

---

# Bridging the Microscopic and Macroscopic Realms of Laser Driven Plasma Dynamics

---

Graeme Bart

Thesis submitted to the Faculty of Graduate and Postdoctoral Studies in partial fulfillment  
of the requirements for a  
**Doctorate of Philosophy in Physics**

Department of Physics  
Faculty of Science  
University of Ottawa

## Abstract

The physical processes shaping laser plasma dynamics take place on length scales ranging from the microscopic (1 ångström) to the macroscopic realms ( $\mu\text{m}$ ). Microscopic field fluctuations due to the motions of individual plasma charges evolve on an atomic scale. Collisional effects influencing thermalization and ionization processes depend on the plasma fields on an atomic level. Simultaneously, collective processes such as plasma oscillations take place on a mesoscopic length scale of many-nm. The macroscopic realm is ultimately determined by the laser wavelength, which typically spans hundreds of nm to a few  $\mu\text{m}$ .

Consequently, ab-initio modelling of laser plasma dynamics requires the resolution of length scales from 1 Å to multiple  $\mu\text{m}$ . As such, in order to bridge the microscopic and macroscopic length scales of light-matter interaction, it is necessary to account for the individual motions of up to  $\sim 10^{11}$  particles. This is a not an insignificant undertaking.

Until recently, approaches to numerical modelling of light-matter interactions were limited to MD and PIC, each with their own limitations. MicPIC has been developed to fill the gap left by MD and PIC but so far has not been adapted for scalable parallel processing on large distributed memory machines. Thus, its full potential was not able to be fully realized until now.

This thesis presents the massively parallel MicPIC method capable of bridging the micro- and macroscopic realms. A wide range of applications that have heretofore not been accessible to theory or, at best, had limited applicability are now open for thorough investigation. Among these are nonlinear nanophotonics, quantum nanophotonics, laser machining, ab-initio dynamics of strongly coupled plasmas, high-harmonic generation, electron and x-ray sources, and optical switching. Two of the first applications of parallel MicPIC to a selection of such problems are shown and discussed below, demonstrating the applicability of the method to a wide variety of newly accessible strong field laser-plasma physics phenomena.

# CONTENTS

<b>Abstract</b>	<b>ii</b>
<b>Table of Contents</b>	<b>iii</b>
<b>Preface</b>	<b>v</b>
<b>List of Publications</b>	<b>vi</b>
<b>List of Figures</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Basic Principles . . . . .	3
1.1.1 Plasma Parameters and Screening . . . . .	4
1.1.2 Kinetic Theory of Plasma . . . . .	7
1.1.3 Ionization Processes and High Harmonic Generation in Atoms and Molecules . . . . .	10
1.1.4 Ionization and Photoemission in Solids . . . . .	12
<b>2 Microscopic Particle-in-Cell</b>	<b>15</b>
2.1 Introduction . . . . .	15
2.1.1 Microscopic Particle-in-Cell . . . . .	16
2.1.2 The MicPIC Algorithm . . . . .	19
2.2 Finite Difference Time Domain . . . . .	19
2.2.1 Solving Maxwell's Equations . . . . .	19
2.2.2 Introducing Radiation: The Total-Field Scattered-Field Technique . . . . .	22
2.2.3 Radiation at the Boundaries: The Uniaxial Perfectly Matched Layer . . . . .	26
2.3 Current Density and Long Range Force Weighting . . . . .	29
2.4 Short Range Force and Microscopic Corrections . . . . .	31
2.5 Solving the Equations of Motion . . . . .	33
2.6 Generating new particles: Ionization . . . . .	34
2.6.1 Impact Ionization . . . . .	34
2.6.2 Tunnel Ionization . . . . .	36

<b>3</b>	<b>Massively parallel microscopic particle-in-cell</b>	<b>37</b>
3.1	Introduction . . . . .	37
3.2	The Method of Parallel MicPIC . . . . .	39
3.2.1	Parallel domain decomposition . . . . .	39
3.3	Theoretical scaling analysis . . . . .	46
3.4	Tested performance . . . . .	49
3.4.1	Optimum Cell Size . . . . .	50
3.4.2	Parallel Speedup and Efficiency . . . . .	54
3.4.3	Scalability for massive-scale jobs . . . . .	60
3.5	Conclusion . . . . .	61
<b>4</b>	<b>Collision Dynamics in Strongly Coupled Plasmas</b>	<b>63</b>
4.1	Introduction . . . . .	63
4.2	Theoretical Background . . . . .	64
4.3	Numerical Methods . . . . .	65
4.4	Simulation Results . . . . .	66
4.4.1	Particle Velocities . . . . .	67
4.4.2	Variations in the Electric Field . . . . .	71
4.4.3	Harmonic Spectrum . . . . .	72
4.5	Conclusions . . . . .	73
<b>5</b>	<b>Photoemission of Electrons from Metal Nanotips</b>	<b>76</b>
5.1	Introduction . . . . .	76
5.2	Theoretical Background . . . . .	77
5.3	Implementation . . . . .	77
5.3.1	Determining the turning points . . . . .	78
5.3.2	Trilinear Interpolation and Calculation of the Electric Potential . . . . .	78
5.3.3	Short Range Field Effects . . . . .	80
5.3.4	Removing the Self-Action . . . . .	81
5.3.5	Parallelization . . . . .	82
5.3.6	Tunnelling . . . . .	83
5.4	Simulation . . . . .	84
5.5	Results . . . . .	85
5.6	Conclusion . . . . .	87
<b>6</b>	<b>Conclusions</b>	<b>90</b>
	<b>Bibliography</b>	<b>91</b>

# PREFACE

This thesis is dedicated to the memory of James Hopkins, a friend of many years who helped kindle my fascination for how the universe worked.

Thank you to my family for the support over the years, both moral and financial, without which I would not have been able to make it this far, with special thanks to my loving partner Ashton Waghorn for her patience, understanding, and support. I also wish to thank my supervisor Thomas Brabec for the opportunity and for his guidance, wisdom, trust, and seemingly infinite patience. Additionally, my gratitude goes out to many friends and colleagues, including but not limited to Joel Cox, Chris Racknor, Ali Hatef, Charles Varin, Chris McDonald, Christian Peltz, Gianfranco Orlando, Rhys Emms, Shudipto Amin, Ingmar Schubert, Jeremiah O’Neil, Nakib Protik, and Michael Nesrallah. Thank you everyone for all the assistance learning my way around MicPIC (especially Charles), the conversations, the collaboration, the encouragement, and the good times.

This thesis is composed of mostly original material, except for portions of chapters two and three which have been adapted in part from the paper I authored which was published in the journal *Computer Physics Communications* (G. Bart, C. Peltz, N. Bigaouette, T. Fennel, T. Brabec, C. Varin, *Massively parallel microscopic particle-in-cell*, *Comp. Phys. Comm.* 219 (2017) 269–285). This work was a collaboration between myself and colleagues Charles Varin and Christian Peltz, my supervisor Thomas Brabec, and Thomas Fennel. The majority of the text and figures were written and generated by myself, with additional text and selected figures prepared by CV. CP, TB, and TF provided valuable feedback regarding the manuscript, which would not have been possible without the foundational software developed by CV, NB, and TF.

# LIST OF PUBLICATIONS

## Peer-Reviewed Publications

- [1] M. Nesrallah, A. Hakami, G. Bart, C. R. McDonald, C. Varin, and T. Brabec, “Measuring the Kerr nonlinearity via seeded Kerr instability amplification: conceptual analysis”, *Opt. Express* 26, 7646-7654 (2018)
- [2] M. Nesrallah, G. Vampa, G. Bart, P. B. Corkum, C. R. McDonald, and T. Brabec, “Theory of Kerr instability amplification”, *Optica* 5, 271-278 (2018)
- [3] C. Varin, R. Emms, G. Bart, T. Fennel, T. Brabec, “Explicit formulation of second and third order optical nonlinearity in the FDTD framework”, *Comp. Phys. Comm.* 222 (2017) 70–83
- \*[4] G. Bart, C. Peltz, N. Bigaouette, T. Fennel, T. Brabec, C. Varin, “Massively parallel microscopic particle-in-cell”, *Comp. Phys. Comm.* 219 (2017) 269–285
- [5] C. Varin, G. Bart, R. Emms, T. Brabec, “Saturable Lorentz model for fully explicit three-dimensional modelling of nonlinear optics”, *Optics Express* 23 (2015) 2686-2695

## Presentations

- [6] G. Bart, “MicPIC: Modeling Laser-Driven Plasma Dynamics with HPC”, HPCS 2015, Montreal, QC, Canada
- [7] G. Bart, “Modelling Optical Processes with MicPIC”, ICN 2013, London, ON, Canada

# LIST OF FIGURES

1.1	Graphical representation of the various plasma parameters and how they are related. The lines of constant parameters are shown in blue, red, and green for the quantum degeneracy parameter, the classical coupling parameter, and the Debye number, respectively. Dashed lines indicate points where the relevant energy scale is comparable to the rest mass of the electron. $\Theta > 1$ above the blue line, while $\Gamma$ increases above the red line, denoting the strong coupling regime. Weak coupling occurs for conditions well below the red line ( $\Gamma \ll 1$ ) . . . . .	5
1.2	Schematic representation of the three ionization processes discussed in the text. Multiphoton emission is shown in a), where the bound electron absorbs multiple photons from the laser field and is liberated from its parent atom. The case where the electron absorbs more photons than is necessary to escape the atom or ion is above-threshold ionization (b). Lastly, if the laser field is strong enough, it can lower the atomic potential seen by the electron to a point where electronic probability begins to leak into the surroundings, increasing the likelihood of tunnel ionization (c). . . . .	10
2.1	Decomposition of the effective inter-particle force in MicPIC for typical interpolation and correction parameters ( $n_{pts} = 7$ , $r_{cut} = 3w_{PIC}$ , and $w_{PIC} = 1.15\Delta x$ ). Curves correspond to the PIC force $\vec{f}_i^{PIC}$ in blue [Eq. (2.5)], the Mic force $\vec{f}_i^{Mic}$ [Eq. (2.9)] in green, and the total MicPIC force $\vec{f}_i = \vec{f}_i^{Mic} + \vec{f}_i^{PIC}$ [Eq. (2.6)] in red. In particular, it is seen that the Mic force is non-zero only within a few-cell distance from the origin. The fact that $\vec{f}_i^{Mic}(r_{cut}) \simeq 0$ ensures a smooth, seamless transition between the short and long range forces calculated by MD (white area) and PIC (shaded grey area), respectively. All forces are in units of $e^2/4\pi\epsilon_0$ and the inter-particle distance is in cell units ( $\Delta x$ ). For the demonstration, the effective particle size $w_0$ was set to $0.5\Delta x$ , but we emphasize that it can take any value, depending on the problem, with short-range contribution increasing in strength with decreasing $w_0$ . In the limit $w_0 \rightarrow 0$ , MicPIC converges to the Coulomb force (full black line). . . . .	18
2.2	Schematic representation of the main MicPIC algorithm in terms of its major components. Each of the light blue segments with a solid border is treated in Sections 2.2–2.5. Light blue with a dashed border indicates optional additional main-loop components of MicPIC. Ionization is discussed in Section 2.6, while photoemission is a more detailed topic and is presented in detail in Chapter 5. Diagnostics and file system access are not considered part of the main loop, since they depend on the experiment and are not necessarily performed every step. . . . .	20

2.3	Schematic illustration of the current density (a) and local field (b) weighting in parallel MicPIC. In (a), the charge $q$ located at the position $x$ and moving with velocity $\vec{v}$ generates a current density $\vec{J} = \rho\vec{v}$ which is sampled with a finite number of points equally spaced and matching the electric field locations on the FDTD grid. In (b), the force the charge feels is due to the interpolated values of the $\vec{E}$ and $\vec{H}$ fields. In the example above, the $x$ -component of the force is calculated from $H_x$ , $E_y$ , and $E_z$ values at integer cell indices and $E_x$ , $H_y$ , and $H_z$ values at half-integer cell indices. . . . .	30
2.4	Schematic representation of the energies involved in the impact ionization process. An arbitrary potential due to the external environment $V_{env}$ is shown in red. This potential exerts an effect $\sigma$ on the ionization potential apparent to the electron in combination with the energy required for bring a new electron from the impactor’s current position (circle, dashed vertical) to the continuum $E^*$ . This shift of the ionization potential results in a new effective potential $IP^*$ . In MicPIC, the new electron is given only enough energy ( $IP^*$ ) to propagate from the position of the ion to the point at which the impacting electron triggered its release. The “bare” ionic potential in MicPIC is shown in blue for hydrogen, with the green curve representing the Coulomb potential. . . . .	35
3.1	A schematic illustration of the MicPIC algorithm showing the location of functions related to communication between processes. Execution proceeds from top to bottom. Pink boxes (left side) indicate communication operations and blue boxes represent functions related to computation. Numbered lines indicate when time measurements are taken during a step (see text for details). Short range communications overlap with short range force and boundary corrections (see text below, 3.2.1). Currents (3.2.1) and particles (3.2.1) are sent and received between processes in a similar manner. . . . .	40
3.2	Schematic representation of a typical 2D [XY] domain decomposition, viewed down the $z$ -axis, showing the sub-grids involved in communication of data between neighbouring processes. The main grid of process 14 is bounded by magenta boxes ( $i-iv$ ), which represent current data from neighbours’ particles extending across boundaries. Light green boxes ( $I-IV$ ) are related to cells from which process 14 requires EM field data in order to properly update the weighted local field of particles near one or more boundaries. Particles whose short range correction radius intersects a boundary exchange their position and charge data across those boundaries in order to correct the PIC force of particles in the vicinity of a boundary. Edge buffers $a-d$ and $A-D$ (light blue, dashed) operate on the same principles. . . . .	42
3.3	Graphical depiction of a parallel impact ionization event. An electron (blue) on rank $r_1$ is made aware of the presence of the ion (pink) with charge state $S^+$ on rank $r_2$ , which is the nearest ion. A successful event caused by the electron entering the ionic radius $r_{ion}$ (beige circle) with sufficient kinetic energy triggers the instantiation of <code>Impact_Event</code> object $IE(r_2)$ on the host process of the ion, which facilitates the exchange of information between the two interacting particles. A new electron is created at the position of the ion with a random velocity $\vec{v}_{el,new}$ and kinetic energy equal to the shifted ionization potential discussed in Section 1.1.3. Due to the inelastic nature of the process, total kinetic energy is not conserved, though the impacting electron loses the ion-relative kinetic energy to the release of the new electron. Momentum conservation yields the final velocities of the two original particles. . . . .	46

3.4	Graphical depiction of parallel tunnel event. A successful tunnel event at ion $i_1$ (left, pink) on rank $r_2$ triggers the generation of a new electron (blue). The ramping begins and moves the electron to its release position (dashed arrow) by spatial step size $\delta\vec{x}$ while the two remain linked by <code>Tunnel_Event</code> object $TE(r_2)$ (curved arrows). During the ramping, the electron crosses into the subdomain of rank $r_1$ (middle); meanwhile the ion's memory location can change due to list operations. The two exchange updated identifications via the <code>Tunnel_Event</code> , which accompanies the electron to $r_1$ . Finally, the electron is unlocked (released) at the final ramp position $\Delta\vec{x}$ and an <code>Unlock.Call</code> is sent to the parent ion so that it can participate in future tunnel ionization events. . . . .	47
3.5	Support sketches for theoretical scaling analysis and parallel decomposition scenarios. In (a), a 2D cut of a 3D subdomain showing the geometry of the boundary layer used for inter-process communications. Also shown in (b) and (c), four-process examples of 2D parallelization approaches corresponding to (b) strong and (c) weak scaling situations. In (b), the calculation of a certain physical problem contained in a volume $V = L^3$ with $N$ particles is split in 4 subdomains, each handled by different processes. As a result of the domain decomposition, each process has a fourth of the total workload. In (c), an initial simulation of a volume $V$ containing $N$ particles is duplicated to simulate a total volume four times as big. The two parallel strategies are typically used: (b) to reduce the global computation time and (c) to benefit from the large memory capacities of distributed simulation clusters to study big material volumes. . . . .	48
3.6	Execution times for the most intensive components of parallel MicPIC for 48 process test series over a range of cell sizes and atomic density $\rho = 22\text{nm}^{-3}$ (see text for more details). All times are averaged over all processes. The top subplots show (a) the total execution times and (b) the time spent on communication functions, including synchronization and overhead related to counting, buffering, etc. Long-range force, short range corrections, current density, and field propagation computation times appear in (c). Totals are shown as large black dots, that are fit using the scaling equations found in Sec. 3.3 [see, in particular, Eqs. (3.2) and (3.4)]. The shaded grey areas outlined by a darker grey curve effectively correspond to (c) $f_{\text{comp}} = 1.38\Delta x^2 + 92.5/\Delta x + 140.0/\Delta x^4$ , (b) $f_{\text{comm}} = 2.4 + 0.09\Delta x + 4 \times 10^{-4}\Delta x^3 + 2.5 \times 10^{-4}\Delta x^4 + 18.5/\Delta x^2 + 0.8/\Delta x^3$ , and (a) $f_{\text{tot}} = f_{\text{comp}} + f_{\text{comm}}$ . Optimal values of $\Delta x$ for fastest execution were found with downhill simplex minimization and are shown as orange stars. In (a), including both computation and communications, the optimum is found to be $\Delta x \simeq 0.35 \text{ nm}$ . . . . .	51
3.7	Performance of MicPIC for 48-process test runs with varying cell size for different materials, represented as plasmas with a single free electron per ion. The simulation volume was held constant and particle number fixed according to the specific atomic density associated with gold (Au, $58.85 \text{ nm}^{-3}$ ), silicon (Si, $49.88 \text{ nm}^{-3}$ ), quartz ( $26.38 \text{ nm}^{-3}$ ), and gaseous nitrogen ( $\text{N}_2$ , $0.05 \text{ nm}^{-3}$ ). a) Arithmetic mean execution time in milliseconds for each step. Violin plots show the distribution of times for a step, each of which is the arithmetic mean over all 48 processes. b) Total execution time required to simulate 1.117fs in hours for each material. Note the distinction between a) and b): the former is per step while the latter is the total time required to complete the number of steps needed to simulate the fixed time window. . . . .	52

3.8	Execution times of the main components of the parallel MicPIC algorithm vs number of processes for a test system of $1.04 \times 10^8$ particles. Each test is shown as a violin plot of the distribution of median execution time of a step across all processes. (a) Particle mover time. (b) FDTD field propagation. (c) Top to bottom: short range force, long range force, and current density. (d-h) Respectively, particle, FDTD, current, short range, and long range communication times. All times are in milliseconds. . . . .	56
3.9	Parallel speedup (a) and efficiency (b) of Parallel MicPIC for a $20\text{nm} \times 138\text{nm} \times 200\text{nm}$ system of $\sim 10^8$ particles. In (b), as the size of the subdomains becomes smaller with increasing number of processes, there is a corresponding drop in the parallel efficiency caused by an increase in the relative fraction of the execution times of the parallel functions (c). Solid lines correspond to Eq. (3.10). (d) and (e) show the largest mean persistent allocations of memory for computation and communication. Some process counts were tested for multiple domain decompositions (max of 4), which are summarized as bars (extrema of the samples) and hollow markers. In a), b), d), and e) markers correspond to the arithmetic means of the speedup and efficiency values, whereas in c) they correspond to the geometric means of the fractions of computation and communication time. This mean is suggested as a way to summarize ratios [1]. . . . .	57
3.10	Execution times of the main components of the parallel MicPIC algorithm vs number of processes for a test system of $5.02 \times 10^9$ particles. Each test is shown as a violin plot of the distribution of median execution time of a step across all processes. The medians of the distributions shown are visible as ticks within the violins. (a) Particle mover time. (b) FDTD field propagation. (c) Top to bottom: short range force, long range force, and current density. (d-g) Respectively, particle, current, short range, and field communication times. All times are in milliseconds. . . . .	58
3.11	Parallel speedup (a) and efficiency (b) of Parallel MicPIC for a $52\text{nm} \times 343\text{nm} \times 2.4\mu\text{m}$ system of $5.02 \times 10^9$ particles. Similar to what is seen in Fig. 3.9, both quantities tend to decrease as the sizes of the subdomains become smaller with increasing number of processes. (c) and (d) display the mean persistent allocations of memory on a process for the grid, microscopic corrections, and particles. . . . .	59
3.12	Results from the BlueGene/Q scalability tests. The base-10 logarithm of the execution time in hours is plotted against the base-2 logarithm of the number of processes. The test series begins with $1.4 \times 10^9$ particles on 1024 processes and doubles in size up to $4.5 \times 10^{10}$ particles on 32768 processes. It appears that using twice as many processes to run a simulation twice as big has a minimal impact on the global computation and communications execution time. . . . .	62
4.1	The laser is propagating from the lower left corner in all figures at a $45^\circ$ angle to the upper right. Polarization is in the same plane ( $x$ increases to the right, $y$ increases upward). a) Global view looking down the $z$ -axis (into the page) of the PIC and MicPIC simulations at step 10050 (indicated), at a time when the laser begins exciting surface plasmons. At this step, the pulse has not yet arrived. Electrons have otherwise not travelled more than a few nanometres from the surface and still exhibit mostly collective behaviour. b) Closer view of the sample in the same direction as a), focused on the far left end of the sample. The weaker particle interactions in PIC permit surface plasmons, whereas microscopic effects in MicPIC prevent electrons from reaching high speeds.. See text for discussion. . . . .	67

