particle-Mesh (TreePM) method [221–223]. Here, only close encounters are dealt with via a tree code. Several trees are integrated into a mesh, where FOURIER transform or other grid-based methods are used for determining far-field contributions. This finally leads to a forest-of-trees approach.

While work on single-processor performance optimization explicitly dedicated to multipole methods is scarce, see e.g. [231–234], many parallelized implementations have been published. A non-exhaustive selection is given in Table 3.2. In the early years of tree code development, efficient vectorization of the traversal played an important role [192, 204, 212, 235, 236] but soon has lost importance with the advent of large-scale distributed-memory compute clusters. However, the use of the computing power of Graphics Processing Units (GPU) provided new impetus to vectorized traversals and the tree code development. Simpler implementations only offload the force computation onto the GPU but thus already allow for utilization of GPU clusters since the traversal is still performed on the CPU, see e.g. [237]. More advanced techniques transfer the complete tree construction and traversal onto a single GPU [238–240]. Finally, multiple-walk technologies have been developed that allow the extension of GPU-only codes onto GPU clusters again [215, 241, 242] and allow for simulations of up to 3.3 billion particles on 576 GPUs [243].

Driven by applications such as stellar dynamics with high density contrasts, an advanced dynamic timestepping approach has been developed by Springel [217]. Using proximity information from the tree, it can identify nearby particles, adjust their individual timestep, e.g. as given in the seminal publication by Aarsett [244], and dynamically

update only modified portions of the tree. This allows for significant better resolution of close collisions or binary systems embedded in large particle clouds.

The use of tree codes in studies on properties of bulk matter leads to the need for periodic boundary conditions. These are usually implemented using the so-called tabulated Barnes-Hut-Ewald (tBHE) method [160, 176, 245, 246], see also Section 4.2.3. Finally, for simulations with DIRICHLET or VON NEUMANN boundary conditions, the tree code based boundary element method has been developed by Christlieb et al. [247].

3.5 Summary

Among the developments listed in the previous section there are numerous promising approaches for reducing time-to-solution and improving applicability to many special use cases. As seen in Table 3.2, PEPC in its original implementation as described in this chapter has already been one of the tree codes with best parallel scalability and among those with largest particle numbers. However, the solid scaling limit beyond 16 k processors and memory restrictions that prevent simulations with more than 256 million particles were already seen to result from conceptual issues. For any progress to be made, new concepts are necessary at this stage.

Concentrating on simulations of large open and periodic boundary systems in this thesis, our primary interests with respect to code development are twofold. First, we want to further improve PEPC's scalability with respect to numbers of processors as well as its capability to deal with larger particle systems. To this end we will show in Chapter 5, that with the new hybrid parallelization developed here efficient scaling of the code with 2 billion particles on up to 300 k processors can be achieved. This is to the best of our knowledge the first successful demonstration of a classical Barnes-Hut tree code on these scales. Secondly, we want to include a new approach for periodic boundary conditions, which is more precise, reliable and computationally less expensive than the tBHE approach used before. This will be addressed in Chapter 4.

4 Periodic boundary conditions

Despite PEPC's unique capability of efficiently handling configurations consisting of tremendous numbers of particles as will be demonstrated in Chapter 5, it is still far away from simulating truly macroscopic systems. As already mentioned when introducing the general N-body problem in Chapter 2, independent of the chosen algorithm, these systems can and will not be treatable with today's or tomorrow's computing and storage resources due to the vast amount of $\mathcal{O}(10^{23})$ particles involved. In fact, the present code with the modifications that will be described in Chapter 5 is able to compute forces between $\mathcal{O}(10^9)$ particles. This corresponds to a bulk cube with 10^3 particles per edge, i.e. a volume of approximately $\mathcal{O}(1\,\mu\mathrm{m}^3)$. For overdense plasmas where the external field only penetrates into the matter with the characteristic skin depth $l_{\rm skin} = c/\omega_{\rm pl} \lesssim$ 10 nm [16] and a laser focus of $10 \,\mu\text{m}^2$ this is already sufficient. However, there is also lots of interesting physics going on in the high-frequency/underdense regime, where matter is transparent for the laser and the full system volume is penetrated and accordingly larger volumes have to be studied. In addition, considering open simulation systems introduces surfaces that are not existing in reality. These have to be eliminated as they can dominate the system dynamics and modify its properties inadvertently. Hence, the simulation region has to be increased at least virtually to infinity for such bulk simulations. This is by default performed in MD simulations using periodic boundary conditions for particle movement as well as for the field and potential. To this end, the simulation region is expanded by mirror images of itself into all relevant spatial directions, leading to an infinite-sized simulation without increasing the actual number of particles. Finite-size effects and the artificial surfaces that are introduced with open boundaries and do not have a correspondence in real bulk matter are thus removed.

4.1 Periodic movement constraint

The necessary movement constraint for all particles is trivial. When using a simulation box with dimensions $(0 \dots L^{(1)}) \times (0 \dots L^{(2)}) \times (0 \dots L^{(3)})$, a simple wraparound via

$$\vec{r}_i^{(k)} = \begin{cases} \vec{r}_i^{(k)} + L^{(k)} & \text{for} & \vec{r}_i^{(k)} < 0\\ \vec{r}_i^{(k)} & \text{for} & 0 \le \vec{r}_i^{(k)} \le L^{(k)}\\ \vec{r}_i^{(k)} - L^{(k)} & \text{for} & L^{(k)} < \vec{r}_i^{(k)} \end{cases}$$
(4.1)

for all coordinate indices k and particles i is sufficient. A generalization to non cuboid box shape is straightforward. Clearly, (4.1) modifies physical variables, such as particle-particle correlations. Consequently, these have to be evaluated with caution.

4.2 Periodic forces and potential

A periodic extension of forces and potential is not as trivial. Here, (2.20) and (2.22) are augmented with an additional summation over all image boxes:

$$\Phi(\vec{r}, \{\vec{r}_1 \dots \vec{r}_N\}) = \sum_{\vec{n} \in \mathbb{Z}^3} \sum_{i=1}^N \frac{q_i}{4\pi\varepsilon_0} \frac{1}{|\vec{r} - (\vec{r}_i + \vec{n} \cdot L)|} + \Phi_{\text{ext}}(\vec{r}) , \qquad (4.2)$$

$$\vec{F}(\vec{r_i}, \{\vec{r_1} \dots \vec{r_N}\}) = \frac{q_i}{4\pi\varepsilon_0} \sum_{\vec{n}\in\mathbb{Z}^3} \sum_{i\neq i \text{ if } \vec{n}=0} q_j \frac{\vec{r_i} - (\vec{r_j} + \vec{n} \cdot L)}{|\vec{r_i} - (\vec{r_j} + \vec{n} \cdot L)|^3} + \vec{F}_{\text{ext}}(\vec{r_i}) , \qquad (4.3)$$

where \vec{n} is an index vector for all boxes. For convenience, a 3D cubic simulation region $L=L^{(1)}=L^{(2)}=L^{(3)}$ is assumed. Again, the generalization to non-cubic and even non-cuboid box shape and 2D or 1D pseudo-periodicity is straightforward. In this case, $\vec{n} \cdot L$ has to be written as the vector matrix product of \vec{n} and the matrix that is constructed from the three lattice basis vectors.

Clearly, the additional sum over \vec{n} cannot be executed directly, since it contains an infinite number of contributions. The effective total number of interactions grows linearly with the number of included mirror boxes and hence cubic with the effective simulation region size. Furthermore, due to the long-range character $\propto 1/r$ of the COULOMB potential, it converges slowly and cannot be truncated arbitrarily. However, simplifications and transformations for this sum are possible and were implemented to reduce the computational expense. These are discussed in the following.

4.2.1 Next-neighbor periodicity

The most trivial approach for emulating periodic boundary conditions is to cut the lattice cell sum in (4.2) and (4.3) after a number of contributions. For example, when restricting it to $|\vec{n}| \leq n_{\text{cut}} = \sqrt{3}$ in the 3D case, the sum only contains B = 27 terms. In a naïve implementation, this increases the necessary computational work by this factor B. However, in the tree code algorithm, the complete tree traversal can simply be performed for every box shift vector individually. Since most neighbor boxes lie farther away than the furthermost particle inside the central box, most of their interactions will involve low-level tree nodes. In the BARNES-HUT algorithm, for every particle $\mathcal{O}(\log N)$ interactions have to be evaluated. Thus, through increasing the number of cells in the

sum over \vec{n} by a factor B, the total number of interactions scales with $\mathcal{O}(N \log(B \cdot N))$, since only forces on the N particles in the central box are of interest.

The precision of this approach can be improved by simply increasing $n_{\rm cut}$, which finally results in an infinite sum for $n_{\rm cut} \to \infty$. Furthermore, due to asymmetry, the resulting potential and fields are not strictly periodic, which violates conservation of energy and momentum.

4.2.2 Nearest-image periodicity

The asymmetry that results from next-neighbor periodicity can be circumvented by using a real-space cutoff

$$\vec{r}_i - (\vec{r}_j + \vec{n} \cdot L) \le L$$
, $i, j = 1, 2, 3$, (4.4)

that has to be performed component-wise. Figure 4.1 shows this cutoff region for a sample particle.

With this technique, momentum conservation can be achieved while the potential is still not periodic due to the incomplete sum. Hence, energy conservation is still violated. Furthermore, the approach is not conveniently applicable in the BARNES-HUT algorithm, since the latter relies in particle-multipole instead of particle-particle interactions. For particle clusters, that happen to lie either side of the cutoff limit, some of their constituents would have to be included, others not. This demands resolving all tree nodes on the boundary. Since its position varies with every particle under consideration, all nodes would have to be resolved leading to an $\mathcal{O}(N^2)$ scaling.

4.2.3 Ewald summation – The tabulated Barnes-Hut-Ewald technique

The most well-known and probably most-used technique for including periodic boundary conditions in Coulomb N-body MD-simulations was originally developed by EWALD for computing crystal lattice energies [248]. It is based on splitting (4.2) into two rapidly converging series in coordinate and momentum space:

$$\Phi(\vec{r}) = \Phi_{\rm cs}(\vec{r}) + \Phi_{\rm ms}(\vec{r}) \tag{4.5}$$

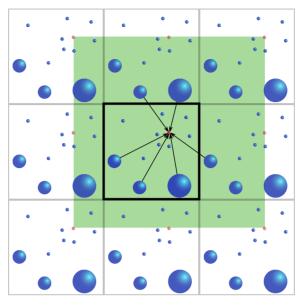


Figure 4.1: When using nearest-image periodicity, interaction outside a real-space cutoff region, here L in every spatial direction (shown as green box for the
red particle), is not allowed. This results in symmetry of included interactions for all particles inside the central simulation cell. Only some of
the included particle-cluster interactions are shown by arrows. The actual
decision for resolving clusters into particles is shown for the central box in
this picture to exaggerate the impression of periodicity. In reality it has
to be performed for the image cells separately.

with

$$\Phi_{cs}(\vec{r}) = \sum_{\vec{n} \in \mathbb{Z}^3} \sum_{j=1}^N q_j \cdot \frac{\operatorname{erfc}(\alpha | \vec{r} - (\vec{r}_j + \vec{n} \cdot L)|)}{|\vec{r} - (\vec{r}_j + \vec{n} \cdot L)|},$$
(4.6)

$$\Phi_{\rm ms}(\vec{r}) = \frac{1}{\pi L} \sum_{\vec{h} \in \mathbb{Z}^3 \setminus 0} \sum_{j=1}^N q_j \cdot \frac{e^{-\frac{\pi^2 |\vec{h}|^2}{\alpha^2 L^2}}}{|\vec{h}|^2} \cdot \cos\left(\frac{2\pi}{L} \vec{h} \cdot (\vec{r} - \vec{r}_j)\right) . \tag{4.7}$$

The dimensionless arbitrary parameter α determines the relative convergence rate of the two series. Typically, $\alpha = 2/L$ is chosen, which ensures both series to converge equally fast. When truncating the sums over \vec{n} and \vec{h} at \vec{n} , $\vec{h} > 3\sqrt{3}$, the last summand is already smaller than first one by a factor of 10^{-12} , which is sufficient for computations with limited machine precision. However, due to the vectorial character of the indices, even when truncating the sums that early, a large number of terms have to be evaluated.

Consequently, an evaluation of (4.5) is not feasible at $\vec{r} = \vec{r_i}$ for all particles i = 1...N. It can be avoided by tabulating the contribution of the image-cell particles as a correction to the classical COULOMB interaction in the form [160, 176, 246]

$$\Phi(\vec{r}) = \underbrace{\Phi(\vec{r}) - \Phi_{\text{COUL}}(\vec{r})}_{\Phi_{\text{EwCorr}}(\vec{r}) : \text{ tabulated treated as usual}} + \underbrace{\Phi_{\text{COUL}}(\vec{r})}_{\Phi_{\text{EwCorr}}}.$$
(4.8)

Here, $\Phi_{\text{COUL}}(\vec{r})$ is the regular COULOMB potential, that can be evaluated with the usual BARNES-HUT tree code algorithm. The computationally expensive correction $\Phi_{\text{EwCorr}}(\vec{r})$ is instead evaluated on a regular equidistant grid that covers the complete simulation region. Using standard linear interpolation, it is evaluated at the instantaneous particle positions $\vec{r_i}$. The electrical field is treated analogously.

This so-called tabulated Barnes-hut-Ewald (tBHE) approach is sufficiently faster than computing (4.5) and the corresponding expression for the field for every particle individually. However, the interpolation procedure introduces additional potential and field errors that undermine the periodic boundary conditions and may compromise energy conservation. Additionally, with the parameter α and the series (4.6) and (4.7) that have to be truncated depending on desired precision, the overall convergence and accuracy are rather difficult to control. Furthermore, although much theoretical work is already available [249–254], the extension of the Ewald-technique to 2D- and 1D-periodic systems requires extensive modifications. These extensions are necessary, for example, to simulate slab and wire geometries.

4.2.4 FMM-approach for periodic boundary conditions

To avoid such restrictions, in this work a parameter-free real-space summation via a renormalization approach is chosen. Therefore, an elegant approach borrowed from the fast multipole method [159, 165, 255] is adopted here. It is based on the bipolar expansion of the inverse distance in terms of LEGENDRE polynomials, already introduced in (2.43), which can be recast into

$$\frac{1}{|\vec{r}_1 - (\vec{r}_2 + \vec{n})|} = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{j=0}^{\infty} \sum_{k=-j}^{j} (-1)^j \mathcal{O}_l^m(\vec{r}_1) \mathcal{M}_{l+j}^{m+k}(\vec{n}) \mathcal{O}_j^k(\vec{r}_2)$$
(4.9)

with the multipole coefficients

$$\mathcal{O}_{l}^{m}(\vec{r} = [r, \theta, \varphi]) = \frac{r^{l}}{(l+m)!} P_{lm}(\cos \theta) e^{-im\varphi}$$
(4.10)

and the Taylor coefficients

$$\mathcal{M}_{l}^{m}(\vec{r} = [r, \theta, \varphi]) = \frac{(l-m)!}{r^{l+1}} P_{lm}(\cos \theta) e^{im\varphi} . \tag{4.11}$$

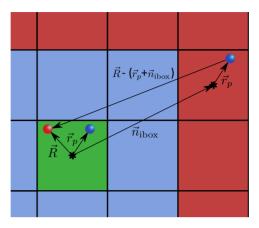


Figure 4.2: Notation for evaluating the far-field contribution with (4.17). To be able to reorder the summation, (4.18) has to be respected, which restricts the application to the red region. The near-field that is to be treated separately is shown in blue, the original (central) simulation cell as green square.

The associated Legendre polynomials $P_{lm}(x)$ differ from the previously used definition by a factor

$$P_{lm}(x) = (-1)^m P_l^m(x) , (4.12)$$

see [163] for details. Often, the Taylor coefficients are also referred to as *local expansion* since they represent an expansion around a center inside the particle cloud while the multipole coefficients correspond to an expansion in the cloud's far-field.

In addition to \mathcal{O}_l^m and \mathcal{M}_l^m , integral coefficients denoted by ω_l^m and μ_l^m , respectively, will be used in the following. These are sums over the charge-weighted coefficients for a set of particles with

$$\omega_l^m = \sum_{p=1}^N q_p \, \mathcal{O}_l^m(\vec{r}_p) \qquad \qquad \mu_l^m = \sum_{p=1}^N q_p \, \mathcal{M}_l^m(\vec{r}_p)$$
 (4.13)

and hence do not have a coordinate argument. Being connected to the real physical expansions, they implicitly carry their expansion center, though. Using the FMM operators M2M (4.31), M2L (4.32), and L2L (4.29) that will be introduced later, they can be manipulated, e.g. shifted to another expansion center or converted from multipole to Taylor expansion and vice versa. Due to linearity the operators can be applied to the pristine expansions \mathcal{O}_l^m and \mathcal{M}_l^m as well as to the integral coefficients ω_l^m and μ_l^m .

For the transformations which follow, the order of summation has to be modified. This requires absolute convergence of the series, which is only guaranteed if

$$|\vec{n}| > |\vec{a}_1| + |\vec{a}_2| \ . \tag{4.14}$$

Then, using the notation in Figure 4.2, the lattice (far-field) contribution to the potential can be written as

$$\Phi^{\text{lat}}(\vec{R}) = \sum_{\vec{n} \in \mathbb{Z}^3} \sum_{p=1}^N \frac{q_p}{\left| \vec{R} - (\vec{r_p} + \vec{n}) \right|}$$
(4.15)

$$= \sum_{\vec{n} \in \mathbb{Z}^3} \sum_{p=1}^N q_p \sum_{l=0}^\infty \sum_{m=-l}^l \sum_{j=0}^\infty \sum_{k=-j}^j (-1)^j \mathcal{O}_l^m(\vec{R}) \mathcal{M}_{l+j}^{m+k}(\vec{n}) \mathcal{O}_j^k(\vec{r_p})$$
(4.16)

$$= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathcal{O}_{l}^{m}(\vec{R}) \sum_{j=0}^{\infty} \sum_{k=-j}^{j} (-1)^{j} \underbrace{\sum_{\vec{n} \in \mathbb{Z}^{3}} \mathcal{M}_{l+j}^{m+k}(\vec{n})}_{\mathcal{L}_{l+j}^{m+k}} \underbrace{\sum_{p=1}^{N} q_{p} \mathcal{O}_{j}^{k}(\vec{r}_{p})}_{\omega_{j}^{k}} . \tag{4.17}$$

The condition for absolute convergence

$$|\vec{n}| > |\vec{R}| + |\vec{r_n}| \tag{4.18}$$

that is required for these transformations can only be guaranteed if all covered image cells that are indexed by $\vec{n} \in \mathbb{Z}^3$ lie well-separated from the central one. For example in the cubic case this means, that they may not share common box edges, as visualized in Figure 4.3. This concept of well-separatedness is of fundamental importance for FMM theory [256]. Using the well-separation parameter ws, a measure for the number of neighbor shells that are excluded from the FMM formalism for the sake of (4.18) is given. Assuming cubic box shape for now, we will use ws=1 in the following. Consequently, the edge length of the near-field region for every box – the blue region in Figure 4.3 – is 2ws+1=3. Naturally, other choices ws>1 are possible to allow for non-cubic box shapes, but in favor of the simple 3 as near-field edge length, the 2ws+1 will be omitted in the following. Our implementation uses this factor correctly to keep the formalisms flexibility, but carrying it would render the explanations unnecessarily clumsy.

The reordering of the summation in (4.17) already indicates the recipe for calculating the far-field contribution. All geometry-dependent terms are absorbed into the lattice coefficients

$$\mathcal{L}_{l+j}^{m+k} = \sum_{\vec{n} \in \mathbb{Z}^3} \mathcal{M}_{l+j}^{m+k}(\vec{n}) . \tag{4.19}$$

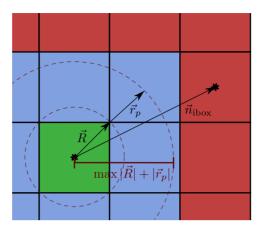


Figure 4.3: For the approach to be applicable, (4.18) has to be respected, which restricts it to the red region. The near-field that is to be treated separately is shown in blue, the original (central) simulation cell as green square.

This was already noted in [159] as a convenient approach of separating geometry parameters from particle positions. However, the lattice coefficients still demand evaluation of an infinite sum, which can be done by an efficient renormalization approach as will be shown later. Further ingredients are the multipole expansion of the central simulation box

$$\omega_j^k = \sum_{p=1}^N q_p \, \mathcal{O}_j^k(\vec{r}_p) \tag{4.20}$$

and the Taylor expansion of all far-field boxes around the simulation box center

$$\mu_l^{m,\text{cent}}(0) = \sum_{j=0}^{\infty} \sum_{k=-j}^{j} (-1)^j \mathcal{L}_{l+j}^{m+k} \,\omega_j^k \,, \tag{4.21}$$

that is composed of both. Then, the far-field contribution to the potential is given as

$$\Phi^{\text{lat}}(\vec{R}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathcal{O}_{l}^{m}(\vec{R}) \mu_{l}^{m,\text{cent}} \approx \sum_{l=0}^{p} \sum_{m=-l}^{l} \mathcal{O}_{l}^{m}(\vec{R}) \mu_{l}^{m,\text{cent}} . \tag{4.22}$$

Computing \mathcal{L}_{l+j}^{m+k} will be the computationally most demanding term for this technique since it involves many iterations over long multipole series. However, as it only depends on all lattice vectors \vec{n} and is independent of any particle positions, it can be completely pre-calculated at the beginning of the simulation and reused in every timestep or – for certain special cases such as a cubic simulation region – even be pre-tabulated. The multipole expansion ω_i^k and Taylor coefficients $\mu_l^{m,\text{cent}}$ can be computed once per

timestep and only (4.22) has to be evaluated for every individual particle. Hence, this approach shows an algorithmic scaling of $\mathcal{O}(N \cdot p)$, where p is the truncation limit for the infinite sums and hence the maximum multipole order to be included. Being the only free parameter here, it can be used to trade computational speed and memory demands against precision. Higher values of p deliver better accuracy, but demand more storage for the higher-order multipole and lattice coefficients and more computational effort for evaluating (4.22) and the necessary intermediate expressions before. However, since the BARNES-HUT tree code already uses low-order multipole expansions for approximating interactions inside the central cell, the choice of p is only of secondary relevance. Naturally, p=2 can be chosen for the evaluation of the far-field to be consistent with the approximation in the central box. However, the renormalization approach which is presented in Section 4.2.5 involves iterating over a series of multipole contributions using the standard FMM operators introduced later. Since their precision is constrained by the choice of p at least during calculation of the lattice coefficients, the multipole order has to be higher, e.g. p = 30, which yields more than 8 digits of precision [257].

The electrical field can be deduced from (4.22) by straightforward differentiation [255], that yields

$$-\vec{E}(\vec{R}) = \frac{\partial}{\partial \vec{R}} \Phi^{\text{lat}}(\vec{R}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left[\frac{\partial}{\partial \vec{R}} \mathcal{O}_{l}^{m}(\vec{R}) \right] \mu_{l}^{m,\text{cent}}$$
(4.23)

and after shifting all coordinates $\vec{r} \rightarrow \vec{r} - \vec{R}$

$$-\vec{E}(\vec{R}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left[\frac{\partial}{\partial \vec{R}} \mathcal{O}_{l}^{m}(\vec{R}) \right]_{\vec{R}=0} \mu_{l}^{m,\text{shift}} . \tag{4.24}$$

Now, the expression in brackets after differentiation contains terms of the form $|\vec{R}|^{l-1}$ that vanish for $\vec{R} = 0$ except for l = 1, so we may write

$$-\vec{E}(\vec{R}) = \sum_{m=-1}^{1} \left[\frac{\partial}{\partial \vec{R}} \mathcal{O}_{1}^{m}(\vec{R}) \right]_{\vec{R}=0} \mu_{1}^{m,\text{shift}} , \qquad (4.25)$$

which after expanding the sum and subsequently making use of the relations $\mathcal{O}_l^{-|m|} = (-1)^{|m|} \mathcal{O}_l^{|m|^*}$ and $\mu_l^{-|m|} = (-1)^{|m|} \mu_l^{|m|^*}$ between the expansion coefficients and their complex conjugates leads to

$$-\vec{E}(\vec{R}) = 2 \left[\operatorname{Re}\{\mathcal{O}_{1}^{1'}(0)\} \operatorname{Re}\{\mu_{1}^{1,\text{shift}}\} - \operatorname{Im}\{\mathcal{O}_{1}^{1'}(0)\} \operatorname{Im}\{\mu_{1}^{1,\text{shift}}\} \right] + \mathcal{O}_{1}^{0'}(0)\mu_{1}^{0,\text{shift}}.$$

$$(4.26)$$

The prime denotes the gradient, that reads

$$\begin{pmatrix} \partial_x \\ \partial_y \\ \partial_z \end{pmatrix} \left(\mathcal{O}_1^0(0), \mathcal{O}_1^1(0) \right) = \begin{pmatrix} 0 & 1/2 \\ 0 & -i/2 \\ 1 & 0 \end{pmatrix}$$

$$(4.27)$$

and finally for the electrical field

$$\vec{E}(\vec{R}) = -\begin{pmatrix} \operatorname{Re}\{\mu_1^{1,\text{shift}}\} \\ \operatorname{Im}\{\mu_1^{1,\text{shift}}\} \\ \operatorname{Re}\{\mu_1^{0,\text{shift}}\} \end{pmatrix}. \tag{4.28}$$

The superscript index shift still denotes, that an expansion around $-\vec{R}$ is used instead the expansion around the cell origin, that results from (4.21). Fortunately, the standard FMM formalism delivers the *local-to-local* (L2L) operator [165, 257]

$$\mu_l^m([\vec{r} - \vec{b}]) = \sum_{j=0}^{\infty} \sum_{k=-j}^{j} \mathcal{O}_{j-l}^{k-m}([\vec{0}], \vec{b}) \mu_j^k([\vec{r}]) , \qquad (4.29)$$

that translates a TAYLOR expansion to a different center, that is given in square brackets. Here it can be applied via

$$\mu_l^{m,\text{shift}} = \sum_{j=0}^{\infty} \sum_{k=-j}^{j} \mathcal{O}_{j-l}^{k-m}(-\vec{R})\mu_l^{m,\text{cent}} . \tag{4.30}$$

It replaces (4.22) and has to be evaluated together with (4.28) for every particle and l = 1 and m = 0,1 to compute the electrical field.

4.2.5 Renormalization approach for the lattice coefficients

The derived expressions (4.22) and (4.28) allow the far-field contribution to the electric potential and field to be computed conveniently and – with only an $\mathcal{O}(N)$ overhead – in an efficient way. However, the lattice coefficients (4.19) as the only geometry-dependent parameters are still to be determined. Being composed of an infinite sum over all far-field lattice cells, they cannot be evaluated directly.

Using an EWALD-like approach, CHALLACOMBE et al. did this for the first time in the FMM framework [256]. However, a similar formulation by Nijboer that is also based on spherical multipole expansions already existed before the Fast Multipole Method has been invented [258]. Kudin and Scuseria followed and extended these works. They proposed a renormalization scheme for evaluating the emerging infinite sums and provided a formulation for arbitrary cell geometry for force and stress tensor [259, 260].

First developments on a pure real-space renormalization approach for the summation over the infinite number of lattice sites using recurrence relations for the multipole coefficients were performed by Berman and Greengard [261] and rigorously extended by Kudin and Scuseria [257] Their efficient renormalization approach, is also utilized in the implementation of periodic boundary conditions in PEPC that is described in the following.

The standard FMM operators multipole-to-multipole (M2M)

$$\omega_l^m([\vec{a} + \vec{b}]) = \sum_{j=0}^l \sum_{k=-j}^j \mathcal{O}_{l-j}^{m-k}([\vec{0}], \vec{b}) \omega_j^k([\vec{a}]) =: \mathcal{O}([\vec{0}], \vec{b}) \triangleleft \omega([\vec{a}]) , \qquad (4.31)$$

and multipole-to-local (M2L)

$$\mu_l^m([\vec{a} - \vec{r}]) = \sum_{j=0}^{\infty} \sum_{k=-j}^{j} \mathcal{M}_{l+j}^{m+k}([\vec{0}], \vec{r}) \omega_j^k([\vec{a}]) =: \mathcal{M}([\vec{0}], \vec{r}) \otimes \omega([\vec{a}]) , \qquad (4.32)$$

where the previously omitted expansion centers are now given in square brackets again, are used to shift multipole expansions to a different center and to transform multipole coefficients into local Taylor coefficients. Naturally, the infinite sums in (4.29) and (4.32) can be truncated at j = p if p is the maximum order of multipoles to be used. Only for the M2L operator, does this introduce an additional error, since higher-order expansions are necessary anyway due to the sums in the indices.

The formal binary operator notation defined on the right hand side of (4.31) and (4.32) will be used later to simplify notation and improve readability. This notation as well as the algorithm that is based on it and will be deduced here loosely follows the one given in [257].

In this section, a generic lattice with linearly independent basis vectors $\vec{t_1}, \vec{t_2}, \vec{t_3} \in \mathbb{R}^3$ will be considered. With the index vector $\vec{k} \in \mathbb{Z}^3$, the shift vectors to the lattice cells are

$$\vec{q} = \sum_{c=1}^{3} k^{(c)} \cdot \vec{t}_c ,$$
 (4.33)

where as before the superscript indices signify vector components. Clearly, periodicity in less than all three spatial dimensions can be handled by fixing the respective coordinates in \vec{k} and hence in \vec{q} to zero. Now, the TAYLOR expansion around the origin due to the cell's contribution at \vec{k} can be deduced by shifting it with the previously introduced M2L operator (4.32) as

$$\mu_{\vec{k}} = \mathcal{M}(\vec{k}) \otimes \omega , \qquad (4.34)$$

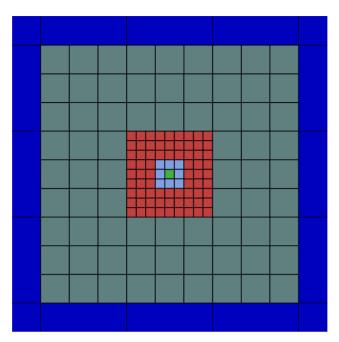


Figure 4.4: Regions to be considered in the derivation of the lattice coefficient algorithm. The original (central) simulation cell **OC** is drawn as a green square, its near-field region **NF** in light blue. The first far-field layer **FF**' is composed of multipole expansions of the central $3^0 \times 3^0 \times 3^0$ cell **OC** that are shifted to appropriate positions. The second far-field layer **FF**" is built out of replicas of the central $3^1 \times 3^1 \times 3^1$ block **OC** \cup **NF** and the next far-field layer **FF**" comprises $3^1 \times 3^1 \times 3^1$ blocks.

where ω is the multipole expansion of the central cell. Since in this expression the center of expansion is fixed at the origin, multiple such terms can be summed to get the expansion for a group of replica cells. In the following, this will be done in a very efficient way to ideally cover the complete periodically filled spatial domain.

Figure 4.4 shows several colored regions around the central cell at $\vec{k} = (0,0,0)$ which will be used in the derivation of an iterative hierarchical process for approximating the infinite lattice sum. The TAYLOR expansion of the far-field region, that is shown as red and teal domains but also lies beyond these, is given by

$$\mu = \sum_{\vec{k} \notin (OC \cup NF)} \mathcal{M}(\vec{k}) \otimes \omega . \tag{4.35}$$

The regions OC and NF denote the original cell and its near-field. Colors correspond to those in Figure 4.4. Due to linearity of the M2L operator, ω can be factored out and only

$$S = \sum_{\vec{k} \notin (OC \cup NF)} \mathcal{M}(\vec{k}) \tag{4.36}$$

has to be computed as the geometry dependent part.

As the starting point for the hierarchy of super-cells that will be used, all cells are considered where $|\vec{k}| \geq 2$ and $k^{(i)} \leq 4$ for i=1,2,3. This region that lies in the far-field of the central cell $\vec{k}=(0,0,0)$ but in the near-field of the central $3^1\times 3^1\times 3^1$ super-cell, will be denoted as **FF**' and is shown in red color in Figure 4.4. With the definition of the translational moments

$$\mathcal{M}^{\star} := \sum_{\vec{k} \in \mathbf{FF}'} \mathcal{M}(\vec{k}) , \qquad (4.37)$$

the Taylor expansion in the sense of (4.35) is

$$\mu_0 = \mu' = \mathcal{M}^* \otimes \omega' \tag{4.38}$$

with

$$\omega' = \omega , \qquad (4.39)$$

that now includes contributions of the $3^2 \times 3^2 \times 3^2$ super-cell without the central (OC \cup NF). It can be seen as a lowest-order approximation to the far-field expansion.

The cells lying in the near-field of **FF**' will be denoted as **FF**". They are shown in teal in Figure 4.4 and cover a $3^3 \times 3^3 \times 3^3$ region with the central $(OC \cup NF \cup FF')$ cut out. As **FF**" is essentially constructed of sub-blocks which are comparable to $(OC \cup NF)$, it can be built out of $3^1 \times 3^1 \times 3^1$ contributions of the form

$$\omega'' = \mathcal{O}^{\star} \lhd \omega' \,, \tag{4.40}$$

with

$$\mathcal{O}^{\star} := \sum_{\vec{k} \in (\text{OC} \cup \text{NF})} \mathcal{O}(-\vec{k}) . \tag{4.41}$$

Using the (with the operator $\mathcal{U}_{\mathcal{M}}()$) that will be defined later) rescaled translational moments from before, the Taylor expansion of **FF**" is

$$\mu'' = \mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star}) \otimes \omega'' \tag{4.42}$$

and together with the contribution of **FF**' yields

$$\mu_1 = \mu_0 + \mu'' \tag{4.43}$$

$$= \mathcal{M}^{\star} \otimes \omega' + \mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star}) \otimes \omega'' . \tag{4.44}$$

The next level of super-cells (that already lies outside the scope of Figure 4.4) can be included by also rescaling \mathcal{O}^* and applying it twice to get the multipole expansion

$$\omega''' = \mathcal{U}_{\mathcal{O}}(\mathcal{O}^{\star}) \triangleleft \omega'' \tag{4.45}$$

of a $3^2 \times 3^2 \times 3^2$ super-cell. It has to be shifted further outward again, to yield

$$\mu''' = \mathcal{U}_{\mathcal{M}}(\mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star})) \otimes \omega''' \tag{4.46}$$

and together with the previous contributions

$$\mu_2 = \mu_1 + \mu''' \tag{4.47}$$

$$= \mathcal{M}^{\star} \otimes \omega' + \mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star}) \otimes \omega'' + \mathcal{U}_{\mathcal{M}}(\mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star})) \otimes \omega''' . \tag{4.48}$$

Since the multipole and TAYLOR coefficients $\mathcal{O}(\vec{r})$ and $\mathcal{M}(\vec{r})$ are polynomials in $|\vec{r}|$, see (4.10) and (4.11), the scaling operators, increasing all lengths by a factor of (2ws+1), become trivial multiplications

$$\mathcal{U}_{\mathcal{O}}(\mathcal{O}_l^m) := \mathcal{O}_l^m(\vec{r} \to (2ws+1)\vec{r}) = (2ws+1)^l \cdot \mathcal{O}_l^m , \qquad (4.49)$$

$$\mathcal{U}_{\mathcal{M}}(\mathcal{M}_{l}^{m}) := \mathcal{M}_{l}^{m}(\vec{r} \to (2ws + 1)\vec{r}) = \frac{\mathcal{M}_{l}^{m}}{(2ws + 1)^{l+1}}.$$
 (4.50)

The series $\mu_0, \mu_1, \mu_2, \ldots$ rapidly converges towards μ since more and more larger portions of the complete far-field are included. However, for the geometry-dependent coefficients to be pre-calculated, it is necessary to factor out the multipole expansion ω of the central cell. It depends on all particle positions and hence changes from one simulation timestep to the next. Using (4.39), (4.40), and (4.45) in (4.38), (4.42), and (4.46), the multipole expansion ω of the central cell always appears rightmost:

$$\mu' = \mathcal{M}^* \otimes \omega \tag{4.51}$$

$$\mu'' = \mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star}) \otimes [\mathcal{O}^{\star} \lhd \omega] \tag{4.52}$$

$$\mu''' = \mathcal{U}_{\mathcal{M}}(\mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star})) \otimes [\mathcal{U}_{\mathcal{O}}(\mathcal{O}^{\star}) \triangleleft [\mathcal{O}^{\star} \triangleleft \omega]]$$
(4.53)

and hence can be factored out in the sense of (4.36). Due to linearity of the operators, the identities

$$\mathcal{M} \otimes [\mathcal{O}_1 \triangleleft \mathcal{O}_2] = [\mathcal{M} \otimes \mathcal{O}_1] \otimes \mathcal{O}_2 \tag{4.54}$$

$$\mathcal{O}_1 \triangleleft [\mathcal{O}_2 \triangleleft \mathcal{O}_3] = [\mathcal{O}_1 \triangleleft \mathcal{O}_2] \triangleleft \mathcal{O}_3 \tag{4.55}$$

$$\mathcal{U}_{\mathcal{M}}(\mathcal{M}) \otimes \mathcal{U}_{\mathcal{O}}(\mathcal{O}) = \mathcal{U}_{\mathcal{M}}(\mathcal{M} \otimes \mathcal{O}) \tag{4.56}$$

hold [257, 262]. These can be applied – after some reorganization – to yield expressions of the form $\mu_n = S_n \otimes \omega$ with

$$S_0 = \mathcal{M}^* \tag{4.57}$$

$$S_1 = \mathcal{M}^* + [\mathcal{U}_{\mathcal{M}}(\mathcal{M}^*) \otimes \mathcal{O}^*] \tag{4.58}$$

$$= \mathcal{M}^{\star} + [\mathcal{U}_{\mathcal{M}}(S_0) \otimes \mathcal{O}^{\star}] \tag{4.59}$$

$$S_2 = \mathcal{M}^* + [\mathcal{U}_{\mathcal{M}}(\mathcal{M}^*) \otimes \mathcal{O}^*] + [[\mathcal{U}_{\mathcal{M}}(\mathcal{U}_{\mathcal{M}}(\mathcal{M}^*)) \otimes \mathcal{U}_{\mathcal{O}}(\mathcal{O}^*)] \otimes \mathcal{O}^*]$$
(4.60)

$$= \mathcal{M}^{\star} + [\mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star} + \mathcal{U}_{\mathcal{M}}(\mathcal{M}^{\star}) \otimes \mathcal{O}^{\star}) \otimes \mathcal{O}^{\star}]$$

$$(4.61)$$

$$= \mathcal{M}^{\star} + [\mathcal{U}_{\mathcal{M}}(S_1) \otimes \mathcal{O}^{\star}]. \tag{4.62}$$

Clearly, this series of partial sums can be recast into the recursive scheme

$$S_0 = \mathcal{M}^* \tag{4.63}$$

$$S_{n+1} = \mathcal{M}^* + [\mathcal{U}_{\mathcal{M}}(S_n) \otimes \mathcal{O}^*] \tag{4.64}$$

that eventually converges to the lattice sum

$$S = \lim_{n \to \infty} S_n \ . \tag{4.65}$$

The resulting lattice coefficients finally can be inserted into (4.35), or (4.21) with $\mathcal{L}=S$, respectively. Informally, this renormalization approach to compute the lattice coefficients can be interpreted as an iterative process that scales the current box expansion, shifts it to every super-block position at the next level and adds the inner **FF**' ring again.