4.2	The laser is propagating from the lower left corner in all figures at a $45^\circ$ angle to the upper right. Polarization is in the same plane ( $x$ increases to the right, $y$ increases upward). a) Global view looking down the $z$ -axis (into the page) of the PIC and MicPIC simulations at step 24300 (indicated), shortly before the arrival of the peak of the laser pulse. Waves of electrons are visible propagating through the sample and a wake of electrons begins to leave the medium. b) Closer view of the sample in the same direction as a), focused on the far end of the sample. Surface plasmons exhibit similar features in both, while the small amplitude waves propagating within tend to be more distinct in the case of PIC. See text for discussion.	68
4.3	The laser is propagating from the lower left corner in all figures at a $45^\circ$ angle to the upper right. Polarization is in the same plane ( $x$ increases to the right, $y$ increases upward). a) Global view looking down the $z$ -axis of the PIC and MicPIC simulations at step 31800 (indicated), at a time when the laser has traversed the majority of the length of the sample ( $x$ increases to the right, $y$ increases upward). Many electrons have left the medium, forming a wake, and those that are returned have excited numerous short period plasma waves inside the material. b) Closer view of the sample in the same direction as a), focused on the far right end of the sample. Surface plasmons reflect off the discontinuity at the boundary in $x$ . Small amplitude waves are more visible in the case of PIC, where the attenuation of short period plasma waves is less pronounced. See text for discussion. . . . .	69
4.4	The laser is propagating from the lower left corner in all figures at a $45^\circ$ angle to the upper right. Polarization is in the same plane ( $x$ increases to the right, $y$ increases upward). a) Global view looking down the $z$ -axis of the PIC and MicPIC simulations at step 40050 (indicated), just before the peak of the laser pulse passes beyond the end of the sample ( $x$ increases to the right, $y$ increases upward). In both simulations, many fast electrons escape the medium, either due to field enhancement at the end of the sample or by accelerating in the wake of the laser. b) Closer view of the sample in the same direction as a), focused on the far right end of the sample. Interfering short period plasma waves are more clear in the PIC case, another example of which is shown to the left of the middle in c). See text for discussion. . . . .	70
4.5	Comparison of the distributions of electron speeds in PIC and MicPIC. Red curves and markers indicate PIC data; MicPIC data are represented in blue. Top: In both cases, magnitudes of electron velocities were sampled into 250 speed bins, shown as markers on the curves, spanning ranges of velocities specific to the data from a given timestep. The data are shown up to $10^7$ m/s, beyond which the particle counts become negligible compared to the peaks and interfere with visibility. Bottom: Plot of only the results from step 10050. . . . .	72
4.6	Comparison of components of electric fields in PIC (a) and MicPIC (b) simulations and the difference between the two (c). Sample is visible in the middle of the subdomain, primarily in plots of $E_y$ . TFSF interfaces are visible as discontinuities in field components near the peripheries. . . . .	74
4.7	The harmonic spectrum for the electric field taken at observation points on the incident and transmitted sides of the sample. In this case, the probe point is taken at a location 373 nm from the left end of the sample and approximately 182 nm from either side in the $y$ -direction.	75
4.8	The harmonic spectrum for the electric field taken at observation points on the incident and transmitted sides of the sample. In this case, the probe point is taken at a location 120 nm from the left end of the sample and approximately 182 nm from either side in the $y$ -direction.	75

5.1	Schematic depiction of a parallel photoemission event. Left: electron 1 approaches the bounding surface of some material (shaded region, left). Components of velocity and electric field in the direction normal to the surface determine whether potential difference relative to the classical turning point is calculated (bottom left, actual MicPIC result) until final point $\vec{r}_f$ (red star). Green shaded region corresponds to Eq. (5.15); any point within $r_{cut}$ of other particles (red cross) is corrected via Eq. (5.14). Right: schematic representation of photoemission procedure. Checks are made for location and turning points; path integrations are performed conditionally. Every path starts on its owner (pink region, right). Any path reaching a boundary is passed to workers for further integration, recursively (yellow region). Some paths may be short; these are stored locally (green region) until full paths are reassembled after communication of segments from workers to owners. Tunnel probability calculations are made on owners via Eqs. (1.49) and (1.50). More detail in text. . . . .	82
5.2	Electric field components $E_x$ (a and c) and $E_y$ (b and d) in a slice at $x = 620$ or $5$ nm beyond the end of the tip. a) and c) show the deviation of the field in the polarization direction as the laser intensity increases. b) and d) show similar behaviour in the propagation direction. $y$ and $z$ are increasing to the right and top of the figures, respectively. See text for more detail.	84
5.3	A slice of the $x$ (a) and $z$ (b) components of the electric field bisecting the tip in the propagation direction at step 8000. $x$ and $z$ are increasing to the right and top of the figures, respectively.	86
5.4	A slice of the $x$ (a) and $z$ (b) components of the electric field bisecting the tip in the propagation direction at step 24000. $x$ and $z$ are increasing to the right and top of the figures, respectively.	87
5.5	Slices of the magnitude of the electric field bisecting the tip in the propagation direction at steps 8000 (a), 12000 (b), and 20000 (c). $x$ and $z$ are increasing to the right and top of the figures, respectively. The pattern in the field slowly moves into the interior of the sample and is does not appear to be linked to the carrier frequency. The small distortions near the end of the tip are individual electrons. . . . .	88
5.6	a) Photoemission paths shown exaggerated by a factor of 10 for visibility. Darker green arrows indicate earlier events; conversely, lighter arrows correspond to events at later times. b) Comparison of photoemission times with the incident laser pulse amplitude. The peaks of the photoemission tend to correspond with the peak amplitudes of the corresponding cycle. .	89

---

## CHAPTER

# ONE

---

## INTRODUCTION

The word *plasma* entered the physics lexicon in 1928 by way of Irving Langmuir who, having been inspired by the main component of blood [2], used it to describe a region in which “an ionized gas contains ions and electrons in about equal numbers so that the resultant space charge is very small” [3]. This simple definition encompasses a wide majority of regular matter in the visible universe. Plasmas exist over many orders of magnitude in both temperature and density, ranging from very exotic conditions to the commonplace.

Many naturally occurring plasmas are extraterrestrial in origin and indeed plasma physics has provided insight into many problems in astrophysics. The interstellar and intergalactic media are very low density examples and are nearly ubiquitous. The interiors of giant planets like Jupiter consist of hot, dense plasmas and the electromagnetic fields they generate interact with their surroundings to produce plasma rich environments [4]. From birth to death, stars are another area of study to which the physics of plasmas has been applied to increase our understanding of nature. Their nurseries, interiors, atmospheres, winds, and even remains are all examples of this “fourth state” of matter [5, 6, 7]. Plasmas do indeed exist in more familiar environments as well. Numerous atmospheric phenomena serve as examples of plasma, including lightning and aurorae, and more unusual ones such as ELVES, sprites, and blue jets [8, 9]. Finally, metals can be considered a form of electron plasma, governed by quantum mechanical effects at low temperatures due to Pauli exclusion effects and subject to Fermi-Dirac statistics [10]. Thus, understanding the behaviour of plasmas is understanding the world around us.

The study of the physics of plasmas is evidently a large and varied discipline, a sizable branch of which is laser plasma physics. Laser plasma physics is itself a varied field of study pertaining to the interaction of intense radiation with charge matter and having benefited greatly from and in turn helped to advance the development of laser technology. Recent decades have seen the convergence of nonlinear optics, plasma physics, and materials science into the study of femto- and attosecond science [11, 12, 13].

Continuing advances in experimental techniques and equipment provide the means to generate sub-cycle near-infrared (NIR) and visible frequency pulses at relativistic intensities ( $> 10^{18}$  W/cm<sup>2</sup>), allowing access to study the dynamics of electrons on extremely short time-scales ( $< 10^{-15}$  s) with field strengths comparable to those found within an atom ( $\gtrsim 10^9$  V/m) [14, 15, 13]. Historically, this has developed from mode-locked solid state sources like ruby, Nd:Glass, and Nd:YAG coupled with organic saturable absorbers provided the means to generate picosecond pulses by the end of the 1960s [16]. Over the next two decades, the state of

the art progressed to a point where it was possible to generate intense femtosecond (fs) pulses of NIR and visible light using newer tunable vibronic crystal laser media like Ti:Sapphire [17, 18] and for the first time it became possible to probe chemical bonds and the dissociation of molecules in real time [19, 20]. Further advances provided the means to achieve significantly higher intensities by increasing the quality and power of the laser signal by reducing noise and using new optical components [14, 11, 15]. Finally, in the latter half of 2001 attosecond pulses were generated and measured for the first time via the process of *High-order Harmonic Generation* (HHG) [21].

HHG is a nonlinear phenomenon first observed in the 1970s with nanosecond infrared CO<sub>2</sub> lasers and solid targets [22, 23, 24]. In the succeeding years, experiments with rare gases demonstrated production of harmonic orders beyond the 100th [25, 26]. It has also been observed in dielectrics and metals [27, 28], as well as semiconductor crystals [29, 30] and from nanoantenna arrays [31]. Whereas a femtosecond is the time scale of vibration and rotation of atomic nuclei, attoseconds are the time scale on which electrons move *inside* atoms and ions. This combined with the fact that HHG has been shown to yield up to XUV and even x-ray frequencies, thus offering very fine spatial resolution, has given rise to the possibility of resolving intra-atomic electronic processes both spatially and temporally. Consequently, it has come within the reach of experiment to be able to take snapshots of important chemical and physical processes on the spatial and temporal scale of its electrons using table-top sources [32].

Closely tied to these advances in laser radiation sources are a number of plasma physics phenomena. Perhaps the most successful semi-classical model for HHG in atomic and molecular gases comes from a plasma perspective: the three step model [33, 34], wherein an electron in an atom or solid is able to tunnel through the potential energy barrier formed by the combination of laser and ionic potential. The free electron is accelerated by the laser and can return to the parent ion during the subsequent cycle with a kinetic energy equivalent to many photons at the carrier frequency. However, it is important to note that HHG is not restricted to low density plasmas. Solid density plasmas have been known to generate high harmonics of the incident field since those early experiments with CO<sub>2</sub> lasers [22]. The advent of intense femtosecond sources has pulled back the veil on a rich area of strong field plasma dynamics, where laser-matter interactions occur on timescales orders of magnitude shorter than the thermal expansion of the sample [35]. Harmonics can be generated by fast-moving electron bunches and intense surface plasmons interacting with non-uniform density profiles and material defects on planar surfaces and thin films [36, 37, 38, 35, 39, 40]. Solid state physics and quantum many-body theory provide an alternate theoretical framework for HHG in crystalline materials. In these situations, harmonics are generated by inter- and intra-band dynamics [41, 42, 43], interactions with the crystal lattice or impurities [44], or by resonantly driven localized plasmons, in the case of graphene ribbons [45].

There exist many other examples of ultrafast nonlinear optics and plasma physics of experimental, industrial, and medical utility. Laser wakefield acceleration [46, 47, 48] is a method for generating compact sources of high energy particles (eg. tens of GeV electrons) from dense plasmas with intensities beyond 10<sup>16</sup> W/cm<sup>2</sup>. Additionally, ablation of metallic, dielectric [49, 50] and organic [51, 52] materials [53] using fs-laser pulses creates localized plasmas [54] and has applications in laser micromachining [55, 56], the production of thin films [57], nanostructures [58, 59], waveguides [60], and medicine [61, 62, 63].

The femtosecond and attosecond control of electrons now afforded to researchers is crucial in the field of electron microscopy in order to gain insight into the structure and behaviour of matter [64, 65]. Sources of these photoelectrons are typically metallic nanorods or nanotips [66, 67]. Moreover, these structures also provide a mechanism for HHG [13]. Although metallic, the energy delivered by intense ultrashort pulses very rapidly raises the temperature of a material sample, thereby surpassing the quantum degeneracy (see below)

common in many solid density materials (ie. metals). This permits, to a good approximation, a classical approach to a theoretical treatment of many of these processes. As an additional example, nanoclusters and nanodroplets exhibit metal-like properties and comprise a broad field of laser plasma physics in their own right [68].

Many of the topics discussed so far involve solid density plasmas, where the coupling between individual charges within plays an important role in the dynamics of laser-plasma interactions. Modelling these phenomena is difficult when short range, many-body effects must be taken into account. Often, theories neglect collisional effects [69, 40], while utilizing numerical models developed primarily for the electrostatic or the weak-coupling regime (see below) [47, 70], where collisional effects are handled via statistical models and rate equations [71, 72, 73].

The current lack of methods for modelling both radiation propagation and collisional effects is an impediment towards progress in many of the area mentioned above. This thesis describes progress made in closing the gap between theory and experiment in strong-field laser plasma physics. Until recently, there has been no method for simulating strong field phenomena in strongly coupled plasmas which is able to include short range, collisional dynamics on length scales ranging from ångströms to micrometers with sub-attosecond temporal resolution. Below, this method is discussed, from its theoretical background, to its implementation on a massively parallel processing scale, and its application to a selection of interesting problems in the field of ultrafast nonlinear laser-plasma physics.

The discussion is organized as follows. The remainder of this chapter introduces some of the relevant background theory for understanding strongly coupled plasmas and some of the mechanisms that can be used to produce plasmas. Next, in Chapter 2, the aforementioned method for numerical simulation of strongly coupled plasmas is presented, including specific theoretical and implementation details. Chapter 3 takes the discussion of the MicPIC method to the massively parallel scale, focusing on how this is achieved and profiling how the resulting software performs with test simulations on some large compute clusters. The first two applications of MicPIC to a selection of interesting physical processes are presented in Chapters 4 and 5. Preliminary results in Chapter 4 demonstrate how collisions affect the dynamics of a hot dense plasmas in a long but finite thin film of fused silica interacting with a non-relativistic femtosecond laser pulse. The harmonic spectrum at different locations along the length of the sample is also investigated. Simulations of photoemission from metal nanotips treated as strongly coupled plasmas are presented in Chapter 5, where the photoelectron yield and dynamics are compared with recent experimental results. Finally, overall conclusions are drawn and the summary of this discussion appear in Chapter 6, where some ideas for future work are briefly mentioned.

## 1.1 Basic Principles

This section aims at introducing the reader to the main theoretical concepts behind much of the material presented in later chapters. First, the concept of plasma coupling is introduced in Section 1.1.1 and the plasma parameters used to quantify this coupling are introduced. Following this, a brief introduction to the kinetic theory of plasmas is provided in Section 1.1.2. The various limits required to yield tractable analytic descriptions of plasmas are discussed, with emphasis on those which are important in numerical models. Section 1.1.3 introduces the ionization processes which yield plasmas when matter is exposed to intense radiation. A brief introduction to relevant models for tunnel and impact ionization are presented. Lastly, Section 1.1.4 outlines a semi-classical theoretical description of the tunneling rates of electrons from the

material surfaces. It is worth noting that the choice of units in this discussion is SI (MKS) unless otherwise specified, although some older literature and texts on nonlinear optics and plasma physics (and astrophysics) favour cgs or atomic units.

### 1.1.1 Plasma Parameters and Screening

The simplest plasma to consider is one that is homogeneous, infinite in extent, and composed of charged particles intermingling with an oppositely charged background [10]. This is the one component plasma (OCP). Since plasmas exist over such a wide range of temperatures and densities, it is useful to be able to classify them based on some characteristic parameter(s), which can be used as a guide for choosing the correct theoretical approach (see Fig. 1.1). One can characterize a classical OCP by comparing the average Coulomb binding energy per particle to the thermal energy of the plasma. For a plasma composed of particles with charge  $q$ , with a number density  $n$ , temperature  $T$ , the classical plasma coupling parameter is

$$\Gamma = \frac{q^2}{4\pi\epsilon_0 k_B T a} \quad (1.1)$$

where  $a = (\frac{4}{3}\pi n)^{-1/3}$  is the Wigner-Seitz radius,  $\epsilon_0$  is the permittivity of free space, and  $k_B$  is Boltzmann's constant. The parameter  $a$  can be thought of as a characteristic radius for the components of the plasma. A classical OCP is said to be weakly coupled when  $\Gamma \ll 1$ , while for  $\Gamma \gg 1$  it is said to be strongly coupled. Thus, a solid density plasma ( $n \sim 10^{28} \text{ m}^{-3}$ ) is strongly coupled for temperatures below  $5.8 \times 10^4 \text{ K}$ . Such temperatures can routinely be found in stellar interiors and situations involving fs-laser pulses incident upon dielectrics and metals.

In the case of metals at low temperatures (eg. room temperature) or extreme environments like the interior of a gas giant planet, quantum mechanical effects cannot be ignored. The characteristic parameter for degenerate plasmas is the ratio of the thermal energy  $k_B T$  of the particles to the Fermi energy  $E_F$  of the material:

$$\Theta = \frac{k_B T}{E_F} \quad (1.2)$$

where

$$E_F = \frac{\hbar^2}{2m} (3\pi^2 n)^{2/3} \quad (1.3)$$

Thus, a plasma is said to be degenerate if the parameter  $\Theta$  is less than unity, which can occur for high temperatures or low densities. A plasma with the same density as that in the previous example, then, is degenerate for temperatures below  $\sim 2.0 \times 10^4 \text{ K}$ .

The fact that plasmas are composed of mobile charges has the consequence that the component charges can readily rearrange themselves to screen out external electric fields. This suggests another parameter to characterize the distance over which this screening occurs. Following the approach of Ichimaru [10], consider a point charge  $q_0$  in a uniform OCP, whose potential  $\phi_0$  is given by the Poisson equation:

$$\nabla^2 \phi_0 = -\frac{q}{\epsilon_0} \delta(\vec{r}) \quad (1.4)$$

where  $\delta(\vec{r})$  is the Dirac delta function. This charge perturbs the positions of the other charges  $q$  in the neighbourhood of  $\vec{r}$ , resulting in the mean density variation  $\langle \delta n \rangle$ . This yields the overall potential  $\phi(\vec{r})$ , which satisfies

$$\nabla^2 \phi = -\frac{1}{\epsilon_0} [q_0 \delta(\vec{r}) - q \langle \delta n \rangle] \quad (1.5)$$

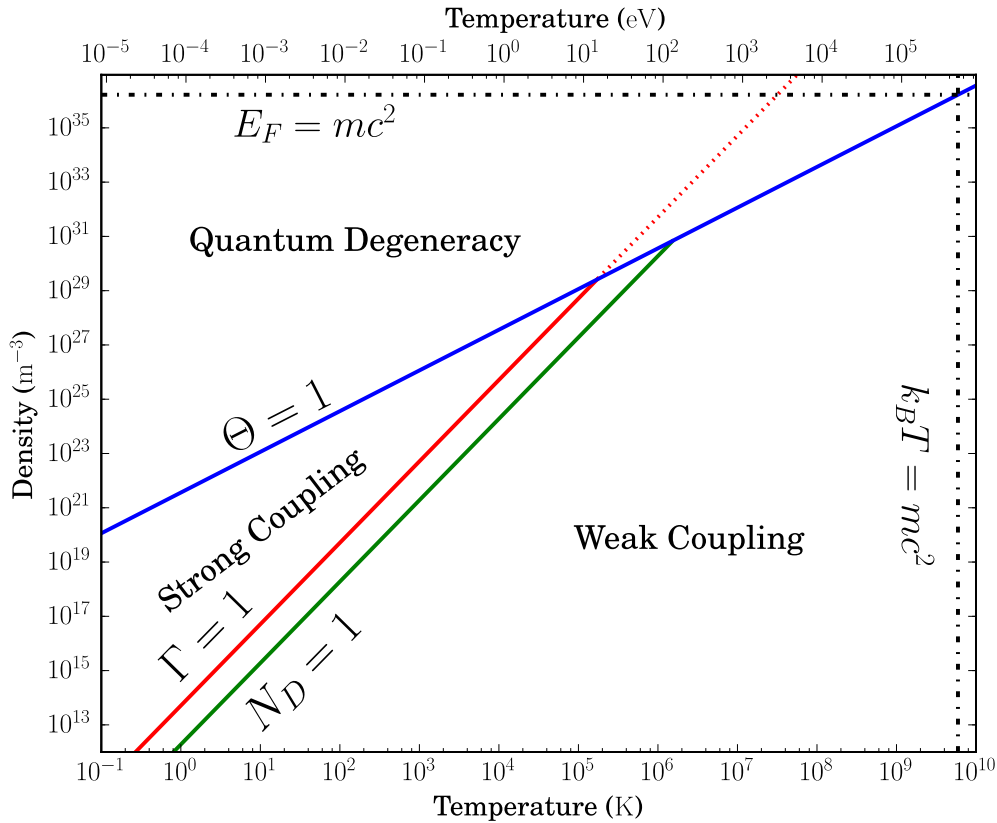


Figure 1.1: Graphical representation of the various plasma parameters and how they are related. The lines of constant parameters are shown in blue, red, and green for the quantum degeneracy parameter, the classical coupling parameter, and the Debye number, respectively. Dashed lines indicate points where the relevant energy scale is comparable to the rest mass of the electron.  $\Theta > 1$  above the blue line, while  $\Gamma$  increases above the red line, denoting the strong coupling regime. Weak coupling occurs for conditions well below the red line ( $\Gamma \ll 1$ )

Provided that the plasma exists in a state of thermal equilibrium, the second term in Eq. 1.5 is expressed in terms of the Boltzmann distribution and the equilibrium number density  $n$  of the OCP:

$$\langle \delta n \rangle = n \left( e^{-\frac{q\phi}{k_B T}} - 1 \right) \quad (1.6)$$

Note that the exponent is the ratio of electrostatic potential energy to thermal energy and is equal to the coupling parameter  $\Gamma$  for  $\|\vec{r}\| = a$ . Then, if the plasma is weakly coupled ( $\Gamma \ll 1$ ), the exponent can be expanded in powers of  $q\phi/k_B T$ , yielding

$$\langle \delta n \rangle \approx -\frac{qn}{k_B T} \phi \quad (1.7)$$

Inserting this result into Eq. 1.5 gives the differential equation

$$\nabla^2 \phi + \frac{qn}{\epsilon_0 k_B T} \phi = -\frac{q_0}{\epsilon_0} \delta(\vec{r}) \quad (1.8)$$

which has a solution of the form

$$\phi(\vec{r}) = \frac{q_0}{r} e^{-r/\lambda_D} \quad (1.9)$$

The characteristic length in the exponential is known as the Debye length and is defined as

$$\lambda_D = \left( \frac{\epsilon_0 k_B T}{q^2 n} \right)^{1/2} \quad (1.10)$$

From this result it is straightforward to calculate the number of plasma particles within a sphere of radius  $\lambda_D$ . This parameter, known as the Debye number or plasma parameter, is a way to quantify how collectively the plasma particles cooperate [74]. Too small a value and the idea of Debye screening ceases to be a rigorously valid concept. Using Eq. 1.10, this parameter is expressed as

$$N_D = \left( \frac{4}{3} \pi n \right) \lambda_D^3 = (3\Gamma)^{-3/2} \quad (1.11)$$

It quickly becomes apparent that the number of particles which cooperate to screen out external fields becomes less than unity for  $\Gamma \gg 1$ . Consider a solid density ( $n = 10^{28} \text{ m}^{-3}$ ) electron plasma at a temperature of  $10^5 \text{ K}$ . Such a plasma will have a Debye length of only  $2.2 \text{ \AA}$  and therefore only 0.435 particles, on average, would be expected within this radius from any given plasma particle.

The ion-sphere model [75] is an intuitive way to characterize strongly coupled plasmas. While it does not provide a succinct parameter akin to the Debye number (Eq. 1.11), it does provide a means to define the internal energy of such a plasma [76]. Here, the negative charge  $-q$  is spread uniformly throughout a spherical volume of radius  $a$ , yielding the charge density

$$\rho_0 = -\frac{3q}{4\pi a^3} \quad (1.12)$$

where  $a$  is again the Wigner-Seitz radius, and the ion is a point charge  $q$  located at  $\vec{r}$ . The radial electrostatic potential  $\phi(r)$  of the negative charge distribution is

$$\begin{aligned} \phi(a) - \phi(r) &= - \int_{\|\vec{r}'\| \leq a} E(r) dr \\ &= \frac{q}{4\pi\epsilon_0 a^3} \int_0^a r dr \end{aligned} \quad (1.13)$$

Noting that the potential at the surface of the sphere is  $\phi(a) = -q/4\pi\epsilon_0 a$  leads, via Eq. 1.13, to an expression for the potential within the radius  $r \leq a$ :

$$\phi(r) = -\frac{q}{8\pi\epsilon_0 a} \left( 3 - \frac{r^2}{a^2} \right) \quad (1.14)$$

The total potential energy is the sum of the potential energy of the ionic charge  $q$  at position  $\vec{r}$ ,  $U_{+q}(\vec{r})$ , and the self-energy of the negatively charged sphere,  $U_{-q}$ . In the case of the former,

$$\begin{aligned} U_{+q}(r) &= q\phi(r) \\ &= -\frac{q^2}{8\pi\epsilon_0 a} \left( 3 - \frac{r^2}{a^2} \right) \end{aligned} \quad (1.15)$$

with the self-energy

$$\begin{aligned} U_{-q} &= \frac{1}{2} \int \rho_0 \phi(\vec{r}) d^3 \vec{r} \\ &= \frac{3q^2}{20\pi\epsilon_0 a} \end{aligned} \quad (1.16)$$

Thus, taking Eq. 1.1 into account, the total potential energy is

$$U_{total} = \left( \frac{r^2}{a^2} - 0.9 \right) \Gamma k_B T \quad (1.17)$$

One can think of this system as a harmonic oscillator, due to the appearance of the  $r^2$  term in Eq. 1.17 [10]. Since the average potential energy in this system is  $\langle U \rangle = \frac{3}{2} k_B T$ , the total internal energy of the plasma becomes

$$\frac{E}{k_B T} = -0.9\Gamma + \frac{3}{2} \quad (1.18)$$

This result demonstrates that for a weakly coupled plasma, the internal energy is approximately that of an ideal gas, whereas a strongly coupled plasma has an internal energy which deviates from this and is very tightly bound.

## 1.1.2 Kinetic Theory of Plasma

Plasmas are dynamic media whose coarse-grained, macroscopic behaviour can be described as arising from the mechanics of large numbers of microscopic particles. Describing such systems exactly, aside from those which are sufficiently small, is an exercise in futility: the numbers of particles to consider is too large. It is necessary to turn to statistical mechanics in order to treat the problems of plasma physics in an analytic fashion.

Those systems for which the degeneracy parameter  $\Theta \lesssim 1$  must invariably incorporate quantum mechanical effects into any model that seeks to characterize them. In a semi-classical sense (where the laser field is treated classically), this amounts to finding a solution to the von Neumann equation

$$\frac{\partial \hat{\rho}}{\partial t} = -\frac{i}{\hbar} \{ \hat{\mathcal{H}}, \hat{\rho} \} \quad (1.19)$$

where  $\hat{\mathcal{H}}$  is the  $N$ -particle Hamiltonian,  $\hat{\rho}$  is the density operator for the system, and  $\{ \hat{A}, \hat{B} \} = \hat{A}\hat{B} - \hat{B}\hat{A}$  defines the commutator of operators  $\hat{A}$  and  $\hat{B}$ .

From statistical mechanics [10, 77], Liouville's theorem states that the phase space probability density  $\rho(\Gamma; t)$  is a conserved quantity, where  $\Gamma = \{ \vec{q}, \vec{p} \}$ , and therefore behaves like an incompressible fluid. That is, at some time  $t + \delta t$ , the probability density function  $\rho$  satisfies the condition  $\rho(\vec{q}', \vec{p}'; t + \delta t) = \rho(\vec{q}, \vec{p}; t)$ . This is equivalent to

$$\begin{aligned} \frac{d\rho}{dt} &= 0 \\ \frac{\partial \rho}{\partial t} + \sum_{i=0}^N \left[ \frac{\partial \rho}{\partial \vec{x}_i} \frac{d\vec{x}_i}{dt} + \frac{\partial \rho}{\partial \vec{p}_i} \frac{d\vec{p}_i}{dt} \right] &= 0 \\ \frac{\partial \rho}{\partial t} + \sum_{i=0}^N \left[ \frac{\partial \rho}{\partial \vec{x}_i} \frac{\partial \mathcal{H}}{\partial \vec{p}_i} - \frac{\partial \rho}{\partial \vec{p}_i} \frac{\partial \mathcal{H}}{\partial \vec{x}_i} \right] &= 0 \\ \frac{\partial \rho}{\partial t} + \{ \rho, \mathcal{H} \} &= 0 \end{aligned} \quad (1.20)$$

where Hamilton's equations have been used to express the time derivatives of  $\vec{x}_i$  and  $\vec{p}_i$  in terms of the  $N$ -particle Hamiltonian and

$$\{ A, B \} = -\{ B, A \} = \sum_{i=1}^N \left[ \frac{\partial A}{\partial \vec{x}_i} \frac{\partial B}{\partial \vec{p}_i} - \frac{\partial B}{\partial \vec{p}_i} \frac{\partial A}{\partial \vec{x}_i} \right] \quad (1.21)$$

is the Poisson bracket of  $A$  and  $B$ . Eq. 1.20 is the Liouville equation for a classical system of particles.

The probability density function  $\rho$  is an ensemble average over an arbitrary number  $d\mathcal{N} = \rho\mathcal{N}d\Gamma$  of individual microstates  $\mu_i$  within the phase space volume element  $d\Gamma$ . Each microstate encompasses an exact description of the configuration of a system of  $N$  point particles via the *Klimontovich distribution function* [78]

$$n(\Gamma; t) = \sum_{i=1}^N \delta(\Gamma - \Gamma_i) \quad (1.22)$$

Inserting Eq. 1.22 into Eq. 1.20 with  $\rho \rightarrow n$  yields

$$\frac{\partial n}{\partial t} + \dot{\vec{x}} \cdot \frac{\partial n}{\partial \vec{x}} + \dot{\vec{p}} \cdot \frac{\partial n}{\partial \vec{p}} = 0 \quad (1.23)$$

where

$$\dot{\vec{p}} = q \left[ \vec{E}_{mic}(\vec{x}) + \vec{E}_{ext}(\vec{x}) + \frac{\vec{p}}{m} \times \left( \vec{B}_{mic}(\vec{x}) + \vec{B}_{ext}(\vec{x}) \right) \right] \quad (1.24)$$

is the Lorentz force, since the system under consideration is a OCP, and  $\{\vec{E}_{mic}, \vec{B}_{mic}\}$  and  $\{\vec{E}_{ext}, \vec{B}_{ext}\}$  are the microscopic and external fields, respectively.

In cases where particle motion is non-relativistic, substituting Eq. 1.22 in Maxwell's equations will yield solutions for the microscopic fields in the electrostatic approximation. This modified expression for the Lorentz force will transform Eq. 1.23 into the *Klimontovich equation* [10, 78]:

$$\left\{ \underbrace{\frac{\partial}{\partial t} + \frac{\vec{p}}{m} \cdot \frac{\partial}{\partial \vec{x}} + q \left[ \vec{E}_{ext} + \frac{\vec{p}}{m} \times \vec{B}_{ext} \right] \cdot \frac{\partial}{\partial \vec{p}}}_{\mathcal{P}(\Gamma)} - \int d\Gamma' \underbrace{\frac{q^2}{4\pi\epsilon_0} \left[ \frac{\partial}{\partial \vec{x}} \frac{1}{\|\vec{x} - \vec{x}'\|} \cdot \frac{\partial}{\partial \vec{p}} \right]}_{\mathcal{V}(\Gamma, \Gamma')} n(\Gamma') \right\} n(\Gamma) = 0 \quad (1.25)$$

This result contains within it an exact description of the plasma. The first term describes the time variation of the distribution function within a fixed  $\Gamma$ -element. The next two terms constituting the operator  $\mathcal{P}(\Gamma)$  are convective terms for the motion of the individual charges from their inertia and due to the action of an external electromagnetic field. Finally, the integral term describes the two-particle Coulomb interactions within the plasma, where  $\mathcal{V}(\Gamma, \Gamma')$  is a two-particle interaction operator.

The goal now becomes finding a way to connect this microscopic picture with a coarse-grained macroscopic one. In order to do this, it is necessary to compute one or more averages over some set of phase space coordinates  $\Gamma'_1, \dots, \Gamma'_m$ . For some macroscopic quantity  $\mathcal{O}$ , which is a function of the coordinates  $\Gamma_1, \dots, \Gamma_n$ , given the probability density function  $\rho$ , the average  $\langle \mathcal{O} \rangle$  is expressed as

$$\langle \mathcal{O}(\Gamma_1, \dots, \Gamma_n) \rangle = \int d\Gamma'_1 \dots d\Gamma'_m \rho(\Gamma'_1, \dots, \Gamma'_m) \mathcal{O}(\Gamma_1, \dots, \Gamma_n) \quad (1.26)$$

The one-particle distribution represents the likelihood of finding any one particle at the phase space point  $\Gamma$ . By extension, the two- and three-particle distribution functions represent the likelihood of finding two and three particles at the phase space coordinates  $(\Gamma_1, \Gamma_2)$  and  $(\Gamma_1, \Gamma_2, \Gamma_3)$ , respectively. These functions can be expressed as [10]

$$f_1(\Gamma_1) = \langle n(\Gamma_1) \rangle \quad (1.27a)$$

$$f_2(\Gamma_1, \Gamma_2) = \langle n(\Gamma_1)n(\Gamma_2) \rangle - \delta(\Gamma_1 - \Gamma_2)f_1(\Gamma_1) \quad (1.27b)$$

$$f_3(\Gamma_1, \Gamma_2, \Gamma_3) = \langle n(\Gamma_1)n(\Gamma_2)n(\Gamma_3) \rangle - \delta(\Gamma_1 - \Gamma_2)f_1(\Gamma_1) - \delta(\Gamma_1 - \Gamma_2)f_2(\Gamma_2, \Gamma_3) - \delta(\Gamma_2 - \Gamma_3)f_2(\Gamma_3, \Gamma_1) - \delta(\Gamma_3 - \Gamma_1)f_2(\Gamma_1, \Gamma_2) \quad (1.27c)$$

Taking note of the fact that the terms in the braces in the Klimontovich equation (Eq. 1.25) will commute with the averaging operation defined in Eq. 1.26, it is possible to express Eq. 1.25 itself in terms of distribution functions. That is,

$$\left[ \frac{\partial}{\partial t} + \mathcal{P}(\Gamma_1) \right] f_1(\Gamma_1) = \int d\Gamma_2 \mathcal{V}(\Gamma_1, \Gamma_2) f_2(\Gamma_1, \Gamma_2) \quad (1.28)$$

This approach is readily extended to include more particles. For a two-particle distribution the operators  $\mathcal{P}$  and  $\mathcal{V}$  are expanded to include both particles, yielding

$$\left[ \frac{\partial}{\partial t} + \mathcal{P}(\Gamma_1) + \mathcal{P}(\Gamma_2) - \mathcal{V}(\Gamma_1, \Gamma_2) - \mathcal{V}(\Gamma_2, \Gamma_1) \right] f_2(\Gamma_1, \Gamma_2) = \int d\Gamma_3 [\mathcal{V}(\Gamma_1, \Gamma_3) + \mathcal{V}(\Gamma_2, \Gamma_3)] f_3(\Gamma_1, \Gamma_2, \Gamma_3) \quad (1.29)$$

Indeed, for the  $s$ -particle distribution function  $f_s(\Gamma_1, \dots, \Gamma_s)$ , one finds for a combined distribution  $\prod_{i=0}^s N(\Gamma_s)$

$$\left[ \frac{\partial}{\partial t} + \sum_{i=1}^s \mathcal{P}_i(\Gamma_i) + \sum_{\substack{i,j=0 \\ i \neq j}}^s \mathcal{V}(\Gamma_i, \Gamma_j) \right] f_s(\Gamma_1, \dots, \Gamma_s) = \int d\Gamma_{s+1} \sum_{i=0}^s \mathcal{V}(\Gamma_i, \Gamma_{s+1}) f_{s+1}(\Gamma_1, \dots, \Gamma_{s+1}) \quad (1.30)$$

The expression in Eq. 1.30 is well-known in statistical mechanics as the *Born-Bogoliubov-Green-Kirkwood-Yvon (BBGKY) hierarchy*, so called because of the dependence of the flow of the  $s$ -particle distribution function on the left-hand side on the possibility of a collision between any one of those particles with any one of the remaining  $N - s$  particles (right-hand side). The  $s$ -particle distribution function  $f_s$  is related probability density function  $\rho$  by the relation [77]

$$f_s(\Gamma_1, \dots, \Gamma_s) = \frac{N!}{(N-s)!} \int \prod_{i=s+1}^N d\Gamma_i \rho(\Gamma_1, \dots, \Gamma_N) \quad (1.31)$$

This hierarchy of equations need not continue arbitrarily. It is sufficient to truncate the it for a number of situations where the higher-order particle interactions are negligible. For example, truncating the BBGKY hierarchy after the second equation, which is equivalent to including only two body collisions, leads to the *Boltzmann equation*:

$$\left[ \frac{\partial}{\partial t} + \mathcal{P}(\Gamma_1) \right] f_1(\Gamma_1) = \int d\Gamma_2 \mathcal{V}(\Gamma_1, \Gamma_2) f_2(\Gamma_1, \Gamma_2) \equiv \left( \frac{\partial f_1}{\partial t} \right)_{coll} \quad (1.32)$$

The term on the right in Eq. 1.32 is a collision term expressed in terms of one-particle distributions, based on the assumption that particles are not correlated before collisions, and a differential scattering cross section dependent upon the two-particle interaction potential  $\mathcal{V}$  [77]. This approximation to the BBGKY hierarchy works well in cases involving plasmas where the electrostatic approximation is valid and the Debye number is small ( $N_D \ll 1$ ). If one excludes collisions entirely from Eq. 1.30, one arrives at the well-known *Vlasov equation*

$$\left[ \frac{\partial}{\partial t} + \mathcal{P}(\Gamma) \right] f_1(\Gamma) = 0 \quad (1.33)$$

This approximation holds in cases where the plasma is dilute and collisions are unlikely, and is valid mainly for weakly-coupled plasmas in the fluid limit [10].




Figure 1.2: Schematic representation of the three ionization processes discussed in the text. Multiphoton emission is shown in a), where the bound electron absorbs multiple photons from the laser field and is liberated from its parent atom. The case where the electron absorbs more photons than is necessary to escape the atom or ion is above-threshold ionization (b). Lastly, if the laser field is strong enough, it can lower the atomic potential seen by the electron to a point where electronic probability begins to leak into the surroundings, increasing the likelihood of tunnel ionization (c).