During the process, it can be noted that the monopole and dipole contributions in S do not converge. The reason for this is a physical one: For the monopole, summing over an infinite number of charged cells leads to an infinite total potential, while the electrical field remains unchanged since only a global potential shift is experienced. While the approach can in principle handle such systems by simply setting the monopole contribution to zero, the resulting lattice energy will be wrong (though not needed in molecular dynamics simulations) and thus — being unphysical anyway — these cases should be avoided. The dipole sum however, is conditionally convergent. It is corrected for by a special procedure, that is explained in more detail in Section 4.2.7.

Using this technique, after n iterations

$$N_n^{\text{FF}} = \left[3^{n \cdot \dim} - 1 \right] (2ws + 1)^{\dim} \tag{4.66}$$

far-field boxes are included in the lattice sum, where dim is the number of spatial dimensions where periodicity is desired. Thanks to this rapid growth, already after n=32 iterations with a maximum multipole order of p=20 and maximum length of Taylor expansions 2p=40, most digits of double precision are converged and the

high-precision results for cubic lattices of [256, 257, 262] are reproduced with sufficient precision. As already mentioned, the approach is completely general due to the free choice of the lattice basis $\vec{t}_{1,2,3}$. It can be applied to non-cubic and even non-cuboid lattices as long as (4.18) is fulfilled, which can be achieved by increasing ws if necessary. Since the calculation of lattice coefficients only has to be performed once per simulation at a cost of few seconds runtime, it has no significant impact on the total runtime and has not been parallelized. The code section, that finally translates the far-field TAYLOR expansion to every particle for evaluating (4.22) and (4.28) is inherently parallel, since with the exception of the already pre-calculated μ_0^{cent} , no global information is necessary.

4.2.6 Near-field

Considering again Figure 4.4, the original cell \mathbf{OC} is handled with the classical Barnes-Hut tree code and the far-field $\mathbf{FF'} \cup \mathbf{FF''} \cup \mathbf{FF'''} \cup \cdots$ is treated via the FMM-renormalization approach. This still leaves the near-field region \mathbf{NF} , that – due to the constraint for absolute convergence (4.18) – cannot be included into the far-field. In the implementation in PEPC, it is simply dealt with as explained in Section 4.2.1 on next-neighbor periodicity through inclusion of these cells into the central cell's tree traversal. All features of this approach that were already mentioned still hold true, but with the FMM far-field technique, this next-neighbor periodicity is no longer an auxiliary construction but the brick that fills the gap between near-field and far-field. Naturally, a too non-cubic cell violates the convergence criterion for the far-field, as can be seen from Figure 4.3. However, this flaw can be mitigated by increasing ws and hence the near-field itself appropriately.

4.2.7 Dipole correction

As shown in [263, 264] and further extended in [265], the electrostatic potential in a lattice can be written as a sum of two contributions

$$\Phi(\vec{r}, C, L) = \Phi^{\text{int}}(\vec{r}) + \Phi^{\text{ext}}(\vec{r}, C, L) ,$$
(4.67)

where C and L denote the shape, volume and other specifications of the unit cell and the complete macroscopic crystal under consideration, respectively. Being only a function of position \vec{r} , the *intrinsic potential* Φ^{int} is a periodic function, while the *extrinsic potential* Φ^{ext} is non-periodic and strongly depends on the choice of the unit cell and the surrounding region, i. e. the position of the cell under consideration inside the crystal. Explicit expressions for Φ^{ext} can be given for some special cases, e.g. an ellipsoidal region or a cubic or spherical region, that will be important for the implementation of

periodic boundary conditions here. For the extrinsic potential, the expression

$$\Phi^{\text{ext}}(\vec{r}, C, L) = \frac{4\pi}{3V}(\vec{r} - \vec{r}_0) \sum_{k=1}^{N} q_k \vec{r}_k - \frac{2\pi}{3V} \sum_{k=1}^{N} q_k |\vec{r}_k|^2$$
(4.68)

can be derived from the general ellipsoidal case. The sums run over all particles in the central cell and hence represent the cell's dipole and trace of its quadrupole moment, \vec{r}_0 denotes the center of the region under consideration and V is the central cell's volume. Note that the expression given in [256] lacks the explicit volume factor, since unit cells are assumed there.

When performing molecular dynamics simulations, periodic boundary conditions are applied to emulate an infinitely extended bulk system to avoid discontinuities in the potential when applying periodic movement constraints, i.e. when inserting particles that left the simulation region on the opposite side. This is necessary to ensure conservation of total energy and momentum. As shown in [263, 264], the standard EWALD procedure as given in Section 4.2.3 provides an intrinsic potential and hence does not show any dependence on the overall shape of the crystal (that in fact does not come into play at all here). This also holds true for the minimum image convention from Section 4.2.2 – see [265].

The pure real-space sum for supporting periodic boundary conditions implemented here, clearly assumes a certain shape of the surrounding crystal matter. Being extended by an equal number of replicas into every periodic direction it is an upscaled version of the unit cell. Although the surfaces lie far away from the central cell, the truncation of the iteration (4.64) in fact defines this shape. Consequently, the method delivers the sum of extrinsic and intrinsic contributions. For it to be usable in molecular dynamics simulations with periodic boundary conditions, the extrinsic potential has to be subtracted so that

$$\Phi^{\text{lat}} = \Phi^{\text{far-field}} - \Phi^{\text{ext}} , \qquad (4.69)$$

where $\Phi^{\text{far-field}}$ is the result of (4.22). The correction for the electrical field can be found as the gradient of (4.68):

$$\vec{E}^{\text{lat}} = \vec{E}^{\text{far-field}} + \frac{4\pi}{3V} \sum_{k=1}^{N} q_k \vec{r}_k \tag{4.70}$$

with $\vec{E}^{\text{far-field}}$ from (4.28). This technique corrects for the conditionally convergent dipole sum in the lattice coefficient algorithm.

As already mentioned, (4.68) is only valid for cubic boxes. An extension for cuboid cell shapes can be derived from the ellipsoid expression given in [265]. A generalization to non-cuboid cells cannot be performed by this approach, though. For applications

with non-orthogonal unit-cells, such as triclinic systems, an approach for eliminating the conditionally convergent dipole contribution is proposed in [259, 260]. It is based on adding a number of particles that compensate for the cell's dipole moment without modifying overall physical properties. Therefore, additional charges $q_{1,2,3,4}$ are placed at the (in the periodic system) translationally equivalent points

$$\vec{p_1} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \vec{p_2} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \vec{p_3} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \vec{p_4} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \quad (4.71)$$

$$q_1 = \frac{\vec{d}^{(1)}}{L^{(1)}}, \qquad q_2 = \frac{\vec{d}^{(2)}}{L^{(2)}}, \qquad q_3 = \frac{\vec{d}^{(3)}}{L^{(2)}}, \qquad q_4 = -(q_1 + q_2 + q_3).$$
 (4.72)

Here, $\vec{d}^{(i)}$ is the *i*th component of the cell dipole moment and $L^{(i)}$ are the respective box lengths. The coordinates of the virtual particles are given in the lattice basis system and obviously correspond to unit cell corners. While these particles indeed annihilate the dipole moment of the central cell and even modify the overall field inside, after periodic continuation, they in fact compensate for each other. In a completely periodic system they reside in the same spot and show a total charge of zero. This approach can also be used for 2D and 1D periodic systems, where again (4.68) cannot be applied directly.

4.3 Verification of correctness and computational overhead

In order to verify correctness of our implementation of the approach for periodic boundary conditions, we return to the melting NaCl crystal that was also used in Section 3.1.3 for analyzing the influence of different choices for the multipole acceptance parameter ϑ . In Figure 4.5 we show the same graph as before in Figure 3.4a but with periodic boundary conditions now. Since a direct evaluation of the sum over periodic mirror cells is not possible, we compare with data obtained using a standard EWALD implementation that is part of the ScaFaCoS library [156, 180]. Obviously, the overall behavior of the error is unchanged and it goes down to $< 10^{-6}$ for $\vartheta = 0$. The approach for periodic boundary conditions does not have any free parameters besides the maximum multipole order p and hence is not depending on ϑ . Accordingly, the primary contributions at $\vartheta > 0$ still result from the multipole approximation that is performed for the central box and the near-field region here, as the far-field contribution to the errors should be visible in the plot as a shift that is globally applied to the curves when compared to Figure 3.4. For $\vartheta = 0.3$, all errors are still bound below 1.5% which is acceptable for our simulations. This results seems to be general as for all test systems, from a very regular Madelung crystal to randomly arranged periodic systems, from comparisons

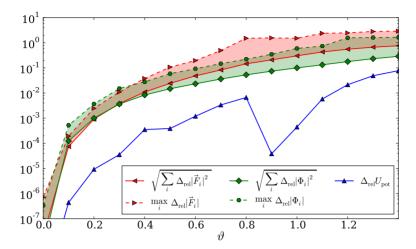


Figure 4.5: Precision comparison for a periodic system and different multipole acceptance parameters ϑ . The graphs show the maximum and root-mean-square relative errors for potential, field and total potential energy compared between PEPC and results from an EWALD approach. The example system considered for this comparisons is the same 12,960 particles open-boundary extract of a melting NaCl crystal as in Figure 3.4 but is now periodically extended to infinity. Reference values were produced using a standard EWALD implementation from the ScaFaCoS library [156, 180]. For $\vartheta=0.0$ the relative errors reach 10^{-14} which is expected as the numerical limit for double precision computations.

with literature values and other methods, e.g. a high-precision FMM [262], it is seen that the main error in potential and forces results from the BARNES-HUT tree code algorithm. Errors introduced from the far-field are usually smaller by two orders of magnitude.

Especially in simulations of periodic systems over a long physical time, the errors introduced can accumulate if they have a dominant direction. This is for example the case if the dipole correction from Section 4.2.7 is omitted. Eventually, such an error accumulation will lead to violating the conservation of total energy. However, seen in Figure 4.6, no significant energy drift in fully periodic systems is observed for $\vartheta \leq 0.5$. This is not the case for the trivial nearest-image method which underlines the importance of a precise treatment of periodicity in our molecular dynamics simulations.

The computational overhead introduced by the approach for periodic boundary conditions implemented here is moderate. The lattice coefficients only have to be computed

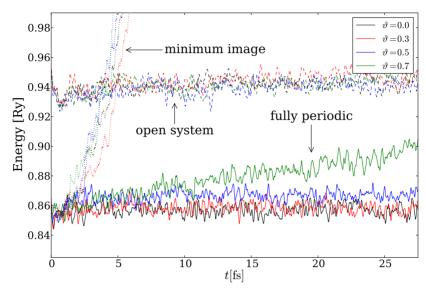


Figure 4.6: Verification of energy conservation with periodic boundary conditions. For different values of the multipole acceptance criterion ϑ , the evolution of the total energy for a cubic system of $N_{\rm el}=N_{\rm ion}=1,000$ charges in equilibrium is studied. For the open and the fully periodic system with the approach presented in this chapter, energy conservation is confirmed for small enough multipole acceptance parameters $\vartheta \leq 0.5$. In contrast, the system that only comprises nearest-image periodicity, see Section 4.2.2, suffers from significant heating due to the discontinuous forces experienced by the particles that are wrapped to the opposite side of the simulation box by means of the periodic movement constraint, Section 4.1. Data for the open, i. e. non-periodic, system is included for getting an estimate of the energy fluctuations that are evident due to the tree code's multipole approximation.

once per simulation and are independent of the number of particles N. Thus, we can omit them in this analysis. The evaluation of (4.20) and (4.21) have an algorithmic scaling of $\mathcal{O}(p \cdot N)$ and $\mathcal{O}(p^3)$, respectively because they are only evaluated once per timestep. The computation of the far field contribution to the force (4.28) with (4.30) has to be performed for every individual particle and accordingly scales as $\mathcal{O}(p^2 \cdot N)$. All these algorithmic scalings can be neglected in comparison to the $\mathcal{O}(N \log N)$ scaling of the Barnes-Hut tree code part that is used for the central cell and the near-field. In the 3D-case, the near-field consists of 26 boxes, so that the number of interaction partners grows by a factor if 27 when including periodic boundaries. Accordingly, the central-

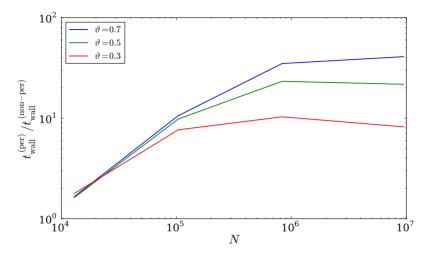


Figure 4.7: Ratio of runtime for periodic and non-periodic setup. Again, the melting NaCl crystal from Figure 3.4 is used – now with different total particle numbers. Starting for $N \geq 10^6$ the runtime factor introduced due to periodic boundary conditions starts to saturate in accordance to the algorithmic scaling of the approach for periodic boundaries. Consequently, for $N \to \infty$, the additional runtime needed is negligible.

cell and near-field contributions yield an $\mathcal{O}(N\log(27N)) \propto \mathcal{O}(N\log(N) + N\log 27)$ which is an additional linear term that vanishes in comparison to the $N\log N$ part for $N\to\infty$. Thus, the overall scaling of the BARNES-HUT tree code is not compromised by the method. This can also be seen in Figure 4.7 that shows the ratio of the runtimes for periodic and non-periodic setups. There, for $N\geq 10^6$ saturation of the additional factor to the runtime of the non-periodic system occurs as the $\mathcal{O}(N\log N)$ -term begins to dominate the overall runtime again.

4.4 Summary

With the renormalization approach for the lattice coefficients and their FMM-based utilization, an efficient technique for including periodic boundary conditions into the parallel BARNES-HUT tree code PEPC has been implemented. It overcomes the limitations of the tBHE approach, that has been implemented experimentally in an older version of the code [246] and even adds further features. Being able to consistently deal with 1D-, 2D-, and 3D-periodic systems in cubic, non-cubic and even non-cuboid ge-

ometries, it prepares the code for simulations of wire-, slab-, and bulk-shaped systems. It is free of any arbitrary parameters and can be driven up to machine precision. However, since the tree code itself currently employs expansions up to quadrupole order in the central cell and its near-field, precision of the far-field is not too critical for our purposes. Since the order of included multipole expansions p is variable in our approach, though, it is also prepared for backends with higher expansions for the central cell in the future.

As the code is now able to perform bulk simulations with virtually an infinite number of particles without increasing the actual particle number, it is applicable for the studies on optic and dielectric properties in solids that will be topic of Chapter 7.

While all effects of artificial surfaces and finite simulation size are mitigated using periodic boundary conditions, partially open system still have to be handled with care. For example, when studying laser-particle acceleration in a slab geometry using 2D periodic boundaries, the periodicity also replicates the laser beam, which is usually not the desired behavior. Here, still the only solution is to increase the system size by increasing the number of simulation particles. However, since such setups will not be of interest in this work, their specialties will not be discussed any further here.

Based on the implementation developed in this work, currently a version for 2D-COULOMB interaction

$$\Phi \vec{r} = \sum_{i=1}^{N} q_i \log \left(\frac{|\vec{r} - \vec{r}_i|}{L} \right) \tag{4.73}$$

with $\vec{r}, \vec{r}_i \in \mathbb{R}^2$ and a metric constant L is being developed [266]. The potential (4.73) represents the interaction of charged parallel rods that perpendicularly penetrate the x-y plane. This can, for example, be used as a model for particles that are confined to external magnetic field lines. Their rapid movement along these lines is averaged to form a charged rod that moves slowly perpendicular to the field lines.

5 Multi-level parallelism

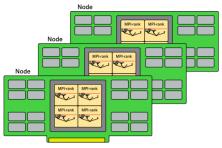
5.1 Hybrid parallelization

Most of today's supercomputers incorporate at least two levels of parallelism: They consist of large numbers of distributed memory nodes, that in turn provide a number of processors sharing common main memory. Currently, the tendency is for the total number of cores to increase, while the amount of memory per node and the total number of nodes as well as the performance of each individual processor have stagnated. For example, the IBM Blue Gene/P machine JUGENE in Jülich had 73,728 nodes, each consisting of 4 processor cores and 2 GB of main memory [267]. Its successor, the Blue Gene/Q supercomputer JUQUEEN consists of 28,672 nodes and 16 processors with doubled clock rate per node, which are theoretically most efficient if they are oversubscribed by a factor of four [268]. This means that every processor should be occupied by four concurrent program threads to obtain maximum performance. Thus, while the compute power per node increased by a factor of approximately

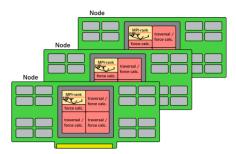
$$2 \text{ (clock ratio)} \times 4 \text{ (processor number ratio)} \times 4 \text{ (optimal threads per processor)} = 16$$
, (5.1)

the amount of memory per node only increased by a factor of 8 to 16 GB. Accordingly, it is necessary to make use of the possibility of processors on the same node to share common data instead of duplicating it or even exchanging information via the message passing system. On such a hybrid machine it is clearly inefficient to spawn 16 or even up to 64 MPI ranks per compute node. Given these constraints it is natural to also add another level of parallelism to the tree code that makes as efficient use of the available hardware parallelism as possible. This can be achieved by only using one MPI rank per compute node and starting 1...4 sub-threads per processor – Figure 5.1.

Besides avoiding duplicate storage, the idea of hybrid parallelization offers a multitude of further advantages. First of all, the total number of MPI ranks is reduced since there is only one rank per node instead of one rank per processor. This reduces the number of communication partners in global as well as in point-to-point message exchange. Furthermore, due to the locality and self-similarity of the space-filling curve that has been used for domain decomposition, the particles that reside on neighboring processors have similar interaction lists anyway. Hence, they request the same remote tree



(a) pure-MPI approach: one rank per core, lots of identical data on shared memory nodes



(b) hybrid: one rank and accordingly one hash table and tree per sharedmemory-node, several worker-threads per rank

Figure 5.1: In contrast to the previous pure-MPI approach, the hybrid parallelization avoids duplicate storage and makes use of the available shared memory parallelism.

nodes. By having only one MPI rank and a tree common to all local processors, those requests only have to be sent once instead of on a per-processor-basis, which reduces the total number of messages that have to be exchanged drastically. Secondly, the reduced number of MPI ranks leads to fewer but larger particle chunks to be distributed during the domain decomposition. This simplifies load balancing across the nodes and postpones load balancing issues to the shared-memory parallelism level, where they can be corrected on the fly as will be shown below.

The newly developed algorithmic techniques and implementation details described in this chapter have partially been published in [269].

5.1.1 An MPI + Pthreads tree traversal

To make use of the two-level parallelism available in hardware, the tree traversal as the most time-consuming part of the code has been parallelized accordingly using a hybrid approach of MPI and Pthreads. Distributed memory parallelism using the *Message Passing Interface* (MPI) was already implemented before and follows the latency hiding scheme that has been presented in Section 3.2. However, besides the already mentioned advantages of a hybrid parallelization, it was also intended to allow for a native overlap of computation and communication during the tree traversal. Therefore, the exchange of tree nodes had to be separated from the tree traversal with its MAC evaluations and the calculation of the actual interactions. Whereas it was performed in a separate subsequent step previously, in the new implementation the interactions are evaluated

immediately, allowing further communication latency to be hidden behind the expensive floating point computations.

Intra-node parallelism is implemented using the standard *POSIX Threads* (Pthreads) facilities, that are available on any POSIX-based operating system and thus maintaining the excellent portability of PEPC. The so-called *threads* represent different execution paths within a single application, executed in parallel if shared memory parallelism is available. Otherwise, the operating system's scheduler assigns them processor time slices to virtualize this parallelism. In contrast to sub-processes, threads share a common address space with each other and their calling program. Despite implementation difficulties with calling C-library functions from Fortran code, these were chosen in favor of OpenMP, a more standard approach for shared memory parallelism. In contrast to OpenMP, Pthreads offer a great deal of better flexibility, especially when very different tasks are to be handled by the different threads. A similar feature has only been included in the most recent version 3.0 of OpenMP.

In Figure 5.2, a flowchart shows the technical details of the approach, a pseudo-code overview is given in Algorithms 5.1 and 5.2. After some initialization, the tree traversal, that is pictorially called walk here, spawns an adjustable number of worker-threads. These distribute the local particle chunk among themselves and perform the tree traversal for their individual particle subsets in parallel. Therefore, they also administrate their own share of the global todo-list, that contains the deferred nodes etc. for all particles. At some point, during their traversals the threads need tree node properties from remote processors. As in the latency-hiding scheme in Algorithm 3.2, the respective parent key is put onto a request-list and the particle traversal is deferred until the necessary data is available. Furthermore, a flag in the hash-table is set for the respective parent node to avoid duplicate requests by other threads. The request-queue itself is shared across the worker threads. It is continually processed by the main thread of the process, responsible for all communication tasks, working as the *communicator thread*. It sends requests for child node properties from the local request queue to remote MPI ranks and answers their requests with data from the local hash table. Upon reply on its own requests, the communicator thread receives tree node data from remote ranks, inserts it into the local hash table and sets the *children-available*-flag for the respective parent node.

After processing other particles from its allocated chunk, the worker thread that originally posted the request, will eventually return to the deferred particle. Since the children-available flag for the previously unavailable node is now set, it can proceed with the traversal here and – depending on the MAC – perform the interaction or request further nodes.

An excerpt of the traversal is additionally shown in Figure 5.3. It illustrates the path taken by an individual request over time. During its alternating task of tree traversal and force computation (green and red), thread 2 hits an unavailable node that it needs in

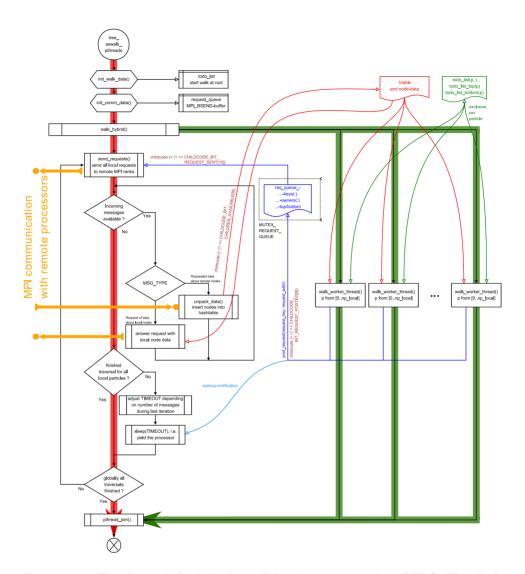


Figure 5.2: Flowchart of the hybrid parallelized tree traversal in PEPC. The thick red arrow represents the communicator (main) thread, green arrows represent worker threads. Information about the request queue length and the amount of received messages during the previous sleep periods of the communicator are used to adjust its suspend period length.

Algorithm 5.1 Hybrid parallelized traversal: Worker-threads

```
while still particles unfinished do
   {\rm take}\ K \leq {\tt MaxParticlesPerThread}
   for k = 1 to K do
       pick tree node from particle's todo list
       evaluate MAC
       if MAC ok then
          call Interact (with node)
       else
          if node's children locally available then
              put all children to front of particle's todo list
          else
              call PostRequest(for child nodes)
              put node on end of particle's todo list
          end if
       end if
   end for
end while
```

Algorithm 5.2 Hybrid parallelized traversal: Communicator-thread

```
while not globally all traversals finished do
   while incoming requests in queue do
      answer with requested data from hash table
   end while
   if not locally all traversal finished then
      while answer for own requests in queue do
          insert received data into hash table
          notify worker about new data
      end while
      while local request list not empty do
          send requests to remote communicators
      end while
      adjust timeout
      sleep
   end if
end while
```

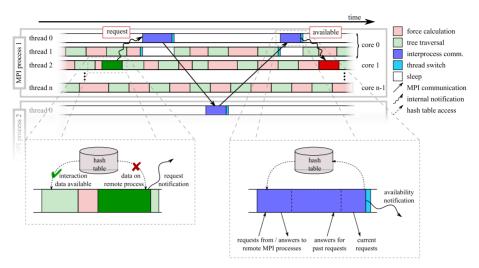


Figure 5.3: Schematic flow of the hybrid parallelized PEPC. For every MPI rank, there is one communicator thread (thread 0) that handles all inter-node communication via MPI and n several worker threads which perform tree traversal and force computation. As the communicator is not permanently occupied, it can share its processor core with one worker thread. Therefore, an advanced scheme for thread switching has been implemented to reduce answer latency caused by sleeping communicators. Figure taken from [269].

order to proceed. It notifies the communicator thread 0 about the missing information and resumes traversal and force computation with some other particle. In the meantime, the communicator forwards this request to the recipients MPI rank and receives the corresponding data at a later time. After insertion into the local hash table, the worker threads are informed by setting the appropriate flag. Now, thread 2 can continue the traversal for the deferred particle.

As demonstrated in Figure 5.4, the implementation of this hybrid parallelization approach works just as expected: All four cores are permanently occupied – three of them with tree traversal and force computation and the other one with communication and computation alternately. The latter is done since there is not enough work in the communicator thread to keep a single computer core permanently occupied. Hence it is expedient, that the communicator's core is shared with another worker. This exactly resembles the projected behavior from Figure 5.3. Additionally, by funneling all MPI communication through the primary thread, no thread-safe MPI implementation is necessary.

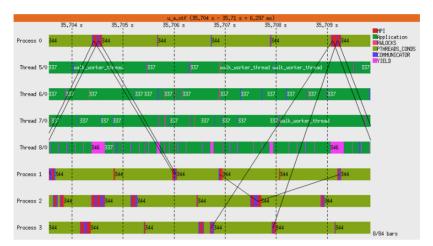


Figure 5.4: Execution trace of the hybrid parallelized PEPC. Only for MPI rank 0, all five threads that are running on the four available cores are shown. For ranks 1 to 3, only the respective communicator thread is visible. Communication (red blocks and black lines) nicely overlap. Furthermore, the alternating cycles of sleeping (olive and magenta blocks) and working (other colors) of the communicator and the last worker thread are visible. This execution trace was obtained using VAMPIR and VAMPIRTRACE [190].

Not only has an additional layer of parallelism been successfully added to the code, but the complete communication scheme benefits from the implemented modification. Before, a globally synchronized bulk-transfer of requests and answers was used. Now, communication is performed in a completely asynchronous point-to-point scheme. Requests are sent and answered as they occur, answers are received and incorporated into the hash table as soon as they are received. Instead of all ranks having to be synchronized, only two communication partners are participating and short latency is, though optimized for, not strictly necessary. The advantage of this scheme for load balancing will be addressed later. However, the approach also poses new difficulties on the internode networking hardware, that is regularly optimized for the often-appearing global communication patterns or only nearest-neighbor communication. Here, fast random point-to-point connections with near and far compute nodes are necessary. Architectures like the Blue Gene/P and Blue Gene/Q with their torus networks already allow for fast communication paths also to far-away compute nodes, and thus are well suited for this technique.

5.1.2 Performance results for the new approach

In the previous section, the newly developed completely asynchronous and hybrid parallelized tree traversal for the classical BARNES-HUT tree code has been presented as a replacement for the previously implemented latency-hiding scheme of WARREN and SALMON. With it, the processors on shared-memory nodes can make use of common administrative and particle data. Through our approach, a native overlap of communication and computation is obtained. Furthermore, the code's memory demands are reduced significantly, since

- 1. interaction lists are not necessary due to the immediate force evaluation,
- 2. all processors on a shared memory node use a common tree,
- 3. fetched data from remote processors is only stored once per node, and
- 4. further administrative data is shared between the processors.

Besides these clear advantages over the MPI-only implementation with one rank per processor, the implementation of the hybrid parallelized tree traversal made an extensive rewrite of the code necessary. This significantly improved code stability and allowed for a substantial restructuring, grouping many different interaction backends and application frontends in a monolithic framework. These extensions will be discussed later in Section 5.3.

In order to test the algorithm, code extensive scaling experiments were performed on the IBM Blue Gene/P machine JUGENE at Jülich Supercomputing Centre on up to 294,912 processors and with up to N=2,048,000,000 particles. Previously, the code was not able to run on these large scales due to prohibitive memory demands. With the new approach, this became possible, see Figure 5.5. Moreover, as depicted in Figure 5.6, the new implementation exhibits high efficiency for homogeneous (Figure 5.6a) as well as inhomogeneous (Figure 5.6b) particle configurations. The previous scaling limit beyond 8,192 processors from Figure 3.8 has been removed. Now, force computation for extremely large systems with hundreds of millions of particles in reasonable wall-clock times is in reach. The evaluation of the mutual forces between 2 billion particles, that would have been taking almost 15 minutes on 4,096 JUGENE cores (if it was possible at all with the previous implementation), now needs less than 15 seconds on the whole machine.

Since the combined tree traversal and force computation routine is the most time-consuming and demanding part of the algorithm as regards scalability, Figure 5.6 concentrates just on this part. As the scaling results are shown to be comparable for the different configurations, we study performance data for the homogeneous setup in more detail. While the tree traversal shows good parallel efficiency, the complete tree code also includes the previously mentioned steps of domain decomposition and local as well as global tree construction with the previous branch node exchange. Figure 5.7 ana-

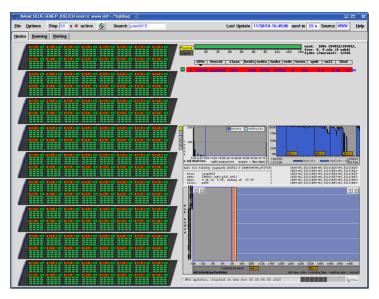
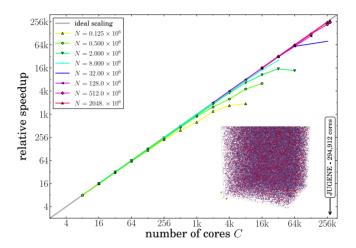


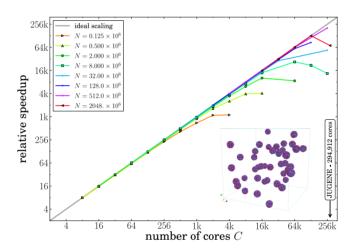
Figure 5.5: LLView [270] screenshot with PEPC running on all 294,912 processors of JUGENE. Before implementing the hybrid parallelization, this was not possible at all.

lyzes the two primary contributions to the overall runtime of a single simulation step for the homogeneous setup with different particle numbers N. While the tree traversal shows the good scaling already demonstrated, the exchange of branch nodes starts to dominate the total runtime as soon as the number of particles per core drops below N/C = 2,000. For simulations with few hundred cores, the exchange is completely negligible. However, the number of branch nodes per rank has empirically been seen to be constant, see [188, 269], which leads to a linear growth $\mathcal{O}(P)$ of their total number with the number of MPI ranks P, compare also Section 5.2. This issue that is intrinsic to the parallel tree code has been significantly mitigated by the hybrid parallelization because there is only one MPI rank per compute node instead of per compute core now. Hence, the overall number of branch nodes and consequently their impact on runtime has been reduced with the developments described in this chapter. To prevent the collective branch node exchange from finally dominating the runtime for very large parallel computers that comprise hundred thousands of distributed-memory nodes, some additional improvements were tested and implemented. They will be discussed shortly in Section 5.2.

To estimate the parallelization overhead, the weak scaling properties of PEPC are shown in Figure 5.8 with especially small workloads per core, i. e. small particle numbers. Since



(a) Setup with a homogeneously filled, overall charge neutral cube.



(b) The inhomogeneous setup is composed of 42 randomly positioned neutral spheres.

Figure 5.6: Strong scaling of the tree traversal and force computation part of the algorithm for different total particle numbers N and compute cores C for the homogeneous (top) and inhomogeneous (bottom) particle distributions. The insets show the test particle setups used for the performance analysis, from [269].

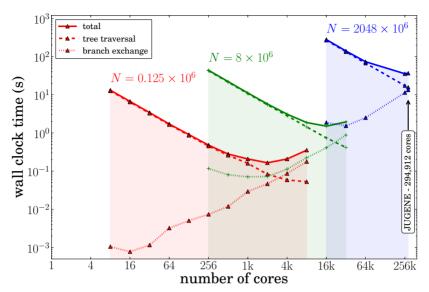


Figure 5.7: Total runtime for the homogeneous setup with different particle numbers. Besides the total, also figures for the two dominating contributions are given, from [269].

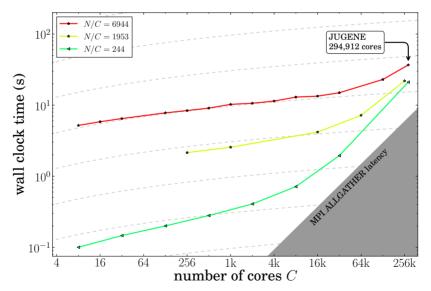


Figure 5.8: Weak scaling of the full code with a homogeneous setup and different numbers of particles per processor, from [269].

the number of particles has to be increased linearly with the number of processors here, an algorithm with nonlinear complexity – as the Barnes-Hut tree code, that scales as $\mathcal{O}(N\log N)$ – cannot yield constant runtimes in this analysis. Instead, the expected behavior is drawn with dashed lines in Figure 5.8. At 6,944 particles per processor, ideal scaling up to 16 k cores and acceptable scaling up to C=288 k cores is achieved. At the lower extreme with 244 particles per core, the surprisingly good scaling already saturates around C=8 k cores, i.e. P=2 k MPI ranks. Again, this has to be attributed to the $\mathcal{O}(P)$ contributions from the remaining collective operations that appear during the tree-build phase. However, as already indicated, fresh developments have been started to mitigate this issue – see Section 5.2.

Despite the non-ideal efficiency for small systems, with its impressive scaling up to hundreds of thousands of processors with particle numbers up to 2 billion, PEPC is well positioned to make efficient use of upcoming supercomputer generations with increasing numbers of cores on every shared memory node.

5.1.3 Intra- and inter-node load balancing

As already mentioned in Section 3.3, PEPC already included a sophisticated load balancing mechanism to evenly distribute the actual work across the MPI ranks. This is performed during the domain decomposition step by using the number of interactions from the previous simulation timestep as the weight of each particle. Instead of evenly distributing the particles across the processors, their total weight per rank is equalized. The new hybrid parallelized implementation especially profits from this approach, since traversal and force computation (which is in fact balanced by construction) are now performed in a monolithic step. In contrast to the earlier version [175], the alternating stages of traversal and force computation do not have to be balanced individually. Furthermore, due to the hybrid parallelization, the total number of MPI ranks is reduced, which simplifies the inter-node load balancing due to a reduced number of particle chunks.

To summarize, the benefit from this already implemented load balancing is actually enhanced with the hybrid parallelization. The overall effect can be seen in Figure 5.9, where for the homogeneous setup with 2 billion particles, the traversal time is reduced by approximately 8%. Even better gain up to 15% is reached for other configurations.

By adding the second level of parallelism, an additional level of load balancing has to be included: The workload for the different processors on every shared-memory node must be balanced. This is achieved by a straightforward competition-technique among the worker threads. They are not allowed to grab all particles at once but must fill their short todo-list successively. Starting from todo-lists of equal length, threads that finish the traversal for a particle may take a new one from the shared pool. Thus, threads that process particles with less work automatically grab more of them, balancing the total work on the node so that all local worker threads finish synchronously.

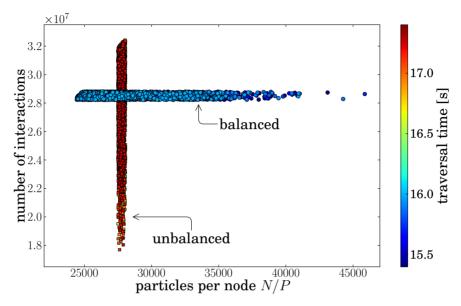


Figure 5.9: Number of computed interactions for a random subset of the 73,728 MPI ranks with 288 k processors simulating 2,048,000,000 homogeneously distributed particles for unbalanced and balanced workload. Colors encode the respective traversal time, including force calculation. Balancing the number of interactions instead of the number of particles yields a reduction of traversal time by approximately 8%, from [269].