### 1.1.3 Ionization Processes and High Harmonic Generation in Atoms and Molecules

Illumination of materials with intense laser radiation will ionize the constituent atoms, creating a plasma. The charges in the plasma are subsequently able to move under the action of the applied laser field, resulting in the absorption of energy and rapid heating of the material. The nature of the ionization process depends on laser and material parameters, and is important in determining the precise dynamics of the light-matter interaction.

Early theoretical work on ionization focused on atoms in static electric fields (eg. [79]) but expanded to include periodic fields with the advent of laser technology in the 1960s and the new intensity regimes to which it provided access. Certainly some of the most important contributions to the topic of ionization in laser fields came from Keldysh [80] and Perelomov et al. [81], who put forth a model for the ionization of atoms and within solids (as interband transitions) in terms of tunnelling probabilities. The two main regimes are characterized by the Keldysh parameter

$$\gamma = \frac{\omega_0 \sqrt{2m_e I_p}}{eE_0} \sim \frac{\omega_0 \sqrt{2m_r E_g}}{eE_0} \quad (1.34)$$

where  $I_p$  ( $E_g$ ) is the ionization potential (band gap energy),  $E_0$  is the intensity of the laser field,  $\omega$  is the laser frequency, and  $m_e$  ( $m_r$ ) is the electron mass (reduced electron-hole mass  $m_r^{-1} = (m_e^{-1} + m_h^{-1})^{-1}$ ) in a gas (solid). The tunnel regime is characterized by  $\gamma \ll 1$  or, in other words, when the driving frequency of the field is much smaller than the frequency corresponding to the ionization or band gap energy (the adiabatic approximation). The field of the laser varies slowly and distorts the atomic potential, resulting in a penetrable barrier through which the electron can tunnel. The opposite case,  $\gamma \gg 1$ , is the multiphoton regime, wherein the electron absorbs an integer number  $n$  of photons in order to be ionized. In such cases, the ionization probability exhibits resonances at frequencies corresponding to integer numbers of photons, while the ionization rate scales with the intensity as  $I^n$ . Further efforts extended and revised the theory of Keldysh to include more complex atoms, ions, and molecules, as well as higher order terms [82, 83].

In particular, the theory of Ammosov, Delone, and Krainov (ADK) provides analytic expressions for the tunnelling rates of an atom or ion in an oscillating field in an arbitrary state. For an atom or ion with a charge state  $Z$  and orbital and magnetic quantum numbers  $l$  and  $m$ , respectively, the ADK tunnel rate is, in atomic units [83],

$$w_{ADK} = C_{n^*l^*} A_{lm} E_{IP} \left( \frac{2\kappa^3}{E} \right)^{2n^* - |m| - 1} e^{-\frac{2\kappa^3}{3E}} \quad (1.35)$$

where

$$C_{n^*l^*}^2 = \frac{2^{2n^*}}{n^* \Gamma(n^* + l^* + 1) \Gamma(n^* - l^*)} \quad (1.36a)$$

$$A_{lm} = \frac{(2l+1)(l+|m|)!}{2^{|m|} (|m|)! (l-|m|)!} \quad (1.36b)$$

$n^* = Z/\kappa$  and  $l^* = n^* - 1$  are effective quantum numbers,  $E_{IP}$  is the ionization potential of the state, and  $E$  is the applied electric field. The expression in Eq. (1.35) is very nonlinear and depends strongly on the relationship between ionization potential, laser electric field, and charge state of the ion.

In a strongly coupled plasma, free electrons move freely and regularly encounter ions. When they do, it becomes possible for the incoming electron to scatter inelastically off the ion, liberating another electron, in a process referred to as *electron impact ionization*. These many-body quantum mechanical interactions are naturally difficult to model ab initio, which has given rise to empirical formulas being proposed to fit experimental scattering cross section values [84]. The scattering cross section of an electron impacting an atom or ion is equal to the sum of the individual cross sections of each electronic shell  $\sigma_{nl}$ :

$$\sigma_{tot} = \sum_{i=0}^n f_i \sigma_{ei} \quad (1.37)$$

where  $\sigma_{ei}$  is the cross section of each electron in shell  $i$  whose occupation number is  $f_i$  and  $i$  begins at the outermost shell and counts inward. Lotz ([84]) proposed a parametric expression for the cross section of the state with principal and orbital quantum numbers  $n$  and  $l$ , respectively, of the form

$$\sigma_{nl} = a_{nl} f_{nl} \frac{\ln E_{im}/E_{nl}}{E_{im} E_{nl}} \left[ 1 - b_{nl} e^{-c_{nl} \left( \frac{E_{im}}{E_{nl}} - 1 \right)} \right] \quad (1.38)$$

where  $E_i$  is the ionization potential of the shell,  $E_{im}$  is the kinetic energy of the impacting electron, and  $a_{nl}$ ,  $b_{nl}$ ,  $c_{nl}$  are parameters. This expression can be simplified, when considering only single ionization from the ground state of the ion, to

$$\sigma_{nl} = a_{nl} q_{nl} \frac{\ln E_{im}/E_{nl}}{E_{im} E_{nl}} \sim a_0 q_{nl} \frac{\ln E_{im}/E_{nl}}{E_{im} E_{nl}} \quad (1.39)$$

where  $a_0$  can be determined from experimental data [84]. However, it is important to note that in solid density matter, microscopic variations in the energy landscape can shift the ionization potentials and therefore modify the scattering cross sections. Specifically, the ionization energy  $E_{nl}$  of a state can be adjusted due to the influence of the surrounding plasma environment according to  $E_{nl}^* = E_{nl} - \Delta V_{env}$ , where  $E_{nl}^*$  is an effective ionization potential for the orbital and  $\Delta V_{env}$  is a shift in the potential energy from the environment, provided the latter is known at the time of impact [85].

High intensity laser radiation incident upon atomic systems is capable of triggering a multiphoton absorption process in which the electron absorbs more photons than is required to surpass the ionization potential of its host atom. This is *above-threshold ionization (ATI)* [86], first observed in the 1970s by measuring the energy spectrum of photoelectrons from atomic and molecular gases and noticing that these electrons had absorbed higher numbers of photons than would be required for ionization [87, 88, 89]. The signature of ATI

is the appearance of multiple peaks in the electron energy spectrum at  $E_p = (m + s)\hbar\omega - I_p$  [90], where  $s$  is the excess number of photons beyond the  $m$  required to overcome the ionization potential  $I_p$ .

An important case to consider is that of an electron liberated by an ionization process returns to its parent ion during the subsequent half-cycle. This is the basis of the three-step or simple-man model for high harmonic generation (HHG) [91, 33, 34]. Here, the electron leaves the atom or molecule by tunnelling through a barrier produced by the suppression of the ionic potential by the laser, or by a multiphoton process. Once free, it is accelerated by the laser field, possibly gaining many multiples of the incident photon energy. Up to this point, the process is reminiscent of the aforementioned MPI or ATI. However, if the electron is returned to its parent ion at some later time when the laser polarization has reversed, recombination can occur. The energy released by this process is the sum of the ionization potential and the energy the electron gained from the driving field; that is,  $n_{max}\hbar\omega = I_p + 3.17U_p$ , where  $n_{max}$  is the highest harmonic order and  $U_p = q^2 E_0^2 / 4m_e\omega^2$  is the ponderomotive energy gained by its excursion outside the ion. The dependence of the ponderomotive energy on the laser intensity and frequency implies that  $U_p$  can become quite large; indeed, harmonic orders of over 100 have been reported for more than twenty years [26, 25].

Solid materials provide an additional avenue for HHG. The routine use of femtosecond laser pulses provides a means to create hot dense plasmas, the electron dynamics of which can evolve on timescales where the ions are essentially static. Nonlinear plasma density waves in solid density plasmas are a promising route to production of HHG [92]. Metallic and dielectric nanostructures, such as silicon nanoantennas [31] and tungsten nanotips [66], are able to utilize strong field enhancement to create electric fields beyond the damage threshold of their component materials. Electrons in these cases are free to move freely within the bulk of the material and either tunnel through potential barrier near the surface of the object or overcome the work function by multiphoton absorption before returning to the material after gaining energy in the following half-cycle of the laser [13].

### 1.1.4 Ionization and Photoemission in Solids

Photoemission of electrons from metal surfaces is analogous to ionization in atomic or molecular gases and research into electron emission in strong fields even predates the laser by several decades. Fowler and Nordheim [93] developed a theory of electron emission in strong static fields “in the new [quantum] mechanics”, noting that Oppenheimer [79] did not pursue the matter. A Keldysh-type theory was put forth later to model tunnelling and multiphoton photoemission from metal surfaces in the presence of an electromagnetic wave [94]. The metal is modelled as a triangular potential of the form

$$V(x) = \begin{cases} -V_0 & x < 0 \\ eE_0x \sin \theta \sin \omega t & x \geq 0 \end{cases} \quad (1.40)$$

where  $V_0$  is equivalent to the work function of the metal,  $E_0$  is the peak laser field,  $\omega$  is the laser frequency, and  $\theta$  is the angle of incidence of the beam. The parameter

$$\gamma = \frac{\omega\sqrt{2m_eV_0}}{eE_0} \quad (1.41)$$

determines the limiting cases for photoemission and is analogous to Eq. 1.34: multiphoton photoemission occurs in the  $\gamma \gg 1$  limit, whereas tunnelling is the primary mode of photoemission for  $\gamma \ll 1$ .

The models of both [93] and [94] utilize the triangular potential, which, although it simplifies the problem of calculating the tunnelling probability, is something of an idealization and neglects the possibility of more

realistic barrier shapes and short range forces from microscopic field variations. More recently, the Fowler-Nordheim theory has been reformulated to clarify some aspects of the mathematics and to include different types of barriers [95].

Taking a semiclassical approach to solving the problem permits the treatment of metal nanostructures and nanoparticles interacting with strong fields as strongly coupled plasmas. The electron plasma temperature is high enough that the plasma is non-degenerate ( $\Theta \gg 1$ ) and the de Broglie wavelength  $\lambda = h/p$  is small compared to variations in the potential barrier. The *Wentzel-Kramers-Brillouin (WKB) approximation* is well-suited to solving the problem of a hot electron impinging upon a potential barrier that is not necessarily triangular.

Consider the Schrödinger equation for an electron with total energy  $E$  interacting with some potential  $V(x)$

$$\psi(x) = -\frac{\hbar^2}{2m_e} \frac{\partial^2 \psi}{\partial x^2} + [V(x) - E] \psi = 0 \quad (1.42)$$

where  $V(x)$  is subject to the condition that it changes slowly with position:

$$\frac{\partial V}{\partial x} \ll 2\lambda^{-1} [V(x) - E] \quad (1.43)$$

The one-dimensional plane wave solution for the wave function  $\psi = e^{\frac{i}{\hbar} p x}$  suggests that there exists some approximate solution  $\psi$  in the region containing  $V(x)$  of the form

$$\psi = e^{\frac{i}{\hbar} S(x)} \quad (1.44)$$

where

$$S(x) = S_0(x) + \hbar S_1(x) + \frac{\hbar^2}{2} S_2(x) + \dots \quad (1.45)$$

is a complex power series expansion in powers of  $\hbar$ . Inserting this approximate solution into the Schrödinger equation yields, to second order in  $\hbar$ ,

$$-\frac{\hbar^2}{2m_e} \frac{\partial^2 \psi}{\partial x^2} + [V(x) - E] \psi = 0 \quad (1.46a)$$

$$\left( \frac{\partial S_0}{\partial x} \right)^2 + [V(x) - E] + \hbar \left( 2 \frac{\partial S_0}{\partial x} \frac{\partial S_1}{\partial x} - i \frac{\partial^2 S_0}{\partial x^2} \right) + \hbar^2 \left[ \left( \frac{\partial S_1}{\partial x} \right)^2 + \frac{\partial S_0}{\partial x} \frac{\partial S_2}{\partial x} - i \frac{\partial^2 S_1}{\partial x^2} \right] = 0 \quad (1.46b)$$

Expressions for each term in Eq. 1.45 can be found by requiring that the term(s) in powers of  $\hbar$  in Eq. 1.46 be independently zero. This condition results in the solutions

$$S_0 = \pm \int \sqrt{2m_e[V(x) - E]} dx \quad (1.47a)$$

$$S_1 = \frac{i}{2} \ln \frac{\partial S_0}{\partial x} \quad (1.47b)$$

$$S_2 = \frac{m_e}{2} \frac{\partial V}{\partial x} [2m_e(V - E)]^{-3/2} - \frac{m_e^2}{4} \int \left( \frac{\partial V}{\partial x} \right)^2 [2m_e(V - E)]^{-5/2} dx \quad (1.47c)$$

Note that the solution for  $S_2(x)$  in Eq. 1.47c is negligible, provided that the aforementioned condition that the difference  $V(x) - E$  holds and that the spatial variation of the potential is small, since both terms are roughly the same order of magnitude. Substituting these solutions back into Eq. 1.44 via Eq. 1.45 leads

to an expression for the wave function including both forward and backward propagating components:

$$\begin{aligned}
\psi(x) &= \psi_0^{(+)} e^{i(\frac{1}{\hbar}S_0+S_1)} + \psi_0^{(-)} e^{-i(\frac{1}{\hbar}S_0+S_1)} \\
&= \frac{\psi_0^{(+)}}{[2m_e(V(x) - E)]^{1/4}} \exp \left\{ \frac{i}{\hbar} \int \sqrt{2m_e [V(x) - E]} dx \right\} \\
&\quad + \frac{\psi_0^{(-)}}{[2m_e(VU(x) - E)]^{1/4}} \exp \left\{ -\frac{i}{\hbar} \int \sqrt{2m_e [V(x) - E]} dx \right\}
\end{aligned} \tag{1.48}$$

Suppose now that an electron in a metal is incident from the left (region **I**) upon a smoothly varying potential barrier, which is sufficiently thick to allow the WKB approximation to hold within the classically forbidden region (region **II**; a few nm), and has some probability of tunnelling into the space on the other side (region **III**). The wave function decays exponentially in region **II**, where  $S(x)$  becomes complex, while the oscillating wave function in region **III** is restricted to a right-propagating component only. It can be shown that by connecting the expressions for the wave function in the three regions, the transmission coefficient takes the form [95, 96]

$$T = P e^{-G} \tag{1.49}$$

where  $P$  is a scaling factor typically on the order of unity and  $G$  is the WKB exponent, defined as

$$G = g_e \int_{x_b}^{x_a} M^{1/2}(x) dx \tag{1.50}$$

The factor  $g_e = \frac{2}{\hbar} \sqrt{2m_e}$  preceding the integral is known as the WKB factor and the limits of integration  $x_a$  and  $x_b$  are the classical turning points of the particle, between which the wavefunction is evanescent. The function in the integrand is the difference between the potential barrier and the total energy of the electron,  $V(x) - E$ .

---

CHAPTER

TWO

---

# MICROSCOPIC PARTICLE-IN-CELL

## 2.1 Introduction

The interaction of light with solid matter at high intensities produces strongly coupled plasmas which have highly correlated nonlinear dynamics. Understanding these physical processes is interesting from a theoretical [49, 39] standpoint and relevant for technological applications, including laser micromachining [60, 55], the realization of new short wavelength radiation [97, 36] and electron sources [48, 98], and in surgical procedures [61, 62]. However, obtaining a theoretical description of these complex processes at the microscopic level is difficult due to their highly collisional nature.

A truly complete microscopic picture of ultrafast light-matter interactions demands precise knowledge of the dynamics of individual atoms, charge carrier propagation and interference, scattering processes, and wave propagation effects. Such an amalgamation of atomic and plasma physics, solid-state theory, and quantum electrodynamics would present a significant challenge. Fortunately, high intensity femtosecond laser pulses create hot dense plasmas, overcoming the electron degeneracy (Eq. 1.2) and permitting a classical approach.

Numerical modelling is able to provide insight into experiments or test theoretical predictions but is still confronted with the problem of scale. In situations where the objects of interest are much smaller than the laser wavelength, propagation effects can be ignored and the electrostatic approximation is sufficient. Larger-scale simulations must properly account for radiation and typically neglect or approximate microscopic effects, either as a result of weak coupling or of linear material response to the driving field. Typical wavelengths encountered in ultrafast processes range from the extreme ultraviolet (100 nm) to the mid infrared ( $\sim 10 \mu\text{m}$ ), while even a modest volume of  $(100 \text{ nm})^3$  material at solid density ( $10^{27} - 10^{28} \text{ m}^{-3}$ ) contains roughly  $10^7$  particles. The dynamics of such a system is no small challenge from a computational perspective.

Two primary numerical methods exist to model these many-body, ultrafast processes. The first, and simpler of the two, is *Molecular Dynamics (MD)*. MD is restricted to the aforementioned class of problems wherein the electrostatic and dipole approximations are sufficient to describe the particle dynamics. Particles experience a spatially uniform, applied external electric or electromagnetic field, while individually further experiencing the net Coulomb force resulting from every other charged particle present. This method works well, mainly for small noble gas cluster nanoplasmas [99, 85] in the nonrelativistic regime, but requires more

advanced techniques, such as tree algorithms or collision models, to exceed counts of  $10^5$  particles [100]. However, MD becomes prohibitively slow for the larger systems mentioned above, since the force calculation scales as  $\mathcal{O}(N^2)$ , where  $N$  is the number of particles in the simulation, and more advanced optimization algorithms are still restricted to the electrostatic regime.

Larger systems requiring wave propagation effects, can utilize the *Particle-in-Cell (PIC)* approach [70, 101]. PIC sacrifices some detail in order to provide access to length scales significantly larger than those available in MD. Particles are typically treated as superparticles: ensembles of particles exhibiting collective behaviour over timescales which are small compared to the collision time of the plasma particles and length scales which are small compared to the wavelength of the laser. Radiation is propagated by explicitly solving the inhomogeneous Maxwell's equations, where the particles or superparticles act as sources, according to some weighting scheme, and the *Finite Difference Time Domain (FDTD)* method provides the solution to the wave equation on a coarse grid each timestep. The spatial extent of the particles requires the electromagnetic field to be interpolated at each of their positions from this grid in order to provide the Lorentz force and integrate the equations of motion.

The superparticle treatment in PIC is adequate for modelling plasmas that are weakly coupled such that the Vlasov equation (Eq. 1.33) is sufficient to describe the plasma dynamics. This coarse-grained approach smooths out microscopic fluctuations and  $N$ -body effects, which cannot be neglected in strongly coupled plasmas, where collisions are much more frequent. While this can be amended to a certain degree by using statistical models to simulate collisions [73, 102, 103], they are primarily suited to handling cool or low density plasmas, amounting to solving the Boltzmann equation (Eq. 1.32).

### 2.1.1 Microscopic Particle-in-Cell

A marriage of the two methods PIC and MD, *Microscopic Particle-in-Cell (MicPIC)* [104, 105] models long range effects, like wave propagation and retardation, with a PIC approach, while utilizing MD in order to perform the  $N$ -body calculations required to more accurately apply forces at short range. It is inspired by the *particle-particle/particle-mesh (P<sup>3</sup>M)* approach, which interpolates long-range forces on a coarse-grained mesh and later corrects the individual inter-particle forces for separations less than some cutoff value to shape the effective short-range interactions [70]. Developed to solve the Poisson equation, P<sup>3</sup>M applies to situations where the electrostatic approximation is justified and is therefore equivalent to MD but scales linearly with particle number. Generalization to fully relativistic short-range corrections is possible and subject to future research.

Particles in MicPIC correspond to individual atoms and electrons but are represented by Gaussian charge density distributions with identical physical width  $w_0$ , ensuring smooth particle interactions by minimizing numerical noise. Explicitly, the charge density for the  $i$ -th MicPIC particle at position  $\vec{r}_i$  is

$$\rho_i(\vec{r}, w_0) = q_i g(\|\vec{r} - \vec{r}_i\|, w_0) \quad (2.1a)$$

$$g(\vec{r}, w) = \frac{1}{\pi^{3/2} w} \exp\left(-\frac{\vec{r} \cdot \vec{r}}{w^2}\right) \quad (2.1b)$$

where  $q_i$  is its charge,  $g(\vec{r}, w)$  is the normalized Gaussian shape function, and  $w_0$  is the width of the charge density, the last of which is intended to soften the Coulomb singularity associated with a point particle. As emphasized previously (see, e.g., [106, 104, 105, 107]), for charges in close encounters the physical width parameter  $w_0$  emulates the screening of Coulomb interactions associated with the effective shielding of Coulomb singularities by quantum uncertainty and the finite width of the particle wavefunctions. The force this particle

experiences is simply [105]

$$f_i = \int \rho_i \left[ \vec{E} + \vec{v} \times \vec{B} \right] d^3 r_i \quad (2.2)$$

A solution to the integral in Eq. 2.2 for every particle is equivalent to solving the Klimontovich equation (Eq. 1.25) in the limit of vanishing  $w_0$ .