5.2 Branch nodes

Naturally, the overall performance of the parallel tree code strongly depends on the overhead that is introduced due to construction and distribution of the tree data structure. As already seen in Figure 5.7, the hybrid parallelized tree traversal – which is the main part of the algorithm – shows a better parallel efficiency than the branch node exchange as soon as the number of particles per core reaches too small numbers. Further analysis shows that with increasing numbers of MPI ranks P, the position of the minima $^{N}/c$ increases linearly. This is to be expected, since with increasing P, the branch nodes move to higher tree levels and hence their total number increases. Since all of them have to be exchanged globally, the respective MPI_ALLGATHERV() call must transfer more and more data to more and more endpoints. This essentially prevents efficient scaling at a certain point.

5.2.1 A-priori branch node estimation

The branch nodes are intrinsic part of the parallel tree code concept – their global exchange cannot be circumvented. Consequently, it is necessary to define an optimal set of branch nodes to reduce their total number. This can be performed in the hashed octtree scheme by comparing the local minimum and maximum particle key on every rank with the highest key of the left and the lowest key of the right neighbor, respectively. The first digit where the keys differ, determines the tree level of the respective branch nodes on both processors as was shown in [188, 269]. With this concept of branch nodes, optimality with respect to tree level is guaranteed: Their level and hence their global number is as small as possible. This could not be proven with other branch finding algorithms that introduced many high-level branches if particles only occupied certain corners of the simulation box.

Empirical tests showed that using this new branch finding approach the number of branch nodes per rank is essentially constant with different number of ranks P – Figure 5.10. Consequently, their total number grows linearly with P and their global exchange leads to the stagnation of parallel efficiency in Figures 5.7 and 5.8. However, with the provable optimality of our branch concept, this scaling constraint cannot be mitigated any further and new approaches are necessary. One of them is sketched in the next section.

5.2.2 Hierarchical branch node clustering

One approach to circumvent the scaling bottleneck that results from further increasing global branch number with increasing P makes use of the hierarchical structure of the exchanged data. In the oct-tree, a maximum of eight branch nodes can contribute to a common parent node in the tree. The idea behind the hierarchical branch node clustering is to find all ranks that contribute to common ancestor nodes on a certain predefined tree level blev. This clustering can efficiently be performed using MPI_COMM_SPLIT() calls with extracts of the branch keys as splitting property ("color" in the MPI interface definition) to create a set of sub-communicators that correspond to the respective ancestor nodes. Their members are then allowed to exchange their contributing branch nodes and to compute the common super branch node.

The number of communication partners and hence the amount of exchanged data in this step is limited by blev. Finally, all super branches are globally exchanged to construct the global part of the tree on all ranks. The number of super branches is significantly smaller than the total number of branches and is – with fixed blev – limited again. Consequently, with proper choice of blev, this mitigates the scaling limit that results from the branch node exchange.

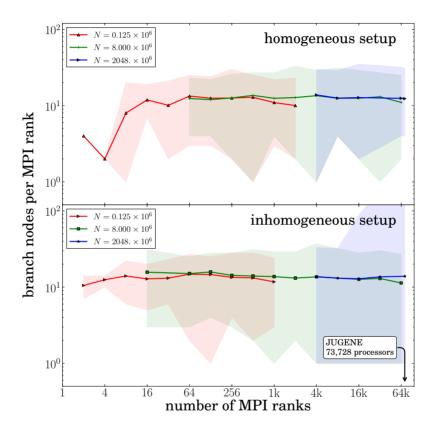


Figure 5.10: Average number of branch nodes per MPI rank and respective minimum and maximum regions (shaded) for the homogeneous and inhomogeneous setup with different particle numbers. For details on the particle configurations see Figure 5.6. From [269].

It is even possible to sweep blev from the maximum tree level towards root to allow for a completely parallel hierarchical tree construction. A first experimental implementation [271] of this new approach shows promising results and will soon be included into PEPC.

5.3 Modularity, applications, and further prospects

In this chapter, only the most important cornerstones of PEPC's evolution during the last three years could be outlined. However, these were not performed in isolation from other applications besides plasma physics. On the contrary, the code's structural redesign and a modular concept developed as a byproduct here provided new impetus to explore a much broader field of applications. Some details are given in Section 5.3.1 before Sections 5.3.2 to 5.3.4 discuss different paths of development that can be pursued in future.

5.3.1 Modularity and further applications

As already indicated, together with the hybrid parallelization of PEPC, an extensive rewrite of the code was undertaken. In the course of this revision process, a rigorous modular concept was developed that cleanly separates the tree code algorithm from the physically relevant parts such as force law, units, particle configurations, diagnostics, etc. The idea is outlined in Figure 5.11. Through specification of mandatory interface routines and data types, the tree code algorithm is kept separate from the *interaction specific backend routines* and the *frontend applications*.

Since the multipole-acceptance criterion is part of the backend, data-dependent variants can be implemented and the algorithm itself can even be exploited to identify the n nearest neighbors of any particle as it is necessary for $smooth\ particle\ hydrodynamics$ (SPH). The tree code itself does not need any knowledge about the nature of the particles apart from their position in space. All physical features such as the creation of the multipole expansion, its propagation through the tree and the interaction routines are provided by the backends. The frontend developer can select a certain backend—for example one that contains all COULOMB-relevant features—or create a new one and develop his frontend based on the supplied data types. When invoked from the frontend, the tree code kernel transparently calls the interaction routines as needed and behaves like a large-scale parallel N-body solver.

These modifications that make the code more of a developer's framework than a library, also allowed a number of development branches of PEPC to be reunified with the main trunk. Currently, it supports five different backends, that range from COULOMB

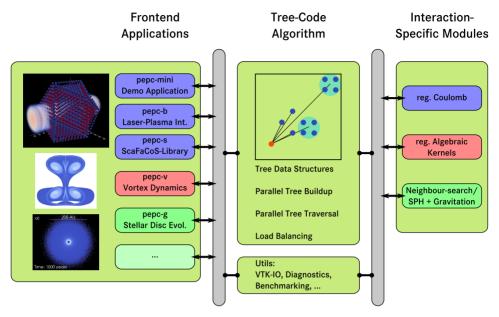


Figure 5.11: Modularity of PEPC. Mandatory interface routines and datatypes (gray boxes) separate the tree code from the interaction-specific backends and the frontend applications. Furthermore, several utilities e.g. for convenient VTK output and tree diagnostics are supplied.

potential via regularized kernels, that are used in fluid simulations with the *vortex* particle method [179], to nearest-neighbor search kernels for SPH studies of stellar disc evolution [219, 220]. Especially the latter application demonstrates the versatility of the modular concept. It makes use of the possibility to include different force laws and multipole acceptance criteria within a single backend and uses the tree traversal for identifying partners for gravitational interaction as well as neighbor particles for the hydrodynamical near-field interaction. Therefore, the MAC is modified in a clever way to gather a list of the closest neighbor particles instead of interaction partners.

Among the currently used 14 frontends are applications for electrostatics and strongly coupled Coulomb systems [104, 176], laser-ion acceleration [169, 171, 272], inclusion of self-generated magnetic fields via a Darwin model [273, 274], plasma-wall interaction in Tokamak fusion vessels [275, 276], instabilities in magnetized plasmas [277], the fluid vortex-method [179] and SPH [219, 220] studies and finally library interfaces for inclusion of the code in third-party libraries. Furthermore, PEPC has become part of the highly scalable Coulomb solver library ScaFaCos [156, 180], that is currently being developed within a BMBF project together with partners from Jülich, Chemnitz,

Bonn, Stuttgart, and Bielefeld and is recently being used together with parallel-in-time approaches for even better scalability [278, 279], see Section 5.3.4.

The approach for including periodic boundary conditions for COULOMB systems presented in Chapter 4 is another good example of the code's modularity. Being implemented solely for the COULOMB backend, it does not interfere in any way with the other backends and immediately with its completion, all frontends using this particular backend were able to make use of it. For example, the fusion-vessel application performs its simulations on plasma-wall interaction in box that is elongated in one direction and periodically continued in the other two.

5.3.2 Projected technical optimizations

While the code is optimized for convenient usability and capability to support many different applications as well as good scalability, *single-core optimizations* were bypassed during this development stage. In particular, the efficient use of hardware features, such as the different memory caches and specialized hardware instructions for rapid floating point evaluation offers great potential for further performance gain [231–234, 280].

Especially the tree traversal with its large number of floating point operations during force computation and lots of random memory access would benefit from such advancements. Clearly, other parts of the code can be improved, too. For example, a hybrid Pthreads/OpenMP+MPI parallelization of the tree assembly phase including the computation of the multipole properties can allow for a more efficient use of the available hardware. In addition, a task-based parallelization using the new programming concept SMPSs (SMP Superscalar) based on the hybrid parallelized tree traversal is currently under development within the framework of the EU TEXT (Towards EXaflop applicaTions) [281] and Mont-Blanc [282] projects. The latter aims at porting supercomputer applications to extremely energy efficient ARM processors. With the task-based approach, the programmer defines code sections and data dependencies between them and a runtime scheduler decides when and even where (e. g. on the main processor, specialized accelerators such as many-integrated core (MIC) modules or available graphics processing units (GPU), or specialized energy efficient processors) a task is executed. This will allow for even better adaptivity to heterogeneous architectures.

5.3.3 Further hybrid parallelism and task-based approaches

As already indicated in Section 5.2.2, by performing a hierarchically parallelized tree construction, the required amount of communication during this stage can be reduced significantly. In addition, a thorough study of the properties of the hashing function (3.17) might reveal that there are better choices available [182]. These can reduce the number

of collisions and hence access latency and improve memory utilization by avoiding large unused regions in the address space. Furthermore, additionally allowing cluster-cluster interactions and then shifting the result towards all child particles instead of only performing particle-cluster interactions can further reduce the necessary number of force evaluations [198, 202, 203]. In combination with a more adaptive MAC such a modification can make the tree code competitive with the fast multipole method with respect to scaling in total number of particles N if only limited accuracy is requested. This was shown in [203] for a tree code, that was – in contrast to PEPC – only parallelized using OpenMP and thus limited to shared memory parallel computers. Ultimately, such a modification could lead to a two-sided MAC, where either local or remote clusters are refined depending on the error incurred by the different approximations. While this will move the algorithm more in the direction of the fast multipole method, it will still keep all the advantages of the tree code, such as dynamic load balancing, arbitrary cell geometry, exchangeable interaction kernel. Dropping the latter feature, an additional Coulomb backend using spherical expansions instead of the Cartesian multipoles can allow for arbitrary multipole order which will be a significant advantage for high-precision computations, where currently ϑ has to be chosen inconveniently small. Finally, an efficient prefetching during traversal can further reduce communication latency: Based on a worst-case MAC estimation, the receiver of a tree node request can already decide to send the requested child nodes and possibly some higher-level additional nodes below that will definitely be needed in a later step of the traversal anyway.

5.3.4 Additional directions of parallelization

Even when implementing all of the improvements proposed above, every particle-based simulation code will reach an intrinsic scalability limit when it comes to very small particle numbers per processor. Then, the amount of work to be done locally is vastly exceeded by the additional overhead due to parallelization. One approach to make use over further available computing capabilities and get beyond the weak scaling limit has been developed during the last few years under the label *Parallel-In-Time* computation.

Here, in addition to the spatial decomposition that reached its saturation point, another direction of parallelism is added by distributing the different timesteps to different processors as demonstrated in Figure 5.12, that extends Figure 5.1b in a natural way. Therefore, the *Parallel Full Approximation Scheme in Space and Time* (PFASST) can be employed [283, 284]. Based on interlaced iterations of Parareal [285] and a Spectral Deferred Correction (SDC) scheme [286], it is a parallel-in-time solver for initial value problems

$$\frac{\partial u}{\partial t} = f(t, u), \qquad u(t = 0) = u_0, \qquad (5.2)$$

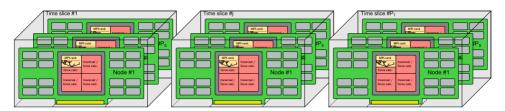


Figure 5.12: The combination of PEPC with PFASST uses $P_T \times P_S$ nodes. Spatial decomposition for each of the P_T time slices is performed by PEPC using P_S nodes within each PEPC-communicator (depicted as one box in the figure). Within PEPC, one MPI-rank per node is used to act as data and communication management thread, while the other cores perform the traversal of the tree using Pthreads, see Section 5.1. For PFASST, this structure is duplicated P_T times to create independently running instances of PEPC. PFASST connects the *i*th node of each box to one new MPI communicator, which results in P_S separated PFASST-communicators for the temporal decomposition, from [278].

which can be applied for the molecular dynamics problem with Newtonian dynamics. The method considers every timestep as a separate initial value problem that gets its initial value as the result of the previous timestep. These problems are solved in parallel starting from approximative solutions to the previous timestep's problem. Therefore a coarse time propagator \mathcal{G} is used to generate approximative solutions at later points in time and a fine propagator $\mathcal F$ is utilized to correct these. While $\mathcal G$ essentially has to be evaluated in serial, \mathcal{F} is computed in parallel for all timesteps. Hence, good scalability can be obtained if \mathcal{G} is computationally cheap and \mathcal{F} is expensive. With its multipole acceptance parameter ϑ – see Section 3.1.3 – the BARNES-HUT tree code offers the possibility to conveniently choose between fast and precise computation. With the corresponding choice $\vartheta(\mathcal{G}) \gg \vartheta(\mathcal{F})$, this combination of the parallel-in-space tree code PEPC and the parallel-in-time method PFASST has already been demonstrated successfully, see [278, 279] for details. Actually performing different physics for both propagators, e.g. a separation of electron and ion dynamics or using hybrid implicitexplicit integration schemes will go beyond this simple approach and has the potential to revolutionize parallel molecular dynamics simulations as the otherwise strongly separated time scales can be studied in a consistent model within one simulation.

5.4 Summary

As shown in this chapter, the hybrid parallelization and modular re-implementation of the Pretty Efficient Parallel Coulomb Solver that was conducted in this work has lead to major advancements in scalability and usability of the tree code algorithm for a multitude of applications. Many minor improvements that evolved out of this way

lead to further improvements for efficiency and reliability. The code has evolved from a monolithically closed parallel Coulomb solver to a parallel Barnes-Hut tree algorithm framework that can cope with billions of particles and efficiently make use of several hundred thousand processors. As a result, it is to the best of our knowledge the first successful demonstration of utilizing a classical Barnes-Hut tree code on a petascale parallel compute cluster. Previously, it was not known whether this was possible at all. With its wide portability, the two-layer parallelism and sophisticated inter- and intra-node load balancing, the code is well prepared for upcoming many-core and even heterogeneous architectures. Furthermore, the support for periodic and pseudo-periodic boundary conditions with arbitrary unit-cell shape that has been included during this work offers the opportunity to bridge the gap from open-boundary microscale systems to bulk matter within a single simulation code.

Naturally, development is not being stopped at this stage. As already indicated, there is still space for future prospects. Especially in algorithm sections that played – as the tree traversal is the dominant part – a secondary role in optimization until now, a hybrid parallelization approach is required to make use of the available shared-memory parallelism. Furthermore, the time-parallelism promises a whole dimension for new developments. However, in the following chapters of this thesis, we will concentrate on applications. We will first use PEPC to study collective phenomena in confined nano plasmas in Chapter 6, where the code's capabilities to efficiently deal with large, open-boundary systems are particularly advantageous. In Chapter 7, we make us of the extension to accurate periodic boundaries from Chapter 4 to simulate collisional heating in bulk matter.

6 Collective electronic properties in nano clusters

If light is interacting with metallic nanostructures, it does not only couple to the collective plasma oscillation as in bulk matter. In fact, there exists a multitude of additional oscillation modes of the free electrons in the metal which can be excited. These result from either volume oscillations with selected wave numbers or surface excitations. In addition, when going from an isolated atom via atomic dimers and clusters of more and more particles, the quasi-continuous band structure of the metal only evolves gradually [287]. The distinct energy levels of the atom split into two components at the transition to the dimer and further fragment to a discrete spectrum that is evident for a finite-size cluster of atoms [288]. Thus, single-electron properties as well as collective features are modified in nanostructures when comparing to either isolated atoms and bulk matter. In this respect, nano structures are more than only large molecules and also significantly different to simple pieces of bulk [287].

Naturally, this difference also appears in their optical properties that can vary significantly from those known for infinite systems or solitary atoms. One impressive example of exploiting these new features is the work of Yu et al. [289, 290]. Using nanostructured obstacles, they are able to focus the emitted light from a semiconductor laser more efficiently than before. To do this, they apply a nanostructured, optically thick gold layer on a dielectric substrate to the emitting facet of the semiconductor. The laser radiation couples to surface plasmons that – due to a special geometry – start emitting radiation with high directionality by themselves. The active region of a semiconductor laser has a size in the order of one wavelength, leading to strong diffraction of the laser beam when leaving the active zone. With their novel nano-antenna array, Yu et al. are able to reduce beam divergence from 63° to 2.4°. Occurring effectively lossless [289], such effects have to potential to revolutionize conventional optics.

During the last few years, a broad field of similar exciting applications for nanostructures in optics has evolved under the term *plasmonics*, see [291–294] for recent topical reviews. In this chapter we study nano clusters as representatives for finite-sized systems that exhibit interesting spectral features which eventually lead to the development of the field of plasmonics.

6.1 Introduction

Motivated by previous studies of RAITZA, REINHOLZ, et al. [132–140], we concentrate on nano clusters in this chapter. These are defined as particle agglomerates with $N=2\dots 10^7$ atoms, in the simplest case of a single species [288]. Extensive reviews on optical properties of such metal nano particles are available in [287, 288, 295]. In the view of this work, they can be seen as representatives for a multitude of differently shaped nano- to microscale structures that can be simulated with our highly efficient parallel tree code PEPC, see Chapter 5. We utilize it for solving the N-body problem and integrating the trajectories for the cluster electrons and ions as an open-boundary system with large particle number.

In their analysis, RAITZA et al. evaluate the total and spatially resolved electron momentum auto correlation spectrum, identify the MIE and bulk plasma resonance and find a cluster size dependent shift of the oscillation spectra. Their simulations are performed with clusters of up to $N_{\rm el}=N_{\rm ion}=1,000$ electrons and ions. In this chapter, we follow and extend their experiments to 300 times larger systems and propose a considerably simpler approach for identifying spatially resolved oscillation patterns that correspond to additional resonances found in the spectra. To this end, we describe the simulation setup and several fundamental cluster properties in Section 6.2. In the subsequent Section 6.3 we show and discuss the different spectral features observed in our simulations. Our spatially resolved diagnostic is demonstrated in Section 6.4, where also the obtained data is discussed in comparison to the findings of Section 6.3.

Results presented in this chapter have been partially published in [296].

6.2 Numerical simulation setup

In order to compare our simulation experiments with work previously done by RAITZA et al. [134], they are performed with an analogous configuration. We consider icosahedral, closed-shell neutral sodium clusters, compare Figure 6.1. These are composed of $N_{\rm el} = N_{\rm ion} = 55\dots 293,789$ electrons and ions that interacts via the PLUMMER/soft core potential (2.16) [104, 149]

$$\Phi^{(\text{Plummer})}(r) = \frac{1}{4\pi\varepsilon_0} \cdot \frac{q}{\sqrt{r^2 + \alpha^2}} \ . \tag{6.1}$$

In this way we avoid numerical heating, compare Section 2.2.3. Furthermore, we use the parameter α to model ionization processes. Therefore, we initially place the electrons in close proximity to their respective ion and adjust the potential cutoff parameter $\alpha = 5.336\,\mathrm{a_B}$. For vanishing electron-ion distance, this gives the ground state ionization

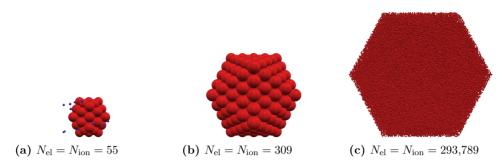


Figure 6.1: Icosahedral structure of the studied clusters with different numbers of electrons and ions (not to scale). For the smallest cluster, some electrons that left the surface are shown.

energy for sodium

$$V_{\text{e-i}}(r \to 0) = \lim_{r \to 0} \frac{1}{4\pi\varepsilon_0} \cdot \frac{q_{\text{el}}q_{\text{ion}}}{\sqrt{r^2 + \alpha^2}} = -5.1 \,\text{eV} \ .$$
 (6.2)

Thus, we mimic the dissociation $Na \rightleftharpoons Na^+ + e^-$. Higher ionization states are not considered.

The particle trajectories are integrated for physical times $t=0\dots 290$ fs using the highly scalable parallel Barnes-Hut tree code PEPC as a solver for the underlying N-body problem, see Chapters 3 and 5. Intrinsically supporting open boundaries, the tree algorithm is perfectly suitable for simulation of isolated, finite-sized particle configurations. In contrast to bulk systems that can be studied in computer experiments using relatively small setups that are extended to infinity by means of periodic boundaries, the cluster systems here require open boundaries and large particle numbers. PEPC is able to simulate configurations from small dimers up to billions of particles in a consistent, fully-kinetic approach, see Chapter 5. With this capability, closing the gap between nano- and bulk scale simulations is now possible.

Our calculations were performed on the high-scalability IBM Blue Gene/Q installation JUQUEEN [268] and the multi-purpose cluster JuRoPa [297] at Jülich Supercomputing Centre.

We start our simulations with icosahedrally arranged ions and marginally displaced electrons at solid density $n_{\rm ion}=n_{\rm el}=28\times10^{21}\,{\rm cm^{-3}}$. The corresponding bulk plasma frequency is $\omega_{\rm pl}=9.43\,{\rm fs^{-1}}$. Initially, the particles are frozen at $T_{\rm ion}=T_{\rm el}=0\,{\rm K}$ but are heated by a linearly polarized laser pulse with \sin^2 envelope. A moderate laser intensity of $I_{\rm L}=1\times10^{13}\,{\rm W/cm^2}$ and a wavelength of $\lambda=436\,{\rm nm}$ in the optical range are chosen. The laser pulse length is set to $t_{\rm pulse}=100\,{\rm fs}$ to allow for sufficient heating of the system.

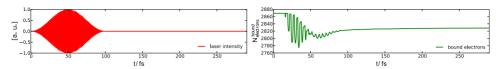


Figure 6.2: Representative evolution of the simulation experiment for a cluster with initially $N_{\rm ion}=N_{\rm el}=2,\!869$ ions and electrons. During $t_{\rm pulse}=100\,{\rm fs}$, it is heated by a laser pulse, that pulls electrons away and hence starts to charge the cluster. Due to the resulting positive net charge, the ion sphere expands. At $t=t_{\rm pulse}$, the ions are fixed. Diagnostics are performed during the following time $t=100\dots290\,{\rm fs}$ without any external field, with fixed ionic background and stable ionization degree.

Due to the laser heating, several electrons leave the cluster's surface. This effectively leads to a charged cluster as shown in Figure 6.2. With its increasing positive charge from the excess of ions, the remaining electrons are confined to the inner volume of the cluster. Only the bound electrons are considered in the present studies on collective electronic properties inside the nano clusters. In agreement with [136], the cluster size is defined via the ionic root mean square radius

$$R_{\text{cluster}}^{(\text{RMS})} = \sqrt{\frac{1}{N_{\text{ion}}} \frac{5}{3} \sum_{i=1}^{N_{\text{ion}}} r_{\text{ion}}^{(i)^{2}}}.$$
 (6.3)

Here, $r_{\text{ion}}^{(i)}$ is the distance of the *i*th ion from the cluster's center of ion mass. Now, electrons can be considered as free if

$$r_{\rm el} > R_{\rm cluster}^{\rm (RMS)}$$
 (6.4)

and

$$E_{\rm el}^{(\rm tot)} = E_{\rm el}^{(\rm kin)} + E_{\rm el}^{(\rm pot)} > 0$$
 (6.5)

Firstly, this means they have left the cluster since their distance $r_{\rm el}$ from its center is larger than the cluster radius. Secondly, they are not moving on a closed orbit as indicated by non-negative total energy. A detailed empirical study of this cluster charging process was performed by Broda [298, 299].

Another aspect of the cluster's positive net charge is its expansion during the heating stage. Since some compensating electrons are missing, the ion charges cannot be screened completely, which renders their mutual repulsion to drive them to flying apart. Naturally, the cluster expansion and electron spill-out modify the average electron density $n_{\rm el}$ and thus the plasma frequency $\omega_{\rm pl}$. This has to be taken into account in our simulations.

In our simulations, beam propagation is not considered. This restricts us to the long-wavelength limit of a homogeneous external field. Such a simplification might be questionable for the large clusters in this study with a diameter of up to $2 \cdot R_{\rm cluster}^{\rm (RMS)} = 27.94\,\mathrm{nm} \approx \lambda/_{15}$. However, we only study plasma properties after laser heating when electrons have reached temperatures of $T_{\rm el} = 1\dots 3\,\mathrm{eV}$. The actual way of reaching this heated state is not of primary interest here as the laser is switched off before any physical diagnostics are performed.

To be able to perform diagnostics of the electronic motion with fixed ionic background, we freeze the ionic drift at $t=t_{\rm pulse}$. This approach corresponds to the restricted molecular dynamics scheme of [134]. During the subsequent diagnostics stage $t=100\,{\rm fs}\dots290\,{\rm fs}$, the electrons interact with each other and with the static ionic background. With the exception of electron/electron and electron/ion collision, they can move freely and thus thermalize rapidly as shown in [134].

During the diagnostics stage, we will measure the electrons' total momentum autocorrelation function as well as the total potential and force experienced by test particles on a spherical grid inside the cluster. All simulations are repeated 128 times with different initial electron positions. With the laser heating that emphasizes perturbations to the initial conditions we are thus able to perform an ensemble averaging for improved statistics of our measured data. The following two Sections 6.3 and 6.4 present the particular diagnostics in more detail and discuss simulations results from either method.

6.3 Total momentum autocorrelation function

The total momentum autocorrelation function is closely related to macroscopic properties such as the dielectric function and the dynamic conductivity in the DRUDE model. Thus, it allows for a connection of microscopic simulation observables to the material properties that are experimentally accessible and were discussed in Section 1.2. In Section 6.3.1 we summarize this connection to prepare the fundamental basis for our results and analyses in Section 6.3.2. Finally, Sections 6.3.3 and 6.3.4 discuss the data obtained.

6.3.1 The dynamical conductivity and its connection to simulation observables

We study size effects on optical and transport properties in finite nanoplasmas. Here, the electrons are – due to the excess of ions – confined to a limited region inside the cluster's volume. Generalizing DRUDE theory for bulk matter, see Section 1.3, the

dielectric function in the long-wavelength limit

$$\varepsilon(\omega) = 1 - \frac{\omega_{\rm pl}^2}{\omega^2} \frac{1}{1 + i\frac{\nu(\omega)}{\omega}}$$
(6.6)

is considered. Via the collision frequency $\nu(\omega)$ it is connected to the internal conductivity [94]

$$\sigma(\omega) = \frac{\varepsilon_0 \,\omega_{\rm pl}^2}{\nu(\omega) - i\omega} \ . \tag{6.7}$$

As discussed in Section 1.4, a number of theoretical approaches for the collision frequency are available. However, most of them are limited to the weakly coupled regime. Molecular dynamics simulations do not suffer from such restrictions and thus can help building a solid data basis also for $\Gamma \gtrsim 1$.

In our simulations, we use a connection of the external conductivity to system observables via the fluctuation-dissipation theorem [71, 81]. In particular, the external conductivity can be introduced via

$$\sigma_{\rm ext}(\omega) = \frac{\Omega_0}{k_{\rm B}T} \cdot \left\langle \vec{J}_{\rm el}; \vec{J}_{\rm el} \right\rangle_{\omega} = \frac{1}{k_{\rm B}T} \frac{n_{\rm el}e^2}{m_{\rm el}^2} \cdot \left\langle \vec{P}_{\rm el}; \vec{P}_{\rm el} \right\rangle_{\omega} , \qquad (6.8)$$

with the normalization volume Ω_0 . Here, $\langle \vec{J}_{\rm el}; \vec{J}_{\rm el} \rangle_{\omega}$ and $\langle \vec{P}_{\rm el}; \vec{P}_{\rm el} \rangle_{\omega}$ are the Laplace transforms

$$\langle A; A \rangle_{\omega} = \int_0^{\infty} e^{i\omega t} (A(0); A(t)) dt$$
 (6.9)

of the autocorrelation functions

$$(A(0); A(t)) = \int_{-\infty}^{\infty} \overline{A(u)} \cdot A(u+t) du$$
 (6.10)

for the total bound electron current \vec{J}_{el} and momentum

$$\vec{P}_{\rm el} = \sum_{k=1}^{N_{\rm el}^{\rm (bound)}} \vec{p}_{\rm el}^{(k)}$$
 (6.11)

where $\vec{p}_{\rm el}^{(k)}$ is the momentum of the kth bound electron. In the following, we use the abbreviation

$$K(\omega) := \left\langle \vec{P}_{\text{el}}; \vec{P}_{\text{el}} \right\rangle_{\omega} \tag{6.12}$$

for the momentum autocorrelation spectrum. For straightforward evaluation of $K(\omega)$, we use the autocorrelation theorem from FOURIER theory [300]. For the function h(t) and its autocorrelation function (h(0); h(t)) it reads

$$\mathcal{F}\left\{|H(\omega)|^2\right\} = (h(0); h(t)) \tag{6.13}$$

and accordingly

$$\mathcal{F}^{-1}\left\{|h(t)|^2\right\} = (H(0); H(\omega)) \tag{6.14}$$

if h(t) and $H(\omega)$ are related via the FOURIER transform

$$h(t) = \mathcal{F} \{ H(\omega) \} := \frac{1}{2\pi} \int_{-\infty}^{\infty} H(\omega) e^{i\omega t} d\omega , \qquad (6.15)$$

$$H(\omega) = \mathcal{F}^{-1}\left\{ (h(t)\right\} := \int_{-\infty}^{\infty} h(t)e^{i\omega t} dt .$$
 (6.16)

We exploit the fact that we are interested in the LAPLACE transform of the momentum autocorrelation function, which is identical to the FOURIER transform since the integrand function vanishes for negative arguments. Thus, omitting all normalizing prefactors and using arbitrary units from here, for K(w) the theorem yields

$$K(\omega) = \left| \mathcal{F} \left\{ \vec{P}_{\text{el}} \right\} \right|^2 . \tag{6.17}$$

Since $\vec{P}_{\rm el} \in \mathbb{R}^3$, its FOURIER transform has to be performed component-wise and is in general a complex vector $\mathcal{F}\left\{\vec{P}_{\rm el}\right\} \in \mathbb{C}^3$.

As evident from (6.8) with (6.12),

$$K(\omega) \propto \sigma_{\rm ext}(\omega)$$
 (6.18)

For a bulk plasma in the long-wavelength limit, it can be shown that the external conductivity has a LORENTZian shape [94]

$$\sigma_{\text{ext}}(\omega) = \frac{\varepsilon_0 \omega_{\text{pl}}^2 \,\omega}{\nu(\omega) \,\omega - i(\omega^2 - \omega_{\text{pl}}^2)} \tag{6.19}$$

that will also be acquired by $K(\omega)$ and

$$\operatorname{Re}\{K(\omega)\} \propto \frac{\varepsilon_0 \frac{\omega_{\text{pl}}^2}{\omega \cdot \nu(\omega)}}{1 + \left(\frac{\omega^2 - \omega_{\text{pl}}^2}{\omega \cdot \nu(\omega)}\right)^2} =: \frac{K}{1 + \left(\frac{\omega^2 - \omega_{\text{pl}}^2}{\gamma}\right)^2} . \tag{6.20}$$

Accordingly, in a bulk plasma $\text{Re}\{K(\omega)\}\$ has a distinct resonance at $\omega = \omega_{\text{pl}}$ with a width γ that is related to the bulk collision frequency $\nu(\omega)$.

In the following section that shows our simulation results on finite plasma systems, we will see that several additional resonances appear in the spectrum $K(\omega)$. We will use an expansion motivated by (6.20) to fit and evaluate their positions. Furthermore, we deduce an explanation for observed resonance shifts.

6.3.2 Simulation results and fitting procedure for the momentum autocorrelation function

As discussed in the previous section, in our simulations the total momentum autocorrelation spectrum $K(\omega)$ is calculated from the total electron momentum $\vec{P}_{\rm el}(t)$ via the autocorrelation theorem (6.17). In Figures 6.3 and 6.4, we show Re $\{K(\omega)\}$ for differently sized nano clusters from $N_{\rm el}=N_{\rm ion}=55\dots293,789$ electrons and ions. The data was obtained by performing a Jackknife averaging procedure [301, 302] over the spectra of all 128 different simulations runs per cluster size. As mentioned before, due to the initial expansion and electron spill-out, the effective electron density $n_{\rm el}$ and hence the plasma frequency $\omega_{\rm pl}$ varies for the different cluster sizes. The position of the bulk plasmon resonance as deduced from the measured densities is highlighted in the figures as the right dashed line.

Being the highest excited frequency in all systems, the rightmost resonance can be identified as plasma oscillation. With the exception of the two largest systems that clearly behave bulk-like in this respect, its position is shifted to lower frequencies with respect to the bulk plasma frequency $\omega_{\rm pl}$. Furthermore, several additional resonances appear. Their number grows with increasing system size. The lowest-frequency resonance in all systems is predicted from standard MIE theory [288, 295] at

$$\omega_{\text{Mie}} = \frac{\omega_{\text{pl}}}{\sqrt{3}} \tag{6.21}$$

for spherical particles, marked as the left dashed line in the graphs. It is often referred to as surface plasmon since it results from interaction of the electron system with the cluster surface. It can be derived by considering two hard-sphere distributions of electrons and ions that oscillate with respect to each other with a small amplitude. Details as well as theoretical studies also for non-spherical, e.g. rod-shaped, clusters can be found in [288, 295]. The theoretical position of the surface plasmon (6.21) is included in Figures 6.3 and 6.4. Again, the lowest-frequency resonances that correspond to the surface plasmon are consistently found at smaller frequencies than expected in our simulation data for the smaller systems. This redshift will be further analyzed in Section 6.3.3.

Starting from intermediate cluster size $N_{\rm el}, N_{\rm ion} \geq 309$, between these well-known resonances, additional peaks are evident in the total momentum autocorrelation spectrum $K(\omega)$. For sizes from $N_{\rm el}, N_{\rm ion} \geq 28,741$ an additional low-frequency resonance appears. To account for these additional resonances, we fit $K(\omega)$ with a linear combination of modified LORENTZian resonance curves

$$\operatorname{Re}\{K(\omega)\} \approx \sum_{p=1}^{N_{\text{peaks}}} K_p \cdot \left[1 + \left(\frac{\omega - \omega_p}{\gamma_p}\right)^2\right]^{-\frac{3}{2}}.$$
 (6.22)

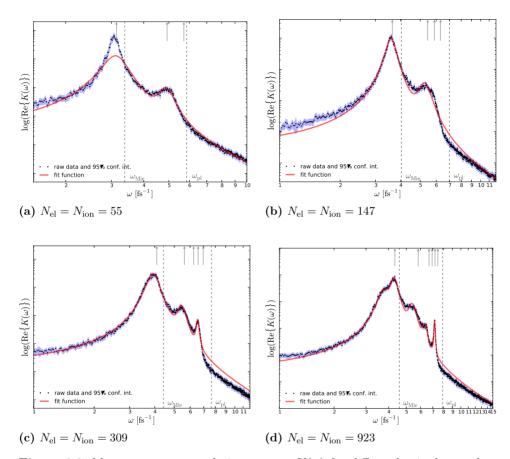


Figure 6.3: Momentum autocorrelation spectra $K(\omega)$ for differently sized nanoclusters. The dots represent the average over 128 distinct runs per cluster size. Shaded areas correspond to the respective 95% confidence interval. Red lines are fitted graphs for the spectra. Arrows correspond to resonances found in the spatially resolved study, see Section 6.4. The position of the bulk plasma frequency $\omega_{\rm pl}$ and the surface plasmon $\omega_{\rm Mie}$ are shown as dashed lines. Graphs for larger clusters are given in Figure 6.4.

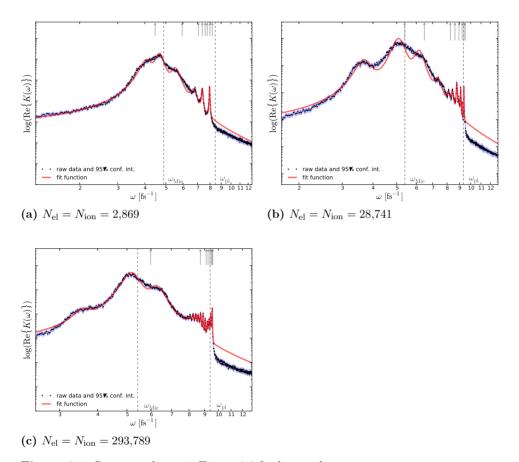


Figure 6.4: Same graphs as in Figure 6.3 for larger clusters.

$\begin{array}{c ccc} N_{\rm ion} = 55 & K_{\rm pl} = 8.72 \\ \hline K_p/K_{\rm pl} & 14.77 & 1.00 \\ \omega_p[\rm fis] & 3.11 & 4.87 \\ \gamma_p[\rm fis] & 0.760 & 0.757 \\ \hline \end{array}$	$\begin{aligned} N_{\rm ion} &= 147 K_{\rm p} \\ K_p/K_{\rm pl} &= 28.35 \\ \omega_p [{\rm fs}] &= 3.58 \\ \gamma_p [{\rm fs}] &= 0.470 \end{aligned}$	$ \begin{array}{c} 1 = 3.44 \\ 1.00 \\ 5.25 \\ 0.867 \end{array} $	$N_{\text{ion}} = 3$ K_p/K_{pl} $\omega_p[\text{fs}]$ $\gamma_p[\text{fs}]$	3.92	3.12 5.43	1.00 6.53 0.255	
$\begin{array}{c cccc} N_{\rm ion} = 923 & K_{\rm pl} = 0.19 \\ \hline K_p/K_{\rm pl} & 13.46 & 39.29 & 4.09 \\ \omega_p[{\rm fs}] & 3.75 & 4.28 & 5.38 \\ \gamma_p[{\rm fs}] & 0.577 & 0.475 & 0.615 \\ \hline \end{array}$	0.61 1.00 6.32 7.16 0.433 0.125						
$\begin{array}{c cccc} N_{\rm ion} = 2,\!869 & K_{\rm pl} = 0.05 \\ \hline K_p/K_{\rm pl} & 18.87 & 37.71 & 7.42 \\ \omega_p [\rm fs] & 4.07 & 4.67 & 5.53 \\ \end{array}$	0.60 0.74 6.70 7.36	1.00 7.96					
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$							
$ \begin{array}{c cccc} \omega_p[\dot{\text{fs}}] & 3.68 & 5.08 & 6.15 \\ \gamma_p[\dot{\text{fs}}] & 0.655 & 0.682 & 0.646 \\ \end{array} $	7.34 8.12 0.594 0.278	8.44 8.77 0.201 0.106	8.94 0.116		9.22 9.	.38 .031	
$\begin{array}{c cccc} N_{\rm ion} = 293,789 & K_{\rm pl} = 0.00022 \\ \hline K_p/K_{\rm pl} & 0.49 & 23.34 & 5.96 \\ \omega_p[\rm fs] & 3.58 & 5.17 & 6.26 \\ \gamma_p[\rm fs] & 0.568 & 0.769 & 0.849 \\ \hline \end{array}$	0.13 0.22 7.53 8.02 0.559 0.358	0.29 0.34 8.25 8.42 0.126 0.126	0.32 8.58 0.123	8.74	8.90 9.	.25 .04 .061	
$K_p/K_{ m pl} \ \omega_p[{ m fs}] \ \gamma_p[{ m fs}]$		1 220	0.17 9.19 0.081	0.25 9.32	0.42 1. 9.44 9.	.00 .54 .031	

Table 6.1: Fit parameters for Figures 6.3 and 6.4. The notation follows (6.22), each parameter triple corresponds to an identified resonance. Instead of K_p , the ratio $K_p/K_{\rm pl}$ is given, where $K_{\rm pl}$ is the amplitude of the highest-frequency resonance.

This ansatz is motivated by (6.20) and a similar approach in [134]. However, we modified the exponent in (6.22) to better reproduce the asymptotic behavior

$$\operatorname{Re}\{K(\omega \to \infty)\} \propto \omega^{-3}$$
 (6.23)

that was observed in our simulations. The necessity for this modification hints at a modified high-frequency behavior of the collision frequency due to the choice of a soft-core potential instead of the Coulomb interaction [83, 91]. The respective fit curves are also included in Figures 6.3 and 6.4 and satisfactorily approximate the measured data for all system sizes. The parameters found by the least-squares fit are given in Table 6.1. Due to applying the modified Lorentzian, the width γ has lost its direct connection to the collision frequency. However, since we are more interested in the resonance positions here, this is acceptable.

Since the additional resonances hint to further excitation modes besides the bulk and surface plasmon, a spatially resolved study is performed in Section 6.4. Hitherto, Sec-

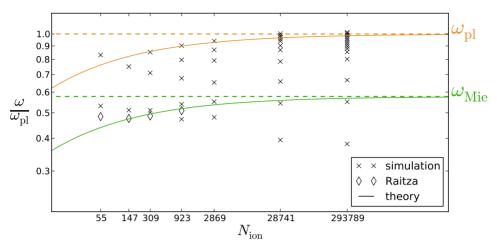


Figure 6.5: Position of the observed resonances normalized to plasma and MIE frequency in dependence of cluster size (crosses ×). The dashed lines denote the position of the classical plasma and MIE frequency. Solid lines correspond to the theoretically expected position (6.36). Diamonds ⋄ are values by RAITZA [136] for small clusters and an identical setup as in our simulations.

tion 6.3.3 examines the observed shift of the resonance spectra towards lower frequencies for smaller cluster sizes quantitatively.

6.3.3 Analysis and theory for the resonance shift

As seen in Figures 6.3 and 6.4, the resonance spectra for the small clusters are consistently redshifted. Figure 6.5 shows the identified resonance positions normalized to the respective plasma frequency. Here it is clearly visible that the surface and bulk plasmon are shifted by the same factor which depends on cluster size.

To find a theoretical estimation for the resonance frequencies, we follow [136] and consider an electron distribution that oscillates in front of a fixed ionic background. Without loss of generality, linear movement along the z-direction is assumed. Then, the one-dimensional equation of motion for the displacement z of the electron cloud with respect to its position of rest is

$$m_{\rm el} N_{\rm el} \frac{\mathrm{d}^2}{\mathrm{d}t^2} z(t) = F_z(z(t)) ,$$
 (6.24)

where F_z is the third component of the restoring force (the x- and y-components vanish and the respective positions are constant) and $m_{\rm el}N_{\rm el}=M_{\rm el}$ is the total mass of the electron cloud that consists of $N_{\rm el}$ particles with mass $m_{\rm el}$. Assuming harmonic oscillations $z(t)=z\cdot {\rm e}^{{\rm i}\omega t}$ and writing the force as gradient of the potential, we find

$$m_{\rm el}N_{\rm el}\omega^2 z = \frac{\partial U(z)}{\partial z}$$
, (6.25)

which after differentiating with respect to z yields

$$m_{\rm el} N_{\rm el} \omega^2 = \frac{\partial^2 U(z)}{\partial z^2} \ . \tag{6.26}$$

The potential can be written as

$$U(z) = \int d^3 \vec{r} \, n_{\rm el}(r) V_{\rm ext}(|\vec{r} - z\vec{e}_z|) \tag{6.27}$$

with the electron density profile $n_{\rm el}(r)$ and the external potential of the ion cloud $V_{\rm ext}(r)$. Both electron and ion distributions are assumed to be isotropic. Hence, the dependency on position could be replaced by the distance argument r. Insertion into (6.26) yields

$$\omega^2 = \frac{1}{m_{\rm el} N_{\rm el}} \frac{\partial^2}{\partial z^2} \int d^3 \vec{r} \, n_{\rm el}(r) V_{\rm ext}(|\vec{r} - z\vec{e}_z|) . \qquad (6.28)$$

The differentiation in (6.28) is only performed with respect to the solitary variable z and hence can be drawn into the integration. Under the assumption of small amplitudes $z \ll r$, it can be shown that

$$\frac{\partial^2}{\partial z^2} V_{\text{ext}}(|\vec{r} - z\vec{e}_z|) \approx \frac{\mathrm{d}^2}{\mathrm{d}r^2} V_{\text{ext}}(r) \cdot t^2 + \frac{1}{r} \frac{\mathrm{d}}{\mathrm{d}r} V_{\text{ext}}(r) \cdot (1 - t^2)$$
 (6.29)

with $t = \cos \vartheta$ and $\vartheta = \sphericalangle(\vec{r}, \vec{e}_z)$. Insertion of (6.29) into (6.28) and integration over both angular parameters in spherical coordinates yields

$$\omega^2 = \frac{4\pi}{3m_{\rm el}N_{\rm el}} \int_0^\infty \int \mathrm{d}r \, n_{\rm el}(r) \left[r^2 \frac{\mathrm{d}^2}{\mathrm{d}r^2} V_{\rm ext}(r) + 2r \frac{\mathrm{d}}{\mathrm{d}r} V_{\rm ext}(r) \right] . \tag{6.30}$$

The external potential of the ion cloud can be written as

$$V_{\text{ext}}(\vec{r}) = \int d^3 \vec{R} \, n_{\text{ion}}(|\vec{R}|) \Phi(|\vec{R} - \vec{r}|)$$
 (6.31)

with the interaction potential $\Phi(r)$ between individual electrons and ions, which is the PLUMMER potential (6.1) in our case.

We now assume hard-sphere ion and electron distributions

$$n_{\rm ion}(r) = \begin{cases} n_{\rm ion} & \text{for } r \le R_{\rm ion} \\ 0 & \text{else} \end{cases}, \qquad n_{\rm el}(r) = \begin{cases} n_{\rm el} & \text{for } r \le R_{\rm el} \\ 0 & \text{else} \end{cases}$$
 (6.32)

with radii R_{ion} and R_{el} , respectively. Insertion of (6.31) with the potential (6.1) and the densities (6.32) into (6.30) yields after tedious calculations

$$\omega^2 = \frac{\omega_{\text{MiE}}^2}{2R_{\text{el}}^3} \left\{ (R_+ - R_-)(R_{\text{el}}^2 + R_{\text{ion}}^2 + \alpha^2) - (R_+ + R_-)R_{\text{el}}R_{\text{ion}} \right\}$$
(6.33)

with

$$R_{+} := \sqrt{(R_{\rm el} + R_{\rm ion})^2 + \alpha^2} ,$$
 (6.34)

$$R_{-} := \sqrt{(R_{\rm el} - R_{\rm ion})^2 + \alpha^2} \ .$$
 (6.35)

This is the surface plasmon resonance frequency a hard-sphere electron distribution that harmonically oscillates in front of an immobile hard-sphere ion cloud with the Plummer interaction potential.

For $R_{\text{ion}} = R_{\text{el}} = R$, i. e. equally sized spheres, (6.33) reduces to

$$\omega^2 = \frac{\omega_{\text{MIE}}^2}{2R^3} \left\{ \sqrt{4R^2 + \alpha^2} (R^2 + \alpha^2) - \alpha (3R^2 + \alpha^2) \right\} , \qquad (6.36)$$

which for $\alpha = 0$, i.e. Coulomb instead of Plummer potential, yields

$$\omega = \omega_{\text{Mir}} \,, \tag{6.37}$$

as expected for consistency with the well-known theory. The factor $1/\sqrt{3}$ between surface and bulk plasmon frequency is a geometry-related constant [288, 295]. Since we also simulate almost-spherical nano particles, despite the resonance shift, this factor is expected to stay valid. Therefore, we expect the modification of the surface plasmon that is given by (6.36) also to apply for the bulk plasmon.

In Figure 6.5, the solid lines were computed using (6.36) for ω_{Mie} and ω_{pl} . Both are consistent with the measured data and the resonances found by RAITZA [136]. Only for the smallest nano cluster with $N_{\text{ion}} = N_{\text{el}} = 55$, an additional blue-shift appears. Obviously, the approximation of small oscillation amplitudes in comparison to cluster size is not valid any more.

6.3.4 Conclusions

In this section, we measured the total electron autocorrelation spectra for differently sized nano clusters. Besides several additional features, they show the expected bulk plasmon and MIE resonances. These are redshifted due to the Plummer/soft core potential that is used in our simulations to avoid numerical heating and mimic ionization

effects. An analogous shift was also observed by RAITZA et al. who used and error function potential

$$\Phi^{(\text{erf})}(r) = \frac{1}{4\pi\varepsilon_0} \cdot \frac{q}{r} \operatorname{erf}\left(\frac{r}{\lambda}\right) \tag{6.38}$$

where the parameter λ was also adjusted to yield the ground state ionization energy of sodium for $r \to 0$. We can explain this shift of the resonance spectrum consistently using a simple model of oscillating spherical charge clouds as long as the oscillation amplitude is small compared to the cluster size.

Besides either plasmon resonances, the simulation data and fits to the spectra in Figures 6.3 and 6.4 show a multitude of additional excitations. Especially, the larger clusters exhibit an extensively structured spectrum of resonances between $\omega_{\rm MIE}$ and $\omega_{\rm pl}$ and even below $\omega_{\rm MIE}$. Since they do not correspond to the classical surface and bulk plasmon, in the following section we perform a spatially resolved study of collective phenomena inside the cluster's electronic subsystem.

6.4 Spatially resolved studies of the electronic resonances

As seen in the previous section, a large number of additional resonances appear in the total momentum autocorrelation function $K(\omega)$ for increasing cluster size. These are not explained by the standard, long-wavelength MIE theory. The assumption, that the resonances between ω_{MIE} and ω_{pl} result from additional excitation modes with finite wave number seems natural. To actually localize the origin of the different spectral features, we performed a spatially resolved study on collective phenomena inside the cluster. The fundamental idea of our diagnostics is described in Section 6.4.1 before in Section 6.4.2 the data obtained is discussed.

6.4.1 Spatially resolved field spectra

In [138–140], Raitza et al. perform a spatially resolved analysis for a spherical nanocluster with $N_{\rm el}=N_{\rm ion}=1{,}000$ electrons and ions. Therefore, they study bilocal momentum correlation spectra

$$K_{a;a'}(\omega) = \int_0^\infty dt \, K_{a;a'}(t) e^{i\omega t}$$
(6.39)

with the correlation matrix

$$K_{a;a'}(t) = \frac{1}{V_{a}, V_{a'}} \left(\vec{J}_{el}(a,0); \vec{J}_{el}(a',t) \right) .$$
 (6.40)

In this notation, the indices a,a' correspond to cells of a spherical grid that covers the whole cluster volume. The volume-normalized total electron current $J_{el}(a,t)/v_a$ in every cell and at every instant in time is used as a simulation observable. This is done in full analogy to the formalism in Section 6.3. The eigenvectors of the matrix $K_{a;a'}(\omega)$ can be considered as excitation patterns. They are further analyzed in [138–140] by means of a spherical FOURIER transform to separate the different oscillation modes.

In our simulations, this diagnostic showed significant sensitivity to noise and limited spatial resolution. In particular, for the smaller clusters grid cells are often unoccupied. Then, $\vec{J}_{\rm el}(a,t)$ cannot be determined reliably for them. This is why we use a much simpler and more direct diagnostic here.

Clearly, the self-consistent electric field experienced by each individual particle is the driver for any collective or non-collective motion inside the cluster. Hence it is natural to exploit it as the simulation observable. In particular, we use the electric potential instead of the field since it is – as an integrated variable – expected to be even less susceptible to noise. To this end, in addition to evaluating the forces from electrons and ions onto all electrons, we also compute the electric potential $\Phi(\vec{r}_a)$ on the nodes \vec{r}_a of a spherical grid that is embedded into the cluster. This can conveniently be done using our tree code implementation by performing an additional tree traversal for the grid node positions through the tree that is constructed from the electrons and ions anyway, see Chapters 3 and 5.

We are interested in time correlations, i.e. frequency resolved observables. These are obtained by performing a FOURIER transform

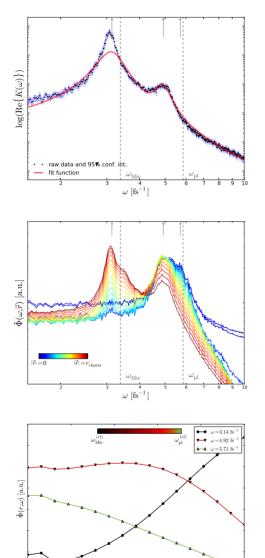
$$\tilde{\Phi}(\omega, \vec{r}_a) = \text{Re}\{\mathcal{F}\{\Phi(t, \vec{r}_a)\}\}$$
(6.41)

for every grid node separately to keep the locality information. This spatially resolved potential spectrum is averaged over all 128 individual simulation runs per cluster size again, yielding the graphs shown in the central plot of the left column in Figures 6.6 to 6.12 that will be discussed in the next section. For the purpose of preparing these graphs, an angular average of $\tilde{\Phi}(\omega, \vec{r_a} = (r_a, \theta_a, \phi_a))$ for fixed cell distances r_a from the cluster's center in the form

$$\tilde{\Phi}(\omega, r_a) = \frac{1}{N_{\theta} N_{\phi}} \sum_{i_a=1}^{N_{\theta}} \sum_{i_a=1}^{N_{\phi}} \tilde{\Phi}(\omega, r_a, \theta(i_{\theta}), \phi(i_{\phi}))$$

$$(6.42)$$

has been performed. Note that the normalizing factor $\sin \theta$ is missing as it did not yield a significant difference in our case. The individual, differently colored lines in the plots correspond to different, fixed values for r_a . For blue lines, $r_a=0$, while for dark red lines $r_a=r_{\rm cluster}$. thus, peaks of blue lines correspond to resonant movement in the inner cluster region while resonances at its surface show as features in the dark red lines. Dashed lines mark the position of the bulk and surface plasmon again. For better



 $r/r_{
m cluster}$

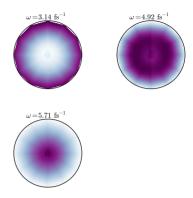
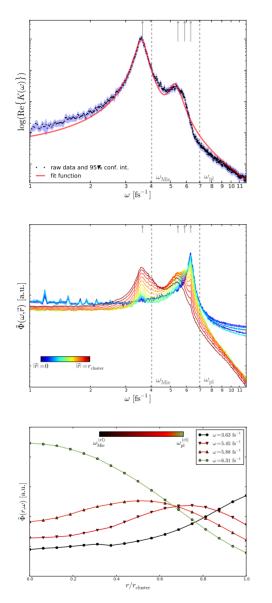


Figure 6.6:

Results of the simulations for a cluster of initially $N_{\rm el} = N_{\rm ion} = 55$ electrons and ions. the graphs to the left show from top to bottom: the total momentum autocorrelation function $K(\omega)$ from Figure 6.3 for comparison; the potential spectrum (6.42) for fixed radii, identified resonance peaks are marked with arrows; the radial dependence of the potential spectrum for the identified resonance fixed frequencies. On the right, the spatially resolved, ϕ -averaged excitation patterns for the resonance frequencies are shown. The color coding and further details are given in the text. The dodecagonal shape of the excitation pattern graphs is an artifact of limited grid resolution $N_{\theta} = 6$.



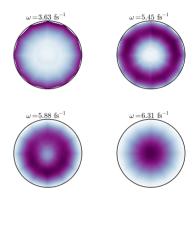
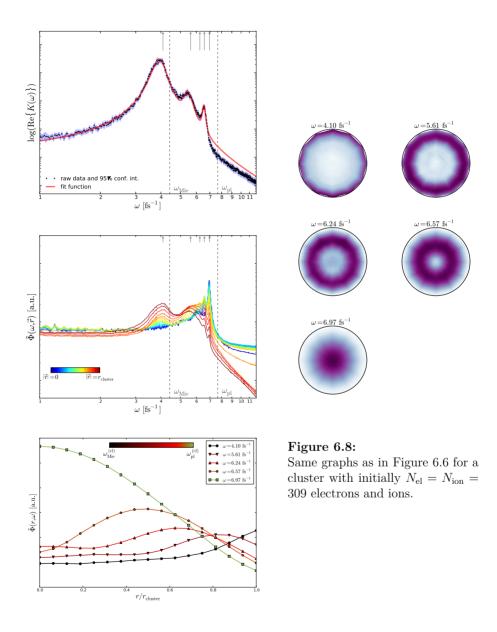
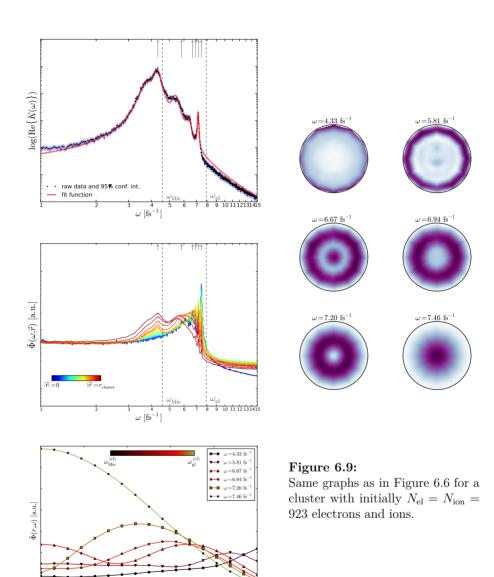
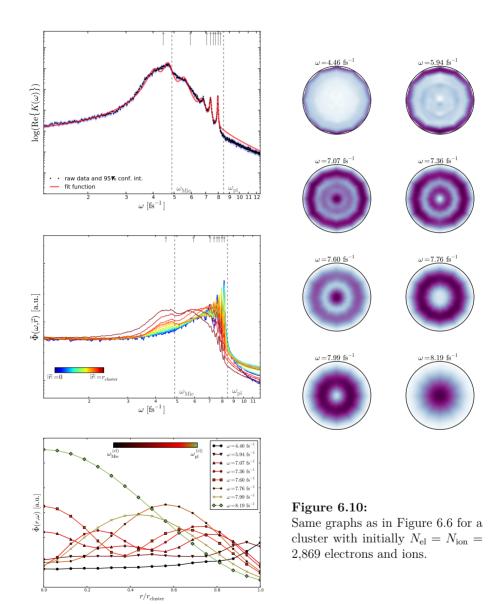


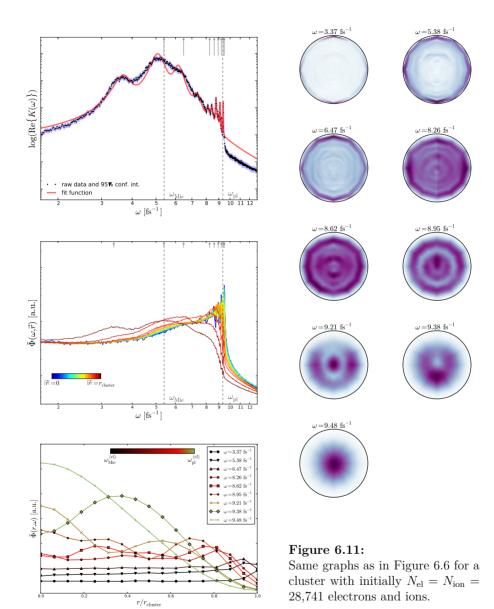
Figure 6.7: Same graphs as in Figure 6.6 for a cluster with initially $N_{\rm el}=N_{\rm ion}=147$ electrons and ions.



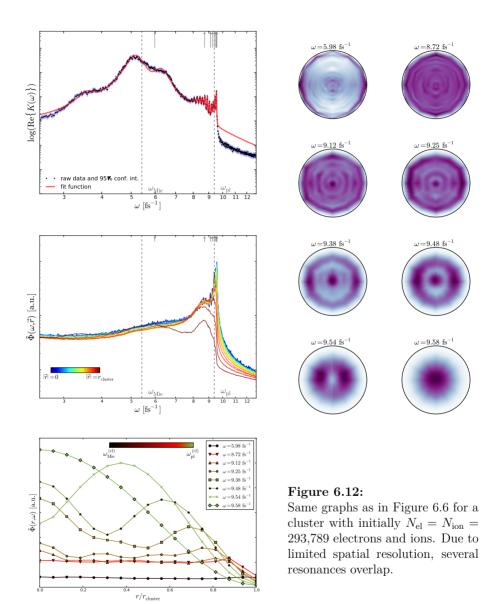
 $r/r_{
m cluster}$







126



comparability, the momentum autocorrelation spectra from Figures 6.3 and 6.4 are also included as the top left graphs in Figures 6.6 to 6.12.

Since with both diagnostics different observables with different normalization are evaluated, a direct comparison of peak heights is not possible. However, their position already yields enough information for an extensive discussion and comparison in the next section. The resonance positions that are identified in the spectra from the spatially resolved study are indicated by gray arrows at the upper edge of both graphs in the figures.

6.4.2 Analysis of the spatially resolved spectra

In the spatially resolved study, numerous resonances are evident that can directly be related to those found in Section 6.3. In particular, the MIE and bulk plasmon can be identified immediately. Many peaks between both resonances can also be found in the graphs for both studies. However, the spectra of $\tilde{\Phi}(\omega, r_a)$ reveal a large amount of additional peaks that are not visible in the spectra for $K(\omega)$. Clearly, these have to correspond to oscillation modes with vanishing total momentum and hence $K(\omega) = \langle \vec{P}_{\rm el}, \vec{P}_{\rm el} \rangle_{\omega} = 0$.

The line colors in the graphs of $\tilde{\Phi}(\omega, r_a)$ already indicate a strong connection between resonance frequency and distance from the cluster center. This is becoming even more evident when considering the lower left plots in Figures 6.6 to 6.12. There, the radial dependence of (6.42) is shown for fixed values of ω that correspond to the identified resonance positions (gray arrows in the other graphs). In these plots, black lines correspond to $\omega \approx \omega_{\rm MIE}$, olive lines to $\omega \approx \omega_{\rm pl}$. As expected, for all cluster sizes MIE contributions result from the cluster's surface and oscillation with $\omega_{\rm pl}$ is observed in its central, bulk-like region. Intermediate frequencies are mostly related to collective oscillation in simple shell regions, i.e. for fixed distance from the center as particularly prominent for $N_{\rm el} = N_{\rm ion} = 55 \dots 309$.

Starting from $N_{\rm el}, N_{\rm ion} \geq 923$, some oscillation modes exhibit a multi-peak structure. For example, there are two peaks visible for $\omega = 6.67\,{\rm fs^{-1}}$ in the lower left graph of Figure 6.9. This is evidence for oscillations with finite wave number. Going towards larger systems, more and more of such excitations appear. This results from more discrete wave numbers to fit into the nano plasma region that is confined to the cluster.

The angular resolved excitation patters

$$\tilde{\Phi}(\omega, r_a, \theta_a) = \frac{1}{N_{\phi}} \sum_{i_{\phi}=1}^{N_{\phi}} \tilde{\Phi}(\omega, r_a, \theta_a, \phi(i_{\phi}))$$
(6.43)

are shown in the right columns of Figures 6.6 to 6.12 for the same frequencies as in the lower-left graphs and indicated by gray arrows. There, collective behavior corresponds to violet color, off-resonant oscillation is shown in light blue to white.

The excitation patterns support the hypothesis of higher, non-trivial excitation modes. However, they cannot be explained a higher modes in a classical fluid model with appropriate boundary conditions as in that case the dispersion relation $\omega^2 = \omega_{\rm pl}^2 + s^2 k^2$ would – in analogy to the BOHM-GROSS case (1.15) – suggest resonances with $\omega > \omega_{\rm pl}$ [303]. The observed resonances feature complex and (mostly) isotropic oscillation that can be seen as radial plasma waves with different wave numbers. In addition, few non-isotropic modes appear, e.g. at $\omega = 9.21\,{\rm fs^{-1}}$ and $\omega = 9.38\,{\rm fs^{-1}}$ for $N_{\rm el} = N_{\rm ion} = 28,741$ and $\omega = 9.54\,{\rm fs^{-1}}$ for $N_{\rm el} = N_{\rm ion} = 293,789$. They have a droplet or dumbbell shape and break the intrinsic symmetry of the system. Their origin is still unclear. However, they can simply originate from the non-spherical cluster shape. Especially for the very large systems, the radial symmetry is strongly broken by the flat faces and sharp edges of the icosahedral particle setup. Here, comparisons with spherical configurations and hence preserved radial symmetry are necessary.

Finally, another interesting feature can be understood using the spatially resolved data. For the larger cluster, in the total momentum autocorrelation spectrum a solitary peak was observed far below the MIE frequency, see Figure 6.5. There, its origin could not be explained. This resonance also appears in the study of $\tilde{\Phi}(r_a,\omega)$, for example for $N_{\rm el}=N_{\rm ion}=28,741$ in Figure 6.11 at $\omega=3.37\,{\rm fs^{-1}}$. The respective excitation pattern exhibits only weak resonance regions at the very outer edge of the cluster where the diagnostics region ends. Apparently, an electron that orbits outside the system influences a potential oscillation here. Since such electrons are considered as bound to the cluster, they are included in the total momentum that finally yields $K(\omega)$ and appear as resonances there.

6.5 Summary

In this chapter, we analyzed collective oscillations in finite nano cluster systems with sizes of $N_{\rm el}=N_{\rm ion}=55\dots293,789$ electrons and ions by means of molecular dynamics simulations using our tree code PEPC. We evaluated the total momentum autocorrelation spectrum for the electronic subsystem and applied a new technique to identify the spatial origin of resonance phenomena in nanoclusters. The surface and bulk plasma resonances were identified in using either method and their redshift for lower cluster sizes has been explained consistently. Furthermore, the spatially resolved potential spectra were utilized to accurately pinpoint the origin of the observed additional resonances for the first time. They essentially reproduced the resonance positions from the autocorrelation spectra but also showed some additional peaks which are interpreted as

vibrational modes with vanishing total momentum that by construction cannot be seen using the former method. As verified via inspection of the spatially resolved oscillation patterns for the resonance frequencies, the additional peaks correspond to excitation of distinct shells inside the cluster and finite wave number oscillation in the confined nanoplasma. Furthermore, oscillation patterns that break the spherical symmetry were found and attributed to the non-spherical shape of the particle configurations in this study.

The results from this section should provide guidance and additional impetus to further analyze the absorption spectra of metallic nano clusters by means of theory and experiment. In addition, a comparison with the latest spatially resolved studies for small systems by Raitza et al. [139, 140] and comparative computations using the recently developed MicPIC method [304, 305] are desirable. Furthermore, a detailed analysis of the excitation patterns could help constructing a spatially dependent dispersion relation inside the cluster that is be relevant for understanding optical and transport properties of isolated clusters as well as finite nanoplasmas in colloidal suspensions and on surfaces. Nanostructured targets for laser-matter interaction experiments are already available and nanoelectronics is rapidly developing towards the scales studied here. There, similar novel exciting applications as the application of plasmonics in optics in the introductory example by Yu et al. [289] can be expected.

7 Collision frequency in bulk plasmas

As summarized in Chapter 1 there exist numerous models on Inverse Bremsstrahlung heating in bulk matter, but experimental data that can be directly compared to theoretical results is still scarce. Molecular dynamics simulations can close this gap to some extent by providing data from computer experiments which can be tuned to assumptions and parameter ranges covered by the theories. The effective collision frequency is a key parameter when studying optical and transport properties such as absorption coefficient and conductivity, as shown in Section 1.3. It will be studied in this chapter using classical molecular dynamics simulations for high-intensity fields in strongly coupled bulk matter. To do this, in Section 7.1, we formulate the connection from the plasma heating rate as a simulation observable to the collision frequency. The details on simulation setup and workflow are given in Section 7.2. Since we calculate the heating rate for fixed sets of plasma conditions $(n_{\rm el}, T_{\rm el})$, it is difficult but necessary to keep temperature constant while measuring the heating rate in these simulations. Appropriate methods to achieve this are discussed in Section 7.3 before we present our simulation results in Section 7.4.

7.1 The plasma heating rate

In order to retrieve the dynamical collision frequency for the electron-ion system in the simulation, a connection to the plasma heating rate is required. For completeness, the cornerstones of the derivation are given here, which borrows freely from the argumentation of [51].

As shown in Section 1.3, the ac-conductivity is given in the DRUDE model via

$$\sigma(\omega) = \frac{\varepsilon_0 \omega_{\rm pl}^2}{-i\omega + \nu_{\rm ei}} \,, \tag{7.1}$$

so that the dielectric function reads

$$\varepsilon(\omega) = 1 - \frac{\sigma(\omega)}{i\omega\varepsilon_0} = 1 - \frac{\omega_{\rm pl}^2}{\omega(\omega + i\nu_{\rm ei})} , \qquad (7.2)$$

where the electron-ion collision frequency ν_{ei} has originally been defined via the electron-ion relaxation time $\tau_{ei} = 1/\nu_{ei}$ as in Sections 1.3 and 1.4, see [43–45, 47, 306].

We start from the dispersion relation (1.32)

$$k^2 c^2 = \omega^2 \varepsilon(\vec{k}, \omega) . \tag{7.3}$$

Neglecting non-local effects in the following, the long-wavelength limit of the dielectric function $\varepsilon(\omega) = \lim_{|\vec{k}| \to 0} \varepsilon(\vec{k}, \omega)$ will be considered. Then, with (7.2)

$$\omega^2 = k^2 c^2 + \omega_{\rm pl}^2 \left(1 - i \frac{\nu_{\rm ei}}{\omega} \right) \tag{7.4}$$

is obtained in the high-frequency limit $\omega \gg \nu_{\rm ei}$. Rewriting (7.4) with a complex frequency

$$\omega = \omega_r - i\frac{\nu}{2} \,, \tag{7.5}$$

the imaginary part is a damping factor for the wave with the damping rate

$$\nu_{\rm E} = \frac{1}{\langle u_{\rm L} \rangle} \frac{\partial}{\partial t} \langle u_{\rm L} \rangle \tag{7.6}$$

and the field energy density $u_{\rm L}$. This can be seen by considering the energy density

$$u_{\rm L} = \frac{1}{2} \varepsilon_0 \left(|\vec{E}|^2 + c^2 |\vec{B}|^2 \right) \tag{7.7}$$

with $|\vec{E}| = c|\vec{B}|$ for an electromagnetic wave [35], the complex frequency (7.5) and a plane wave (1.31), so that

$$u_{\rm L} = \varepsilon_0 |\vec{E}|^2 = \varepsilon_0 E_{\rm L}^2 \cdot e^{2i(\vec{k} \cdot \vec{r} - \omega t)} \cdot e^{-\nu \cdot t}$$
(7.8)

and

$$\nu_{\rm E} = -\frac{1}{\langle u_{\rm L} \rangle} \frac{\partial}{\partial t} \langle u_{\rm L} \rangle = -\frac{1}{\langle |\vec{E}|^2 \rangle} \frac{\partial}{\partial t} \langle |\vec{E}|^2 \rangle = \nu . \tag{7.9}$$

Insertion of (7.5) into (7.4) again in the high-frequency limit yields

$$\omega_r^2 - i\nu_E\omega_r = k^2 c^2 + \omega_{\rm pl}^2 - i\frac{\nu_{\rm ei}\omega_{\rm pl}^2}{\omega_r} , \qquad (7.10)$$

which after separation of real and imaginary part gives

$$\omega_r = \sqrt{k^2 c^2 + \omega_{\rm pl}^2} \,, \tag{7.11}$$

$$\nu_{\rm E} = \frac{\omega_{\rm pl}^2}{\omega_{\rm z}^2} \nu_{\rm ei} \ . \tag{7.12}$$

Clearly, the energy loss of the external field due to electron-ion collisions leads to an increase of system energy, since the electrons' oscillatory trajectories are converted to randomized motion due to collisions. Thus, the increase of total system energy $U_{\rm sys}$ is

$$\frac{\mathrm{d}}{\mathrm{d}t}U_{\mathrm{sys}} = -\frac{\mathrm{d}}{\mathrm{d}t}\langle \int u_{\mathrm{L}} \mathrm{d}V \rangle , \qquad (7.13)$$

where the integral of the laser energy density is performed over the complete system. From (7.9) with constant system volume,

$$\frac{\mathrm{d}}{\mathrm{d}t}U_{\mathrm{sys}} = \int \nu_{\mathrm{E}} \langle u_{\mathrm{L}} \rangle \mathrm{d}V \ . \tag{7.14}$$

Then making use of (7.12) and assuming spatial homogeneity, the expression

$$\frac{\mathrm{d}}{\mathrm{d}t}U_{\mathrm{sys}} = \frac{\omega_{\mathrm{pl}}^2}{\omega_r^2}\nu_{\mathrm{ei}} \cdot \varepsilon_0 \langle |\vec{E}|^2 \rangle \tag{7.15}$$

is obtained. The mean value $\langle \ldots \rangle$ represents averaging over a cycle, which – for harmonic waves (1.31) in the long-wavelength limit – gives

$$\langle |\vec{E}|^2 \rangle = \frac{1}{2} E_{\rm L}^2 \,.$$
 (7.16)

In the simulated system of N particles,

$$U_{\text{sys}} = \sum_{i=1}^{N} U_i = N \cdot U \tag{7.17}$$

with the average particle energy U, so that (7.15) can be recast into the form

$$\nu_{\rm ei} = \frac{\omega^2}{\omega_{\rm pl}^2} \cdot \frac{2}{\varepsilon_0 E_{\rm L}^2} \cdot n \cdot \frac{\mathrm{d}}{\mathrm{d}t} U , \qquad (7.18)$$

where n = N/V is the particle density. Using

$$v_{\rm osc} = \frac{e}{m} \frac{E_{\rm L}}{\omega} \,, \tag{7.19}$$

$$v_{\rm therm} = \sqrt{\frac{3k_{\rm B}T}{m}} = \sqrt{\frac{2E_{\rm therm}}{m}} , \qquad (7.20)$$

another useful expression

$$\nu_{\rm ei} = \frac{v_{\rm therm}^2}{v_{\rm esc}^2} \cdot \frac{1}{E_{\rm therm}} \frac{\mathrm{d}}{\mathrm{d}t} U \tag{7.21}$$

can be derived. Both expressions (7.18) and (7.21) are ready-to-use tools for deducing the electron-ion collision frequency through measuring the increase of average thermal energy per particle during the simulation.

Since averaging over a full cycle is necessary when computing the particles' kinetic energy or temperature, any laser- or configuration-induced drift will have to be subtracted appropriately. The latter may arise from non-uniform ion placement in the simulation box that induces large-scale oscillations with spatial extent $\gg \lambda_D$. Drift correction can be performed for example via

$$\frac{3}{2}k_{\rm B}T = \frac{m}{2}\left(\langle |\vec{v}|^2 \rangle - |\langle \vec{v} \rangle|^2\right) , \qquad (7.22)$$

where averaging over all particles of the considered single species – here electrons – is performed.

It is worth noticing, that via POYNTINGs theorem [35], the energy conversion rate is also connected to the microscopic current \vec{j} via

$$\nu_{\rm E} = \frac{\langle \vec{j} \cdot \vec{E} \rangle}{\varepsilon_0 \langle |\vec{E}|^2 \rangle} \ . \tag{7.23}$$

This is the usual approach for microscopic theories for $\nu_{\rm ei}$.