The first level of the MicPIC method is identical to the PIC method. For efficient long-range force calculations, numerical particles are represented on a coarse mesh by a wide charge density distribution of the form  $\rho_i^{(PIC)}(\vec{r}) = q_i g(\vec{r} - \vec{r}_i, w_{PIC})$ , where the width is chosen such that  $w_{PIC} > w_0$ . The motions of all  $N$  particles contribute to the global current density field

$$\vec{j}(\vec{r}, t) = \sum_{i=1}^N \rho_i^{(PIC)}(\vec{r}, t) \vec{v}_i \quad (2.3)$$

where  $v_i$  is the velocity of the  $i$ -th particle. This current density then acts as a source term in the microscopic Maxwell equations. In MicPIC, the meshed electromagnetic fields obey the microscopic Faraday and Maxwell-Ampère equations:

$$\nabla \times \vec{e}^{(PIC)} = -\frac{\partial}{\partial t} \vec{b}^{(PIC)} \quad (2.4a)$$

$$\nabla \times \vec{b}^{(PIC)} = \mu_0 \vec{j}^{(PIC)} + \frac{1}{c^2} \frac{\partial}{\partial t} \vec{e}^{(PIC)} \quad (2.4b)$$

These fields ultimately contribute the the Lorentz force  $f_i$  at the position  $\vec{r}_i$  of each particle according to

$$f_i^{(PIC)} = \int \rho_i^{(PIC)} \left[ \vec{e}^{(PIC)} + \vec{v}_i \times \vec{b}^{(PIC)} \right] d^3 r \quad (2.5)$$

whereby the  $i$ -th particle feels the effect of all the other particles through the electromagnetic field induced by the total particle current  $\vec{j}^{PIC}$ . The numerical integration of Eqs (2.3)-(2.5) is self-consistent and includes important collective wave propagation effects like scattering, interference, retardation, and absorption. Stopping here would yield a PIC method with atomic-scale resolution, capable of fully electro-dynamically dealing with plasmas where  $\Gamma \lesssim 1$ , but would consistently and systematically underestimate the effect of the short range particle interactions for cases where  $\Gamma \gg 1$ .

The approach of MicPIC is to build upon the method outlined above with a second level, corresponding to the ‘‘Mic’’ in MicPIC, which accounts for short-range interactions in the following way (see also [105] for more details). The linearity of Maxwell’s equations permits the decomposition into the PIC force and the residual microscopic correction force  $\vec{f}_i^{(MIC)}$ , intended to account for the short range influence of the neighbours of the  $i$ -th particle.

$$\vec{f}_i = \underbrace{\int \left[ \rho_i \left( \vec{e} + \vec{v}_i \times \vec{b} \right) - \rho_i^{(PIC)} \left( \vec{e}^{(PIC)} + \vec{v}_i \times \vec{b}^{(PIC)} \right) \right] d^3 r}_{\vec{f}_i^{(Mic)}} + \underbrace{\int \rho_i^{(PIC)} \left( \vec{e}^{(PIC)} + \vec{v}_i \times \vec{b}^{(PIC)} \right) d^3 r}_{\vec{f}_i^{(PIC)}} \quad (2.6)$$

The first term written in Eq. (2.6),  $\vec{f}_i^{Mic}$ , emphasizes the numerical correction operation performed in MicPIC, i.e., modifying the coarse, PIC-calculated force  $\vec{f}_i^{PIC}$  with only the missing, underestimated contribution to the theoretically expected, screened force in Eq. (2.2). A numerical example of the MicPIC force decomposition is given in Fig. 2.1.

The expression in Eq. (2.6) is a formal decomposition that does not involve any approximation, and as such, yields the exact total force on the  $i$ -th particle. However, its full numerical implementation is cumbersome. In practice, it has been observed that  $\vec{f}_i^{Mic}$  only contributes in any significant way up to

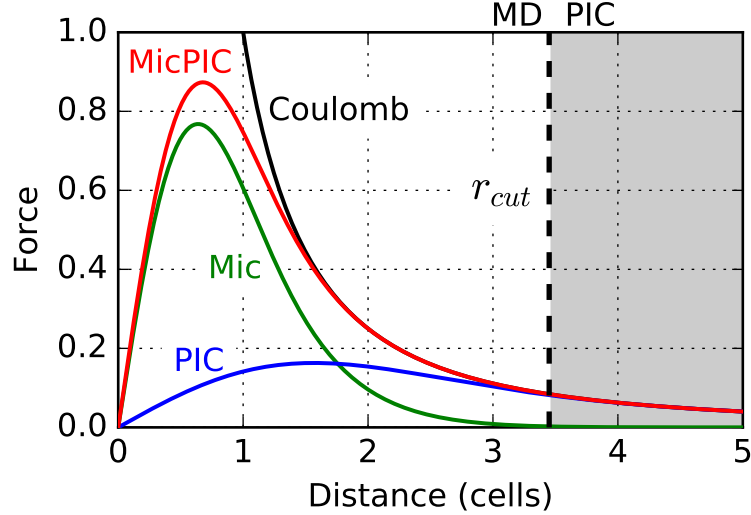


Figure 2.1: Decomposition of the effective inter-particle force in MicPIC for typical interpolation and correction parameters ( $n_{pts} = 7$ ,  $r_{cut} = 3w_{PIC}$ , and  $w_{PIC} = 1.15\Delta x$ ). Curves correspond to the PIC force  $\vec{f}_i^{PIC}$  in blue [Eq. (2.5)], the Mic force  $\vec{f}_i^{Mic}$  [Eq. (2.9)] in green, and the total MicPIC force  $\vec{f}_i = \vec{f}_i^{Mic} + \vec{f}_i^{PIC}$  [Eq. (2.6)] in red. In particular, it is seen that the Mic force is non-zero only within a few-cell distance from the origin. The fact that  $\vec{f}_i^{Mic}(r_{cut}) \simeq 0$  ensures a smooth, seamless transition between the short and long range forces calculated by MD (white area) and PIC (shaded grey area), respectively. All forces are in units of  $e^2/4\pi\epsilon_0$  and the inter-particle distance is in cell units ( $\Delta x$ ). For the demonstration, the effective particle size  $w_0$  was set to  $0.5\Delta x$ , but we emphasize that it can take any value, depending on the problem, with short-range contribution increasing in strength with decreasing  $w_0$ . In the limit  $w_0 \rightarrow 0$ , MicPIC converges to the Coulomb force (full black line).

a few PIC cells away from each individual particle (see Fig. 2.1). Consequently, it is enough to account for microscopic corrections only within a certain cutoff radius  $r_{cut}$ , set to a value large enough such that  $\vec{f}_i^{Mic}(r_{cut}) \ll \vec{f}_i^{PIC}(r_{cut})$ . Effectively, it is observed that  $r_{cut} \approx 3w_{PIC}$  is sufficient for most applications and presents a good balance between computational efficiency and accuracy [105].

The only approximation made formally in the theory of MicPIC is that the short-range interactions within a correction sphere of radius  $r_{cut}$  are electrostatic and instantaneous. With this the microscopic force reduces to

$$\vec{f}_i^{Mic} = -\nabla_{\vec{r}_i} \sum_j \int \int \frac{\rho_i(\vec{r}) \rho_j(\vec{r}') - \rho_i^{PIC}(\vec{r}) \rho_j^{PIC}(\vec{r}')}{4\pi\epsilon_0 \|\vec{r} - \vec{r}'\|} d^3r d^3r'. \quad (2.7)$$

With the Gaussian charge distributions of arbitrary width  $w$  (Eqs. (2.1a) and (2.1b)), integration leads to the interaction potential energy between particles  $i$  and  $j$  of the form

$$V_{ij}(r_{ij}, w) = \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}} \operatorname{erf}\left(\frac{r_{ij}}{w}\right), \quad (2.8)$$

which yields the expression for the microscopic correction to the PIC force

$$\vec{f}_i^{Mic} = -\sum_j \nabla_{\vec{r}_i} [V_{ij}(r_{ij}, w_0) - V_{ij}(r_{ij}, w_{PIC})]. \quad (2.9)$$

where the reader is reminded that  $w_0$  is the effective, physical particle size and  $w_{PIC}$  is its size on the PIC level.

In keeping with the electrostatic approximation employed to derive the analytic forms for the microscopic corrections, MicPIC applies the correction as an adjustment to the electric field experienced by a particle from PIC. Thus, in practice, Eqs. (2.8) and (2.9) become

$$\delta V_{ij}^{(mic)}(r_{ij}) = \frac{q_j}{4\pi\epsilon_0 r_{ij}} \left[ \operatorname{erf}\left(\frac{r_{ij}}{w_0}\right) - \operatorname{erf}\left(\frac{r_{ij}}{w_{\text{PIC}}}\right) \right] \quad (2.10a)$$

$$\begin{aligned} \delta \vec{E}_{ij}(r_{ij}) &= -\nabla \delta V_{ij}(r_{ij}) \\ &= \frac{q_j}{4\pi\epsilon_0} \left\{ \frac{1}{r_{ij}^2} \left[ \operatorname{erf}\left(\frac{r_{ij}}{w_{\text{PIC}}}\right) - \operatorname{erf}\left(\frac{r_{ij}}{w_0}\right) \right] - \frac{2}{\sqrt{\pi} r_{ij}} \left[ \frac{1}{w_{\text{PIC}}} e^{r_{ij}^2/w_{\text{PIC}}^2} - \frac{1}{w_0} e^{r_{ij}^2/w_0^2} \right] \right\} \hat{r}_{ij} \quad (2.10b) \end{aligned}$$

## 2.1.2 The MicPIC Algorithm

The MicPIC algorithm consists of four major components: long range force, short range force correction, particle updates, and field updates (see Fig. 2.2). Execution proceeds in the following manner. A step first involves computation of the long range PIC force (Eq. (2.5)), which is followed by the short range force corrections (Mic force, Eqs. (2.9) and (2.10b)). Next, particle positions and velocities are updated, which allows the current density to be determined. Any particles crossing out of a process' sub-domain and into another are removed from the old and carried over into the new one. The final update is that of the electromagnetic field, propagated via the Finite Difference Time Domain method and using the newly updated current density as a source.

The remainder of this chapter is dedicated to a more detailed discussion of MicPIC and the function of its primary component parts; it is organized in the following manner. Section 2.2 includes an in-depth look into the FDTD algorithm and how it deals with the introduction of external sources (eg. laser) in addition to radiating fields at the simulation domain boundaries. The next section (2.3) focuses on the weighting of the current densities onto the grid and interpolation of the long range PIC force to which each particle is subjected. Section 2.4 deals with the short range corrections to the force, which constitute the Mic stage of MicPIC. The final section includes a brief discussion of the algorithm used to integrate the particle equations of motion and how it is implemented.

## 2.2 Finite Difference Time Domain

The strength of MicPIC comes from its ability to maintain an accurate picture of the dynamics of strongly coupled plasmas interacting with intense electromagnetic fields in a fully self-consistent electrodynamic fashion. In order to implement radiative processes in MicPIC, a numerical implementation of the wave equation is required. The *Finite Difference Time Domain (FDTD)* method is one of, if not the most, popular numerical solution to Maxwell's equations and is suitable for investigating a wide range of electromagnetic problems. The remainder of this section draws from portions of the discussion of the implementation of FDTD found in [108], which is an authoritative text on many aspects of computational electromagnetism.

### 2.2.1 Solving Maxwell's Equations

From a most general standpoint, FDTD solves the wave equation in the time domain via integration of the two curl equations expressing Faraday's Law and the Maxwell-Ampère Law. The only assumption at this

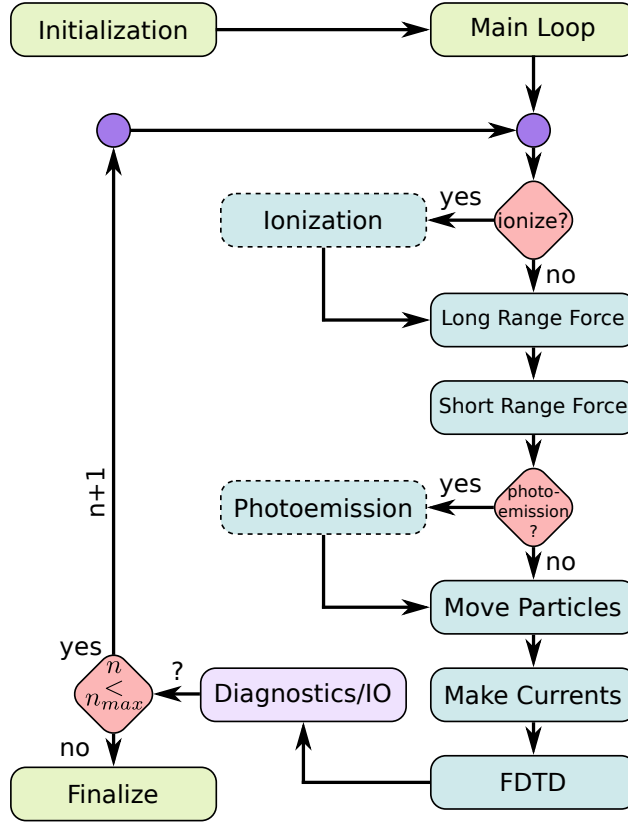


Figure 2.2: Schematic representation of the main MicPIC algorithm in terms of its major components. Each of the light blue segments with a solid border is treated in Sections 2.2–2.5. Light blue with a dashed border indicates optional additional main-loop components of MicPIC. Ionization is discussed in Section 2.6, while photoemission is a more detailed topic and is presented in detail in Chapter 5. Diagnostics and file system access are not considered part of the main loop, since they depend on the experiment and are not necessarily performed every step.

point is that there are no free charges (or magnetic monopoles) and that the electric and magnetic flux density fields  $\vec{E}$  and  $\vec{B}$  are divergenceless. Explicitly,

$$\nabla \cdot \vec{E} = 0 \quad (2.11a)$$

$$\nabla \cdot \vec{B} = 0 \quad (2.11b)$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \quad (2.11c)$$

$$\nabla \times \vec{H} = \vec{J} + \frac{\partial \vec{D}}{\partial t} \quad (2.11d)$$

where  $\vec{D} = \varepsilon_0(1+\chi_e)\vec{E}$  is the induced electric displacement field, the magnetic field  $\vec{H}$  is such that  $\vec{B} = \mu_0(1+\chi_m)\vec{H}$ , and  $\chi_e$  and  $\chi_m$  are the electric and magnetic susceptibilities, respectively. Thus, it is straightforward to simulate bulk responses like dispersion and loss by specifying the appropriate set of empirical material parameters. In MicPIC, the material response is intended to be an emergent behaviour arising from the underlying microscopic particle dynamics and therefore only simulates the vacuum response, where  $\vec{D} = \varepsilon_0\vec{E}$ ,  $\vec{B} = \mu_0\vec{H}$ , and  $\vec{J}$  is due to the particle motions.

Rearranging the curl equations Eqs. 2.11c and 2.11d to solve for the time derivatives and expanding into individual components results in a system of six coupled partial differential equations:

$$\frac{\partial E_x}{\partial t} = \frac{1}{\varepsilon_0} \left[ \frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} - J_x \right] \quad (2.12a)$$

$$\frac{\partial E_y}{\partial t} = \frac{1}{\varepsilon_0} \left[ \frac{\partial H_x}{\partial z} - \frac{\partial H_z}{\partial x} - J_y \right] \quad (2.12b)$$

$$\frac{\partial E_z}{\partial t} = \frac{1}{\varepsilon_0} \left[ \frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} - J_z \right] \quad (2.12c)$$

$$\frac{\partial H_x}{\partial t} = \frac{1}{\mu_0} \left[ \frac{\partial E_y}{\partial z} - \frac{\partial E_z}{\partial y} \right] \quad (2.12d)$$

$$\frac{\partial H_y}{\partial t} = \frac{1}{\mu_0} \left[ \frac{\partial E_z}{\partial x} - \frac{\partial E_x}{\partial z} \right] \quad (2.12e)$$

$$\frac{\partial H_z}{\partial t} = \frac{1}{\mu_0} \left[ \frac{\partial E_x}{\partial y} - \frac{\partial E_y}{\partial x} \right] \quad (2.12f)$$

The numerical approach to Eqs. (2.12a)–(2.12f) is to replace them with a system of finite difference equations, with field values at equally spaced instants in time being sampled at sets of equally spaced points in space, which constitute a *grid* occupying the simulation volume (*domain*). For simplicity, let the value of any given field component sampled at a point on this grid can be represented as  $F_\alpha(i\Delta x, j\Delta y, k\Delta z; n\Delta t) = F_\alpha|_{i,j,k}^n$ , where  $i, j, k$  are grid cell indices and  $n$  is an integer. Furthermore, the derivative of this quantity in time, to second order accuracy, is  $\partial_t F_\alpha|_{i,j,k} = (F_\alpha|_{i,j,k}^{n+1/2} - F_\alpha|_{i,j,k}^{n-1/2})/\Delta t$ . With this notation, the finite difference expressions for Eqs. (2.12a)–(2.12c) become

$$\begin{aligned} \frac{1}{\Delta t} \left[ E_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - E_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} \right] &= \frac{1}{\varepsilon_0} \left\{ \frac{1}{\Delta y} \left[ H_z|_{i,j+1,k+\frac{1}{2}}^n - H_z|_{i,j,k+\frac{1}{2}}^n \right] \right. \\ &\quad \left. - \frac{1}{\Delta z} \left[ H_y|_{i,j+\frac{1}{2},k+1}^n - H_y|_{i,j+\frac{1}{2},k}^n \right] + J_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^n \right\} \end{aligned} \quad (2.13a)$$

$$\begin{aligned} \frac{1}{\Delta t} \left[ E_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - E_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n-\frac{1}{2}} \right] &= \frac{1}{\varepsilon_0} \left\{ \frac{1}{\Delta z} \left[ H_x|_{i+\frac{1}{2},j,k+1}^n - H_x|_{i+\frac{1}{2},j,k}^n \right] \right. \\ &\quad \left. - \frac{1}{\Delta x} \left[ H_z|_{i+1,j,k+\frac{1}{2}}^n - H_z|_{i,j,k+\frac{1}{2}}^n \right] + J_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^n \right\} \end{aligned} \quad (2.13b)$$

$$\begin{aligned} \frac{1}{\Delta t} \left[ E_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - E_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n-\frac{1}{2}} \right] &= \frac{1}{\varepsilon_0} \left\{ \frac{1}{\Delta x} \left[ H_y|_{i+1,j+\frac{1}{2},k}^n - H_y|_{i,j+\frac{1}{2},k}^n \right] \right. \\ &\quad \left. - \frac{1}{\Delta y} \left[ H_x|_{i+\frac{1}{2},j+1,k}^n - H_x|_{i+\frac{1}{2},j,k}^n \right] + J_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^n \right\} \end{aligned} \quad (2.13c)$$

and those for Eqs. (2.12d)–(2.12f) become

$$\frac{1}{\Delta t} \left[ H_x|_{i+\frac{1}{2},j,k}^{n+1} - H_x|_{i+\frac{1}{2},j,k}^n \right] = \frac{1}{\mu_0} \left\{ \frac{1}{\Delta z} \left[ E_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - E_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right] - \frac{1}{\Delta y} \left[ E_z|_{i,j+\frac{1}{2},k+\frac{1}{2}}^n - E_z|_{i,j-\frac{1}{2},k+\frac{1}{2}}^n \right] \right\} \quad (2.14a)$$

$$\frac{1}{\Delta t} \left[ H_y|_{i,j+\frac{1}{2},k}^{n+1} - H_y|_{i,j+\frac{1}{2},k}^n \right] = \frac{1}{\varepsilon_0} \left\{ \frac{1}{\Delta x} \left[ E_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - E_z|_{i-\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} \right] - \frac{1}{\Delta z} \left[ E_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - E_x|_{i,j+\frac{1}{2},k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \right\} \quad (2.14b)$$

$$\frac{1}{\Delta t} \left[ H_z|_{i,j,k+\frac{1}{2}}^{n+1} - H_z|_{i,j,k+\frac{1}{2}}^n \right] = \frac{1}{\varepsilon_0} \left\{ \frac{1}{\Delta y} \left[ E_x|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - E_x|_{i-\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} \right] - \frac{1}{\Delta x} \left[ E_y|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - E_y|_{i,j-\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} \right] \right\} \quad (2.14c)$$

In expressing Eqs. (2.13a)–(2.14c), the electric fields have been evaluated at half-integer time steps in terms of the magnetic field taken at integer time steps and integer or half-integer positions on the FDTD grid, depending on the components involved. Similarly, magnetic field values are taken at integer time steps in terms of the electric field values at integer ones. Such a staggering of the field elements in space and time allows them to be updated, one in terms of the other, in a leapfrog scheme. Moreover, the value of each field component at every point is determined by the value of that same component one time step earlier plus a contribution from the other field, which allows a simulation to simply update pre-existing data. This simple yet powerful algorithm for updating the interleaved  $\vec{E}$  and  $\vec{H}$  fields is attributed to Yee [109] and is the basis of FDTD.