7.2 Simulations on the heating rate in hydrogen plasmas

7.2.1 Simulation setup

For measuring the heating rate $\nu_{\rm E}$ due to external radiation, we consider a neutral hydrogen plasma at different densities, temperatures and with varying frequency $\omega_{\rm laser}$ and intensity $I \propto v_{\rm osc}^2$ of the externally applied linearly polarized laser radiation. The latter parameters are chosen from the range

$$\frac{\omega_{\text{laser}}}{\omega_{\text{pl}}} = 0.1 \dots 100.0 \tag{7.24}$$

to sufficiently sample both the low- and high-frequency limits and

$$\frac{v_{\text{osc}}}{v_{\text{therm}}} = 0.1...20.0$$
 (7.25)

to cover a broad intensity range from the DAWSON-OBERMAN limit [56] $v_{\rm osc} \ll v_{\rm therm}$ to the high-intensity limit studied by SILIN [58]. In agreement with the available theoretical data, the force from the external radiation is only applied to the system's electrons.

$n_{\rm el}[{\rm cm}^{-3}]$	$T_{\rm el} [{ m eV}]$	0.1	0.5	1	2	5
1×10^{18}	$\Gamma =$	2.32	0.46	0.23	0.12	0.05
	$\Theta =$	27	137	274	548	1371
1×10^{20}	$\Gamma =$	10.8	2.15	1.08	0.54	0.22
	$\Theta =$	1.27	6.36	12.7	25.5	63.7
1×10^{22}	$\Gamma =$		10.0	5.00	2.50	1.00
	$\Theta =$		0.30	0.59	1.18	2.95
$n_{\rm el}[{\rm cm}^{-3}]$	$T_{\rm el} [{ m eV}]$	7.7	10	33	50	100
1×10^{18}	$\Gamma =$	0.03	0.02	0.01	0.005	0.002
	$\Theta =$	2111	2742	9050	13712	27424
1×10^{20}	$\Gamma =$	0.14	0.11	0.03	0.02	0.01
	$\Theta =$	98.0	127	420	636	1273
1×10^{22}	$\Gamma =$	0.65	0.50	0.15	0.10	0.05
	$\Theta =$	4.55	5.91	19.5	29.5	59.0

Table 7.1: Densities and temperatures used in the simulations in this work.

Ions are not directly affected. This is also justified by the time-scale of ionic motion which, thanks to the mass ratio $m_{ion}/m_{el} = 1,836$, is significantly longer than that of the electronic motion. Although this is similar to an adiabatic approximation, ions are allowed to gain additional energy and momentum due to collisions with electrons during the simulation. Hence, some fraction of the laser energy absorbed by the electron is transferred to ionic motion. For the sake of energy and momentum conservation in the simulation, fixing the ionic background is not an option here. Clearly, this transferred energy also has to be considered when measuring the system heating rate as described in Section 7.3.

The temperature and density values under consideration are given in Table 7.1 together with the respective values for the coupling parameter Γ and the degeneracy parameter Θ , defined by (1.7) and (1.8).

With $\Theta \lesssim 1$ especially the near-bulk density case $n_{\rm el} = 1 \times 10^{22} \, {\rm cm}^{-3}$ requires quantum effects to be taken into account. This is done through utilizing the Kelbg potential, a modification to the classical Coulomb interaction. Details are given in Section 7.2.2. Afterwards, Section 7.2.3 will describe the actually used particle configurations, numerical setup, and simulation workflow.

7.2.2 Kelbg potential

Naturally, the pure COULOMB potential (1.27) is not usable for the simulation of electron-ion interactions in dense systems. Due to the constrained timestep granularity, particles can approach each other closely. Thus, the potential's singularity at $r \to 0$ leads to an unphysical heating that would render any attempt of heating rate measurement futile. Accordingly, the singularity has to be removed. Furthermore, as already seen in Table 7.1, with $\Theta \lesssim 1$, quantum effects start to play a role in the selected parameter region.

Derived from a two-particle Slater sum, the Kelbg potential [307–310]

$$\Phi^{\text{(Kelbg)}}(r_{ij}) = \frac{q_i q_j}{4\pi\varepsilon_0} \cdot \left\{ \frac{1}{r_{ij}} \left(1 - e^{-\frac{r_{ij}^2}{\lambda_{ij}^2}} \right) + \frac{\sqrt{\pi}}{\lambda_{ij}} \left(1 - \operatorname{erf} \frac{r}{\lambda_{ij}} \right) \right\}$$
(7.26)

$$\lim_{r_{ij}\to 0} \Phi^{(\text{Kelbg})}(r_{ij}) = \frac{q_i q_j}{4\pi\varepsilon_0} \cdot \frac{\sqrt{\pi}}{\lambda_{ij}}$$
(7.27)

removes the singularity at $r \to 0$ due to Heisenberg's uncertainty effect (see Figure 7.1) and already includes quantum correlations [4]. Since (7.26) is a perturbation theoretical expansion in the interaction parameter (charge) $e^2 \to 0$, it is in general only valid for low to moderate coupling [4]. However, even in the region $\Gamma \gtrsim 1$, it still delivers an acceptable regularization of the Coulomb potential and is hence also applied there.

The parameter

$$\lambda_{ij} = \frac{\hbar}{\sqrt{\mu_{ij}k_{\rm B}T}} \tag{7.28}$$

is the thermal wavelength with the reduced mass [308]

$$\mu_{ij} = \left[\frac{1}{m_i} + \frac{1}{m_j}\right]^{-1} . {(7.29)}$$

The relevant species in our simulations are electrons and ions/protons. Thus, three combinations $(i,j) \in \{(e,p),(p,p),(e,e)\}$ have to be considered separately.

For $r/\lambda > 2$, the Kelbg potential smoothly approximates the Coulomb interaction, see Figure 7.1. In our simulations, we make use of the fact, that (7.26) is only a near-field correction to the Coulomb law. Instead of performing and including a full multipole expansion of $\Phi^{(\text{Kelbg})}$ into our tree code, we only replace (near) particle-leaf interactions, that were treated with the pure Coulomb expression otherwise, with the

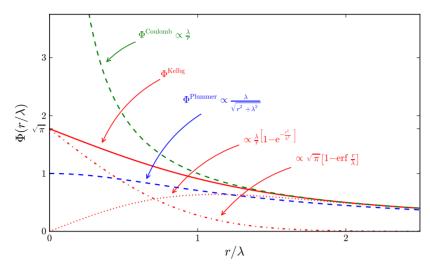


Figure 7.1: Comparison of Kelbg (solid red line) and Coulomb (dashed green line) potential as well as the Kelbg near- (dash-dotted red line) and far-field (dashed red line) contribution separately. For brevity, $\frac{q_iq_j}{4\pi\varepsilon_0}\frac{1}{\lambda}=1$ is set in the plot. Being finite at the origin (with the convention above $\Phi^{(\text{Kelbg})}(r\to 0)=\sqrt{\pi}$) the Coulomb singularity is removed by using the Kelbg expression. In addition, the frequently used Plummer potential $\Phi^{(\text{Plummer})}(r)\propto \frac{\lambda}{\sqrt{r^2+\lambda^2}}$, see also (2.16) is shown as dashed blue line.

Kelbg interaction and its respective expression for the force

$$\vec{F}^{(\text{Kelbg})}(\vec{r}_{ij}) = \frac{q_i q_j}{4\pi\varepsilon_0} \left(1 - e^{-\frac{|\vec{r}_{ij}|^2}{\lambda_{ij}^2}} \right) \frac{\vec{r}_{ij}}{|\vec{r}_{ij}|^3} , \qquad (7.30)$$

$$\lim_{r \to 0} \vec{F}^{(\text{Kelbg})}(r \cdot \vec{e}_{ij}) = \frac{q_i q_j}{4\pi\varepsilon_0} \frac{\vec{e}_{ij}}{\lambda_{ij}^2} \,. \tag{7.31}$$

This introduces a discontinuity of the forces for distances where interactions with particle clusters through their multipole expansion start to appear. However, such discontinuities are evident anyway due to the different levels of accuracy when interaction with clusters at different tree levels are considered and does not lead to an enhanced numerical heating, see Figure 7.2. Also, momentum conservation is not found to be violated by this approximation.

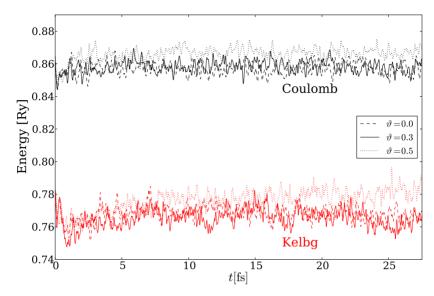


Figure 7.2: Comparison of total energy over time for Kelbg and Coulomb potential to verify energy conservation. See Figure 4.6 for the numerical setup. The different interactions yield different potential energy. Accordingly, the total energy is modified when exchanging the potential. Only for $\vartheta > 0.3$ a significant energy drift is observed for the Kelbg interaction. This results form the fact, that the Kelbg potential used in the near-field is not continuously matched to the far-field Coulomb interaction. Thus, as we are using $\vartheta = 0.3$, see also Section 3.1.3, numerical heating will not spoil our results.

7.2.3 Simulation workflow

Simulations are performed for an overall neutral two-component plasma in a cubic simulation region that is periodically continued as described in Chapter 4. To obtain sufficient statistics even for small heating rates, particle numbers are chosen between $N_{\rm el}=N_{\rm ion}=5,000\ldots2,500,000$ electrons and ions. Depending on system size, the calculations are done on the general purpose cluster JuRoPa [297] and the high-scalability IBM Blue Gene/Q system JUQUEEN [268] on up to 16,384 processor cores.

A typical simulation scenario is shown in Figure 7.3 for a configuration of $N_{\rm el} = N_{\rm ion} = 5,000$ electrons and ions. The method is similar to that of [103]. After being initialized with MAXWELLian velocities at equal temperatures $T_{\rm el} = T_{\rm ion}$, the system is in a non-equilibrium state due to the randomly assigned particle positions. Accordingly, the potential energy is inevitably too high, but the system rapidly adjusts by seeking a

lower potential energy state, accelerating and thus heating the electrons. Due to the mass ratio $m_{\rm p} \approx 1,836 m_{\rm el}$ between protons and electrons, the ions are mainly not affected here. A relaxation to thermal equilibrium is not possible within the next few femtoseconds and a two-temperature plasma is evident during this phase I. In phase II, a thermostat is applied to both species separately to obtain a defined equilibrium state at the desired temperature again. It is deactivated in the next phase III which is used to verify energy conservation before switching on the external laser in phase IV.

After activating the external harmonic field, the electron kinetic energy starts to oscillate with doubled laser frequency due to the laser-induced collective motion. During this oscillation, the electrons undergo collisions with the ions that are – compared to the electron velocity – essentially at rest. Due to these collisions, some fraction of the directed electronic oscillation is converted into random omnidirectional motion. Consequently, the drift-corrected electron kinetic energy $U_{\rm kin,\,w/o\,drift}^{\rm (electrons)}$ that corresponds to the systems temperature (7.22) shows a steady increase. The energy transfer to the ions due to the collisions is negligible. Consequently, their kinetic energy and hence temperature stays approximately constant.

The electronic heating due to collisions is exactly the process to be studied with our simulations. As shown in Section 7.1 using (7.18) or (7.21), the rate of energy transfer is used to determine the electron-ion collision frequency $\nu_{\rm ei}$. Since this should ideally be done at fixed plasma conditions, in particular for $T_{\rm el} = {\rm const}$, the next section describes thermostat algorithms that allow for measuring heating rates at constant temperature and proposes an additional alternative scheme of measurement.

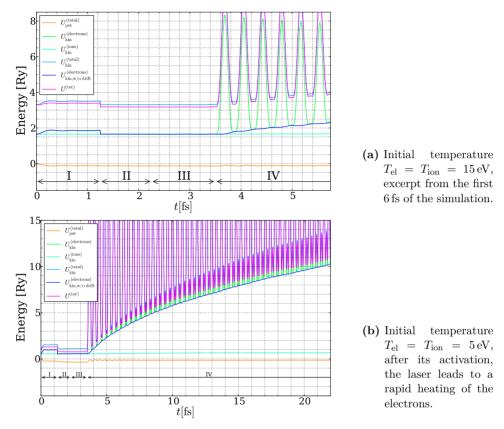


Figure 7.3: Typical simulation scenario as used in this chapter, here for $N_{\rm el}=N_{\rm ion}=5,000$ electrons and ions with identical initial temperatures and densities $n_{\rm el}=n_{\rm ion}=1\times 10^{22}\,{\rm cm^{-3}}$. Due to the non-equilibrium random particle configuration, during phase I until $t=1.25\,{\rm fs}$, a rapid decrease of potential energy and thus heating of the electrons is observed. To reach a single-temperature plasma in equilibrium, in phase II ($t=1.25\dots 2.25\,{\rm fs}$) a thermostat is applied. During phase III until $t=3.5\,{\rm fs}$, it is deactivated to verify energy conservation. At $t=3.5\,{\rm fs}$ the external harmonic field with $v_{\rm osc}/v_{\rm therm}=0.2$ and $\omega/\omega_{\rm pl}=3.0$ is activated leading to a considerable heating of the electrons in the system. Their total kinetic energy oscillates with doubled laser frequency. Due to collisions, their collective oscillation is converted to random motion which shows as an increase of the drift-corrected kinetic energy. The slope of this increase corresponds to the heating rate $\nu_{\rm E}$.

7.3 Method of measurement

As seen in Section 7.1, the electron-ion collision frequency for the particle systems under consideration can be computed from the particle heating rate, which is a measure for the actual energy transfer into the system. Since the collision frequency is known to be strongly dependent on temperature T already in leading order [48, 51]

$$\nu_{\rm ei} \propto \frac{n}{T^{3/2}} \tag{7.32}$$

with the particle density n, measurement of the heating rate during the simulation is a delicate process. As seen in Figure 7.3, the system is considerably heated by the incident laser. Consequently, either the measurement has to be performed instantaneously by computing the value and gradient of the temperature curve at every instant in time, or by keeping the temperature artificially constant through some thermostat algorithm. In the latter case, the energy that is drained from the system through the thermostat has to be measured. Both strategies entail various advantages and drawbacks, that will be summarized in the following.

7.3.1 Constant-temperature simulations

Conserving particle number N, simulation volume V, and system energy E, standard methods of molecular dynamics simulations in general sample the phase space of a microcanonical ensemble. In contrast, macroscopic experiments usually are subject to the canonical ensemble, where instead of total energy the system temperature T is externally fixed, e.g. via a heat bath, or even the isothermal-isobaric ensemble, where volume can vary to keep the pressure P constant. Consequently, numerous different methods have been developed to fix the mean kinetic energy, which is system temperature. For the purpose of fixing the temperature during laser irradiation in the conducted simulation experiment, two of these were used: the Berendsen thermostat that introduces artificial friction and the Nosé-Hoover scheme based on virtual scaling of the physical time, which can also be interpreted as extending the Lagrangian of the system to include a virtual heat bath. Both are briefly summarized in the following paragraphs to evaluate their applicability in the simulations.

From velocity rescaling to artificial friction – The Berendsen thermostat

The most obvious technique for fixing the particle temperature is a linear scaling

$$\vec{v}_i(t) \to \sqrt{\frac{T_0}{T(t)}} \cdot \vec{v}_i(t)$$
 (7.33)

of every particle's velocity to obtain the desired mean kinetic energy. Here, T(t) is the instantaneously measured, uncorrected temperature and T_0 its desired value. This method was first proposed by WOODCOCK [311]. It conserves total linear and angular momentum while strictly fixing the mean kinetic energy per particle. By construction, fluctuations of the temperature are not allowed here. However, in the canonical ensemble, temperature does fluctuate [9, 145]. One way to allow such fluctuations is to introduce a friction term into the equations of motion [312], which leads to the damped force method [313]

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{v}_i = \frac{1}{m_i} \left(\vec{F}_i - \alpha \vec{v}_i \right) \tag{7.34}$$

with particle velocities $\vec{v_i}$ and forces $\vec{F_i}$ and an appropriately chosen constant α . If, with the temperature definition

$$\frac{3}{2}(N-1)k_{\rm B}T = \sum_{i=1}^{N} \frac{m_i}{2} |\vec{v}_i|^2 \tag{7.35}$$

the temperature T is required to remain constant

$$\frac{\mathrm{d}}{\mathrm{d}t}T = 0 , \qquad (7.36)$$

then

$$\alpha = \frac{\sum_{i=1}^{N} \vec{F}_{i} \cdot \vec{v}_{i}}{\sum_{i=1}^{N} \vec{v}_{i} \cdot \vec{v}_{i}}$$
 (7.37)

is obtained. This approach maintains the canonical distribution [312, 314]. Due to the formal constraint of fixed total momentum, the system in (7.35) has 3(N-1) instead of 3N degrees of freedom.

For inclusion of the force damping to fix the system temperature at $T = T_0$, the velocity step of the leap-frog integrator (2.11) is modified as follows [313, 315]. Firstly, an unconstrained velocity half-step

$$\vec{v}'_i(t) = \vec{v}_i(t - \Delta t/2) + \frac{1}{m_i} \vec{F}_i(t) \cdot \Delta t/2$$
 (7.38)

is performed to get an estimate of the projected temperature at time t. Then, with

$$\beta = \frac{1}{1 + \frac{\alpha}{m} \Delta t/2} = \sqrt{\frac{3(N-1)k_{\rm B}T_0}{m\sum_{i=1}^{N} |\vec{v}_i'(t)|^2}}$$
(7.39)

the full velocity step reads

$$\vec{v}_i(t + \Delta t/2) = \vec{v}_i(t - \Delta t/2) \cdot (2\beta - 1) + \frac{\beta}{m} \vec{F}_i(t) \cdot \Delta t$$
 (7.40)

Thus, in contrast to the trivial velocity rescaling, the forces are also rescaled in this damped force approach with a leap-frog integrator.

The choice of (7.39) with (7.38) guarantees the desired temperature at instant t. However, being essentially a linear rescaling process, the temperature compensation tends to overshoot for $t + \Delta t/2$. This leads to an undesired oscillation between too high and too low average kinetic energy at times $t - \Delta t/2$ and $t + \Delta t/2$. Due to this non-continuous behavior, initial temperature perturbations are not damped but carried along as a hidden artifact. This can be circumvented by allowing the system to relax to the desired temperature within a larger time-span instead of forcing the relaxation to happen immediately during the next timestep. The actual way of damping the system temperature towards its desired value is given in [316, 317], where particle velocities are rescaled from \vec{v} to $\lambda \vec{v}$ with

$$\lambda = \sqrt{1 + \frac{\Delta t}{\tau} \left[\frac{T_0}{T(t)} - 1 \right]} \,, \tag{7.41}$$

which in (7.34) corresponds to [318]

$$\alpha = \frac{1}{2\tau} \left[\frac{T_0}{T(t)} - 1 \right] . \tag{7.42}$$

Here, T(t) is the measured uncorrected temperature and τ can be interpreted as a relaxation time, so that

$$\frac{\mathrm{d}}{\mathrm{d}t}T(t) = \frac{T_0 - T(t)}{\tau} \ . \tag{7.43}$$

Clearly, temperature is not constrained immediately in this scheme. Instead, the system slowly develops towards T_0 via corrected temperatures $\mathcal{T}(t)$. Thus, as long as the temperature deviations are small enough, no overshooting appears and without external perturbations, T_0 will be reached while temperature fluctuations are still allowed.

This constant temperature approach can also be used to permanently remove energy from the laser-irradiated system that is heated by the external field. We want to use (7.21) to determine the electron-ion collision frequency. Therefore, the amount of heat, i.e. inner energy that is removed from the system through the thermostat has to be measured. Although $\frac{d}{dt}U$ is directly connected to β and α , respectively, in the current implementation the projected velocities $\vec{v}_i'(t)$, see (7.38), and corrected velocities $\vec{v}_i(t)$ are used to compute the uncorrected temperatures T(t) and their corrected values T(t). These velocities can be determined by performing the integration in half-steps. Using the BERENDSEN approach of slow temperature relaxation, it is important to note that T(t) is in general not equal to T_0 . Then,

$$\frac{\mathrm{d}}{\mathrm{d}t}U \approx \frac{\Delta U}{\Delta t} = \frac{3}{2}Nk_{\mathrm{B}}\frac{T(t) - \mathcal{T}(t)}{\Delta t}$$
(7.44)

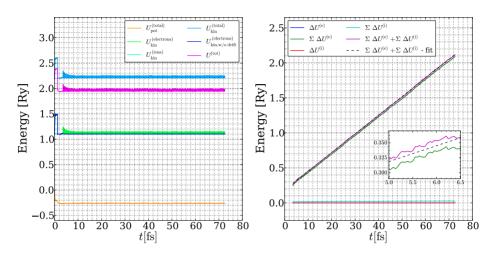


Figure 7.4: System energy per particle for a homogeneous neutral plasma with $N_{\rm el}=N_{\rm ion}=5{,}000$ particles (left). The simulation workflow follows Figure 7.3 with the difference that with activation of the external field at $t=3.5\,{\rm fs}$, a separate Berendsen thermostat is applied to each species. This keeps the electron and ion temperature constant. The amount of energy $\Sigma\Delta U^{(e)}+\Sigma\Delta U^{(i)}$ that is drawn from both subsystems per particle (right) is fit with a linear function (dashed line). The graph of the fit function lies much too close to be distinguishable from the data. The inset that shows an excerpt for times $t=5.0\,{\rm fs}\dots 6.5\,{\rm fs}$ allows for an overview of the actual deviation between fit and simulation data.

is a measure for the amount of energy that is converted per timestep from directed laser-driven drift to undirected motion. Again, it is important that T(t) and T(t) only contain drift-corrected velocity contributions as given in (7.22), since the laser-induced oscillation is clearly not allowed to contribute here.

Being a stochastic property, $\Delta U/\Delta t$ has to be measured over several laser and plasma periods (whichever is longer) to yield reliable results. This can be done conveniently by accumulating ΔU , i.e. collecting all energy that has been removed over time

$$\Sigma \Delta U(k \cdot \Delta t) := \sum_{i=0}^{k} \Delta U(i \cdot \Delta t) . \tag{7.45}$$

The slope of $\Sigma \Delta U$ yields $\frac{d}{dt}U$.

This process is shown for an example setup in Figure 7.4. Here, a homogeneous neutral plasma system with $N_{\rm el}=N_{\rm ion}=5{,}000$ particles and $T_{\rm el}=T_{\rm ion}=10\,{\rm eV}$ is heated by

an external electric field with $v_{\rm osc}/v_{\rm therm}=0.2$. During its first phases, the workflow is identical to Figure 7.3. However, in this case, the thermostat is now active in stage IV, i.e. after the laser has been switched on at $t=3.5\,{\rm fs}$, and the drift-corrected electron kinetic energy remains constant. The laser-induced oscillation is still visible. The amount of energy converted from collective drift to random undirected motion can be seen in the graph of $\Sigma\Delta U_{\rm el}$, that represents the accumulated energy drawn from the electron subsystem.

Though not being directly affected by the laser, due to electron-ion collisions the ionic subsystem is also heated. With their significantly larger mass $m_{\rm ion} \approx 1,836 m_{\rm el}$, this process is considerably slower than for the electrons, but still visible in the simulations. To ensure definite fixed physical conditions, a second thermostat is applied to the plasma ions and their heating is also included in the calculation of the total heating rate.

Using a linear regression for $\Sigma \Delta U_{\rm el} + \Sigma \Delta U_{\rm ion}$ in Figure 7.4, the system heating rate $\frac{\rm d}{{
m d}t}U$ can be identified as the slope of the linear fit. With (7.18) or (7.21), the collision frequency can then be computed. Results – also in comparison with other thermostat and heating rate measurement approaches – are given in Section 7.4.

The Nosé-Hoover thermostat

Besides the velocity rescaling and damped-force approach, another popular and successful method of introducing constraints on thermodynamic observables in molecular dynamics simulations is due to Nosé, who originally formulated it for constant pressure [319, 320] and later for constant temperature dynamics [321, 322]. In this approach, an additional degree of freedom s is introduced, that takes the role of an external heat bath. Its interaction with the physical system is defined via a scaling of the momenta $\vec{p}_i = m_i s^2 \dot{\vec{r}}_i$. With the potential energy $(f+1)k_{\rm B}T_0 \ln s$, where f is the number of degrees of freedom of the original system and T_0 its desired temperature, and the kinetic term $\frac{p_s^2}{2Q}$, the total system Hamiltonian

$$\mathcal{H}_N = \sum_{i=1}^N \frac{\vec{p}_i^2}{2m_i s^2} + U(\{\vec{r}_{i=1\dots N}\}) + \frac{p_s^2}{2Q} + (f+1)k_B T_0 \ln s$$
 (7.46)

is a conserved quantity. With the kinetic term and the virtual momentum p_s , a dynamic equation for s can be constructed. Obviously, the Q parameter is analogous to a classical mass, though it has units of energy \cdot time². It is used to influence the rapidity of heat exchange dynamics between the system and the heat bath. The partition function

$$Z = \frac{1}{N} \int dp_s \int ds \int d^N \vec{p} \int d^N \vec{r} \, \delta \left(\mathcal{H}_N - E^{(\text{ext})} \right)$$
 (7.47)

with the extended system energy $E^{(\text{ext})}$ can be integrated after the substitution $\vec{p}_i/s = \vec{p}_i$ and using $\delta(f(s)) = \frac{\delta(s-s_0)}{|f'(s_0)|}$ if s_0 is the only zero of f(s), see [145, 322–324]. This yields

$$Z = \frac{1}{f+1} \sqrt{\frac{2\pi Q}{k_{\rm B} T_0}} e^{\frac{E(\text{ext})}{k_{\rm B} T_0}} \cdot Z^{(\text{can})}$$
 (7.48)

with the canonical partition function

$$Z^{(\text{can})} = \frac{1}{N!} \int d^N \vec{p} \int d^N \vec{r} e^{\frac{\mathcal{H}(\vec{p}', \vec{r})}{k_B T_0}}$$

$$(7.49)$$

and the physical (real system) Hamiltonian

$$\mathcal{H}(\vec{p}', \vec{r}) = \sum_{i=1}^{N} \frac{\vec{p}_i'^2}{2m_i} + U(\{\vec{r}_{i=1...N}\}) . \tag{7.50}$$

Thus, with (7.48) a canonical ensemble is sampled when deriving dynamic equations from (7.46) under the assumption of ergodicity.

In the original formulation of Nosé, s can be interpreted as a parameter that scales every timestep [321, 322]. This is clearly undesirable in a molecular dynamics simulation that relies on equidistant sampling of the time domain. Using the non-canonical change of variables

$$\vec{p}_i' = \frac{\vec{p}_i}{s}$$
, $dt' = \frac{\mathrm{d}t}{s}$, $\frac{1}{s} \frac{\mathrm{d}}{\mathrm{d}t'} s = \frac{\mathrm{d}}{\mathrm{d}t'} \eta$, $p_s = p_\eta$, (7.51)

and f = 3N - 1 due to the fixed temperature, the equations of motion

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{r}_{i} = \frac{\vec{p}_{i}}{m_{i}},$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{p}_{i} = \vec{F}_{i} - \frac{p_{\eta}}{Q}\vec{p}_{i},$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\eta = \frac{p_{\eta}}{Q},$$

$$\frac{\mathrm{d}}{\mathrm{d}t}p_{\eta} = \sum_{i=1}^{N} \frac{\vec{p}_{i}^{2}}{m_{i}} - 3Nk_{\mathrm{B}}T_{0}$$
(7.52)

can be derived after removing the primes from all variables.

Though being non-Hamiltonian, this set of variables still samples the canonical phase space, but now in the real physical variables *including time*. The newly introduced variable η is general not necessary for the physical evolution. It plays an important role in the analysis of the phase-space distribution [325] and will also appear in the Hamiltonian of the extended system which is used for evaluation of the heating rate.

The reformulation (7.51) of the original approach was introduced by HOOVER [326] and has since become part of the standard toolbox in molecular dynamics. The set of dynamic equations (7.52) are thus labeled NOSÉ-HOOVER equations.

As before, there is an additional friction term in the physical momentum equation in (7.52) that – depending on its sign – can also act as an acceleration. Again, its dynamic is driven by the instantaneous difference between current and desired temperature value. In contrast to the BERENDSEN thermostat, the heat bath is now intrinsically included in the formalism. Thus – despite the temperature rescaling in the physical system – the extended system HAMILTONian that together with (7.51) reads

$$\mathcal{H}_{N}(\{\vec{r}_{i=1...N}\}, \{\vec{p}_{i=1...N}\}, \eta, p_{\eta})$$

$$= \mathcal{H}(\{\vec{r}_{i=1...N}\}, \{\vec{p}_{i=1...N}\}) + \frac{p_{\eta}^{2}}{2Q} + 3Nk_{B}T_{0}\eta, \qquad (7.53)$$

is an integral of motion. Hence its conservation allows for verification of the approach itself as well as numerical correctness in the implementation.

It can be shown that Nosé-Hoover dynamics violates ergodicity for small and stiff systems [326, 327] and several attempts were made to circumvent this problem, e.g. in [325, 328]. The most popular approach by Martyna et al. introduces further extended system variables η_2, p_{η_2} that act as thermostats for $\eta = \eta_1$ and $p_{\eta} = p_{\eta_1}$. This guarantees the latter to be Maxwell-Boltzmann distributed again and finally builds a hierarchical chain of thermostats, since also η_2 and p_{η_2} have to be thermally distributed. Otherwise leading to an infinite so-called Nosé-Hoover chain [145, 324, 325], this hierarchy has to be truncated after adding M thermostat momenta. The choice of M=1 reduces to the original Nosé-Hoover dynamics and will be used within this work, since only large nonstiff systems are studied. Here, an early truncation of the chain is appropriate without considerably violating the canonical ergodic sampling of phase-space.

A formulation of Nosé-Hoover chain dynamics with the leap-frog integrator (2.10) and (2.11) has already been given in [329]. However, relying on an iterative procedure for determination of the particle and extended system momenta, it violates time reversal symmetry and increases computational cost in the integrator. A later approach in [330] solves the former issue, leading to an even more elaborate algorithmic concept. Thus, a much simpler method was chosen for this work. It essentially includes the extended system variables of the Nosé-Hoover chain into the leap-frog scheme [331]. In the

formulation with velocities instead of momenta it reads

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{r}_i = \vec{v}_i \,, \tag{7.54}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\vec{v}_i = \vec{a}_i(\{\vec{r}_{i=1...N}\}) - v_{\eta_1}\vec{v}_i , \qquad (7.55)$$

$$\frac{\mathrm{d}}{\mathrm{d}t}v_{\eta_1} = a_{\eta_1}(\{\vec{v}_{i=1...N}\}) - v_{\eta_2}\vec{v}_{\eta_1} , \qquad (7.56)$$

$$\frac{\mathrm{d}}{\mathrm{d}t}v_{\eta_k} = a_{\eta_k}(v_{\eta_{k-1}}) - v_{\eta_{k+1}}\vec{v}_{\eta_k} , \qquad k = 2, \dots, M ,$$
 (7.57)

where $\vec{a}_i = \vec{F}_i/m_i$ are the usual (force-driven) physical accelerations and the Nosé accelerations are given via

$$a_{\eta_1}(\{\vec{v}_{i=1...N}\}) = \frac{1}{Q_1} \left(\sum_{i=1}^N m_i \vec{v}_i^2 - 3Nk_B T_0 \right) , \qquad (7.58)$$

$$a_{\eta_k}(v_{\eta_{k-1}}) = \frac{1}{Q_k} \left(Q_{k-1} v_{\eta_{k-1}}^2 - k_B T_0 \right) \quad \text{for} \quad k = 2, \dots, M$$
 (7.59)

and the Nosé positions satisfy

$$\frac{\mathrm{d}}{\mathrm{d}t}\eta_k = v_{\eta_k} \ . \tag{7.60}$$

Now, the two sets $S_x = \{\vec{r}_i, \eta_{2k}, v_{\eta_{2k-1}}\}$ and $S_v = \{\vec{v}_i, \eta_{2k-1}, v_{\eta_{2k}}\}$ with i = 1, ..., N and k = 1, ..., M/2 are defined. All members of S_x are computed on full timesteps t and members of S_v on the half timesteps $t \pm \Delta t/2$ in the sense of the leap-frog scheme.

For M=1 as in this work, the NOSÉ-HOOVER leap-frog algorithm according to [331] can be reduced to

$$\vec{r}_{i}(t) = \vec{r}_{i}(t - \Delta t) + \vec{v}_{i}(t - \Delta t/2) \cdot \Delta t$$

$$v_{\eta_{1}}(t) = v_{\eta_{1}}(t - \Delta t) + a_{\eta_{1}}(t - \Delta t/2) \cdot \Delta t$$

$$\vec{v}_{i}(t + \Delta t/2) = \vec{v}_{i}(t - \Delta t/2) \cdot e^{-\Delta t v_{\eta_{1}}(t)} + \vec{a}_{i}(t) \cdot \Delta t \cdot e^{-\Delta t/2 v_{\eta_{1}}(t)}$$

$$\eta_{1}(t + \Delta t/2) = \eta_{1}(t - \Delta t/2) + v_{\eta_{1}}(t) \cdot \Delta t$$

$$a_{\eta_{1}}(t + \Delta t/2) = \frac{1}{Q_{1}} \left(\sum_{i=1}^{N} m_{i} \vec{v}_{i}(t + \Delta t/2)^{2} - 3Nk_{B}T_{0} \right) .$$

$$(7.61)$$

Hence, only a minor modification of the original integrator from Section 2.2.2 is necessary.

It should be mentioned that the algorithm of [331] is not undisputed, see [332]. For certain simulation setups it is not as energy-conserving as the approach of [330]. However, it was shown that this does not pose a problem onto realistic scenarios [333] and

that both approaches can even be derived from a similar factorization of the Liouville propagator [334]. In the present simulations, conservation of total energy is explicitly checked by performing test runs with vanishing laser intensity and was found to be satisfied by (7.61).

As seen before, compare Section 7.1, the system heating rate and thus the change in total energy is the quantity of interest for determining the electron-ion collision frequency. From (7.53) it can be seen that the heating of the system, which results in an increase of total system energy and hence \mathcal{H} , will also lead to an increase in \mathcal{H}_N . Thus, even although the thermostat permanently drains energy from the physical system to keep its temperature constant, all changes in \mathcal{H}_N only reflect the energy transfer from and into the full system including the heat bath.

For computing \mathcal{H}_N , knowledge of all phase-space variables at full timesteps t is necessary. Since members of S_v are only defined at the intermediate half steps $t \pm \Delta t/2$, their values at t have to be computed explicitly. Performing from (7.61) one half step backwards,

$$\vec{v}_i(t) = \frac{1}{2} \left[\vec{v}_i(t + \Delta t/2) \left(1 + e^{\Delta t v_{\eta_1}(t)} \right) - \vec{a}_i(t) \cdot \Delta t e^{\Delta t/2 v_{\eta_1}(t)} \right] , \qquad (7.62)$$

$$\eta_1(t) = \eta_1(t + \Delta t/2) - v_{\eta_1}(t) \cdot \Delta t/2 \tag{7.63}$$

is found. Accordingly, $\mathcal{H}_N(t)$ can conveniently be evaluated after performing the position and velocity update in the simulation.

The only free variable in the Nosé-Hoover thermostat approach is the choice of the relaxation parameters Q_i . According to [325], the optimal choice is

$$Q_1 = 3Nk_{\rm B}T_0 \cdot \tau^2 \,, \tag{7.64}$$

$$Q_{i=2...M} = k_{\rm B} T_0 \cdot \tau^2 \,. \tag{7.65}$$

Here, τ is the characteristic temperature relaxation time scale. Clearly, in addition to being small enough to sufficiently compensate for heating due to the external field and large enough to avoid parasitic oscillation, the relaxation timescale also has to be chosen off-resonant with any internal system frequencies, such as the plasma frequency $\omega_{\rm pl}$, the laser frequency $\omega_{\rm laser}$ and their possible beat modes. Keeping this constraint in mind, in the present simulations,

$$\tau \approx 20 \cdot \Delta t \tag{7.66}$$

was chosen, though, no significant influence on the overall heating rate was found with other choices.

Figure 7.5 shows the same setup with $N_{\rm el} = N_{\rm ion} = 5{,}000$ particles and $T_{\rm el} = T_{\rm ion} = 10\,{\rm eV}$ as in Figure 7.4. Here, instead of the drained energy, the extended system HAMILTONian \mathcal{H}_N that includes the energy transferred to the heat bath is observed. As before,

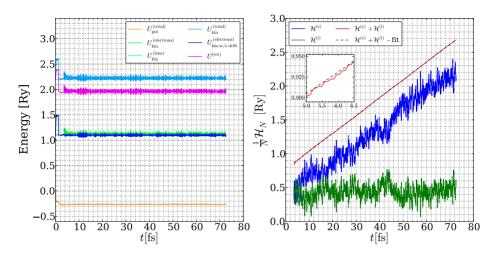


Figure 7.5: System energy (left) and extended system Hamiltonian \mathcal{H}_N (right) per particle for the identical setup as in Figure 7.4 with a Nosé-Hoover thermostat. Again, the (dashed) linear fit function to the total extended system Hamiltonian $\mathcal{H}_N^{(e)} + \mathcal{H}_N^{(i)}$ lies too close to the simulation data to be distinguishable. The inset shows an excerpt for times $t = 5.0 \, \mathrm{fs} \dots 6.5 \, \mathrm{fs}$ to get an idea of their proximity. The slope of the fit is identical to that of Figure 7.4, yielding the same heating rate as with the Berendsen thermostat.

heating primarily occurs in the electronic system. The extended system HAMILTONians for the electronic and ionic subsystem $\mathcal{H}_N^{(\mathrm{e})}$ and $\mathcal{H}_N^{(\mathrm{i})}$ strongly fluctuate. However, their sum $\mathcal{H}_N^{(\mathrm{e})} + \mathcal{H}_N^{(\mathrm{i})}$ shows a nearly linear line as expected for a constant heating rate. Consequently it is presumable that a considerable energy exchange between ions and electrons occurs and influences their kinetic and potential energy. Again, the total extended system HAMILTONian $\mathcal{H}_N^{(\mathrm{e})} + \mathcal{H}_N^{(\mathrm{i})}$ is evaluated for the heating rate measurement. Now, besides the kinetic contributions, it also contains the potential energy. Hence, even a hypothetical increase in correlation energy due to the laser-plasma interaction could be covered by this approach. As before, the slope of a linear fit of $\mathcal{H}_N^{(\mathrm{e})} + \mathcal{H}_N^{(\mathrm{i})}$ yields the overall system heating rate $\frac{\mathrm{d}}{\mathrm{d}t}U$.

Either presented thermostat approach offers the possibility of measuring the heating rate in a canonical ensemble conveniently and with only small modifications to the particle trajectory integrator. Keeping the physical conditions essentially fixed, they furthermore allow for straightforward averaging over a long period of time with constant

 $v_{\rm osc}/v_{\rm therm}$. Additionally, from the linear regression error, an estimate for the statistical uncertainty can be derived.

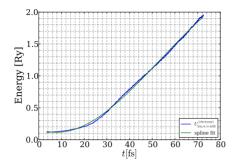
The thermostat approach is an appealing method for making use of ergodicity and replacing ensemble averages by a time average that is realized by means of a linear fit to the drained energy/total \mathcal{H}_N . However, it is limited to a slightly non-intuitive discrete sampling of the temperature domain. For every set of temperatures $T_{\rm el}$, $T_{\rm ion}$, a full simulation run has to be performed. In the next section, an alternative simulation scheme is proposed that allows for sweeping through temperature ranges by making use of the laser-induced heating that is already present anyway.

7.3.2 Instantaneous heating rate measurement

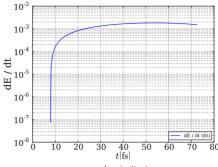
In the previous section, two approaches for simulations at constant temperature were presented. These allow for determination of the collision-driven plasma heating rate at distinct points in the NVT-parameter space. However, this approach itself is quite cumbersome. In fact, it is not necessary to really fix the particle temperature and then determine the heating rate by means of accumulating the energy that is drained from the system. Instead, a natural approach would be to let the system heat up and determine the heating rate $\frac{\mathrm{d}}{\mathrm{d}t}U$ directly from the slope of the temperature curve without artificially cooling the system again. With this approach, that was also proposed in [103] but has not been followed further, a range of system temperatures can be swept during a single simulation. Clearly, this method does not need any modifications to the particle dynamics itself. By not imposing any artificial constraint, it is significantly closer to the real physical system that is simulated.

Figure 7.6 gives a pictorial view of how this is actually implemented. It shows the drift-corrected kinetic energy per electron after laser activation and coefficients that are derived from that for an example simulation with $N_{\rm el}=N_{\rm ion}=5,000$ particles and initial temperatures of $T_{\rm el}=1.0\,{\rm eV}$ and $T_{\rm ion}=1,000\,{\rm K}$ and densities $n_{\rm el}=n_{\rm ion}=1\times10^{22}\,{\rm cm}^{-3}$. The laser frequency is $\omega_{\rm laser}=3\,\omega_{\rm pl}$. These parameters are the same as in simulations performed in [103]. The according data is included for comparison in Figure 7.6d. As we will compare our simulation results to the data of [103] extensively in Section 7.4, we postpone any discussion of the observed data.

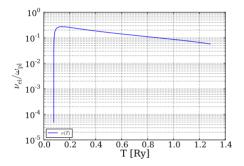
As already seen in Figure 7.3, without a thermostat the system is substantially heated by the external field. To reduce stochastic fluctuations when evaluating the time derivative, the heating curve in Figure 7.6a is fitted by a smooth analytic expression. Therefore, a piecewise cubic spline function $U^{(\text{spline})}(t)$ has proven to be suitable. Figure 7.6b shows its derivative $\frac{d}{dt}U^{(\text{spline})}(t)$ – the heating rate. Using the inverse function of $U^{(\text{spline})}(t)$, the explicit dependence on time t can be replaced by a temperature parameter and using (7.21), the temperature-dependent collision frequency $\nu_{\text{ei}}(T)$ can be computed,



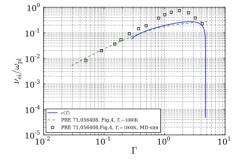
(a) For smoothing stochastic fluctuations, the electron kinetic energy $U_{\rm kin}^{({\rm e})}({\rm t})$ is fit by a piecewise cubic spline function $U^{({\rm spline})}(t)$.



(b) The derivative $\frac{d}{dt}U^{\text{(spline)}}(t)$ delivers the system heating rate.



(c) Using the inverse function of $U^{\text{(spline)}}(t)$, the temperature-dependent heating rate $\frac{\mathrm{d}}{\mathrm{d}t}U(T)$ can be computed that is used to evaluate (7.21) to get $\nu_{\mathrm{ei}}(T)$.



(d) After an additional variable transformation $T \to \Gamma$, the simulation data can for example be compared to results from [103] which are included here as squares and broken line. In contrast to their simulations for distinct temperature values, a temperature range is covered by every single simulation in the instantaneous heating rate scheme.

Figure 7.6: Demonstration of the instantaneous heating rate measurement for a system with $N_{\rm el}=N_{\rm ion}=5{,}000$ particles and initial temperatures of $T_{\rm el}=1.0\,{\rm eV}$ and $T_{\rm ion}=1{,}000\,{\rm K}$ and densities $n_{\rm el}=n_{\rm ion}=1\times10^{22}\,{\rm cm}^{-3}$. These parameters correspond to similar simulations and theoretical considerations in [103]. Data from that publication is included in Figure 7.6d for comparison.

see Figure 7.6c. With a final variable transformation from temperature to the coupling parameter Γ , see (1.8), a comparison with the results from [103] is possible and already shows adequate consistency.

Despite being a very promising approach for determination of the heating rate, this does not come for free. Especially in the region of the upper and lower temperature boundary, the resulting heating rate is not completely reliable. For the very small temperatures directly after the step-like laser activation, the electrons will not immediately move harmonically. Hence reliable heating rates might not be observed during the first few laser periods. This leads to an underestimation of the heating rate during the first time of measurement. Finally, in the high-temperature region, the timestep conditions (compare Section 2.2.3) that are only checked at simulation start-up by setting an appropriately small timestep are approached or even overrun due to the significantly higher particle velocities in the hot system. Thus, also when approaching very high temperatures, care has to be taken when interpreting collision frequencies from such a run. A corresponding artifact can be seen at $\Gamma = 0.3$ in Figure 7.6d, where $\nu_{\rm ei}$ seems to suddenly decrease faster than before. However, both issues can be mitigated by performing simulations with different initial temperatures to cover the low-temperature region and small enough timestep to also treat the high-temperature region correctly. A good compromise is to avoid covering broad Γ ranges with a single simulation, but to bridge them with a number of shorter temperature sweeps instead. In addition, it is expedient to exclude heating rates obtained during the first laser periods.

One more important issue is the question of keeping $v_{\rm osc}/v_{\rm therm}$ constant during one simulation run. Although this is not the real physical picture, since the laser intensity will not follow the system temperature but will stay constant during the heating process, the available data and theoretical models for comparison are usually only available for $v_{\rm osc}/v_{\rm therm} = {\rm const.}$ This is a rather artificial convention compared to real experiments were the material is heating up while laser intensity is kept as constant as possible. However, for comparability we adhere to this convention. Therefore, the laser intensity in simulations with this method is adjusted after every laser cycle appropriately. Then, (7.9) does not have to modified since (7.16) is still valid for every individual cycle and with the instantaneous measurement of the curve's slope, the change of $E_{\rm L}$ between the cycles does not influence the derivative. For preserving validity of the theoretical foundations, we mitigate $v_{\text{osc}}/v_{\text{therm}} = \text{const}$, though, since intensity adjustment is not performed immediately after a temperature change but after the laser cycle is finished. In a consequent implementation, this should be done in every timestep, but then the derivative in (7.9) would also include contributions of change of the amplitude $E_{\rm L}$. Although no significant difference was seen in exploratory simulations for either method, an instantaneous intensity adjustment would spoil the sinusoidal shape of the external field's oscillation and hence make the applicability of the derivation in Section 7.1 at least questionable.

The final drawback with this method is the question of statistical fluctuations, making the measurement less precise and reliable than the long-term average from Section 7.3.1. Clearly, this is true. However, again assuming ergodicity of the system, the time-average can also be replaced by an ensemble-average, which can be assumed to be appropriate for sufficiently large particle numbers. This is precisely what is available using significantly improved tree code method introduced in Section 5.1. With it, fully kinetic simulations with in total up to two billion particles are possible.

Since these mentioned drawbacks of the method have not been thoroughly evaluated until now, both techniques – constant temperature and instantaneous heating rate measurement – are used in this work. The respective results are given in the following section and differences are discussed where appropriate.

7.4 Simulation results

In the first parts of this chapter, we discussed the connection between the collision frequency in warm dense matter and the plasma heating rate, described the simulation setup and workflow, and showed different methods for actually determining the plasma heating rate at defined sets of conditions ($n_{\rm el}, T_{\rm el}$). Together with the implementation of periodic boundary conditions of Chapter 4, these are the basis for the simulation results given in this section. We will concentrate the discussion of our results on three commonly studied dependencies. First, the scaling of the collision frequency with the coupling parameter Γ will be discussed in Section 7.4.1. Afterwards, we evaluate the dependence of $\nu_{\rm ei}$ on laser frequency in Section 7.4.2, and will finally approach scaling with the field amplitude, i. e. $v_{\rm osc}/v_{\rm therm}$, in Section 7.4.3.

7.4.1 From weak to strong correlations – scaling with the nonideality parameter

The scaling of $\nu_{\rm ei}/\omega_{\rm pl}$ as a function of the nonideality parameter Γ is shown in Figure 7.7. There, our simulation results are compared to data and analytical results from different methods published by several authors. Keeping in mind that $\Gamma \ll 1$ corresponds to the weak coupling regime of an ideal gas, here the results of the different theories and simulations have to be consistent. This is the case and can be interpreted as a very fundamental cross-check for validity of our simulations.

Error bars given in the Figure denote the 95% confidence region that was determined from the linear fit to $\mathcal{H}_N^{(e)} + \mathcal{H}_N^{(i)}$, as in Figure 7.5 for each simulation run, i.e. for every data point in the plot, separately. As we have long-running simulations covering up to several hundreds of laser cycles and large particle numbers of up to $N_{\rm el} = N_{\rm ion} = 2.5 \times 10^6$ electrons and ions, we obtain statistics at least an order of magnitude better than previous work to date. These statistical errors are not the only source of uncertainty, though. As we start our simulations with uncorrelated, randomly placed ions, they do not have a distinguished configuration as it would be the case for strongly coupled COULOMB crystals for example. Accordingly, our simulations correspond to essentially setting the ionic structure factor to unity as it is done in many theoretical approaches. To get an estimate of the influence of the actual ionic structure, several simulations for identical plasma parameters but different initial particle positions will have to be performed in future. In addition, a thorough study on the phase transition from the COULOMB liquid to a crystal is indicated. We will take this idea up again briefly in Section 8.2.

When approaching the regime of stronger coupling $\Gamma \gtrsim 1$, theories and simulation results start to differ. The analytic approach of Decker et al. is not applicable here

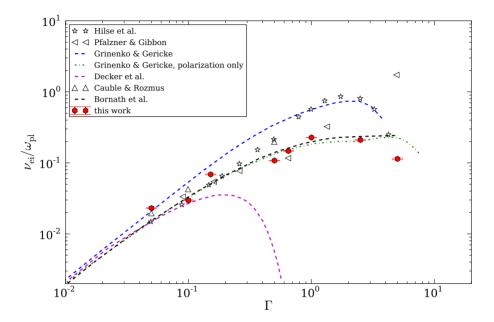


Figure 7.7: Normalized electron-ion collision frequency $\nu_{\rm ei}/\omega_{\rm pl}$ for a fully ionized hydrogen plasma with $n_{\rm el}=n_{\rm ion}=1\times 10^{22}\,{\rm cm}^{-3}$ in a weak harmonic external field $v_{\rm osc}/v_{\rm therm}=0.2$ with frequency $\omega/\omega_{\rm pl}=3.0$. Among our results, published data by several authors for the same plasma parameters is included. The data points correspond to simulations by HILSE et al. [103] and PFALZNER and GIBBON [104] as well as theoretical computations with a HNC structure factor by CAUBLE and ROZMUS [69]. Curves represent analytic results from GRINENKO and GERICKE [105], BORNATH et al. [101], and DECKER et al. [17]. In addition, the polarization contribution as derived in [105] is given separately. Our simulations were performed using a Nosé-Hoover thermostat as described in Section 7.3.1, error bars denote the 95% confidence region from the linear fit to $\mathcal{H}_N^{(e)} + \mathcal{H}_N^{(i)}$ as in Figure 7.5. Due to the good statistics from the simulations runs with large particle numbers up to 5×10^6 , these errors are of the order of 1.5%.

as the Coulomb logarithm (1.48) becomes negative rendering the theory invalid for $\Gamma \gtrsim 1$. Our simulation results are in good correspondence with the analytical results of BORNATH et al. [101] who evaluated the collision frequency based on a quantum kinetic approach. It is fundamentally very similar to the starting point of DECKER et al. but – as expected when going from classical to quantum mechanical expressions - does not need artificial integration cutoffs for the COULOMB logarithm (1.48). Instead, convergence if the k-space integral is ensured by additional exponential factors in the integrand that can be interpreted as quantum diffraction effects, see [101] for details. Similar values for ν_{ei}/ω_{pl} are also reported by GRINENKO and GERICKE [105] if only the polarization contribution is considered. It does not account for hard collision contributions $\nu_{\rm ei}^{\rm (hc)}$ that are an additive term to $\nu_{\rm ei}$ in their theory. These are dealt with an analogy to stopping power theories and yield the difference between the Grinenko curves in the figure. As [105] does not contain extensive details on their approach for $\nu_{\rm ei}^{\rm (hc)}$ but concentrates on the polarization contribution, an evaluation is difficult here. A hint might come from the data of earlier simulations of PFALZNER and Gibbon [104]. In contrast to our present simulations, they used a Plummer/soft core interaction, see (2.16) and Figure 7.1, instead of the quantum-mechanical Kelbg pseudopotential (7.26). Unfortunately, no information on their choice of the interaction cutoff parameter is available. Obviously, collisions are emphasized for $\Gamma \gtrsim 1$ in their simulations when compared to the present data. However, as the Kelbg potential is in general only valid for the weakly coupled case (see Section 7.2.2) it is only used as a regularization for the COULOMB potential with automatic choice of the cutoff parameter as a function of the thermal wavelength (7.28). Hence, it might also – due to considerable potential screening – underestimate collisional effects due to hard collisions. This interpretation matches the coincidence with the data of GRINENKO and GERICKE if hard collisions are omitted in their approach. Clearly, studies in the sensitivity of the strong-coupling dependence of $\nu_{\rm ei}$ on the choice of the potential and potential cutoff parameters would be necessary to clarify this point.

To finalize the discussion in Γ dependence, Figure 7.8 compares the results of the instantaneous heating rate measurement method, Section 7.3.2 to the data points obtained with the constant-temperature simulations via application of a Nosé-Hoover thermostat, Section 7.3.1. The results of the instantaneous measurement mutually coincide for the different simulation runs. However, there are deviation visible at the high-temperature/low- Γ endpoints of the individual curves. There, the collision frequency appears to start increasing again. This can be attributed to numerical heating due to an insufficient simulation timestep. The time discretization is adjusted according to the constraints described in Section 2.2.3 with the initial temperature in our simulations. While the respective limits stay valid during the full run of a constant-temperature simulation, they get violated as the system is heated in the instantaneous measurement scheme. Thus, for improving the low- Γ tails of the curves, an adaptive time stepping

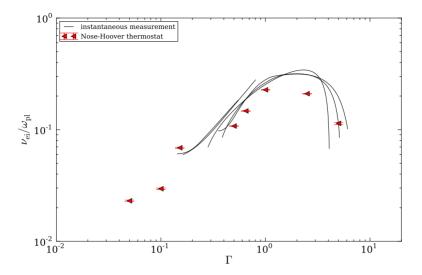


Figure 7.8: Comparison of ν_{ei}/ω_{pl} from instantaneous heating rate measurement as described in Section 7.3.2 with values obtained using the Nosé-Hoover thermostat approach from Section 7.3.1 for the same physical parameters as in Figure 7.7.

would be necessary, which is technically possible, see [217] and Section 3.4, but has not been implemented within this work.

With the exception of these high-temperature endpoints, the continuously measured curves agree with the constant-temperature data. This can primarily be attributed to the very good statistics due to the large particle numbers that renders the time averaging unnecessary. In fact, this allows one to sweep across a temperature range instead of the discrete sampling done until now with conventional molecular dynamics simulations. Thus, the large particle numbers available using the tree code approach open the range for more natural simulations that really follow the experimental interaction scenario and are not restricted to only probe the instantaneous state of an experiment for statistics purposes.

7.4.2 Dynamic behavior of the collision frequency

The dynamic behavior, i. e. frequency dependence of $\nu_{\rm ei}$ from our simulations is shown in Figures 7.9 and 7.10 for two different temperatures $T_{\rm el}=T_{\rm ion}=2.0\,{\rm eV}$ and $T_{\rm el}=T_{\rm ion}=7.7\,{\rm eV}$ and fixed densities $n_{\rm el}=n_{\rm ion}=1\times10^{22}\,{\rm cm}^{-3}$. It is compared to the real

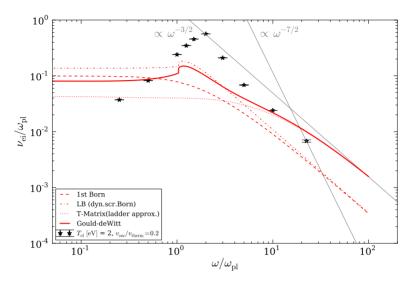


Figure 7.9: Dynamic behavior of the collision frequency for a fully ionized hydrogen plasma with $n_{\rm el} = n_{\rm ion} = 1 \times 10^{22} \, {\rm cm}^{-3}$, $T_{\rm el} = T_{\rm ion} = 2.0 \, {\rm eV}$, $v_{\rm osc}/v_{\rm therm} = 0.2$, $\Gamma = 2.5$, $\Theta = 1.18$, i. e. a moderately coupled, weakly degenerate system. Triangles represent simulation data, error bars denote the 95 % confidence interval from the linear fit for $\mathcal{H}_N^{(e)} + \mathcal{H}_N^{(i)}$. Curves give the theoretical behavior for moderately coupled systems with arbitrary degeneracy from a quantum statistical treatment in different approximations [81, 83, 89, 91, 94]. They straight gray line denotes the asymptote $\propto \omega^{-7/2}$ which is the expected high-frequency scaling of $\nu_{\rm ei}(\omega)$ with the KELBG potential [83] but is not observed here as the rightmost data point is still at a moderate frequency $\omega/\omega_{\rm pl} = 20$.

part of the collision frequency from a quantum statistical approach given by Reinholz, Wierling, Röpke et al. [81, 83, 89, 91, 94]. In Figure 7.9, the different contributions to their Gould-DeWitt scheme are shown separately. The method includes strong collisions via a statically screened T-matrix approach as well as the dynamically screened Born approximation (Lenard-Balescu collision term). With this technique, near and far collisions are treated consistently. However, it is limited to moderate coupling $\Gamma \lesssim 4$ [335] which is fulfilled for both parameter sets in this comparison. The Lippman-Schwinger equation for the T-matrices and integrals for the statically screened Born and dynamically screened Lenard-Balescu contributions were solved numerically with a program set by Millat [83] that was parallelized [84] and has been ported to the Blue Gene/Q system Juquen recently.

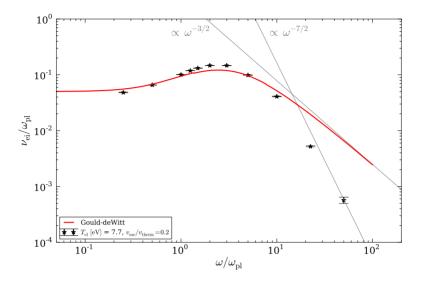


Figure 7.10: Collision frequency from simulations and GOULD-DEWITT technique for $n_{\rm el}=n_{\rm ion}=1\times 10^{22}\,{\rm cm^{-3}},~T_{\rm el}=T_{\rm ion}=7.7\,{\rm eV},~v_{\rm osc}/v_{\rm therm}=0.2,~\Gamma=0.65,~\Theta=4.55,~{\rm i.\,e.}~$ a weakly coupled, classical plasma. For details see Figure 7.9. As the available data points reach $\omega/\omega_{\rm pl}=50$, the high-frequency asymptote is reached in this case.

While Figure 7.10 shows very good agreement of simulation and theory for $\omega/\omega_{\rm pl} \lesssim 20$, results in Figure 7.9 differ significantly. This difference can be attributed to quantum effects that arise as Θ approaches unity. They are included in the theory but are only dealt with via a pseudopotential approach in the simulations. This results in different treatment of potential screening and near-field cutoff and apparently overestimates collisions in this parameter range. Comparably enhanced collision frequencies in comparisons of simulations to this theory were also found for different plasma parameters in other studies [91, 94].

The different behavior of theory and simulation at high frequencies is another artifact of the choice of interaction potential. As shown by MILLAT [83], utilization of the KELBG potential leads to an asymptotic behavior

$$\nu_{\rm ei}(\omega) \propto \omega^{-7/2} \quad \text{for} \quad \omega \to \infty .$$
 (7.67)

For the standard COULOMB potential,

$$\nu_{\rm ei}(\omega) \propto \omega^{-3/2} \quad \text{for} \quad \omega \to \infty$$
 (7.68)

is expected and also reproduced by the analytic expressions [83, 91, 92, 94]. For comparison, an asymptote $\propto \omega^{-7/2}$ is included in Figures 7.9 and 7.10 and well reproduced in the latter where the simulated frequencies reach up to $\omega/\omega_{\rm pl} = 50$.

7.4.3 Dependence on field amplitude

In addition to correlation effects due to strong coupling, i.e. high densities and low temperatures, we are also interested in the field amplitude dependence of the collision frequency. The high-field behavior of optical and transport properties in warm dense matter is is of particular interest today as with the construction and operation of free electron lasers worldwide, this regime becomes accessible for experiments. As seen in Figure 1.2, today's installations are already reaching the range of $v_{\rm osc}/v_{\rm therm} \gtrsim 1$ for high frequencies $\omega/\omega_{\rm pl} \gg 1$ at bulk densities and are in fact able to actively probe the warm dense matter regime. A fundamental theoretical approach including nonlinearity effects due to $v_{\rm osc}/v_{\rm therm} \gg 1$ is given by Silin [58]

$$\nu_{\text{ei}}^{\text{(SILIN)}} = 32\sqrt{2\pi} \frac{n_{\text{el}} q_{\text{el}} q_{\text{ion}}}{m_{\text{el}}^2 v_{\text{therm}}^3} \left(\frac{v_{\text{osc}}}{v_{\text{therm}}}\right)^{-3} \cdot Q\left(\frac{v_{\text{osc}}}{2v_{\text{therm}}}\right) \cdot \mathcal{L}$$
(7.69)

with

$$Q(r) := \int_0^\infty z^2 e^{-z^2} \left[I_0(z^2) - I_1(z^2) \right] dz$$
 (7.70)

where $I_n(z)$ is the modified BESSEL function of first kind [163] and \mathcal{L} is the COULOMB logarithm (1.48). In the low-field limit, this result corresponds to the well-known DAW-SON-OBERMAN collision frequency [56], but for strong fields decreases as

$$\nu_{\rm ei} \left(\frac{v_{\rm osc}}{v_{\rm therm}}\right) \propto \left(\frac{v_{\rm osc}}{v_{\rm therm}}\right)^{-3} \quad \text{for} \quad \frac{v_{\rm osc}}{v_{\rm therm}} \to \infty .$$
 (7.71)

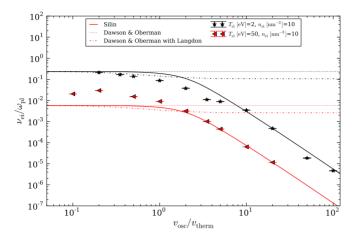
This scaling is clearly reproduced by our simulations, as shown in Figure 7.11a. There, the expressions by DAWSON and OBERMAN and SILIN are plotted in comparison to our simulation results for two parameter sets with nonideality $\Gamma = 0.1$ and $\Gamma = 2.5$, respectively. As the COULOMB logarithm and thus (7.69) becomes invalid when approaching $\Gamma \gtrsim 1$, the analytical curves haven been shifted to overlap with the data in the high-field limit by scaling \mathcal{L} . Accordingly, we have to restrict our discussion to qualitative features such as the curves' overall shape instead of absolute values.

However, we can also interpret our shift factors for the collision frequencies in the simulations at $n_{\rm ion} = n_{\rm el} = 1 \times 10^{22} \, \rm cm^{-3}$ as a scaling that corrects the COULOMB logarithm which would give too low values for the collision frequencies otherwise. The respective factors for $\nu_{\rm ei}^{\rm (Silin)}$ from (7.69) are

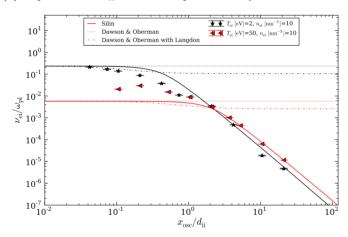
$$\nu_{\rm ei}^{\rm (simulation)} = 12.117 \cdot \nu_{\rm ei}^{\rm (SiLIN)} \quad \text{at} \quad T_{\rm el} = 2 \, {\rm eV} ,$$

$$\nu_{\rm ei}^{\rm (simulation)} = 1.667 \cdot \nu_{\rm ei}^{\rm (SiLIN)} \quad \text{at} \quad T_{\rm el} = 50 \, {\rm eV}$$
(7.72)

$$\nu_{\rm ei}^{\rm (simulation)} = 1.667 \cdot \nu_{\rm ei}^{\rm (Silin)}$$
 at $T_{\rm el} = 50 \,\mathrm{eV}$ (7.73)



(a) Dependence of $\nu_{\rm ei}$ on electron quiver velocity.



(b) Dependence of ν_{ei} on electron quiver amplitude.

Figure 7.11: Collision frequency for a fully ionized hydrogen plasma with $n_{\rm el}=n_{\rm ion}=1\times 10^{22}\,{\rm cm^{-3}}$ and temperatures $T_{\rm el}=2.0\,{\rm eV}$ ($\Gamma=2.5$) and $T_{\rm el}=50.0\,{\rm eV}$ ($\Gamma=0.65$) at $^{\omega}/_{\omega_{\rm pl}}=3.0$ in dependence of quiver velocity and quiver amplitude, respectively. Lines are theoretical curves by Dawson and Oberman [56] (dotted), additional Langdon correction [64] (dashed) and Silin [58] (solid).

and thus renormalize the COULOMB logarithm for the selected set of parameters.

The expected asymptotic behavior in the high-field limit is observed for both parameter sets. Compliance between theory and simulation in the low-field limit is only found for $T_{\rm el}=2.0\,{\rm eV}$, though, which is surprising as the colder plasma corresponds to the stronger coupled regime. However, also in this case there is a significant deviation from the SILIN results towards lower values around $v_{\rm osc}/v_{\rm therm}\approx 2$. In contrast, an enhanced collision frequency is observed for $v_{\rm osc}/v_{\rm therm}\lesssim 2$ in the hotter system $T_{\rm el}=50.0\,{\rm eV}$.

In [104], a reduced collision frequency found around $v_{\rm osc}/v_{\rm therm} \approx 2$ was attributed to the Langdon effect – deformation of the Maxwellian distribution function due to insufficient electron-electron relaxation [64]. This correction, which is deduced from a fit to numerical solution of the Boltzman equation with non-Maxwellian velocity distribution, is a multiplicative factor of the form

$$1 - \frac{0.553}{1 + \left(\frac{0.27}{\alpha}\right)^{0.75}} \quad \text{with} \quad \alpha = Z \left(\frac{v_{\text{osc}}}{v_{\text{therm}}}\right)^2. \tag{7.74}$$

It is also included in Figure 7.11a as a dashed-dotted line that makes the DAWSON-OBERMAN result a function of $v_{\rm osc}/v_{\rm therm}$. However, it does neither explain slope nor value of any of the observed deviations.

A hint to a probable explanation is given in Figure 7.11b that shows the data from Figure 7.11a over a rescaled axis. The abscissa is scaled in x_{osc}/d_{ii} , i.e. electron quiver amplitude over average ion distance, compare Section 1.1 for these parameters. Here, the unexpected features for both parameter sets appear for the same range $0.1 \le x_{\text{osc}}/d_{\text{ii}} \le 1$. If $x_{osc}/d_{ii} > 1$, electrons pass by more than one ion during one quiver period on average and hence suffer from multiple stochastic collisions. In contrast, for $x_{\rm osc}/d_{\rm fi}$ < 1, they oscillate in front of a single ion and are scattered several times at the same center. Thus, they experience correlated collisions in the picture of Decker et al. [17] where such an effect was proposed at scales of $x_{osc}/\lambda_D < 1$. In their interpretation, the effect will lead to enhanced collision frequencies as we observe for $T_{\rm el} = 50.0\,{\rm eV}$ but could not be matched to their own simulation data. As we observe a reduced collision frequency for $T_{\rm el} = 2.0\,{\rm eV}$ that appears to be contradictory, additional effects have to play a role here which is perhaps a hint towards a much more complex picture. To further analyze these features, they first have to be verified by going to lower field amplitudes, where the DAWSON-OBERMAN result has to be reached in both cases for consistency. However, as heating rates are low for weak fields, long-running simulations will be necessary that in addition make use of the good statistics offered by large particle numbers to obtain reliable results.

7.5 Summary

Collisions in a plasma lead to a conversion of the collective oscillatory electronic motion to undirected, random dynamics eventually defining the system's temperature. In the present chapter, this connection which is eponymous for the term *collisional heating* was utilized to derive the electron-ion collision frequency from simulations of a bulk hydrogen plasma across two orders of magnitude of the coupling parameter Γ .

For a weakly coupled, i.e. ideal plasma, good coincidence with the theoretically expected scaling was found. At stronger coupling, we measured significant deviation from data published on earlier simulations and different theories. Agreement of our simulations with analytical results by Bornath et al. [101] and Grinenko and Gericke [105] was found in this region. As the primary cause for deviations from and between the theories and simulations in this comparison, we identified the very different treatment of cutoffs for close encounters – classically with a smoothed potential or quantum mechanically by means of pseudopotentials or with full inclusion of degeneracy in theory. These findings motivate thorough investigations on the collision frequency's strong-correlation limit in future. In particular, its the dependence of the choice of interaction potential as well as the introduced cutoff parameters is of interest when it comes to comparing the different approaches for the collision frequency in strongly correlated Coulomb systems.

Our comparisons on frequency dependence of the collision frequency with analytic expressions from a quantum statistical approach showed good agreement. Deviations in the high-frequency limit were expected – again due to the choice of the Kelbg pseudopotential – and were consistently found in our simulations.

In contrast, the scaling of $\nu_{\rm ei}$ with quiver velocity and amplitude, i. e. with laser intensity, revealed unexpected features around $0.1 \lesssim v_{\rm osc}/v_{\rm therm} \lesssim 1.0$. We found a reduced collision frequency for the colder and enhanced collision frequency for the hotter system when comparing to well-established theoretical results. These could not be explained with modifications of the Maxwellian electron velocity distribution, i. e. the Langbon effect, as it was the case in [104]. Accordingly, other effects will have to be considered. In particular – becoming evident as the electron excursion reaches the inter-ionic separation – correlated collisions as described by Decker et al. are potential candidates for causing the deviations observed. Again, additional studies in the parameter region just below $v_{\rm osc}/v_{\rm therm} \lesssim 1$ are motivated before giving a final evaluation here.

8 Wrapping things up: a conclusion

" Wie es auch die größte Freude macht, in solchem Zeitpunkt des intensivsten Aufschwungs der eigenen wissenschaftlichen Disziplin zu leben, wo es eine Überfülle von Aufgaben gibt, die sich der Bearbeitung darbieten, [...so muß] jeder, der an dieser Entwicklung auch nur zu einem kleinen Teil helfen will, [...] sich die ernste Frage vorlegen, an welcher Ecke des Gebäudes er selbst Hand anlegen kann, falls es nutzbringend zu werden versprechen soll.

€ [1]

As it gives great pleasure to live in such a time of an intense rise of the personal scientific discipline, where there are a plethora of tasks presenting themselves to be dealt with, anyone who wants to help in this development only to a small extent [...] has to seriously ask himself, to which corner of the building he wants to contribute, if it shall promise to become useful. •• [1], translated from the German original.

8.1 Summary

As this thesis opens with a motivating quotation by PAUL DRUDE who contributed some significant share to the foundation of the theories used here, we also conclude with his words. Within this thesis, four corners of the contorted scientific building of plasma physics and its annex molecular dynamics simulations were actually approached.

We started from a parallel implementation of the Barnes-Hut tree code and extended it to periodic boundary conditions and scalability for extremely large particle numbers and hundred thousands of processors. Both developments are not mutually dependent but complementary. Periodic boundary conditions primarily eliminate boundary and surface effects which are introduced artificially in simulations due to limited system size. This is desirable for simulations on properties of bulk matter which are studied in simulations at small extracts of the full system. In contrast, large particle numbers are a necessary prerequisite for dealing with large systems that in fact have open boundaries also in reality. Here, a virtual expansion of system size by means of periodicity is not applicable as the boundary effects to be studied would be removed. The implementation of both features is a milestone in the development of PEPC, our simulation code. Thanks to its modular concept that also has been developed as by-product of

this work, not only the strongly-coupled-plasma part of the simulation suite profits from these advances but all other frontend applications, too.

However, despite interesting developments contributed to the field of numerical algorithms, such a simulation program is only a tool. Its creation and improvement does not end in itself but is the foundation for dealing with real physical problems and studying exciting effects in simulations. In this work, we used our highly scalable parallel tree code for two applications that in fact required the developments made.

First, we studied collective phenomena in the electronic subsystem in confined nanoplasmas. Here, we proposed a new diagnostic technique to track resonance phenomena with spatial resolution. Higher oscillation modes and non-isotropic excitation patterns were found. Being able to simulate with respect to particle number 300 times larger systems than in previous studies, we used our code's potential to bridging from nano to bulk scale. In this transition we have seen that the bulk properties gradually evolve from the small clusters' discrete two-peak spectra via a number of additional resonances towards the dense spectrum of bulk matter where due to infinite system size and only limited by crystal structure virtually all oscillation modes are allowed.

Being also interested in optical and transport properties, in the second physical application of this thesis we dealt with the collision frequency in bulk matter. In the DRUDE model it is the key parameter to experimental observables such as opacity, conductivity, reflectivity to name but a few. This parameter essentially describes the friction experienced by electrons in matter and thus determines their dynamical behavior. We could show that our simulations yield results consistent with different theories and are able to further extend the data basis for collision frequencies in bulk matter. Again, the good statistics that we observed is due to our ability to study large systems over long timescales and hence is geared to the algorithm development done before. However, we also identified an anomalous, hitherto unexplored feature of collisional absorption in the strongly-coupled, high-field regime, namely enhanced and depleted collision rates. Furthermore, we identified a strong dependence of $\nu_{\rm ei}$ on the interaction potential and its parameters.

Thus, coming back to the introductory words, the simulation experiments conducted in this work did not only contribute to new scientific insight but also posed a number of new questions to be addressed in future. Accordingly, our simulations are indeed part of a third scientific corner stone next to theory and experiment. As such, progress is not to be stopped but research will go on on the application as well as the algorithm side. Some prospective tracks will be listed in the next and final section of this work.

8.2 Outlook

With essentially four different topics – two of them algorithm related and two topical physical applications – this work covers a broad range of specific questions. So will also do the overview on possible future prospects that we give here in loose order to conclude this work.

Faster computation with multiple multipoles As seen in Figure 2.2, from the algorithmic point of view the tree code is in a central position between the direct N^2 summation and the Fast Multipole Method for evaluating the self-consistent potential and forces in large sets of particles. Scaling as $\mathcal{O}(N)$ in the number of particles N, the latter outperforms the classical Barnes-Hut tree code if processed particles per second are used as a measure. However, with the cell-cell interactions proposed by Dehnen [198, 202, 203] and his adaptive multipole acceptance criterion, linear scaling is also reported for a tree code. With inclusion of such an approach, we are expecting to considerably reduce time-to-solution as less interactions have to be evaluated. As an intermediate step, we are taking the reuse of interaction lists for nearby particles [204, 205] into account. With these modifications, the tree traversal time will reduce dramatically and other parts of the code could start dominating. These will also have to be reformulated for hybrid parallelization. In particular, for the parallel tree construction this will be a challenging task that will finally lead to even better utilization of the processors available.

Unbounding the boundaries The approach for periodic boundary conditions implemented in this work is based on a spherical harmonics expansion of the Coulomb potential. This expansion relies on the fact, that such an expansion is possible because the potential is the Green's function of the Laplace operator in three dimensions in spherical coordinates. Accordingly, it is not applicable to the class of generalized algebraic kernels introduced in [179] for utilization in the vortex method with our tree code PEPC. However, we expect that an extension to these potentials based on analytic work by Cohl [336, 337] on generalization of spherical expansions for these kernels is possible. This would allow for closed-system vortex fluid simulations that were restricted to open boundaries in [179] until now.