The overall result is a self-consistent, second-order accurate method to solve Maxwell's equations in a vacuum on a grid. In practice, MicPIC uses a uniform, isotropic grid, characterized by the cell size  $\Delta x$  and a time step given by the *Courant condition*:

$$\Delta t = \frac{\Delta x}{\sqrt{3}c} \quad (2.15)$$

The leapfrogging algorithm is numerically stable for all frequency components  $\omega$  propagating with numerical wave vector  $\vec{k}$  which satisfy the numerical dispersion relation [108]

$$\left[ \frac{1}{c\Delta t} \sin\left(\frac{\omega\Delta t}{2}\right) \right]^2 = \left[ \frac{1}{\Delta x} \sin\left(\frac{k_x\Delta x}{2}\right) \right]^2 + \left[ \frac{1}{\Delta y} \sin\left(\frac{k_y\Delta y}{2}\right) \right]^2 + \left[ \frac{1}{\Delta z} \sin\left(\frac{k_z\Delta z}{2}\right) \right]^2 \quad (2.16)$$

## 2.2.2 Introducing Radiation: The Total-Field Scattered-Field Technique

The numerical solution to Maxwell's equations outlined in Section 2.2.1, while complete, would by itself only be useful for simulating a plasma in thermal equilibrium (and not at all for a bulk material). It is necessary, in some way, to introduce an external source of radiation to the simulation to excite any media contained within. It is possible to specify the driving field everywhere, based on some known waveform; this is inefficient. It is also possible to directly apply an external force to the plasma particles in a PIC or MicPIC simulation based on some analytic expression for a known beam waveform and geometry (eg. a Gaussian beam profile); this is cumbersome and evaluating special functions at every particle's position becomes computationally quite expensive for more than a moderate amount of particles. A third option, primarily for the introduction of plane waves to a simulation, is to utilize the *Total-Field Scattered-Field (TFSF)* technique.

The TFSF approach to incident radiation sources in FDTD relies on the linearity of Maxwell's equations. In other words, the total field electric and magnetic fields can be decomposed according to

$$\vec{E}_{tot} = \vec{E}_{src} + \vec{E}_{sca} \quad (2.17a)$$

$$\vec{H}_{tot} = \vec{H}_{src} + \vec{H}_{sca} \quad (2.17b)$$

where the *src*-subscript indicates the incident fields and *sca* indicates the scattered field induced as a response. The linearity also implies that the Yee algorithm applies to each field separately, which permits the simulation volume to be physically decomposed into two regions. The first, the total field region, contains the field solutions for both the incident (source) wave and the scattered fields. The second region only propagates the scattered field solutions and lies external to the total field region. The interface(s) between the two regions require(s), for reasons of consistency, the addition or subtraction of certain field components from the incident wave, which is specified from some known solution on an auxiliary grid.

In the most general case, the total field region is entirely embedded within the scattered field region, where the interface between the two is a closed surface composed of six rectangular faces. Substituting the appropriate components of  $[\vec{E}_{tot}, \vec{H}_{tot}]$  and  $[\vec{E}_{sca}, \vec{H}_{sca}]$  separately into Eqs. (2.13a)–(2.14c) will yield self-consistent solutions to the fields everywhere except the TFSF interfaces. Here, the implementation of a set of conditions is necessary in order to preserve the consistency of the field everywhere. These conditions ultimately stem from Eqs. (2.17a) and (2.17b).

Consider, for the sake of illustration, the one-dimensional example of a wave propagating in the  $z$  direction with electric and magnetic fields reduced to  $\vec{E} \rightarrow E_x$  and  $\vec{H} \rightarrow H_y$ , respectively. At the low  $z$  interface with cell index  $k_0$ , there is an inconsistency in the difference equation for the electric field: for  $k < k_0$  only the scattered field exists. That is,

$$E_{x,tot}|_{k_0}^{n+1} = E_{x,tot}|_{k_0}^n - \frac{\Delta t}{\varepsilon_0 \Delta z} \left[ H_{y,tot}|_{k_0+\frac{1}{2}}^{n+\frac{1}{2}} - H_{y,sca}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}} \right] \quad (2.18)$$

provided there are no sources present. Making note of the relation in Eq. (2.17b), specifically,  $H_{y,tot}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}} = H_{y,src}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}} + H_{y,sca}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}}$  it is clear that Eq. (2.18) would be consistent in the total field if an additional term were added to the right hand side:

$$E_{x,tot}|_{k_0}^{n+1} = E_{x,tot}|_{k_0}^n - \frac{\Delta t}{\varepsilon_0 \Delta z} \left[ H_{y,tot}|_{k_0+\frac{1}{2}}^{n+\frac{1}{2}} - H_{y,sca}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}} \right] + \frac{\Delta t}{\varepsilon_0 \Delta z} H_{y,src}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}} \quad (2.19)$$

This consistency condition assumes the value of  $H_{y,src}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}}$  is known and can be inserted into the field update for  $E_x$ . A similar approach will produce a consistency condition for  $H_{y,tot}|_{k_0-\frac{1}{2}}^{n+\frac{1}{2}}$ .

This principle of these consistency conditions can readily be extended to the case of three dimensions for each of the six possible TFSF interfaces. By adding and subtracting the appropriate values of each of the necessary components on the boundary, the smooth propagation of radiation into and out of the total field region is ensured. In three dimensions, the explicit form of these consistency relations becomes

$$E_y|_{i_0,j,k}^{n+1} = \left\langle E_y|_{i_0,j,k}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{z,src}|_{i_0-\frac{1}{2},j,k}^{n+\frac{1}{2}}, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.20a)$$

$$E_z|_{i_0,j,k}^{n+1} = \left\langle E_z|_{i_0,j,k}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{y,src}|_{i_0-\frac{1}{2},j,k}^{n+\frac{1}{2}}, \quad j_0 \leq j \leq j_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.20b)$$

$$H_y|_{i_0-\frac{1}{2},j,k}^{n+1} = \left\langle H_y|_{i_0-\frac{1}{2},j,k}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{z,src}|_{i_0,j,k}^n, \quad j_0 \leq j \leq j_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.20c)$$

$$H_z|_{i_0-\frac{1}{2},j,k}^{n+\frac{1}{2}} = \left\langle H_z|_{i_0-\frac{1}{2},j,k}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{y,src}|_{i_0,j,k}^n, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.20d)$$

$$E_y|_{i_1,j,k}^{n+1} = \left\langle E_y|_{i_1,j,k}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{z,src}|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}}, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.21a)$$

$$E_z|_{i_1,j,k}^{n+1} = \left\langle E_z|_{i_1,j,k}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{y,src}|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}}, \quad j_0 \leq j \leq j_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.21b)$$

$$H_y|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}} = \left\langle H_y|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{z,src}|_{i_1,j,k}^n, \quad j_0 \leq j \leq j_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.21c)$$

$$H_z|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}} = \left\langle H_z|_{i_1+\frac{1}{2},j,k}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{y,src}|_{i_1,j,k}^n, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.21d)$$

$$E_x|_{i,j_0,k}^{n+1} = \left\langle E_x|_{i,j_0,k}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{z,src}|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}}, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.22a)$$

$$E_z|_{i,j_0,k}^{n+1} = \left\langle E_z|_{i,j_0,k}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{x,src}|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}}, \quad i_0 \leq i \leq i_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.22b)$$

$$H_x|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}} = \left\langle H_x|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{z,src}|_{i,j_0,k}^n, \quad i_0 \leq i \leq i_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.22c)$$

$$H_z|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}} = \left\langle H_z|_{i,j_0-\frac{1}{2},k}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{x,src}|_{i,j_0,k}^n, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.22d)$$

$$E_x|_{i,j_1,k}^{n+1} = \left\langle E_x|_{i,j_1,k}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{z,src}|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}}, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.23a)$$

$$E_z|_{i,j_0,k}^{n+1} = \left\langle E_z|_{i,j_0,k}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{x,src}|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}}, \quad i_0 \leq i \leq i_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.23b)$$

$$H_x|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}} = \left\langle H_x|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{z,src}|_{i,j_1,k}^n, \quad i_0 \leq i \leq i_1, \quad k_0 + \frac{1}{2} \leq k \leq k_1 - \frac{1}{2} \quad (2.23c)$$

$$H_z|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}} = \left\langle H_z|_{i,j_1+\frac{1}{2},k}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{x,src}|_{i,j_1,k}^n, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad k_0 \leq k \leq k_1 \quad (2.23d)$$

$$E_x|_{i,j,k_0}^{n+1} = \left\langle E_x|_{i,j,k_0}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{y,src}|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}}, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad j_0 \leq j \leq j_1 \quad (2.24a)$$

$$E_y|_{i,j,k_0}^{n+1} = \left\langle E_y|_{i,j,k_0}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{x,src}|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}}, \quad i_0 \leq i \leq i_1, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2} \quad (2.24b)$$

$$H_x|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}} = \left\langle H_x|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{y,src}|_{i,j,k_0}^n, \quad i_0 \leq i \leq i_1, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2} \quad (2.24c)$$

$$H_y|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}} = \left\langle H_y|_{i,j,k_0-\frac{1}{2}}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{x,src}|_{i,j,k_0}^n, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad j_0 \leq j \leq j_1 \quad (2.24d)$$

$$E_x|_{i,j,k_1}^{n+1} = \left\langle E_x|_{i,j,k_1}^{n+1} \right\rangle_{inc} - \frac{\Delta t}{\varepsilon_0 \Delta x} H_{y,src}|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}}, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad j_0 \leq j \leq j_1 \quad (2.25a)$$

$$E_y|_{i,j,k_1}^{n+1} = \left\langle E_y|_{i,j,k_1}^{n+1} \right\rangle_{inc} + \frac{\Delta t}{\varepsilon_0 \Delta x} H_{x,src}|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}}, \quad i_0 \leq i \leq i_1, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2} \quad (2.25b)$$

$$H_x|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}} = \left\langle H_x|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}} \right\rangle_{inc} + \frac{\Delta t}{\mu_0 \Delta x} E_{y,src}|_{i,j,k_1}^n, \quad i_0 \leq i \leq i_1, \quad j_0 + \frac{1}{2} \leq j \leq j_1 - \frac{1}{2} \quad (2.25c)$$

$$H_y|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}} = \left\langle H_y|_{i,j,k_1+\frac{1}{2}}^{n+\frac{1}{2}} \right\rangle_{inc} - \frac{\Delta t}{\mu_0 \Delta x} E_{x,src}|_{i,j,k_1}^n, \quad i_0 + \frac{1}{2} \leq i \leq i_1 - \frac{1}{2}, \quad j_0 \leq j \leq j_1 \quad (2.25d)$$

where the terms with  $\langle \dots \rangle_{inc}$  correspond to the inconsistent finite difference equations analogous to Eq. (2.18).

The calculation of the source field is straightforward and is handled separately by a 1D FDTD operating with the same cell size and time step. This auxiliary source field running alongside the main FDTD provides a dynamic look-up table (LUT) from which the source field values can be retrieved. Consider some point of interest located on a TFSF boundary, whose position with respect to the origin is denoted as  $\vec{r}$ . The origin, in this case, is also on the TFSF interface; it is usually chosen to be at some intersection of these faces based on a desired direction of propagation (eg.  $(i_0, j_0, k_0)$ ). For a plane wave with wave vector  $\hat{k}$  to be injected into the simulation via the TFSF technique, as wavefronts propagate a *delay distance*  $d$  through the total field region, they intersect the boundary at a set of points such that  $d = \hat{k} \cdot \vec{r}$ . With this delay distance, the auxiliary field components can be interpolated according to

$$E_{src}|_d^n = (1 - d'_E)E_{src}|_{m_0+\text{int}(d)}^n + d'_E|_{m_0+\text{int}(d)+1} \quad (2.26a)$$

$$H_{src}|_d^{n+\frac{1}{2}} = (1 - d'_H)H_{src}|_{m_0-\frac{1}{2}+\text{int}(d+\frac{1}{2})}^{n+\frac{1}{2}} + d'_H|_{m_0+\frac{1}{2}+\text{int}(d+\frac{1}{2})} \quad (2.26b)$$

where  $d'_E = d - \text{int}(d)$ ,  $d'_H = d + 1/2 - \text{int}(d + 1/2)$ ,  $\text{int}(d)$  returns the largest integer in  $d$ , and  $m_0$  is the "origin" of the 1D auxiliary coordinate system.

In order to supply the interpolated source field values onto the TFSF interface with the correct  $\vec{E}$  and  $\vec{H}$  orientations, it is necessary to specify the direction of the polarization vector of the source  $\hat{E}_{src}$ . The transformation from the 1D auxiliary coordinates to the 3D MicPIC coordinated system, requires a third angle, denoted  $\psi$ , which is determined from the propagation direction  $\hat{k} = (\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta)$  of the source and the direction of polarization via the expression  $\cos \psi = \hat{s} \cdot \hat{E}_{src}$ , where  $\hat{s} = \hat{k} \times \hat{z} / \|\hat{k} \times \hat{z}\|$ . With this final result, the field components interpolated from the 1D source grid can be transformed according to

$$E_{x,src}|_d^n = E_{src}|_d^n (\cos \psi \sin \phi - \sin \psi \cos \theta \cos \phi) \quad (2.27a)$$

$$E_{y,src}|_d^n = -E_{src}|_d^n (\cos \psi \cos \phi + \sin \psi \cos \theta \sin \phi) \quad (2.27b)$$

$$E_{z,src}|_d^n = E_{src}|_d^n (\sin \psi \sin \phi) \quad (2.27c)$$

$$H_{x,src}|_d^{n+\frac{1}{2}} = H_{src}|_d^{n+\frac{1}{2}} (\sin \psi \sin \phi + \cos \psi \cos \theta \cos \phi) \quad (2.27d)$$

$$H_{y,src}|_d^{n+\frac{1}{2}} = H_{src}|_d^{n+\frac{1}{2}} (-\sin \psi \cos \phi + \cos \psi \cos \theta \sin \phi) \quad (2.27e)$$

$$H_{z,src}|_d^{n+\frac{1}{2}} = -H_{src}|_d^{n+\frac{1}{2}} (\cos \psi \sin \theta) \quad (2.27f)$$

This set of equations is the final link between the source and FDTD via the TFSF technique. The one-dimensional source is interpolated and transformed according to position and orientation of propagation and polarization vectors. Each point on the TFSF interface is updated with these sampled data according to Eqs. (2.20a)–(2.25d).

Each source is ultimately represented in MicPIC as an individual 1D FDTD grid, the polarization and propagation directions, and the source function itself, which contains the temporal profile of the wave or pulse. This represents a highly compact and fast method of injecting radiation sources into the simulation, the implementation of which is well-suited to an object-oriented approach. For example, the radiation source is represented as a 176 byte "manager" object in charge of providing sampled source field data to the TFSF functions when queried, evolving the source in time. The source function itself is a polymorphic object that provides the manager with the waveform of the source at each step; for a typical Gaussian pulse this is another 48 bytes. For the 1D grid itself, only 16 bytes per cell are required, so that a source with as many as  $10^4$  cells would altogether only occupy  $\sim 160$  kilobytes. The linearity of Maxwell's equations coupled with this fast and compact numerical radiation source opens the door for including various beam geometries which can themselves be represented as superpositions of plane wave components. For example, it has been

demonstrated that it is possible to inject good quality Gaussian beams into a FDTD simulation with the TFSS technique by superpositioning a set of plane wave sources with a specified set of origins and propagation directions [110].

### 2.2.3 Radiation at the Boundaries: The Uniaxial Perfectly Matched Layer

So far the discussion regarding the FDTD method has focused on propagating radiation and introducing it to a simulation volume, which is, of course, finite in extent. However, the issue of how to deal with radiation at the boundaries of a domain has yet to be addressed. There are, generally speaking, three ways to approach this problem.

The first is to do nothing. The expressions governing the calculation of the electric and magnetic fields in Eqs. (2.13a)–(2.14c) combined with the finite dimensions of the simulation domain have the consequence of having to truncate the leapfrog update. Some of the field components cannot be updated at the boundaries, which results in an incident wave being reflected. This type of boundary condition is the *Perfect Electrical Conductor (PEC)*, because it reflects radiation as though it were a material with infinite conductivity, and is of limited use.

*Periodic Boundary Conditions (PBC)* overcome this limitation by replacing the non-updated cells along a domain boundary with those on the opposing boundary. For example, if  $E_x$  cannot be updated on a  $z = 0$  boundary because  $H_y$  is missing, the application of the PBC will instead use the values of  $H_y$  and  $H_z$  from the  $z = z_{max}$  boundary. The PBC is straightforward to implement on simulations involving a single compute process and treats a simulation as an infinite array of identical copies arranged in the periodic direction. For a more detailed discussion of PBC boundaries including their implementation in parallel process simulations, see Chapter 3.

The rest of the discussion in this section is focused on the third possibility, the *Absorbing Boundary Condition (ABC)*. This type of boundary, ideally, absorbs all incident radiation without reflections, regardless of polarization or frequency, and effectively simulates an extension of the finite domain to infinity. The most widely used implementation of this boundary type is the *Perfectly Matched Layer (PML)* first introduced by Berenger [111]. This method involves placing a non-physical medium adjacent to the domain boundaries, whose material parameters are derived from a split field formulation of Maxwell's curl equations. An improvement upon this mathematically motivated method came with the introduction of the *Uniaxial PML (UMPL)* by Gedney, which is more flexible, efficient, and based directly on Maxwell's equations [112].