Furthermore, as it is in general possible to apply the formulation for periodic boundaries given in Chapter 4 to non-cubic and even triclinic cells, this is not the case for the indispensable extrinsic to intrinsic correction that in the formulation of Section 4.2.7 is only valid for cubic cells and periodicity in three dimensions. Utilization of the compensating charge approach by Kudin and Scuseria [259, 260] will mitigate this issue. Then, the approach can also be applied rigorously for example in simulations

on plasma wall interaction [174, 275] where an elongated simulation cell is periodically extended in two spatial directions to model the plasma sheath in front of the wall.

De-serializing time Parallel-in-time integration as a very promising approach for making even better use of available computer resources and reducing time-to-solution has been discussed briefly in Section 5.3.4. It will allow for studying very different time scales such as electronic and ionic motion in consistent simulations without a compromise concerning resolution. Such an approach can provide a solution to the most demanding problem of molecular dynamics when performed on the level of individual electrons and ions. Due to the mass ratio of $m_{\rm p} \approx 1,836 m_{\rm el}$, the dynamic's time scales vastly differ. Being obstructive in conventional molecular dynamics simulations, this will even be advantageous for the efficiency of the Parallel-in-time method. Thus, it could provide a consistent picture of electronic and ionic relaxation from a single simulation.

Correlation and ionic structure In our simulations on bulk collision frequency, we used a relaxation step from random particle positions to an organized structure that corresponds to the required Γ value. For even nonideality parameters, i.e. higher densities and lower temperatures, crystallization will occur that changes the material's properties dramatically. Unfortunately, with low temperatures this process is rather slow. Studying it in more detail by making use of the extremely good statistics offered by large particle numbers, possibly accelerated by means of the parallel-in-time method, will provide information on this phase transition that is not completely understood up to now. Going towards these higher densities, degeneracy effects will start to dominate. Up to a certain extent, they can for example be dealt with a path-integral molecular dynamics approach [129, 338]. In its present formulation, it can be implemented on top of a classical Coulomb/Kelbg potential based molecular dynamics method such as PEPC.

Towards experiments With its capability of simulating systems with billions of particles in a fully kinetic scheme with self-consistent electrostatics, spatial scales as used in real-world experiments become reachable. Thus, performing simulations with counterparts in laboratory experiment are desirable to make use of PEPC's predictive capabilities. In particular, experiments on transport in nano wires or optical features of nano structures could be the first of such applications. Although we concentrated on optical properties in this work, these are neither limits of the underlying theory nor constraints of simulation capabilities. In fact it is in principle also possible to study transport properties (thermal and electrical conductivity), nonequilibrium and relaxation processes, boundary effects, influence of external magnetic fields, etc. using the

molecular dynamics method. Naturally, fully microscopic *ab initio* simulations of experiments such as the aluminum example [2] that was mentioned in Chapter 1 are still far. For example, a technique for dealing with ionization processes will clearly be necessary at such high intensities. However, with today's computational capabilities and efficient methods such as the Barnes-Hut tree code and its incarnation PEPC chances have never been better to connect results from simulations on digital computers to the analogous world outside.

From collisions back to the observable world As we have seen in Sections 1.4 and 7.4, there are numerous theories on bulk collision frequency. Proposing values that differ by almost an order of magnitude in the strong-field regime, an evaluation of their predictive capabilities is difficult. This is not a result of bad but just of different theories. They include or omit different mechanisms which all might be relevant in the parameter space considered. Hence, a direct comparison might be pointless. For example, as recently shown [85] the collision frequencies from the quantum kinetic approach and linear response theory actually do not represent the same theoretical object. However, they can be related to each other either by tracing them to their theoretical origin as done in [85] or on the level of well-defined physical observables such as the dielectric function [86].

As stated in Chapter 1, theory, experiment, and simulation are inseparably geared to each other. Consequently, they should be compared on the same parametric level. Maybe, the return to experimental observables already known at DRUDE's time will provide the opportunity to for a better unification of these three scientific corner stones in the field of warm dense matter.

Acknowledgments

Being a multidisciplinary study that covers aspects from physics as well as from computer science, this work is an attempt of putting the hand to more than one corner of the scientific house. This would not have been possible without the faith and support of numerous people.

First, I would like to thank Prof. Paul Gibbon for his supervision. He guided me along the borderline between physics and computer science and admitted enough freedom to me for exploring both fields. In addition, I would like to express my sincere gratitude to Prof. Hans-Jörg Kull for his hospitality and advice as well as for the opportunity to work on this thesis as a somehow external part of the research group Laser Physics at RWTH Aachen. Furthermore – representing the Jülich Supercomputing Centre Staff – I would like to thank Prof. Thomas Lippert. Being considerably more than just the place of my desk, JSC is an inspiring location with great possibilities for doing research in fundamental as well as in computer sciences.

Only few parts of the scientific building are as important as its foundation. For me, it has been laid at Rostock University with Prof. Gerd Röpke and PD Heidi Reinholz whom I would like to thank for their guidance and advice towards and also during my time as a Ph.D. student. In addition I am grateful to Thomas Raitza and Igor Morozov who shared many details on their simulation results.

Despite the haste of contemporary science, there has always been some time for not only working at this thesis but also discussing problems that cover a whole bunch of different fields from hydrodynamics across root-soil interaction and astrophysics towards a myriad of non-scientific topics. All these conversations eminently enriched the last three years by offering views and thoughts outside the box of plasma physics and tree codes. For this great time I would like to thank Robert Speck, Natalie Schröder, and Ivo Kabadshow who – first being colleagues – became very good friends. This is also true for Benedikt Steinbusch, Dirk Brömmel and too many others to name all of them.

Finally I would like to thank Tatjana as well as my parents and grandparents. Devoting their love, their strong faith, and their absolute support they gave me strength and patience throughout this project that never would have been completed without them.

Bibliography

- [1] **P. Drude**. "Antrittsrede von Paul Drude im Plenum der Preußischen Akademie der Wissenschaften (Sitzungsberichte der Akademie der Wissenschaften zu Berlin 1906)." In: *Zur Elektronentheorie der Metalle*. Ostwalds Klassiker der exakten Wissenschaften. Frankfurt am Main: Deutsch Verlag, 2006, 56–58.
- [2] **B. Nagler et al.** "Turning solid aluminium transparent by intense soft X-ray photoionization." *Nature Physics* 5 (2009), 693–696.
- [3] I. Langmuir. "Oscillations in Ionized Gases." Proceedings of the National Academy of Science 14 (1928), 627–637.
- [4] W.-D. Kraeft, D. Kremp, W. Ebeling, and G. Röpke. Quantum Statistics of Charged Particle Systems. Akademie-Verlag, Berlin, 1986.
- [5] **H. R. Griem**. "High-density corrections in plasma spectroscopy." *Physical Review* 128 (1962), 997–1003.
- [6] W. Lotz. "Electron-impact ionization cross-sections and ionization rate coefficients for atoms and ions from hydrogen to calcium." Zeitschrift für Physik A Hadrons and Nuclei 216 (1968), 241–247.
- [7] R. Redmer. "Electrical conductivity of dense metal plasmas." Physical Review E 59 (1999), 1073–1081.
- [8] M. R. Zaghloul, M. A. Bourham, and J. M. Doster. "A simple formulation and solution strategy of the Saha equation for ideal and nonideal plasmas." *Journal of Physics D: Applied Physics* 33 (2000), 977–984.
- [9] L. D. Landau and E. M. Lifshitz. Statistical Physics. 3rd ed. Vol. 5. Course of Theoretical Physics. Butterworth-Heinemann, 1996.
- [10] C. Kittel. Introduction to Solid State Physics. 8th. John Wiley and Sons, Inc, 2005.
- [11] E. M. Lifshitz and L. P. Pitaevskii. Physical Kinetics. Reprint. Vol. 10. Course of Theoretical Physics. Butterworth-Heinemann, 1981.
- [12] H.-J. Kull. Plasmaphysik. Lecture Notes, RWTH Aachen University. 2007.
- [13] R. Redmer. Einführung in die Plasmaphysik. Lecture Notes, Rostock University, Institut für Physik. 2008.

- [14] L. Tonks and I. Langmuir. "Oscillations in ionized gases." *Physical Review* 33 (1929), 195–210.
- [15] D. Bohm and E. Gross. "Theory of Plasma Oscillations. A. Origin of Medium-Like Behavior." Physical Review 75 (1949), 1851–1864.
- [16] P. Gibbon. Short Pulse Laser Interactions with Matter. London: Imperial College Press, 2005.
- [17] C. D. Decker, W. B. Mori, J. M. Dawson, and T. Katsouleas. "Nonlinear collisional absorption in laser-driven plasmas." *Physics of Plasmas* 1 (1994), 4043.
- [18] **W. Ackermann et al.** "Operation of a free-electron laser from the extreme ultraviolet to the water window." *Nature Photonics* 1 (2007), 336–342.
- [19] DESY Free Electron Laser Flash. http://flash.desy.de/ [online, accessed February 2013].
- [20] Thomas Jefferson Lab Free Electron Laser Program. http://wwwold.jlab.org/FEL/felspecs.html [online, accessed February 2013].
- [21] **P. Emma et al.** "First lasing and operation of an ångstrom-wavelength free-electron laser." *Nature Photonics* 4 (2010), 641–647.
- [22] **B. Rudek et al.** "Ultra-efficient ionization of heavy atoms by intense X-ray free-electron laser pulses." *Nature Photonics* 6 (2012), 858–865.
- [23] D. Pile. "X-rays: First light from SACLA." Nature Photonics 5 (2011), 456–457.
- [24] **T. Ishikawa et al.** "A compact X-ray free-electron laser emitting in the sub-angström region." *Nature Photonics* 6 (2012), 540–544.
- [25] E. Allaria, C. Callegari, D. Cocco, W. M. Fawley, M. Kiskinova, C. Masciovecchio, and F. Parmigiani. "The FERMI@Elettra free-electron-laser source for coherent x-ray physics: photon properties, beam transport system and applications." New Journal of Physics 12 (2010), 075002.
- [26] E. Allaria et al. "Highly coherent and stable pulses from the FERMI seeded free-electron laser in the extreme ultraviolet." Nature Photonics 6 (2012), 699– 704.
- [27] FERMI @ elletra, Free Electron Laser of Multidisciplinary Investigations. http://www.elettra.trieste.it/FERMI/index.php?n=Main.Parameter [online, accessed February 2013].
- [28] C. Jian-Hui, D. Hai-Xiao, G. Qiang, L. Dong-Guo, W. Dong, Z. Meng, and Z. Zhen-Tang. "Operating the SDUV-FEL with the echo-enabled harmonic generation scheme." *Chinese Physics C* 33 (2009), 706–710.
- [29] **Z. T. Zhao et al.** "First lasing of an echo-enabled harmonic generation free-electron laser." *Nature Photonics* 6 (2012), 360–363.

- [30] B. W. J. McNeil and N. R. Thompson. "X-ray free-electron lasers." Nature Photonics 4 (2010), 814–821.
- [31] P. R. Ribic and G. Margaritondo. "Status and prospects of x-ray free-electron lasers (X-FELs): a simple presentation." *Journal of Physics D: Applied Physics* 45 (2012), 213001.
- [32] Petawatt Hoch- Energie Laser für SchwerIoneneXperimente (PHELIX). http://www.gsi.de/phelix [online, acessed February 2013].
- [33] **J. C. Maxwell**. "A Dynamical Theory of the Electromagnetic Field." *Philosophical Transactions of the Royal Society* 155 (1865), 459–512.
- [34] L. D. Landau and E. M. Lifshitz. Classical Theory of Fields. Revised. Vol. 2. Course of Theoretical Physics. Butterworth-Heinemann, 1987.
- [35] J. D. Jackson. Classical Electrodynamics. 3rd ed. New York: Wiley, 1998.
- [36] E. M. Lifshitz, L. D. Landau, and L. P. Pitaevskii. Electrodynamics of Continuous Media. 2nd. Vol. 8. Course of Theoretical Physics. Butterworth-Heinemann. 1984.
- [37] **J. J. Thomson**. "XXIV. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure." *Philosophical Magazine Series* 6 7 (1904), 237–265.
- [38] J. J. Thomson. "XL. Cathode Rays." Philosophical Magazine Series 5 44 (1897), 293–316.
- [39] **H. A. Lorentz**. "The motion of electrons in metallic bodies I." *KNAW*, *Proceedings* 7 (1905), 438–453.
- [40] **H. A. Lorentz**. "The motion of electrons in metallic bodies II." KNAW, Proceedings 7 (1905), 585–593.
- [41] **H. A. Lorentz**. "The motion of electrons in metallic bodies III." KNAW, Proceedings 7 (1905), 684–691.
- [42] **H.-J. Kull**. Laserphysik Physikalische Grundlagen des Laserlichts und seiner Wechselwirkung mit Materie. Oldenbourg Wissenschaftsverlag, 2010.
- [43] P. Drude. "Zur Elektronentheorie der Metalle." Annalen der Physik 306 (1900), 566–613.
- [44] **P. Drude**. "Zur Elektronentheorie der Metalle; II. Teil. Galvanomagnetische und thermomagnetische Effecte." *Annalen der Physik* 308 (1900), 369–402.
- [45] P. Drude. "Zur Ionentheorie der Metalle." Physikalische Zeitschrift 1 (1900), 161–165.

- [46] **G. Wiedemann and R. Franz**. "Ueber die Wärme-Leitungsfähigkeit der Metalle." *Annalen der Physik und Chemie* 165 (1853), 497–531.
- [47] A. Sommerfeld and H. Bethe. "Elektronentheorie der Metalle." In: Handbuch der Physik. Vol. 24-2. Springer Verlag, Heidelberg, 1933, 333–622.
- [48] L. Spitzer Jr. Physics of Fully Ionized Gases. 2nd. Interscience Tracts in Physics and Astronomy. J. Wiley & Sons, New York and London, 1962.
- [49] L. D. Landau. "Die kinetische Gleichung für den Fall Coulombscher Wechselwirkung." *Phys. Z. Sowjetunion* 10 (1936), 154–164.
- [50] L. Spitzer Jr and R. Härm. "Transport phenomena in a completely ionized gas." Physical Review 89 (1953), 977–981.
- [51] W. L. Kruer. The Physics of Laser Plasma Interaction. Frontiers in Physics. Addison Wesley Publishing Company, 1988.
- [52] **A. Lenard**. "On Bogoliubov's kinetic equation for a spatially homogeneous plasma." *Annals of Physics* 10 (1960), 390–400.
- [53] R. Balescu. "Irreversible Processes in Ionized Gases." Physics of Fluids 3 (1960), 52.
- [54] H. A. Gould and H. E. DeWitt. "Convergent kinetic equation for a classical plasma." *Physical Review* 155 (1967), 68–74.
- [55] V. I. Perel and G. M. Eliashberg. "Absorption of electromagnetic waves in a plasma." *Soviet Physics JETP* 14 (1962), 633–637.
- [56] J. M. Dawson and C. Oberman. "High-Frequency Conductivity and the Emission and Absorption Coefficients of a Fully Ionized Plasma." *Physics of Fluids* 5 (1962), 517–524.
- [57] S. Rand. "Inverse Bremsstrahlung with High-Intensity Radiation Fields." Physical Review 136 (1964), B231–B237.
- [58] V. P. Silin. "Nonlinear high-frequency plasma conductivity." Soviet Physics JETP 20 (1965), 1510–1516.
- [59] P. J. Catto and T. Speziale. "Strong field inverse bremsstrahlung via a Lorentz model." Physics of Fluids 20 (1977), 167.
- [60] G. J. Pert. "Inverse bremsstrahlung absorption in strong radiation fields during binary collisions: 'straight-line path' approximation." Journal of Physics B: Atomic and Molecular Physics 12 (1979), 2755–2769.
- [61] L. Schlessinger and J. Wright. "Inverse-bremsstrahlung absorption rate in an intense laser field." *Physical Review A* 20 (1979), 1934–1945.
- [62] V. P. Silin and S. A. Uryupin. "Absorption of intense electromagnetic radiation in collisions of charged particles." Soviet Physics JETP 54 (1981), 485–492.

- [63] A. Brantov, W. Rozmus, R. Sydora, C. E. Capjack, V. Y. Bychenkov, and V. T. Tikhonchuk. "Enhanced inverse bremsstrahlung heating rates in a strong laser field." *Physics of Plasmas* 10 (2003), 3385.
- [64] A. Langdon. "Nonlinear inverse bremsstrahlung and heated-electron distributions." Physical Review Letters 44 (1980), 575–579.
- [65] R. D. Jones and K. Lee. "Kinetic theory, transport, and hydrodynamics of a high-Z plasma in the presence of an intense laser field." *Physics of Fluids* 25 (1982), 2307.
- [66] S. Pfalzner. "Influence of strong laser fields on the inverse bremsstrahlung collision frequency." Applied Physics B Photophysics and Laser Chemistry 55 (1992), 368–372.
- [67] P. Mulser and A. Saemann. "Dielectric Theory of Collisional Absorption in Intense Laser Fields." Contributions to Plasma Physics 37 (1997), 211–227.
- [68] P. Mulser. "Standard Models of Collisional Absorption and their Equivalence." Contributions to Plasma Physics 41 (2001), 243–246.
- [69] R. Cauble and W. Rozmus. "The inverse bremsstrahlung absorption coefficient in collisional plasmas." Physics of Fluids 28 (1985), 3387.
- [70] **J. S. Rowlinson**. "The equation of state of dense systems." Reports on Progress in Physics 28 (1965), 169–199.
- [71] J. P. Hansen and I. R. McDonald. Theory of simple liquids. 3rd ed. Academic Press, 2006.
- [72] **G. Röpke**. "Dielectric function and electrical dc conductivity of nonideal plasmas." *Physical Review E* 57 (1998), 4673–4683.
- [73] **G. Röpke and A. Wierling**. "Dielectric function of a two-component plasma including collisions." *Physical Review E* 57 (1998), 7075–7085.
- [74] G. Röpke, A. Selchow, A. Wierling, and H. Reinholz. "Lindhard dielectric function in the relaxation-time approximation and generalized linear response theory." *Physics Letters A* (1999), 365–369.
- [75] G. Röpke, R. Redmer, A. Wierling, and H. Reinholz. "Response function including collisions for an interacting fermion gas." *Physical Review E* 60 (1999), R2484.
- [76] A. Wierling, S. Sorge, and G. Röpke. "Using the Dielectric Function for Diagnostics of Laser-Induced Plasmas." Contributions to Plasma Physics 39 (1999), 73–76.
- [77] H. Reinholz, R. Redmer, G. Röpke, and A. Wierling. "Long-wavelength limit of the dynamical local-field factor and dynamical conductivity of a two-component plasma." *Physical Review E* 62 (2000), 5648.

- [78] **G. Röpke**. "Quantum-statistical approach to the electrical conductivity of dense, high-temperature plasmas." *Physical Review A* 38 (1988), 3001–3016.
- [79] A. Esser and G. Röpke. "Debye-Onsager relaxation effect in fully ionized plasmas." *Physical Review E* 58 (1998), 2446.
- [80] A. Esser, R. Redmer, and G. Röpke. "Interpolation formula for the electrical conductivity of nonideal plasmas." *Contributions to Plasma Physics* 43 (2003), 33–38.
- [81] H. Reinholz. "Dielectric and optical properties of dense plasmas." Annales de Physique 30 (2005), 1–187.
- [82] T. Millat. "Dielektrische Funktion stark gekoppelter Plasmen." Diploma thesis. Rostock University, 1999.
- [83] **T. Millat**. "Absorption und dynamische Stoßfrequenz in nichtidealen Plasmen." PhD thesis. Rostock University, 2003.
- [84] M. Winkel. "Theory of the Dynamical Collision Frequency in Dense Plasmas and Application to Reflectivity of Xe-Plasmas." Diploma thesis. Rostock University, 2009.
- [85] **H. Reinholz and G. Röpke**. "Dielectric function beyond the random-phase approximation: Kinetic theory versus linear response theory." *Physical Review E* 85 (2012), 036401.
- [86] M. Veysman, H. Reinholz, M. Winkel, and G. Röpke. "Optical Conductivity from kinetic and quantum statistical approach. Application to aluminum plasmas." [in preparation] (2013).
- [87] L. P. Kadanoff and G. Baym. Quantum Statistical Mechanics Green's Function Methods in Equilibrium and Nonequilibrium Problems. Ed. by D. Pines. Frontiers in Physics A Lecture Note and Reprint Series. New York: W. A. Benjamin Inc, 1962.
- [88] G. Röpke and M. Winkel. Green's functions technique for statistical ensembles. Lecture Notes. 2009.
- [89] A. Wierling, T. Millat, G. Röpke, R. Redmer, and H. Reinholz. "Inverse bremsstrahlung of hot, weakly coupled plasmas." *Physics of Plasmas* 8 (2001), 3810.
- [90] R. Thiele, R. Redmer, H. Reinholz, and G. Röpke. "Using the Gould–De-Witt scheme to approximate the dynamic collision frequency in a dense electron gas." *Journal of Physics A: Mathematical and General* 39 (2006), 4365–4368.
- [91] T. Millat, A. Selchow, A. Wierling, H. Reinholz, R. Redmer, and G. Röpke. "Dynamic collision frequency for a two-component plasma." *Journal of Physics A: Mathematical and General* 36 (2003), 6259–6264.

- [92] A. Wierling, T. Millat, and G. Röpke. "Classical bremsstrahlung in a non-ideal plasma with effective interaction." *Journal of Plasma Physics* 70 (2004), 185–197.
- [93] A. Selchow, G. Röpke, A. Wierling, H. Reinholz, A. Przystawik, and G. Zwicknagel. "Dynamic structure factor for a two-component model plasma." *Physical Review E* 64 (2001), 056410.
- [94] H. Reinholz, I. V. Morozov, G. Röpke, and T. Millat. "Internal versus external conductivity of a dense plasma: Many-particle theory and simulations." *Physical Review E* 69 (2004), 066412.
- [95] I. V. Morozov, H. Reinholz, G. Röpke, A. Wierling, and G. Zwick-nagel. "Molecular dynamics simulations of optical conductivity of dense plasmas." *Physical Review E* 71 (2005), 066408.
- [96] T. Bornath, D. Kremp, and M. Schlanges. "Two-particle problem in a nonequilibrium many-particle system." Physical Review E 60 (1999), 6382–6394.
- [97] M. Moll, M. Schlanges, T. Bornath, and V. P. Krainov. "Inverse bremsstrahlung heating beyond the first Born approximation for dense plasmas in laser fields." *New Journal of Physics* 14 (2012), 065010.
- [98] D. Kremp, D. Semkat, T. Bornath, M. Bonitz, M. Schlanges, and P. Hilse. "Interaction of partially ionized plasmas with electromagnetic fields." *Journal of Physics: Conference Series* 35 (2006), 53–70.
- [99] J. R. Adams, N. S. Shilkin, V. E. Fortov, V. K. Gryaznov, V. B. Mintsev, R. Redmer, H. Reinholz, and G. Röpke. "Coulomb contribution to the direct current electrical conductivity of dense partially ionized plasmas." *Physics of Plasmas* 14 (2007), 62303–62310.
- [100] **D. Kremp, T. Bornath, M. Bonitz, and M. Schlanges**. "Quantum kinetic theory of plasmas in strong laser fields." *Physical Review E* 60 (1999), 4725–4732.
- [101] **T. Bornath, M. Schlanges, P. Hilse, and D. Kremp.** "Nonlinear collisional absorption in dense laser plasmas." *Physical Review E* 64 (2001), 026414.
- [102] T. Bornath, D. Kremp, P. Hilse, and M. Schlanges. "Collisional absorption in strong laser fields." *Journal of Physics: Conference Series* 11 (2005), 180–187.
- [103] P. Hilse, M. Schlanges, T. Bornath, and D. Kremp. "Collisional absorption of dense plasmas in strong laser fields: Quantum statistical results and simulation." *Physical Review E* 71 (2005), 056408.
- [104] S. Pfalzner and P. Gibbon. "Direct calculation of inverse-bremsstrahlung absorption in strongly coupled, nonlinearly driven laser plasmas." *Physical Review E* 57 (1998), 4698–4705.

- [105] A. Grinenko and D. O. Gericke. "Nonlinear Collisional Absorption of Laser Light in Dense Strongly Coupled Plasmas." Physical Review Letters 103 (2009), 065005.
- [106] M. Schlanges, T. Bornath, D. Kremp, and P. Hilse. "Quantum kinetic approach to transport processes in dense laser plasmas." Contributions to Plasma Physics 43 (2003), 360–362.
- [107] T. Bornath, M. Schlanges, P. Hilse, and D. Kremp. "Correlation and quantum effects on collisional absorption in dense laser-produced plasmas." *Journal of Physics A: Mathematical and General* 36 (2003), 5941–5948.
- [108] H.-J. Kull and L. Plagne. "Quantum-mechanical dielectric model of the electron—ion collision frequency in strong laser fields." Physics of Plasmas 8 (2001), 5244.
- [109] H.-J. Kull and V. T. Tikhonchuk. "Fast electrons from electron-ion collisions in strong laser fields." *Physics of Plasmas* 12 (2005), 063301.
- [110] G. Rascol, H. Bachau, V. T. Tikhonchuk, H.-J. Kull, and T. Ristow. "Quantum calculations of correlated electron-ion collisions in a strong laser field." *Physics of Plasmas* 13 (2006), 103108.
- [111] **H.-J. Kull.** "Position–momentum correlations in electron–ion scattering in strong laser fields." *New Journal of Physics* 14 (2012), 055013.
- [112] K. Kim, B. Yellampalle, J. Glownia, a. Taylor, and G. Rodriguez. "Measurements of Terahertz Electrical Conductivity of Intense Laser-Heated Dense Aluminum Plasmas." *Physical Review Letters* 100 (2008), 135002.
- [113] K. Widmann, T. Ao, M. Foord, D. Price, A. Ellis, P. Springer, and A. Ng. "Single-State Measurement of Electrical Conductivity of Warm Dense Gold." *Physical Review Letters* 92 (2004), 125002.
- [114] V. B. Mintsev and Y. B. Zaporozhets. "Reflectivity of dense plasma." Contributions to Plasma Physics 29 (1989), 493–501.
- [115] H. Reinholz, G. Röpke, I. V. Morozov, V. B. Mintsev, Y. B. Zaporozhets, V. E. Fortov, and A. Wierling. "Density profile in shock wave fronts of partially ionized xenon plasmas." *Journal of Physics A: Mathematical and General* 36 (2003), 5991–5997.
- [116] H. Reinholz, Y. B. Zaporozhets, V. B. Mintsev, V. E. Fortov, I. V. Morozov, and G. Röpke. "Frequency-dependent reflectivity of shock-compressed xenon plasmas." *Physical Review E* 68 (2003), 036403.
- [117] Y. B. Zaporozhets, V. B. Mintsev, V. K. Gryaznov, V. E. Fortov, H. Reinholz, T. Raitza, and G. Röpke. "Reflectivity of nonideal plasmas." *Journal of Physics A: Mathematical and General* 39 (2006), 4329–4333.

- [118] T. Raitza. "Reflexion an dichten Plasmen." Diploma thesis. Rostock University, 2006.
- [119] T. Raitza, H. Reinholz, G. Röpke, V. B. Mintsev, and A. Wierling. "Reflectivity in shock wave fronts of xenon." *Journal of Physics A: Mathematical and General* 39 (2006), 4393–4399.
- [120] M. Winkel, H. Reinholz, A. Wierling, G. Röpke, Y. B. Zaporozhets, and V. B. Mintsev. "Dynamical Collision Frequency in Warm Dense Matter: Exemplary Application to Plasma Reflectivity." Contributions to Plasma Physics 49 (2009), 687–691.
- [121] Y. B. Zaporozhets, V. B. Mintsev, V. K. Gryaznov, V. E. Fortov, H. Reinholz, and G. Röpke. "Interaction of explosively driven dense plasmas with a low-intensity laser radiation." *Journal of Physics A: Mathematical and Theoretical* 42 (2009), 214063.
- [122] Y. B. Zaporozhets, V. B. Mintsev, V. K. Gryaznov, V. E. Fortov, H. Reinholz, and G. Röpke. "The Investigation of s- and p-Polarized Reflectivities of Nonideal Plasma." *Contributions to Plasma Physics* 50 (2010), 60–63.
- [123] P. Renaudin, C. Blancard, G. Faussurier, and P. Noiret. "Combined Pressure and Electrical-Resistivity Measurements of Warm Dense Aluminum and Titanium Plasmas." *Physical Review Letters* 88 (2002), 215001.
- [124] A. Mostovych, K. Kearney, J. Stamper, and A. Schmitt. "Measurements of plasma opacity from laser-produced optically thin strongly coupled plasmas." *Physical Review Letters* 66 (1991), 612–615.
- [125] **J. F. Benage**. "Review of electrical resistivity measurements of dense aluminum and comparison to theory." *Physics of Plasmas* 7 (2000), 2040.
- [126] **J. Clérouin et al.** "A database for equations of state and resistivities measurements in the warm dense matter regime." *Physics of Plasmas* 19 (2012), 082702.
- [127] **D. Semkat, R. Redmer, and T. Bornath**. "Collisional absorption in aluminum." *Physical Review E* 73 (2006), 066406.
- [128] V. S. Filinov, P. Thomas, I. Varga, T. Meier, M. Bonitz, V. Fortov, and S. Koch. "Interacting electrons in a one-dimensional random array of scatterers: A quantum dynamics and Monte Carlo study." *Physical Review B* 65 (2002), 165124.
- [129] M. Bonitz et al. "Theory and simulation of strong correlations in quantum Coulomb systems." *Journal of Physics A: Mathematical and General* 36 (2003), 5921–5930.

- [130] V. S. Filinov, M. Bonitz, A. V. Filinov, and V. O. Golubnychiy. "Wigner Function Quantum Molecular Dynamics." In: *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, 41–60.
- [131] **G. Zwicknagel and T. Pschiwul**. "WPMD simulations of a two-component plasma." *Journal of Physics A: Mathematical and General* 39 (2006), 4359–4364.
- [132] H. Reinholz, T. Raitza, and G. Röpke. "Correlations, collision frequency and optical properties in laser excited clusters." *International Journal of Modern Physics B* 21 (2007), 2460–2473.
- [133] H. Reinholz, T. Raitza, G. Röpke, and I. V. Morozov. "Optical and transport properties in dense plasmas collision frequency from bulk to cluster." *International Journal of Modern Physics B* 22 (2008), 4627–4641.
- [134] T. Raitza, H. Reinholz, G. Röpke, and I. V. Morozov. "Collision frequency of electrons in laser excited small clusters." *Journal of Physics A: Mathematical and Theoretical* 42 (2009), 214048.
- [135] T. Raitza, H. Reinholz, G. Röpke, I. V. Morozov, and E. Suraud. "Laser Excited Expanding Small Clusters: Single Time Distribution Functions." Contributions to Plasma Physics 49 (2009), 496–506.
- [136] T. Raitza. "Dynamischer Strukturfaktor von Plasmen in hoch angeregten Clustern." PhD thesis. Rostock University, 2010.
- [137] **T. Raitza, H. Reinholz, and G. Röpke**. "Dynamical Spatially Resolved Response Function of Finite 1-D Nano Plasmas." *International Journal of Modern Physics B* 24 (2010), 4961–4978.
- [138] T. Raitza, G. Röpke, H. Reinholz, and I. V. Morozov. "Spatially resolved dynamic structure factor of finite systems from molecular dynamics simulations." *Physical Review E* 84 (2011), 036406.
- [139] T. Raitza, I. Broda, H. Reinholz, and G. Röpke. "Collective Excitations of Finite-Temperature Nano Plasmas." Contributions to Plasma Physics 52 (2012), 118–121.
- [140] T. Raitza, H. Reinholz, P.-G. Reinhard, G. Röpke, and I. Broda. "Spatially resolved collective excitations of nano-plasmas via molecular dynamics simulations and fluid dynamics." *New Journal of Physics* 14 (2012), 115016.
- [141] L. D. Landau and E. M. Lifshitz. Mechanics. 3rd editio. Vol. 1. Course of Theoretical Physics. Butterworth-Heinemann, 1976.
- [142] M. Griebel, S. Knapek, G. Zumbusch, and A. Calgar. Numerische Simulation in der Moleküldynamik. Springer Berlin Heidelberg New York, 2004.

- [143] J. H. Williamson. "Low-Storage Runge-Kutta Schemes." Journal of Computational Physics 56 (1980), 48–56.
- [144] C. K. Birdsall and A. Langdon. Plasma Physics via Computer Simulation. The Adam Hilger Series on Plasma Physics. IOP Publishing Ltd, 1991.
- [145] D. Frenkel and B. Smit. Understanding Molecular Simulations From Algorithms to Applications. Computational Science Series. Academic Press, New York and London, 2002.
- [146] W. C. Swope. "A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters." The Journal of Chemical Physics 76 (1982), 637.
- [147] L. Verlet. "Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules." *Physical Review* 159 (1967), 98–103.
- [148] **F. H. Harlow and J. E. Welch**. "Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface." *Physics of Fluids* 8 (1965), 2182.
- [149] H. C. Plummer. "On the problem of distribution in globular star clusters." Monthly Notices of the Royal Astronomical Society 71 (1911), 460–470.
- [150] M. Hilbert and P. López. "The world's technological capacity to store, communicate, and compute information." Science (New York, N.Y.)) 332 (2011), 60–5.
- [151] **H.-J. Kull**. Computersimulation von Plasmen. Lecture Notes, RWTH Aachen University. 2003.
- [152] R. W. Hockney and J. W. Eastwood. Computer Simulation Using Particles. Bristol and New York: IOP Publishing Ltd, 1988.
- [153] **S. Colombi**. "Dynamics of the large-scale structure of the universe: N-body techniques." New Astronomy Reviews 45 (2001), 373–377.
- [154] P. Gibbon and G. Sutmann. "Long-Range Interactions in Many-Particle Simulation." In: Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms, Lecture Notes. Ed. by J. Grotendorst, D. Marx, and A. Muramatsu. Vol. 10. NIC Series. John von Neumann Institute for Computing, Jülich, 2002, 467–506.
- [155] G. Sutmann, P. Gibbon, and T. Lippert, eds. Fast Methods for Long-Range Interactions in Complex Systems. IAS-Series. Forschungszentrum Jülich, Zentralbibliothek, 2011.
- [156] A. Arnold et al. "Comparison of Scalable Fast Methods for Long-Range Interactions." *Physical Review E [submitted]* (2013).

- [157] A. W. Appel. "An Efficient Program for Many-Body Simulation." SIAM Journal on Scientific and Statistical Computing 6 (1985), 85–103.
- [158] J. E. Barnes and P. Hut. "A hierarchical O(N log N) force-calculation algorithm." *Nature* 324 (1986), 446–449.
- [159] L. Greengard and V. Rokhlin. "A fast algorithm for particle simulations." Journal of Computational Physics 73 (1987), 325–348.
- [160] S. Pfalzner and P. Gibbon. Many-Body Tree Methods in Physics. Cambridge: Cambridge University Press, 1996, 180.
- [161] H. Heuser. Lehrbuch der Analysis, Teil 1. 15th ed. Vieweg+Teubner, 2003.
- [162] L. Papula. Mathematik für Ingenieure und Naturwissenschaftler Band 1: Ein Lehr- und Arbeitsbuch für das Grundstudium. 13th ed. Vieweg+Teubner Verlag, 2011.
- [163] M. Abramowitz and I. Stegun. Handbook of mathematical functions: with formulas, graphs, and mathematical tables. National Bureau of Standards, 1964.
- [164] M. J. Caola. "Solid harmonics and their addition theorems." Journal of Physics A: Mathematical and General 11 (1978), L23.
- [165] C. A. White and M. Head-Gordon. "Derivation and efficient implementation of the fast multipole method." The Journal of Chemical Physics 101 (1994), 6593.
- [166] I. Kabadshow. Periodic Boundary Conditions and the Error-Controlled Fast Multipole. 2010.
- [167] PEPC The Pretty Efficient Coulomb Solver. http://www.fz-juelich.de/ias/jsc/pepc [online, accessed November 2012].
- [168] **P. Gibbon**. *PEPC*: Pretty Efficient Parallel Coulomb-solver. Tech. rep. Forschungszentrum Jülich GmbH, 2003.
- [169] P. Gibbon, F. N. Beg, E. L. Clark, R. G. Evans, and M. Zepf. "Tree-code simulations of proton acceleration from laser-irradiated wire targets." *Physics of Plasmas* 11 (2004), 4032.
- [170] P. Gibbon, W. Frings, S. Dominiczak, and B. Mohr. "Performance Analysis and Visualization of the N-Body Tree Code PEPC on Massively Parallel Computers." In: Parallel Computing: Current & Future Issues of High-End Computing, Proceedings of the International Conference ParCo 2005. Ed. by G. Joubert, W. Nagel, F. Peters, O. Plata, P. Tirado, and E. Zapata. Vol. 33. John von Neumann Institute for Computing, Jülich, 2006, 367.
- [171] P. Gibbon. "Mesh-Free Simulation of High Intensity Laser Interactions with Collisional Plasmas." In: *NIC Symposium 2006*. Ed. by G. Münster, D. Wolf, and M. Kremer. Vol. 32. John von Neumann Institute for Computing, Jülich, 2006, 333–340.