Consider a region divided into two half spaces, where  $x < 0$  is a vacuum and a UPML is present in the half space  $x \geq 0$ . The material response in the PML region is described by the tensor quantities  $\overleftrightarrow{\epsilon}_2 = \epsilon_0 \overleftrightarrow{\epsilon}$  and  $\overleftrightarrow{\mu}_2 = \mu_0 \overleftrightarrow{\mu}$ , where

$$\epsilon_{11} = a, \quad \epsilon_{22} = b, \quad \epsilon_{33} = b, \quad \epsilon_{i \neq j} = 0 \quad (2.28a)$$

$$\mu_{11} = c, \quad \mu_{22} = d, \quad \mu_{33} = d, \quad \mu_{i \neq j} = 0 \quad (2.28b)$$

The *uniaxial* qualifier in the name UPML is derived from the rotational symmetry of these tensors about an axis (the  $x$ -axis in this case). Let the incident wave be defined in terms of the harmonic functions

$$\tilde{\vec{H}}_i = \tilde{H}_i e^{i(k_{ix}x + k_{iy}y)} \hat{z} \quad (2.29a)$$

$$\tilde{\vec{E}}_i = \frac{\tilde{H}_i}{\epsilon_0 \omega} e^{i(k_{ix}x + k_{iy}y)} (k_{iy} \hat{x} - k_{ix} \hat{y}) \quad (2.29b)$$

Working now in the PML region, where electromagnetic fields satisfy the transformed curl equations in the frequency-momentum domain (where  $\partial_t \rightarrow -i\omega$  and  $\nabla_\alpha \rightarrow i\vec{k}_\alpha$ ;  $\alpha$  corresponds to either region 1 or 2)

$$\vec{k}_r \times \vec{E}_r = \omega \overleftrightarrow{\mu}_2 \vec{H}_r \quad (2.30a)$$

$$\vec{k}_r \times \vec{H}_r = -\omega \overleftrightarrow{\epsilon}_2 \vec{E}_r \quad (2.30b)$$

the transmitted wave in region 2 takes the form

$$\vec{H}_2 \equiv \vec{H}_t = \tau \tilde{H}_i e^{i(k_{tx}x + k_{ty}y)} \hat{z} \quad (2.31a)$$

$$\vec{E}_2 \equiv \vec{E}_t = \frac{\tau \tilde{H}_i}{\epsilon_2 \omega} e^{i(k_{tx}x + k_{ty}y)} \left( \frac{k_{ty}}{a} \hat{x} - \frac{k_{tx}}{b} \hat{y} \right) \quad (2.31b)$$

where  $H_r = \tau H_i$  and  $\tau$  is the transmission coefficient and is equivalent to the total fields  $[\vec{E}_2, \vec{H}_2]$  in that region. The total fields  $[\vec{E}_1, \vec{H}_1]$  in region 1 become

$$\vec{H}_1 = \tilde{H}_i (1 + \Gamma e^{-2ik_{ix}x}) e^{i(k_{ix}x + k_{iy}y)} \hat{z} \quad (2.32a)$$

$$\vec{E}_1 = \frac{\tilde{H}_i}{\epsilon_0 \omega} [k_{iy} (1 + \Gamma e^{-2ik_{ix}x}) \hat{x} - k_{ix} (1 - \Gamma e^{-2ik_{ix}x}) \hat{y}] e^{i(k_{ix}x + k_{iy}y)} \quad (2.32b)$$

where  $\Gamma$  is the reflection coefficient and the factor  $e^{-2ik_{ix}x}$  arises due to  $k_{rx} = -k_{ix}$ . Ensuring the continuity of the electric and magnetic fields at the interface between the two boundaries leads to expressions for the reflection and transmission coefficients, which are

$$\Gamma = \frac{bk_{ix} - k_{tx}}{bk_{ix} + k_{tx}} \quad (2.33a)$$

$$\tau = \frac{2bk_{ix}}{bk_{ix} + k_{tx}} \quad (2.33b)$$

Furthermore, Eqs. (2.30a) and (2.30b) can be combined to form a wave equation for the electric and magnetic fields. Inserting either Eq. (2.31a) or Eq. (2.31b) into the appropriate wave equation, along with Eqs. (2.28a) and (2.28b), will produce a system of equations that relate the tensor components  $a$ ,  $b$ ,  $c$ , and  $d$  (see [108] or [112] for more detail). Doing this and ensuring phase matching at the interface between the two regions yields

$$k_{tx} = bk_{ix} \quad (2.34a)$$

$$k_{ty} = k_{iy} \quad (2.34b)$$

provided that  $a = b^{-1}$ ,  $c = d^{-1}$ ,  $a = c$  (and the corollary  $b = d$ ), and  $k_i^2 = k_t^2 = \epsilon_0 \mu_0 \omega^2$  at  $x = 0$ .

This result shows that by inserting Eq. (2.34a) into Eq. (2.33a), the reflection of the incident wave at the interface is zero for all incident polarizations, angles, and frequencies. Additionally, the relationship between the tensor components indicates that both the permittivity and permeability of the PML in the  $x \geq 0$  region are multiples of the same uniaxial tensor. That is,  $\overleftrightarrow{\epsilon}_2 = \epsilon_0 \overleftrightarrow{s}$  and  $\overleftrightarrow{\mu}_2 = \mu_0 \overleftrightarrow{s}$ , where  $\overleftrightarrow{s}$  has the components  $s_{11} = s_x^{-1}$ ,  $s_{22} = s_x$ ,  $s_{33} = s_x$ , and  $s_{i \neq j} = 0$ .

Choosing  $s_x = 1 + i\sigma_x/\epsilon_0\omega$ , the expressions for the transmitted wave become

$$\vec{H}_2 = \tilde{H}_i e^{k_{ix}x + k_{iy}y} e^{-\sigma_x x \eta_0 \cos \theta} \hat{z} \quad (2.35a)$$

$$\vec{E}_2 = \tilde{H}_i e^{k_{ix}x + k_{iy}y} e^{-\sigma_x x \eta_0 \cos \theta} (s_x \eta_0 \cos \theta \hat{x} - \eta_0 \cos \theta \hat{y}) \quad (2.35b)$$

where  $\eta_0 = \sqrt{\mu_0/\varepsilon_0}$  is the vacuum impedance and  $\theta$  is the angle the wave vector  $\vec{k}$  makes with the  $x$ -axis. The incident wave will propagate across the boundary, retaining the same phase velocity inside the PML, while exponentially decaying in the  $x$  direction, for any angle, frequency, or polarization.

This procedure can be extended to the other two directions to yield a general three dimensional formulation of a PML layer with the material response tensor  $\overleftarrow{\mathbf{s}} = \overleftarrow{\mathbf{s}}_x \otimes \overleftarrow{\mathbf{s}}_y \otimes \overleftarrow{\mathbf{s}}_z$ , where

$$\overleftarrow{\mathbf{s}} = \begin{bmatrix} s_x^{-1} & 0 & 0 \\ 0 & s_x & 0 \\ 0 & 0 & s_x \end{bmatrix} \otimes \begin{bmatrix} s_y & 0 & 0 \\ 0 & s_y^{-1} & 0 \\ 0 & 0 & s_y \end{bmatrix} \otimes \begin{bmatrix} s_z & 0 & 0 \\ 0 & s_z & 0 \\ 0 & 0 & s_z^{-1} \end{bmatrix} = \begin{bmatrix} s_x^{-1}s_y s_z & 0 & 0 \\ 0 & s_x s_y^{-1} s_z & 0 \\ 0 & 0 & s_x s_y s_z^{-1} \end{bmatrix} \quad (2.36)$$

where  $s_\alpha = 1 + i\sigma_\alpha/\varepsilon_0\omega$  for direction  $\alpha$ . In a PML with a response described by this tensor, the Maxwell curl equations become (in Cartesian coordinates)

$$\frac{\partial \tilde{E}_z}{\partial y} - \frac{\partial \tilde{E}_y}{\partial z} = i\omega\mu_0 \frac{s_y s_z}{s_x} \tilde{H}_x \quad (2.37a) \quad \frac{\partial \tilde{H}_z}{\partial y} - \frac{\partial \tilde{H}_y}{\partial z} = -i\omega\varepsilon_0 \frac{s_y s_z}{s_x} \tilde{E}_x \quad (2.38a)$$

$$\frac{\partial \tilde{E}_x}{\partial z} - \frac{\partial \tilde{E}_z}{\partial x} = i\omega\mu_0 \frac{s_x s_z}{s_y} \tilde{H}_y \quad (2.37b) \quad \frac{\partial \tilde{H}_x}{\partial z} - \frac{\partial \tilde{H}_z}{\partial x} = -i\omega\varepsilon_0 \frac{s_x s_z}{s_y} \tilde{E}_y \quad (2.38b)$$

$$\frac{\partial \tilde{E}_y}{\partial x} - \frac{\partial \tilde{E}_x}{\partial y} = i\omega\mu_0 \frac{s_x s_y}{s_z} \tilde{H}_z \quad (2.37c) \quad \frac{\partial \tilde{H}_y}{\partial x} - \frac{\partial \tilde{H}_x}{\partial y} = -i\omega\varepsilon_0 \frac{s_x s_y}{s_z} \tilde{E}_z \quad (2.38c)$$

Ultimately, transforming Eqs. (2.37a)–(2.38c) to the time domain is required to be compatible with the main FDTD method. Fourier transforming these equations numerically can be avoided, thus improving efficiency and accuracy, by defining a set of component-wise constitutive relations, given by

$$\tilde{D}_x = \varepsilon_0 \frac{s_z}{s_x} \tilde{E}_x \quad \tilde{D}_y = \varepsilon_0 \frac{s_x}{s_y} \tilde{E}_y \quad \tilde{D}_z = \varepsilon_0 \frac{s_y}{s_z} \tilde{E}_z \quad (2.39a)$$

$$\tilde{B}_x = \mu_0 \frac{s_z}{s_x} \tilde{H}_x \quad \tilde{B}_y = \mu_0 \frac{s_x}{s_y} \tilde{H}_y \quad \tilde{B}_z = \mu_0 \frac{s_y}{s_z} \tilde{H}_z \quad (2.39b)$$

For clarity, consider only Eq. (2.38a) and the appropriate constitutive relation from Eq. (2.39a), which, after transforming both to the time domain, become

$$\frac{\partial D_x}{\partial t} + \frac{\sigma_y}{\varepsilon_0} D_x = \frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} \quad (2.40a)$$

$$\frac{\partial D_x}{\partial t} + \frac{\sigma_x}{\varepsilon_0} D_x = \varepsilon_0 \left( \frac{\partial E_x}{\partial t} + \frac{\sigma_x}{\varepsilon_0} E_x \right) \quad (2.40b)$$

Now in a form appropriate for the FDTD method, these expressions can be formulated in finite difference notation. After some algebra, Eqs. (2.40a) and (2.40b) yield equations for updating  $E_x$  and  $D_x$  in a given cell at time step  $n+1$ :

$$E_x|_{i+\frac{1}{2},j,k}^{n+1} = \frac{2\varepsilon_0 - \sigma_z \Delta t}{2\varepsilon_0 + \sigma_z \Delta t} E_x|_{i+\frac{1}{2},j,k}^n + \frac{1}{(2\varepsilon_0 + \sigma_z \Delta t) \varepsilon_0} \quad (2.41a)$$

$$D_x|_{i+\frac{1}{2},j,k}^{n+1} = \frac{2\varepsilon_0 - \sigma_y \Delta t}{2\varepsilon_0 + \sigma_y \Delta t} D_x|_{i+\frac{1}{2},j,k}^n + \frac{2\varepsilon \Delta t}{2\varepsilon_0 + \sigma_y \Delta t} \left[ \frac{1}{\Delta y} \left( H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - H_z|_{i+\frac{1}{2},j-\frac{1}{2},k}^{n+\frac{1}{2}} \right) - \frac{1}{\Delta z} \left( H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - H_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right) \right] \quad (2.41b)$$

Similar procedures will yield expressions for the remaining components of the fields. These equations require a two-step approach in practice, where first  $\vec{D}$  ( $\vec{B}$ ) is updated and the result is used to update the  $\vec{E}$  ( $\vec{H}$ ) field immediately afterwards.

A final point of discussion involves the choice of  $\sigma_\alpha$  terms in the update equations for the PML fields. It is theoretically permissible to simply apply a step in the parameter  $\sigma_\alpha$  at the interface the PML region makes with the vacuum, though, in practice, this can result in spurious reflections due to the discrete nature of the FDTD grid [108]. MicPIC instead uses a polynomial grading scheme where  $\sigma_\alpha$  decreases for each cell deeper inside the PML region. For example, for some  $x$  which increases as it goes deeper:

$$\sigma_x(x) = \sigma_{x,0} \left(\frac{x}{d}\right)^m \quad (2.42)$$

where, in the case of MicPIC,  $\sigma_{x,0} = 0.8(m+1)/\Delta x \eta_0$  and  $m = 3$ . These must be set at the cell or half-cell coordinate matching that of the corresponding field value. In the case of Eq. (2.41a), the value of  $E_x$  is taken at  $i + \frac{1}{2}, j, k$ . Therefore,  $\sigma_z$  must be evaluated at integer  $k$  coordinates, while  $\sigma_x$  is evaluated at half integer  $i$  coordinates.

## 2.3 Current Density and Long Range Force Weighting

The source term for the PIC fields in Eq. (2.4b) require the extended charge densities of the particles to be interpolated onto the FDTD grid via Eqs. (2.13a)–(2.13c). This can be accomplished quite easily with the use of so-called “top-hat” [113] or “tent” [114] functions [101, 70] or a more general b-spline function [115]. The simplest function is the top-hat, where the charge density of the particle is a constant value within some specified, often cell-sized, region of the grid. Although simple in its implementation, top-hat particles have charge densities that suddenly “switch on” to some non-negligible value. This stepping of the charge can result in numerical heating and decreased numerical stability, especially when a large fraction of a fast-moving particle enters a new cell. Higher-order functions circumvent this issue by gradually ramping the charge density, improving accuracy and stability, at the expense of additional calculation. In general, the shape function  $g$  of the charge density is normalized to unity and its values are taken at those points in space corresponding to the appropriate cell index coordinates in Eqs. (2.13a)–(2.14c). Particles with greater spatial extents and smoother shape functions will, at least at first glance, require more computation time in order to compute weighted current density values.

For the PIC computation of the long-range interaction forces in MicPIC, a total current density is built from all the individual current densities that are accumulated on the electromagnetic FDTD grid. As shown at Eqs. (2.1a) and (2.1b), MicPIC particles are modelled by rigid, spherical, and three dimensional Gaussian distributions centred at the particle positions and moving through space at the particle velocities. The Gaussian charge shape functions effectively extend to infinity. However, only a finite number of sampling points  $n_{pts}$  for each spatial dimension can be used [105, 104]. Typically, the number of sampling points is an odd number: one point near the peak of the Gaussian and  $(n_{pts} - 1)/2$  points on each side of the near-peak point. For most applications,  $n_{pts} = 7$  is sufficient, for a total of  $7^3 = 343$  points to sample the entire 3D charge distribution in  $x$ ,  $y$ , and  $z$  [105].

The aforementioned middle, near-peak sampling point is defined by the nearest FDTD grid point (NGP) from the peak of the Gaussian distribution. If the position of the  $l$ -th particle is  $(x_l, y_l, z_l)$ , then the NGP coordinate  $(\bar{x}_l, \bar{y}_l, \bar{z}_l)$  is computed using the standard rounding technique found in [101], e.g., for the  $x$  direction,  $\bar{x}_l = \text{int}(x_l/\Delta x + 1/2)\Delta x$ , where  $\text{int}()$  is a typecast from the C/C++ `double` data type to the `int` type, which returns the largest integer in a floating point number. Then the different components of the current density vector for the  $l$ -th particle at the FDTD grid coordinates  $(x_{l,i}, y_{l,j}, z_{l,k})$  covered by sampling

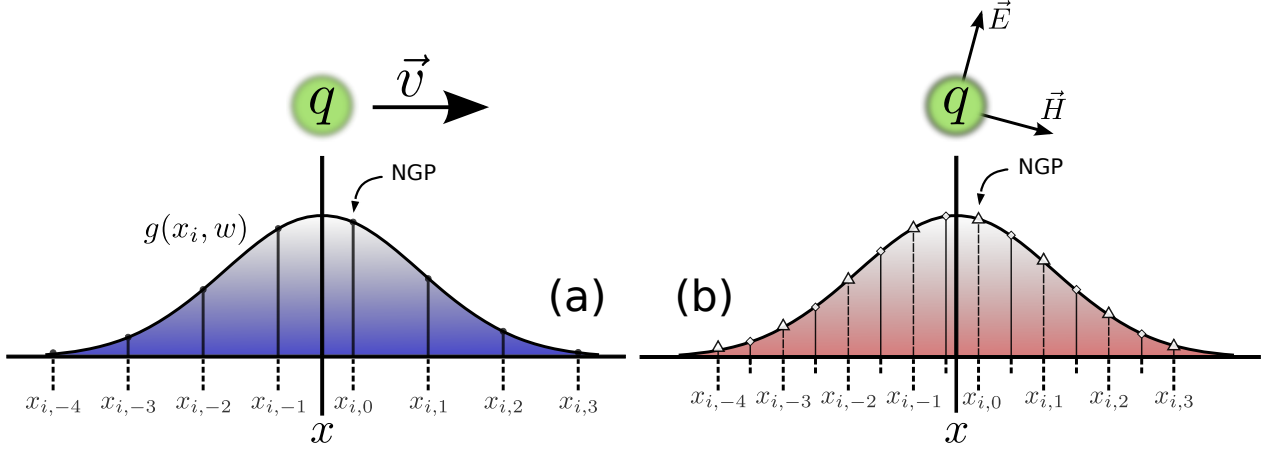


Figure 2.3: Schematic illustration of the current density (a) and local field (b) weighting in parallel MicPIC. In (a), the charge  $q$  located at the position  $x$  and moving with velocity  $\vec{v}$  generates a current density  $\vec{J} = \rho\vec{v}$  which is sampled with a finite number of points equally spaced and matching the electric field locations on the FDTD grid. In (b), the force the charge feels is due to the interpolated values of the  $\vec{E}$  and  $\vec{H}$  fields. In the example above, the  $x$ -component of the force is calculated from  $H_x$ ,  $E_y$ , and  $E_z$  values at integer cell indices and  $E_x$ ,  $H_y$ , and  $H_z$  values at half-integer cell indices.

the Gaussian shape function are defined as

$$J_{l,x}(i, j, k) = q_l v_{l,x} g(\bar{x}_l - x_{l,i} - \Delta x/2) g(\bar{y}_l - y_{l,j}) g(\bar{z}_l - z_{l,k}), \quad (2.43a)$$

$$J_{l,y}(i, j, k) = q_l v_{l,y} g(\bar{x}_l - x_{l,i}) g(\bar{y}_l - y_{l,j} - \Delta x/2) g(\bar{z}_l - z_{l,k}), \quad (2.43b)$$

$$J_{l,z}(i, j, k) = q_l v_{l,z} g(\bar{x}_l - x_{l,i}) g(\bar{y}_l - y_{l,j}) g(\bar{z}_l - z_{l,k} - \Delta x/2). \quad (2.43c)$$

The half step location of each component of the current density is a result of the Yee algorithm used in the FDTD segment of the code. It corresponds to the location of the electric field component in that same direction on the FDTD mesh via the discretization of the Maxwell-Ampère law. 1D current sampling on the FDTD mesh is depicted in Fig. 2.3(a).

The electromagnetic field on the PIC level at a particle's position is interpolated in a similar fashion. Effectively, the field acting on particle  $l$  is the weighted sum of all the field components from the mesh cells in the vicinity of its NGP coordinate  $(\bar{x}_l, \bar{y}_l, \bar{z}_l)$ . The sampling of the weight function again needs to be done at locations that match the Yee-FDTD staggering scheme. Assuming  $n_{pts}$  is odd, and with  $n'_{pts} = (n_{pts} - 1)/2$ , this gives:

$$E_{l,x} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} E_x(i, j, k) g(\bar{x}_l - x_{l,i} - \Delta x/2) g(\bar{y}_l - y_{l,j}) g(\bar{z}_l - z_{l,k}) \quad (2.44a)$$

$$E_{l,y} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} E_y(i, j, k) g(\bar{x}_l - x_{l,i}) g(\bar{y}_l - y_{l,j} - \Delta x/2) g(\bar{z}_l - z_{l,k}) \quad (2.44b)$$

$$E_{l,z} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} E_z(i, j, k) g(\bar{x}_l - x_{l,i}) g(\bar{y}_l - y_{l,j}) g(\bar{z}_l - z_{l,k} - \Delta x/2) \quad (2.44c)$$

and

$$B_{l,x} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} \mu_0 H_x(i,j,k) g(\bar{x}_l - x_{l,i}) g(\bar{y}_l - y_{l,j} - \Delta x/2) g(\bar{z}_l - z_{l,k} - \Delta x/2) \quad (2.45a)$$

$$B_{l,y} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} \mu_0 H_y(i,j,k) g(\bar{x}_l - x_{l,i} - \Delta x/2) g(\bar{y}_l - y_{l,j}) g(\bar{z}_l - z_{l,k} - \Delta x/2) \quad (2.45b)$$

$$B_{l,z} = \sum_{i,j,k=-n'_{pts}}^{n'_{pts}} \mu_0 H_z(i,j,k) g(\bar{x}_l - x_{l,i} - \Delta x/2) g(\bar{y}_l - y_{l,j} - \Delta x/2) g(\bar{z}_l - z_{l,k}) \quad (2.45c)$$

Electromagnetic field interpolation from the FDTD mesh to an individual particle location is depicted in Fig. 2.3(b).