- [172] M. Keldenich. "Optimierung des Multipol-Akzeptanz Kriteriums im Barnes-Hut Baumalgorithmus." JSC Internal Report, FZJ-JSC-IB-2008-04. Fachhochschule Aachen, 2008.
- [173] M. Hofmann, G. Rünger, P. Gibbon, and R. Speck. "Parallel Sorting Algorithms for Optimizing Particle Simulations." In: Cluster Computing Workshops and Posters (CLUSTER WORKSHOPS), 2010 IEEE International Conference on. 2010, 1–8.
- [174] P. Gibbon, R. Speck, A. Karmakar, L. Arnold, W. Frings, B. Berberich, D. Reiter, and M. Mašek. "Progress in Mesh-Free Plasma Simulation with Parallel Tree Codes." *Plasma Science, IEEE Transactions on* 38 (2010), 2367–2376.
- [175] R. Speck, L. Arnold, and P. Gibbon. "Towards a petascale tree code: Scaling and efficiency of the PEPC library." *Journal of Computational Science* 2 (2011), 138–143.
- [176] S. Pfalzner and P. Gibbon. "A 3D hierarchical tree code for dense plasma simulation." Computer Physics Communications 79 (1994), 24–38.
- [177] L. Hernquist. "Hierarchical N-body methods." Computer Physics Communications 48 (1988), 107–115.
- [178] J. K. Salmon and M. S. Warren. "Skeletons from the treecode closet." *Journal of Computational Physics* 111 (1994), 136–155.
- [179] R. Speck. "Generalized Algebraic Kernels and Multipole Expansions for the Vortex Particle Method." PhD thesis. Bergische Universität Wuppertal, 2011.
- [180] BMBF Project 01 IH 08001 A-D: ScaFaCoS Scalable Fast Coulomb Solver. www.scafacos.de, [online, accessed May 2012].
- [181] M. S. Warren and J. K. Salmon. "A parallel hashed Oct-Tree N-body algorithm." Proceedings of the 1993 ACM/IEEE conference on Supercomputing Supercomputing '93 (1993), 12–21.
- [182] D. E. Knuth. The Art of Computer Programming Volume 3: Sorting and Searching. 2nd editio. Addison Wesley Publishing Company, 1998.
- [183] H. Sagan. Space-filling curves. Universitext. Springer Berlin / Heidelberg, 1994.
- [184] M. Bader. Raumfüllende Kurven.
 Begleitendes Skriptum zum entsprechenden Kapitel der Vorlesung "Algorithmen des Wissenschaftlichen Rechnens", Institut für Informatik der Technischen Universität München, Lehrstuhl für numerische Programmierung und Ingenieuranwendungen in der Informatik. 2004.
- [185] **H. Shi and J. Schaeffer**. "Parallel sorting by regular sampling." *Journal of Parallel and Distributed Computing* 14 (1992), 361–372.

- [186] P. Gibbon, R. Speck, L. Arnold, M. Winkel, and H. Hübner. "Parallel Tree Codes." In: Fast Methods for Long-Range Interactions in Complex Systems, Summer School, 6-10 September 2010, Jülich. Vol. 2. Forschungszentrum Jülich GmbH, 2010, 65–84.
- [187] R. Speck, P. Gibbon, and M. Hoffmann. Efficiency and scalability of the parallel Barnes-Hut tree code PEPC. 2009.
- [188] **H. Hübner**. "A priori minimisation of algorithmic bottlenecks in the parallel tree code PEPC." Berichte des Forschungszentrums Jülich, Jül 4339. Fachhochschule Aachen, 2011, 123.
- [189] L. Chenyang, Z. Hong, and W. Nengchao. "Fast N-Dimensional Hilbert Mapping Algorithm." 2008 International Conference on Computational Sciences and its Applications (2008), 507–513.
- [190] M. S. Müller, A. Knüpfer, M. Jurenz, M. Lieber, H. Brunst, H. Mix, and W. E. Nagel. "Developing scalable applications with Vampir, VampirServer and VampirTrace." In: Parallel Computing: Architectures, Algorithms and Applications. Ed. by C. Bischof, M. Bücker, P. Gibbon, G. R. Joubert, T. Lippert, B. Mohr, and F. Peters. Vol. 38. NIC Series. John von Neumann Institute for Computing, Jülich, 2007, 637–644.
- [191] **G. Xue**. "An O(n) time hierarchical tree algorithm for computing force field in n-body simulations." *Theoretical Computer Science* 197 (1998), 157–169.
- [192] J. Makino. "Comparison of Two Different Tree Algorithms." Journal of Computational Physics 88 (1990), 393–408.
- [193] R. J. Anderson. "Tree Data Structures for N-Body Simulation." SIAM Journal on Computing 28 (1999), 1923–1940.
- [194] C.-O. Ahn and S. H. Lee. "A new treecode for long-range force calculation." Computer Physics Communications 178 (2008), 121–127.
- [195] J. K. Salmon. "Parallel hierarchical N-body methods." PhD thesis. California Institute of Technology, Pasadena, California, 1991.
- [196] **J. Dubinski**. "A parallel tree code." New Astronomy 1 (1996), 133–147.
- [197] Q. X. Wang. "Variable order revised binary treecode." Journal of Computational Physics 200 (2004), 192–210.
- [198] E. Gafton and S. Rosswog. "A fast recursive coordinate bisection tree for neighbour search and gravity." Monthly Notices of the Royal Astronomical Society 418 (2011), 770–781.
- [199] J. Waltz, G. Page, S. Milder, J. Wallin, and A. Antunes. "A Performance Comparison of Tree Data Structures for N-Body Simulation." Journal of Computational Physics 178 (2002), 1–14.

- [200] Z.-h. Duan and R. Krasny. "An adaptive treecode for computing nonbonded potential energy in classical molecular systems." *Journal of Computational Chemistry* 22 (2001), 184–195.
- [201] M. S. Warren and J. K. Salmon. "A portable parallel particle program." Computer Physics Communications 87 (1995), 266–290.
- [202] **W. Dehnen**. "A very fast and momentum-conserving tree code." *The Astro-*physical Journal Letters 536 (2000), L39–42.
- [203] W. Dehnen. "A Hierarchical O(N) Force Calculation Algorithm." *Journal of Computational Physics* 179 (2002), 27–42.
- [204] **J. E. Barnes**. "A modified tree code: don't laugh; it runs." *Journal of Computational Physics* 87 (1990), 161–170.
- [205] U. Becciani, V. Antonuccio-Delogu, and M. Gambera. "A Modified Parallel Tree Code for N-Body Simulation of the Large-Scale Structure of the Universe." *Journal of Computational Physics* 163 (2000), 118–132.
- [206] H. A. Boateng and R. Krasny. "Comparison of treecodes for computing electrostatic potentials in charged particle systems with disjoint targets and sources." Journal of computational chemistry (2013).
- [207] L. Hernquist. "Performance characteristics of tree codes." The Astrophysical Journal Supplement Series 64 (1987), 715–734.
- [208] L. Hernquist and J. E. Barnes. "Are some N-body algorithms intrinsically less collisional than others?" *The Astrophysical Journal* 349 (1990), 562–569.
- [209] J. Krek, N. Jelić, and J. Duhovnik. "Grid-free treecode method in diode simulation." *Nuclear Engineering and Design* 261 (2013), 238–243.
- [210] K. Matyash, R. Schneider, R. D. Sydora, and F. Taccogna. "Application of a Grid-Free Kinetic Model to the Collisionless Sheath." Contributions to Plasma Physics 48 (2008), 116–120.
- [211] P. Li, H. Johnston, and R. Krasny. "A Cartesian treecode for screened coulomb interactions." *Journal of Computational Physics* 228 (2009), 3858–3868.
- [212] J. K. Salmon, M. S. Warren, and G. S. Winckelmans. "Fast parallel tree codes for gravitational and fluid dynamical N-body problems." *International Journal of Supercomputer Applications* 8 (1986), 129–142.
- [213] G. S. Winckelmans, J. K. Salmon, M. S. Warren, A. Leonard, and B. Jodoin. "Application of Fast Parallel and Sequential Tree Codes to Computing Three-Dimensional Flows with the Vortex Element and Boundary Element Methods." ESAIM: Proceedings 1 (1996), 225–240.
- [214] J. K. Salmon and M. S. Warren. "Parallel, out-of-core methods for N-body simulation." In: Proceedings of the Eighth SIAM Conference on Parallell Processing for Scientific Computing. 1997.

- [215] T. Hamada, T. Narumi, R. Yokota, K. Yasuoka, K. Nitadori, and M. Taiji. "42 TFlops hierarchical N-body simulations on GPUs with applications in both astrophysics and turbulence." In: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis SC '09. New York, New York, USA: ACM Press, 2009, 1.
- [216] L. Hernquist and N. Katz. "TREESPH A unification of SPH with the hierarchical tree method." The Astrophysical Journal Supplement Series 70 (1989), 419–446.
- [217] V. Springel, N. Yoshida, and S. D. White. "GADGET: a code for collisionless and gasdynamical cosmological simulations." New Astronomy 6 (2001), 79–117.
- [218] J. Stadel, J. Wadsley, and D. Richardson. "High performance computational astrophysics with pkdgrav/gasoline." In: *High Performance Computing Systems and Applications*. The Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers, 2002, 501–523.
- [219] D. Madlener. "Numerische Simulation von protostellaren Scheiben mit Treecodes und Smoothed Particle Hydrodynamics." Diploma Thesis. Cologne University, 2008.
- [220] A. Breslau. "Highly-parallel Smoothed Particle Hydrodynamics modelling of protoplanetary discs." Diploma Thesis. Cologne University, 2011.
- [221] J. S. Bagla. "TreePM: A code for cosmological N-body simulations." Journal of Astrophysics and Astronomy 23 (2002), 185–196.
- [222] J. Dubinski, J. Kim, C. Park, and R. Humble. "GOTPM: a parallel hybrid particle-mesh treecode." New Astronomy 9 (2004), 111–126.
- [223] V. Springel. "The cosmological simulation code GADGET-2." Monthly Notices of the Royal Astronomical Society 364 (2005), 1105–1134.
- [224] M. S. Warren and J. K. Salmon. "Astrophysical N-body simulations using hierarchical tree data structures." *Proceedings Supercomputing '92* (1992), 570–576.
- [225] M. S. Warren and J. K. Salmon. "A parallel, portable and versatile treecode." In: Seventh SIAM Conference on Parallell Computing. 1995, 1–6.
- [226] M. S. Warren, J. K. Salmon, D. Becker, M. P. Goda, T. Sterling, and G. S. Winckelmans. "Pentium Pro Inside: I. A Treecode at 430 Gigaflops on ASCI Red, II. Price/Performance of \$50/Mflop on Loki and Hyglac." In: Supercomputing, ACM/IEEE 1997 Conference. 1997, 61.
- [227] U. Becciani, V. Antonuccio-Delogu, and M. Comparato. "FLY: MPI-2 high resolution code for LSS cosmological simulations." *Computer Physics Communications* 176 (2007), 211–217.

- [228] F. Gioachin, A. Sharma, S. Chakravorty, C. L. Mendes, and L. V. Kalé. "Scalable Cosmological Simulations on Parallel Machines." In: VECPAR'06 Proceedings of the 7th international conference on High performance computing for computational science. 2007, 476–489.
- [229] P. Jetley, F. Gioachin, C. Mendes, and L. V. Kalé. "Massively Parallel Cosmological Simulations with ChaNGa." In: Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on. 2008, 1–12.
- [230] B. Jeon, J. D. Kress, L. A. Collins, and N. Grønbech-Jensen. "Parallel TREE code for two-component ultracold plasma analysis." Computer Physics Communications 178 (2008), 272–279.
- [231] K. Nitadori, J. Makino, and P. Hut. "Performance tuning of N-body codes on modern microprocessors: I. Direct integration with a hermite scheme on x86 64 architecture." New Astronomy 12 (2006), 169–181.
- [232] A. Tanikawa, K. Yoshikawa, T. Okamoto, and K. Nitadori. "N-body simulation for self-gravitating collisional systems with a new SIMD instruction set extension to the x86 architecture, Advanced Vector eXtensions." New Astronomy 17 (2012), 82–92.
- [233] A. Tanikawa, K. Yoshikawa, K. Nitadori, and T. Okamoto. "Phantom-GRAPE: Numerical software library to accelerate collisionless N-body simulation with SIMD instruction set on x86 architecture." New Astronomy 19 (2013), 74–88.
- [234] R. Yokota. "An FMM Based on Dual Tree Traversal for Many-core Architectures." arXiv:1209.3516 [cs.NA] (2012).
- [235] **J. Makino**. "Vectorization of a Treecode." *Journal of Computational Physics* 87 (1990), 148–160.
- [236] L. Hernquist. "Vectorization of tree traversals." Journal of Computational Physics 87 (1990), 137–147.
- [237] P. Jetley, L. Wesolowski, F. Gioachin, L. V. Kalé, and T. R. Quinn. "Scaling Hierarchical N-body Simulations on GPU Clusters." In: 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE, 2010, 1–11.
- [238] E. Gaburov, J. Bédorf, and S. P. Zwart. "Gravitational tree-code on graphics processing units: implementation in CUDA." *Procedia Computer Science* 1 (2010), 1119–1127.
- [239] M. Burtscher and K. Pingali. "An Efficient CUDA Implementation of the Tree-Based Barnes Hut n-Body Algorithm." In: *GPU Computing Gems Emerald Edition*. 2011, 75–92.

- [240] J. Bédorf, E. Gaburov, and S. P. Zwart. "Bonsai: A GPU Tree-Code." arXiv:1204.2280 [astro-ph.IM] (2012), 5.
- [241] T. Hamada, K. Nitadori, K. Benkrid, Y. Ohno, G. Morimoto, T. Masada, Y. Shibata, K. Oguri, and M. Taiji. "A novel multiple-walk parallel algorithm for the Barnes-Hut treecode on GPUs towards cost effective, high performance N-body simulation." Computer Science Research and Development 24 (2009), 21–31.
- [242] W. M. Brown, P. Wang, S. J. Plimpton, and A. N. Tharrington. "Implementing molecular dynamics on hybrid high performance computers short range forces." *Computer Physics Communications* 182 (2011), 898–911.
- [243] T. Hamada and K. Nitadori. "190 TFlops Astrophysical N-body Simulation on a Cluster of GPUs." In: 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis SC'10. IEEE, 2010, 1–9.
- [244] S. J. Aarseth. "Direct integration methods of the N-body problem." Astrophysics and Space Science 14 (1971), 118–132.
- [245] L. Hernquist, F. F. Bouchet, and Y. Suto. "Application of the Ewald method to cosmological N-body simulations." *The Astrophysical Journal Supplement Series* 75 (1991), 231–240.
- [246] **F. Eilenberger**. "Implementation of a Tabulated Barnes-Hut-Ewald Algorithm for Periodic Boundaries in PEPC." In: *Beiträge zum Wissenschaftlichen Rechnen Ergebnisse des Gaststudentenprogramms 2006 des John von Neumann-Instituts für Computing.* Ed. by **R. Esser**. Jülich: Forschungszentrum Jülich GmbH, 2008, 21–40.
- [247] A. J. Christlieb, R. Krasny, J. Verboncoeur, J. Emhoff, and I. Boyd. "Grid-free plasma Simulation techniques." *IEEE Transactions on Plasma Science* 34 (2006), 149–165.
- [248] P. P. Ewald. "Die Berechnung optischer und elektrostatischer Gitterpotentiale." Annalen der Physik 369 (1921), 253–287.
- [249] S. W. D. Leeuw and J. W. Perram. "Electrostatic lattice sums for semiinfinite lattices." Molecular Physics: An International Journal at the Interface Between Chemistry and Physics 37 (1978), 1313–1322.
- [250] J. C. Shelley. "Boundary condition effects in simulations of water confined between planar walls." *Molecular Physics: An International Journal at the Interface Between Chemistry and Physics* 88 (1996), 385–398.
- [251] **E. Spohr**. "Effect of electrostatic boundary conditions and system size on the interfacial properties of water and aqueous solutions." *The Journal of Chemical Physics* 107 (1997), 6342.

- [252] I.-C. Yeh and M. L. Berkowitz. "Ewald summation for systems with slab geometry." *The Journal of Chemical Physics* 111 (1999), 3155.
- [253] J. de Joannis, A. Arnold, and C. Holm. "Electrostatics in periodic slab geometries. II." *The Journal of Chemical Physics* 117 (2002), 2503.
- [254] A. Bródka. "Comment on "Electrostatics in periodic slab geometries. I" [J. Chem. Phys. 117, 2496 (2002)]." The Journal of Chemical Physics 121 (2004), 7032–4.
- [255] I. Kabadshow. The Fast Multipole Method Alternative Gradient Algorithm and Parallelization. Berichte des Forschungszentrums Jülich. Forschungszentrum Jülich GmbH, 2006.
- [256] M. Challacombe, C. A. White, and M. Head-Gordon. "Periodic boundary conditions and the fast multipole method." The Journal of Chemical Physics 107 (1997), 10131.
- [257] K. N. Kudin and G. E. Scuseria. "Revisiting infinite lattice sums with the periodic fast multipole method." *Journal of Chemical Physics* 121 (2004), 2886– 2890.
- [258] **B. Nijboer and F. De Wette**. "On the calculation of lattice sums." *Physica* 23 (1957), 309–321.
- [259] **K. N. Kudin and G. E. Scuseria**. "A fast multipole method for periodic systems with arbitrary unit cell geometries." *Chemical Physics Letters* 283 (1998), 61–68.
- [260] K. N. Kudin and G. E. Scuseria. "Analytic stress tensor with the periodic fast multipole method." *Physical Review B* 61 (2000), 5141–5146.
- [261] C. L. Berman and L. Greengard. "A renormalization method for the evaluation of lattice sums." *Journal of Mathematical Physics* 35 (1994), 6036.
- [262] I. Kabadshow. "Periodic Boundary Conditions and the Error-Controlled Fast Multipole Method." PhD thesis. Bergische Universität Wuppertal, 2012.
- [263] A. Redlack and J. Grindlay. "The electrostatic potential in a finite ionic crystal." Canadian Journal of Physics 50 (1972), 2815–2825.
- [264] A. Redlack and J. Grindlay. "Coulombic potential lattice sums." J. Phys. Chem. Solids 36 (1975), 73–82.
- [265] J. E. Roberts and J. Schnitker. "How the unit cell surface charge distribution affects the energetics of ion-solvent interactions in simulations." The Journal of Chemical Physics 101 (1994), 5024.
- [266] B. Steinbusch. The logarithmic potential Series expansions and periodic systems.
 Presentation at Simulation Laboratory Plasma Physics Group meeting at Jülich Supercomputing Centre. 2013.

- [267] JUGENE website. http://www.fz-juelich.de/ias/jsc/jugene [online, accessed May 2012].
- [268] JUQUEEN website. http://www.fz-juelich.de/ias/jsc/juqueen [online, accessed November 2012]. 2012.
- [269] M. Winkel, R. Speck, H. Hübner, L. Arnold, R. Krause, and P. Gibbon. "A massively parallel, multi-disciplinary Barnes-Hut tree code for extremescale N-body simulations." Computer Physics Communications 183 (2012), 880– 889.
- [270] W. Frings, M. Riedel, A. Streit, and D. Mallmann. "LLview: User-level Monitoring in Computational Grids and e-Science Infrastructures." In: German e-Science Conference. 2007.
- [271] **H. Peschke**. "Hierarchical Tree Construction in PEPC." In: *JSC Guest Student Programme Proceedings 2011*. Ed. by **M. Winkel**. Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, 2011, 21–33.
- [272] **P. Gibbon**. "Resistively enhanced proton acceleration via high-intensity laser interactions with cold foil targets." *Physical Review E* 72 (2005), 026411.
- [273] M. Mašek and P. Gibbon. "A mesh-free Darwin model for non-radiative plasma simulation." ECA 34th EPS Conference on Plasma Phys. Warsaw, 2 6 July 2007 31 (2007), 5–8.
- [274] M. Mašek and P. Gibbon. "Mesh-Free Magnetoinductive Plasma Model." IEEE Transactions on Plasma Science 38 (2010), 2377–2382.
- [275] **B. Berberich**. "Entwicklung und Anwendung eines Tree-Codes in Simulationsszenarios der Plasma-Wand-Wechselwirkung." ger. PhD thesis. Universität Düsseldorf, 2012.
- [276] C. Salmagne. private communication.
- [277] R. D. Sydora. Complex Particle Kinetics and Mesh-Free Methods: Advanced Multi-Scale Models for Kinetic Edge Modeling of Magnetic Fusion Devices. Tech. rep. Forschungszentrum Jülich, 2011.
- [278] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. L. Minion, M. Winkel, and P. Gibbon. "Integrating an N-body problem with SDC and PFASST." In: 21st International Conference on Domain Decomposition Methods [in press]. Rennes, France, 2012.
- [279] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. L. Minion, M. Winkel, and P. Gibbon. "A massively space-time parallel N-body solver." In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. SC '12, Salt Lake City, Utah. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, 92:1–92:11.

- [280] G. Hager and G. Wellein. "Optimization Techniques for Modern High Performance Computers." In: 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, 731–767.
- [281] EU TEXT project Towards EXaflop applicaTions. http://www.project-text.eu/ [online, accessed November 2012].
- [282] EU Mont-Blanc project. http://www.montblanc-project.eu/ [online, accessed November 2012].
- [283] M. L. Minion. "A Hybrid Parareal Spectral Deferred Corrections Method." Communications in Applied Mathematics and Computational Science 5 (2010), 265–301.
- [284] M. Emmett and M. L. Minion. "Toward an efficient parallel in time method for partial differential equations." Communications in Applied Mathematics and Computational Science 7 (2012), 105–132.
- [285] J.-L. Lions, Y. Maday, and G. Turinici. "A "parareal" in time discretization of PDE's." Comptes Rendus de l'Académie des Sciences Series I Mathematics 332 (2001), 661–668.
- [286] A. Dutt, L. Greengard, and V. Rokhlin. "Spectral Deferred Correction Methods for Ordinary Differential Equations." *BIT Numerical Mathematics* 40 (2000), 241–266.
- [287] P.-G. Reinhard and E. Suraud. Introduction to Cluster Dynamics. Weinheim, Germany: Wiley-VCH Verlag GmbH & Co. KGaA, 2004.
- [288] U. Kreibig and M. Vollmer. Optical Properties of Metal Clusters. 1st editio. Berlin: Springer, 1995.
- [289] N. Yu, J. Fan, Q. J. Wang, C. Pflügl, L. Diehl, T. Edamura, M. Yamanishi, H. Kan, and F. Capasso. "Small-divergence semiconductor lasers by plasmonic collimation." *Nature Photonics* 2 (2008), 564–570.
- [290] C. Gmachl. "Plasmonics: A sharper approach." Nature Photonics 2 (2008), 524–525.
- [291] D. K. Gramotnev and S. I. Bozhevolnyi. "Plasmonics beyond the diffraction limit." Nature Photonics 4 (2010), 83–91.
- [292] L. Novotny and N. van Hulst. "Antennas for light." Nature Photonics 5 (2011), 83–90.
- [293] M. I. Stockman. "Nanoplasmonics: past, present, and glimpse into future." Optics Express 19 (2011), 22029.
- [294] M. Kauranen and A. V. Zayats. "Nonlinear plasmonics." *Nature Photonics* 6 (2012), 737–748.

- [295] M. Quinten. Optical Properties of Nanoparticle Systems. Weinheim, Germany: Wiley-VCH Verlag GmbH & Co. KGaA, 2011.
- [296] M. Winkel and P. Gibbon. "Spatially resolved electronic correlations in nanoclusters." *Contributions to Plasma Physics* 53 (2013), 254–262.
- [297] JUROPA/HPC-FF website. http://www.fz-juelich.de/ias/jsc/juropa [online, accessed November 2012].
- [298] I. Broda. "-." Diploma thesis [to be completed in 2013]. Johannes Kepler Universität Linz, 2013.
- [299] H. Reinholz, I. Broda, T. Raitza, and G. Röpke. "Size Dependence of Minimum Charge of Excited Nano-Plasmas." Contributions to Plasma Physics 53 (2013), 263–269.
- [300] J. Kauppinen and J. Partanen. Fourier Transforms in Spectroscopy. Berlin: WILEY-VCH Verlag, 2001.
- [301] W. A. Stahel. Statistische Datenanalyse Eine Einführung für Naturwissenschaftler. vieweg Lehrbuch angewandte Mathematik. Braunschweig/Wiesbaden: Friedr. Vieweg & Sohn Verlagsgesellschaft mbH, 1995.
- [302] M. Knap. "Jackknife und Bootstrap." Bachelor Thesis. Graz University of Technology, 2007, 1–71.
- [303] H.-J. Kull. private communication. 2013.
- [304] C. Varin, C. Peltz, T. Brabec, and T. Fennel. "Attosecond Plasma Wave Dynamics in Laser-Driven Cluster Nanoplasmas." *Physical Review Letters* 108 (2012), 175007.
- [305] C. Peltz, C. Varin, T. Brabec, and T. Fennel. "Fully microscopic analysis of laser-driven finite plasmas using the example of clusters." New Journal of Physics 14 (2012), 065011.
- [306] G. D. Mahan. *Many-Particle Physics*. 2nd. Physics of Solids and Liquids. Plenum Press, New York and London, 1990.
- [307] G. Kelbg. "Theorie des Quanten-Plasmas." Annalen der Physik 467 (1963), 219–224.
- [308] G. Kelbg. "Quantenstatistik der Gase mit Coulomb-Wechselwirkung." Annalen der Physik 467 (1963), 354–360.
- [309] G. Kelbg. "Klassische statistische Mechanik der Teilchen-Mischungen mit sortenabhängigen weitreichenden zwischenmolekularen Wechselwirkungen." Annalen der Physik 469 (1964), 394–403.
- [310] G. Kelbg and H. J. Hoffmann. "Quantenstatistik realer Gase und Plasmen." Annalen der Physik 469 (1964), 310–318.

- [311] L. V. Woodcock. "Isothermal molecular dynamics calculations for liquid salts." Chemical Physics Letters 10 (1971), 257–261.
- [312] W. G. Hoover. "Atomistic nonequilibrium computer simulations." *Physica A: Statistical Mechanics and its Applications* 118 (1983), 111–122.
- [313] **D. Brown and J. Clarke**. "A comparison of constant energy, constant temperature and constant pressure ensembles in molecular dynamics simulations of atomic liquids." *Molecular Physics: An International Journal at the Interface Between Chemistry and Physics* 51 (1984), 1243–1252.
- [314] W. G. Hoover, A. J. C. Ladd, and B. Moran. "High-Strain-Rate Plastic Flow Studied via Nonequilibrium Molecular Dynamics." *Physical Review Letters* 48 (1982), 1818–1820.
- [315] M. P. Allen and D. J. Tildesley. Computer Simulation of Liquids. reprint. Oxford: Oxford University Press, 1991.
- [316] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, a. DiNola, and J. R. Haak. "Molecular dynamics with coupling to an external bath." The Journal of Chemical Physics 81 (1984), 3684.
- [317] B. Hess et al. GROMACS Groningen Machine for Chemical Simulations, User Manual for version 4.5.4. 2010.
- [318] P. H. Hünenberger. "Thermostat algorithms for molecular dynamics simulations." Advanced Polymer Science 173 (2005), 105–149.
- [319] S. Nosé. "A study of solid and liquid carbon tetrafluoride using the constant pressure molecular dynamics technique." The Journal of Chemical Physics 78 (1983), 6928.
- [320] S. Nosé and M. L. Klein. "Constant pressure molecular dynamics for molecular systems." Molecular Physics: An International Journal at the Interface Between Chemistry and Physics 50 (1983), 1055–1076.
- [321] S. Nosé. "A unified formulation of the constant temperature molecular dynamics methods." The Journal of Chemical Physics 81 (1984), 511.
- [322] S. Nosé. "A molecular dynamics method for simulations in the canonical ensemble." Molecular Physics: An International Journal at the Interface Between Chemistry and Physics 52 (1984), 255–268.
- [323] D. C. Rapaport. The Art of Molecular Dynamics Simulation. 2nd(repr.) Cambridge University Press, 2005.
- [324] M. E. Tuckerman. Statistical Mechanics: Theory and Molecular Simulation. Oxford Graduate Texts. Oxford University Press, 2010.
- [325] G. J. Martyna, M. L. Klein, and M. E. Tuckerman. "Nosé-Hoover chains: The canonical ensemble via continuous dynamics." *The Journal of Chemical Physics* 97 (1992), 2635.

- [326] W. G. Hoover. "Canonical dynamics: Equilibrium phase-space distributions." Physical Review A 31 (1985), 1695–1697.
- [327] S. Toxvaerd and O. H. Olsen. "Canonical Molecular Dynamics of Molecules with Internal Degrees of Freedom." *Berichte der Bunsengesellschaft für physikalische Chemie* 94 (1990), 274–278.
- [328] D. Kusnezov, A. Bulgac, and W. Bauer. "Canonical Ensembles from Chaos." Annals of Physics 204 (1990), 155–185.
- [329] G. J. Martyna, D. J. Tobias, and M. L. Klein. "Constant pressure molecular dynamics algorithms." *The Journal of Chemical Physics* 101 (1994), 4177.
- [330] G. J. Martyna, M. E. Tuckerman, D. J. Tobias, and M. L. Klein. "Explicit reversible integrators for extended systems dynamics." *Molecular Physics* 87 (1996), 1117–1157.
- [331] S. Jang and G. a. Voth. "Simple reversible molecular dynamics algorithms for Nose-Hoover chain dynamics." *The Journal of Chemical Physics* 107 (1997), 9514.
- [332] M. E. Tuckerman and G. J. Martyna. "Comment on "Simple reversible molecular dynamics algorithms for Nosé-Hoover chain dynamics" [J. Chem. Phys. 107, 9514 (1997)]." The Journal of Chemical Physics 110 (1999), 3623.
- [333] S. Jang and G. a. Voth. "Response to "Comment on 'Simple reversible molecular dynamics algorithms for Nosé-Hoover chain dynamics'" [J. Chem. Phys. 110, 3623 (1999)]." The Journal of Chemical Physics 110 (1999), 3626.
- [334] **T. Yamamoto**. "Comment on "comment on 'simple reversible molecular dynamics algorithms for Nosé-Hoover chain dynamics' " [J. Chem. Phys. 110, 3623 (1999)]." The Journal of chemical physics 124 (2006), 217101.
- [335] H. Reinholz. private communication.
- [336] H. S. Cohl. "Fourier and Gegenbauer expansions for fundamental solutions of the Laplacian and powers in R^d and H^d." PhD thesis. University of Auckland, 2010.
- [337] **H. S. Cohl.** "On a generalization of the generating function for Gegenbauer polynomials." arXiv:1105.2735 [math.CA] (2013), 1–12.
- [338] V. S. Filinov, M. Bonitz, A. V. Filinov, and V. O. Golubnychiy. "Wigner Function Quantum Molecular Dynamics." In: Lecture Notes in Physics. Ed. by H. Fehske, R. Schneider, and A. Weiße. Vol. 739. Lecture Notes in Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, 41–60.

1. Three-dimensional modelling of soil-plant interactions: Consistent coupling of soil and plant root systems

by T. Schröder (2009), VIII, 72 pages

ISBN: 978-3-89336-576-0 URN: urn:nbn:de:0001-00505

2. Large-Scale Simulations of Error-Prone Quantum Computation Devices

by D. B. Trieu (2009), VI, 173 pages

ISBN: 978-3-89336-601-9 URN: urn:nbn:de:0001-00552

3. NIC Symposium 2010

Proceedings, 24 – 25 February 2010 | Jülich, Germany edited by G. Münster, D. Wolf, M. Kremer (2010), V. 395 pages

ISBN: 978-3-89336-606-4

URN: urn:nbn:de:0001-2010020108

4. Timestamp Synchronization of Concurrent Events

by D. Becker (2010), XVIII, 116 pages

ISBN: 978-3-89336-625-5

URN: urn:nbn:de:0001-2010051916

5. UNICORE Summit 2010

Proceedings, 18 – 19 May 2010 | Jülich, Germany

edited by A. Streit, M. Romberg, D. Mallmann (2010), iv, 123 pages

ISBN: 978-3-89336-661-3

URN: urn:nbn:de:0001-2010082304

6. Fast Methods for Long-Range Interactions in Complex Systems

Lecture Notes, Summer School, 6 – 10 September 2010, Jülich, Germany

edited by P. Gibbon, T. Lippert, G. Sutmann (2011), ii, 167 pages

ISBN: 978-3-89336-714-6

URN: urn:nbn:de:0001-2011051907

7. Generalized Algebraic Kernels and Multipole Expansions for Massively Parallel Vortex Particle Methods

by R. Speck (2011), iv, 125 pages

ISBN: 978-3-89336-733-7

URN: urn:nbn:de:0001-2011083003

8. From Computational Biophysics to Systems Biology (CBSB11)

Proceedings, 20 - 22 July 2011 | Jülich, Germany

edited by P. Carloni, U. H. E. Hansmann, T. Lippert, J. H. Meinke, S. Mohanty,

W. Nadler, O. Zimmermann (2011), v, 255 pages

ISBN: 978-3-89336-748-1

URN: urn:nbn:de:0001-2011112819

9. UNICORE Summit 2011

Proceedings, 7 - 8 July 2011 | Toruń, Poland

edited by M. Romberg, P. Bała, R. Müller-Pfefferkorn, D. Mallmann (2011), iv,

150 pages

ISBN: 978-3-89336-750-4

URN: urn:nbn:de:0001-2011120103

10. Hierarchical Methods for Dynamics in Complex Molecular Systems

Lecture Notes, IAS Winter School, 5 – 9 March 2012, Jülich, Germany edited by J. Grotendorst, G. Sutmann, G. Gompper, D. Marx (2012), vi,

540 pages

ISBN: 978-3-89336-768-9

URN: urn:nbn:de:0001-2012020208

11. Periodic Boundary Conditions and the Error-Controlled Fast Multipole Method

by I. Kabadshow (2012), v, 126 pages

ISBN: 978-3-89336-770-2

URN: urn:nbn:de:0001-2012020810

12. Capturing Parallel Performance Dynamics

by Z. P. Szebenyi (2012), xxi, 192 pages

ISBN: 978-3-89336-798-6

URN: urn:nbn:de:0001-2012062204

13. Validated force-based modeling of pedestrian dynamics

by M. Chraibi (2012), xiv, 112 pages

ISBN: 978-3-89336-799-3

URN: urn:nbn:de:0001-2012062608

14. Pedestrian fundamental diagrams:

Comparative analysis of experiments in different geometries

by J. Zhang (2012), xiii, 103 pages

ISBN: 978-3-89336-825-9

URN: urn:nbn:de:0001-2012102405

15. UNICORE Summit 2012

Proceedings, 30 - 31 May 2012 | Dresden, Germany

edited by V. Huber, R. Müller-Pfefferkorn, M. Romberg (2012), iv, 143 pages

ISBN: 978-3-89336-829-7

URN: urn:nbn:de:0001-2012111202

16. Design and Applications of an Interoperability Reference Model for Production e-Science Infrastructures

by M. Riedel (2013), x, 270 pages

ISBN: 978-3-89336-861-7

URN: urn:nbn:de:0001-2013031903

Schriften des Forschungszentrums Jülich IAS Series

17. Route Choice Modelling and Runtime Optimisation for Simulation of Building Evacuation

by A. U. Kemloh Wagoum (2013), xviii, 122 pages

ISBN: 978-3-89336-865-5

URN: urn:nbn:de:0001-2013032608

18. Dynamik von Personenströmen in Sportstadien

by S. Burghardt (2013), xi, 115 pages

ISBN: 978-3-89336-879-2

URN: urn:nbn:de:0001-2013060504

19. Multiscale Modelling Methods for Applications in Materials Science

by I. Kondov, G. Sutmann (2013), 326 pages

ISBN: 978-3-89336-899-0

URN: urn:nbn:de:0001-20130902

20. High-resolution Simulations of Strongly Coupled Coulomb Systems with a Parallel Tree Code

by M. Winkel (2013), xvii, 196 pages

ISBN: 978-3-89336-901-0

URN: urn:nbn:de:0001-2013091802

Plasma systems that can be experimentally studied today are reaching from hot, low-density plasmas of fusion research to cold dense solids that are dominated by quantum-mechanical effects and strong correlations. Their consistent theoretical description requires a multitude of effects to be considered. In particular, strong correlations pose significant difficulties here. Computer simulations provide a tool for bridging between experiments and theory as they do not suffer from these complications.

The experimentally accessible optical and transport properties in plasmas are primarily featured by the electronic subsystem, such as its collective behavior and interaction with the ionic background, i.e. Coulomb collisions. In this work the collisional behavior of warm dense bulk matter and collective effects in nano plasmas are investigated by means of molecular dynamics simulations. To this end, simulation experiments performed earlier on electronic resonances in metallic nano clusters are extended to significantly larger systems. The observed complex resonance structure is analyzed using a newly introduced spatially resolved spectral diagnostic. As a second field of study, the bulk collision frequency as the key parameter for optical and transport properties in warm dense matter is evaluated in a generalized Drude approach for a hydrogen-like plasma. Here, the combined high-field and strong coupling regime that is only scarcely covered by theoretical models is of primary interest.

To solve the underlying N-body problem for both applications, a highly parallel Barnes-Hut tree code is utilized and considerably extended with respect to functionality, versatility, and scalability. With its new excellent scalability to hundred thousands of processors and simulation setups consisting of up to billions of particles and its support for periodic boundary conditions with an efficient and precise real-space approach it delivers highly resolved results and is prepared for further studies on the warm dense matter regime. Here, its unique predictive capabilities can finally be used for connecting to real-world experiments.

This publication was written at the Jülich Supercomputing Centre (JSC) which is an integral part of the Institute for Advanced Simulation (IAS). The IAS combines the Jülich simulation sciences and the supercomputer facility in one organizational unit. It includes those parts of the scientific institutes at Forschungszentrum Jülich which use simulation on supercomputers as their main research methodology.