Gaussian charge weighting implementation in MicPIC is not strictly charge conserving. However, the error associated with sampling the Gaussian charge distributions with a finite number of points is limited by using enough sampling points and choosing an appropriate particle width ratio  $w_{PIC}/\Delta x$  (see supplemental material of [106] for details). Using 7 sampling points per dimension with  $w_{PIC} = 1.15\Delta x$ , conserves 99.998% of the total particle charge.

We emphasize that splitting the 3D weight function into a multiplication of three 1D weight functions, i.e.,  $g(x, y, z) = g(x)g(y)g(z)$ , is essential for efficient computation. For example, computing  $g(x, y, z)$  using 7 sampling points per dimension results in  $7^3 = 343$  Gaussian function evaluations, while  $g(x)g(y)g(z)$  requires only  $3 \times 7 = 21$ . Note, however, that this still must be done for each particle, which becomes prohibitive when billions of them are involved. This motivates the use of a look-up table to interpolate continuous argument values of the Gaussian function from a discrete set of precomputed points. During MicPIC's initialization phase, a look-up table for the Gaussian function of unit width at 2001 equally spaced points on the interval  $[-20.0, 20.0]$  is computed and stored. Later on, weights at any point in the table interval can be quickly retrieved with a linear interpolation from the closest points. By doing so, the absolute error is no larger than  $10^{-4}$ . A similar optimization technique is used for the efficient evaluation of the short-range interactions that is described in the following section.

## 2.4 Short Range Force and Microscopic Corrections

The update of the long range PIC force with corrections due to the presence of particles at close range is the key component of MicPIC and is what brings its capacity to simulate plasmas definitively into the strongly coupled regime. MicPIC is distinguished from PIC by the short range corrections performed to the force on each particle to enforce the proper microscopic interactions and collisions without having to rely on statistical models based on binary-collisions theories. Such corrections have been shown to only be required within some specified cutoff radius  $r_{cut} \lesssim 3w_{PIC}$  (see [105] and Fig. 2.1). For particle pairs whose inter-particle distance lies within  $r_{cut}$ , the particle interaction is corrected according to Eq. (2.9). This operation modifies the electric field acting on each of the two particles involved (Eq. (2.10b)), while leaving the magnetic field unchanged. This assumption is valid provided that the particles are non-relativistic, which is justified for most light-matter interaction scenarios in the strongly-coupled regime. The evaluation of the interaction energies as of Eq. (2.8) involves the evaluation of two error functions per particle pair. For efficient computation, MicPIC uses a look-up table where precomputed values for the difference in Eq. (2.10a) are stored. Typically, 2000 sampling points over the range  $[0, r_{cut}/\Delta x]$  was shown to be sufficient. A comparison between the effective

MicPIC and Coulomb forces is shown in Fig. 2.1. While any value can be chosen, MicPIC typically uses a value of  $r_{cut} = 3\Delta x$ .

It is prohibitive to loop through the complete list of particles in order to locate those neighbours within the cutoff radius. Doing so would require, in entirety, computation of  $\binom{N}{2} = \frac{1}{2}N(N-1)$  unique particle pair separations, where  $N$  is the total number of particles. In practical terms, the position of each particle is a continuous quantity and there is no direct correspondence between where a particle is within its memory buffer and where it is in the simulation volume. This motivates a method for rapidly organizing and locating those neighbouring particles which are within  $r_{cut}$  of a particle whose force is being corrected. The cell index method [70, 116] is employed as a way to sort and locate particles in MicPIC quickly and with a relatively low memory footprint. The simulation domain is decomposed into a three dimensional grid, which is congruent with but logically separate from that which is used by the FDTD, which allows the user to balance the load between short- and long-range force calculations. Too large a cell reduces the resolution of the field and produces too many neighbours for a particle, thus slowing the short-range calculations. Conversely, too small a cell leads to excessive load on the field update calculations and very large memory consumption (see Section 3.3).

There are two primary goals when organizing this *cell list*: to traverse the main particle list as little as possible and to have all the information available needed to access particles in a given cell and move to the next one with minimal calculation. At the beginning of a simulation ( $t = 0$ ), a portion of the initialization procedure is to allocate this list, every element of which is a C++ `struct`-type object containing the coordinates of the cell in units of  $\Delta x$ , the number of particles occupying the cell at any given time (the *block*), an offset counter for seeking through these particles, and a C++ pointer to the beginning of the organized block. The cell list exists in a layer of abstraction above the actual list of *sorted particle* objects, which is re-populated every time step, and serves as reference. It then becomes a straightforward matter of looping over the cell list elements. For a cell with coordinates  $(i, j, k)$ , the list of neighbours is composed of a subset of the main list whose cells have coordinates  $\{(i - r_{cut}, j - r_{cut}, k - r_{cut}), \dots, (i, j, k), \dots, (i + r_{cut}, j + r_{cut}, k + r_{cut})\}$  and when the dimensions of the mesh are known, it is a simple matter to compose such a list.

Every time step, the main particle list is traversed and particle positions are converted into cell coordinates and the block size of the corresponding cell is incremented. The buffer intended to contain the sorted particles is partitioned once the size of every block is determined, which simply entails traversing the cell list and setting each cell's sorted particle pointer member to the address of the first free cell following the block of the previous cell. Once the organizational aspects of the cell list are determined, it is necessary to loop through the main particle list a second time. This time, the memory addresses of every particle are stored in the sorted particle list. In this way, locating the neighbours required to perform corrections to the short range force on a particle can be located quickly, despite occupying disordered locations in the main list; all that is necessary is a cell index coordinate.

This leads to an essential point. Performing the short range corrections to all  $N$  particles scales as  $\mathcal{O}(N)$ , since three list traversals are required: twice for the partition and sort and once of the cell list to perform corrections on each particle within, and the operations on each particle are constant in time. Using this approach, correcting all the particles scales as  $\propto NN_c r_{cut}^3$ , where  $N_c$  is the average number of particle per cell. This readily emphasizes the influence of the material density  $\rho$  on MicPIC performance. Thus, for a small enough FDTD mesh spacing  $\Delta x$ ,  $N_c = \rho(\Delta x)^3 \ll N$  and MicPIC's microscopic correction routine scales effectively as  $\mathcal{O}(N)$ .

During the correction phase, the execution moves through the cell list in ascending order and the population of each block is read from memory. If there are particles present, a list of its neighbours is determined

based on  $r_{cut}$  and its coordinates. Next, short range corrections are applied cell-by-cell, particle-by-particle, based on whether the neighbour cell in question has already had its own particles corrected. In practice, the cell of interest is corrected by particles in neighbouring cells that follow it in the list; its own particles contributing only if they follow the particle currently being corrected. The local electric field of each, determined by the weighting scheme outlined in Section 2.3, is adjusted by a value for each binary interaction according to Eq. (2.10b). This logic ensures that the short range corrections are applied only once for any given pair of particles.

## 2.5 Solving the Equations of Motion

As a solution to the classical Liouville equation (Eq. (1.20)), MicPIC must integrate the equations of motion of the particles under the action of the Lorentz force forward in time, where the position of the  $i$ -th particle at step  $n$  is updated to  $n + 1$  according to a simple finite difference scheme

$$\vec{r}_i^{n+1} = \vec{r}_i^n + \dot{\vec{r}}_i^{n+\frac{1}{2}} \Delta t \quad (2.46)$$

Under a simple restoring force or gravity, for example, and this would be a straightforward process. However, the velocity dependence of the Lorentz force via the curl operation in the magnetic force, that is,

$$\frac{\dot{\vec{r}}_i^{n+\frac{1}{2}} - \dot{\vec{r}}_i^{n-\frac{1}{2}}}{\Delta t} = \frac{q_i}{m_i} \left[ \vec{E}(\vec{r}_i^n) + \frac{\dot{\vec{r}}_i^{n+\frac{1}{2}} + \dot{\vec{r}}_i^{n-\frac{1}{2}}}{2} \times \vec{B}(\vec{r}_i^n) \right] \quad (2.47)$$

does not allow one to simply disentangle the vector components by factoring, thus requiring more elaborate solutions to this implicit update equation (eg. [115]). To accomplish the integration in a fast and numerically stable manner, MicPIC employs the de facto standard in plasma physics, the Boris algorithm [117].

Boris conceived of a way to turn the implicit expression in Eq. (2.47) into an explicit form by performing a series of partial updates involving intermediate velocities. Consider expressions for the velocities of the form

$$\dot{\vec{r}}_i^{n+\frac{1}{2}} = \vec{v}_+ + \frac{q_i}{2m_i} \vec{E}(\vec{r}_i^n) \quad (2.48a)$$

$$\dot{\vec{r}}_i^{n-\frac{1}{2}} = \vec{v}_- - \frac{q_i}{2m_i} \vec{E}(\vec{r}_i^n) \quad (2.48b)$$

Inserting these expressions into Eq. (2.47) will cancel the electric field term and yield

$$\vec{v}_+ - \vec{v}_- = (\vec{v}_+ + \vec{v}_-) \times \vec{t}^n \quad (2.49)$$

where  $\vec{t}^n = (q_i/2m_i) \vec{B}(\vec{r}_i^{n+\frac{1}{2}})$ . So far, the new velocity is able to be given in terms of some intermediate velocity plus half the contribution from the electric force, which can in turn be defined in terms of another intermediate velocity and the magnetic field, although still implicitly. In order to disentangle the two intermediate velocities, consider a third velocity given by

$$\vec{v}' = \vec{v}_- + \vec{v}_- \times \vec{t}^n \quad (2.50)$$

and, by extension,

$$\vec{v}' = \vec{v}_+ - \vec{v}_+ \times \vec{t}^n \quad (2.51)$$

Now, these two equations can be summed, yielding a result whose cross product with  $\vec{t}^n$  is

$$2\vec{v}' \times \vec{t}^n = (\vec{v}_+ + \vec{v}_-) \times \vec{t}^n - (\vec{v}_- - \vec{v}_+) \|\vec{t}^n\|^2 + [(\vec{v}_- - \vec{v}_+) \cdot \vec{t}^n] \vec{t}^n \quad (2.52)$$

However, it is clear from Eq. (2.49) that  $\vec{v}_- - \vec{v}_+ \perp \vec{t}^n$ , so the final term in Eq. (2.52) vanishes and the first term becomes the difference  $\vec{v}_+ - \vec{v}_-$ . This ultimately leads to the result

$$\vec{v}_+ = \vec{v}_- + \vec{v}' \times \vec{s}^n \quad (2.53)$$

where

$$\vec{s}^n = \frac{2\vec{t}^n}{1 + \|\vec{t}^n\|^2} \quad (2.54)$$

The final set of updates for the velocity of particles in MicPIC is now complete:

$$\vec{v}_- = \dot{\vec{r}}_i^{n-\frac{1}{2}} + \frac{q_i}{2m_i} \vec{E}(\vec{r}^n) \quad (2.55a)$$

$$\vec{v}_+ = \vec{v}_- + (\vec{v}_- + \vec{v}_- \times \vec{t}^n) \times \vec{s}^n \quad (2.55b)$$

$$\dot{\vec{r}}_i^{n+\frac{1}{2}} = \vec{v}_+ + \frac{q_i}{2m_i} \vec{E}(\vec{r}_i^n) \quad (2.55c)$$

Thus, the velocity update due to the Lorentz force acting on a particle can be decomposed into an explicit three-step process, wherein half the electric force is applied before the magnetic force update, ending with the remaining half of the electric force contribution. It is a second-order, explicit algorithm and has the advantage of being computationally inexpensive and accurate, having been shown to have bounded energy error over long time scales due to its phase-space volume preserving nature [118].

## 2.6 Generating new particles: Ionization

When a simulation using MicPIC begins, the charge density everywhere is kept uniformly zero in order to reduce high frequency noise production due to the sudden appearance of non-negligible current densities on the FDTD grid. It is possible to realize this practically either by initializing electrons and ions at common positions, by beginning with only neutral particles, or by using a combination of the two. A system of initially neutral particles is especially useful for simulating the initial formation process of the plasma; for example, in a nanocluster or nanodroplet of atomic gas [104].

MicPIC includes as optional and distinct features both impact and tunnel ionization. Each of these is discussed in the following sections, primarily in terms of implementation details.

### 2.6.1 Impact Ionization

Impact ionization occurs when an incoming electron scatters inelastically with an atom or ion, producing an additional, newly liberated electron in the process. An electron in the vicinity of an ion will have some non-zero probability of initiating an impact event, provided the relative kinetic energy of the electron is sufficient to release another from within the ion. MicPIC in its current state implements sequential ionization events, where only single electrons are released from relaxed atomic or ionic states according to



As an initial step, each particle must be checked to determine whether impact ionization is likely, based on the “size” of the ion. This ionic size is set to half the distance to the nearest neighbour ion, which in turn can be determined, in the case of MicPIC, during the short range correction phase. Additionally, an electron’s distance to its nearest ionic neighbour is used as a criterion, when compared to the size of the




Figure 2.4: Schematic representation of the energies involved in the impact ionization process. An arbitrary potential due to the external environment  $V_{env}$  is shown in red. This potential exerts an effect  $\sigma$  on the ionization potential apparent to the electron in combination with the energy required to bring a new electron from the impactor’s current position (circle, dashed vertical) to the continuum  $E^*$ . This shift of the ionization potential results in a new effective potential  $IP^*$ . In MicPIC, the new electron is given only enough energy ( $IP^*$ ) to propagate from the position of the ion to the point at which the impacting electron triggered its release. The “bare” ionic potential in MicPIC is shown in blue for hydrogen, with the green curve representing the Coulomb potential.

latter, for initiating an impact event. That is, if an electron is “inside” an ion, an impact ionization event will occur.

The relative location and velocity of the impactor determines the cross section of the interaction and, if an event does take place, the velocities of all three particles afterwards. Prior to the event, the new electron liberated by the impact does not exist in the simulation. A hypothetical electron within the ion will see an effective potential due to the residual ionic charge, composed of the nuclear charge and the remaining electrons (ionic potential), and an external one due to the surrounding plasma environment. Computing the potential due to the environment involves the difference in energy  $\sigma = \phi(\vec{x}_{im}) - \phi(0)$ , where  $\phi(0)$  is the potential measured at the position of the ion and  $\phi(\vec{x}_{im})$  is measured at the position  $\vec{x}_{im}$  of the electron relative to the ion. The potential energy  $E^*$  of the hypothetical electron at relative position  $\vec{x}_{im}$  in the ionic potential of a positive ion with a charge state  $n+$  is found by interpolating the MicPIC potential for a particle with charge state  $(n+1)+$  ( $w = w_0$  in Eq. (2.8)) from a look-up table, similar to what is done for the short range force corrections. Thus,  $\sigma$  is simply the potential energy of the hypothetical electron in the environment with respect to the position of the ion.

The impacting electron must have enough kinetic energy relative to the ion to overcome the effective ionization potential  $P^* = P - \Delta_{env}$ , where  $\Delta = E^* + \sigma$  and  $P$  is the ionization potential of the bare atom or ion. The hypothetical electron will be instantiated in the simulation provided the above energy condition is satisfied and that  $\pi b^2 < \sigma_{tot}$ , where  $b = \|\vec{v} \times \vec{x}_{im}\| / \|\vec{v}\|$  is the impact parameter [85],  $\vec{v}$  is the velocity of the impactor, and  $\sigma_{tot}$  is the electron impact cross section. The last of these is calculated using the simplified

empirical Lotz formula [84, 85]

$$\sigma_{tot}(E_{im}) = \left(450 \text{ \AA}^2 \text{ eV}^2\right) \sum_{n,l} f_{nl} \frac{\ln(E_{im}/E_{nl}^*)}{E_{im}E_{nl}} \quad (2.57)$$

where  $E_{im}$  is the kinetic energy of the impacting electron,  $f_i$  are the shell occupation numbers, and the effective ionization potentials  $P_i$  can be provided to MicPIC as input data at runtime. If an impact event is triggered, a new electron is instantiated with a randomized velocity and enough kinetic energy to overcome the shifted ionic barrier  $P^*$ , while the velocity of the impactor is adjusted to ensure that momentum is conserved.

## 2.6.2 Tunnel Ionization

Matter exposed to strong electric fields can be ionized when a bound electron has a non-negligible probability of tunnelling out of its parent atom or ion, resulting from lowering of the atomic/ionic potential by the field. The likelihood of tunneling can be determined by using the well-known theory of Ammosov, Delone, and Krainov (ADK) [83] for atoms and ions exposed to oscillating electromagnetic fields. Modelling tunnel ionization allows MicPIC to simulate the formation of a plasma from an initially neutral target exposed to strong fields, in addition to further ionization of existing ionic species, thus expanding the range of physical scenarios able to be investigated.

When directed to simulate tunnel ionization, MicPIC will seek through the particle list and determine the probability of an event occurring for any atoms or ions it finds. The magnitude of the electric field at the position of the atom or ion and its current ionization state are used to calculate the ADK tunnelling rate via Eq. (1.35), where the ionization potentials are the same as those used in the impact ionization calculations (see Section 2.6.1). A pseudo-random number is used to determine whether an event actually takes place.

It is necessary that a newly created electron does not suddenly appear at the periphery of the ion from which it originated, in order to reduce spurious radiation. Instead, MicPIC employs a type of temporary object (*Tunnel\_Event*) to ramp the positions of a new electron to its classical turning point outside the ion from the position of the ion itself. These objects exist for the lifetime of the ramping process (MicPIC presently uses five time steps) and are used to link the “tunnel” electrons with their ions, allowing the ramping of positions to proceed quickly. Additionally, ions which generated tunnel electrons are “locked” during the aforementioned lifetime to prevent a tunnel ionization cascade, where one tunnel electron triggers a domino effect of tunnel events from a single atom or ion. Every time step, the tunnel event ramps of the electrons are advanced, where any which are completed are released and propagate normally throughout the simulation domain.

---

CHAPTER

THREE

---

# MASSIVELY PARALLEL MICROSCOPIC PARTICLE-IN-CELL

## 3.1 Introduction

When intense laser light interacts with solid-density materials, ionization occurs [81, 119]. In the first (inner) ionization mechanism, bound electrons are promoted to a conduction state where they can freely move throughout the solid. In the (outer) ionization stage, they gain enough energy to overcome the work function and leave the material. While they are accelerated by the laser field, electrons in the conduction state can also free further electrons through laser-assisted impact ionization or share their energy with the lattice ions via electron-ion collisions [120, 121]. This results in a modification and/or destruction of the material, which is of great interest for industrial and clinical applications, such as laser surgery [62, 63, 122], thin-film deposition [57], laser cutting and welding [123, 124], and laser-writing of channels and waveguides in solids [60, 125]. Laser heating of electrons and ions lifts the quantum degeneracy of solids and, to a good approximation, the particle interactions and dynamics can be treated classically.

Modelling the laser-driven plasma dynamics created during laser-induced material modification and machining processes [58, 126, 127, 59] presents researchers with the difficulty of resolving light-matter interaction on both the microscopic and macroscopic length scales. Microscopic phenomena require precise knowledge of the properties and trajectories of individual particles, whereas the macroscopic scale necessitates lengths large enough to account for electromagnetic wave propagation effects. Given that the inter-atomic distance in a solid-density material is on the order of an ångström and that typical laser wavelengths are in the extreme ultraviolet (XUV,  $\sim 100\text{nm}$ ) to the mid-infrared (MIR, a few  $\mu\text{m}$ ) range, bridging the microscopic and macroscopic realms requires resolving roughly four orders of magnitude in length [105]. Furthermore, a  $1\mu\text{m}^3$  volume of material at solid density contains approximately  $10^{10} - 10^{11}$  particles, depending on the material composition. This presents a true challenge for computer modelling. A massively parallel approach to the microscopic particle-in-cell (MicPIC) method is currently the only avenue to meet all of these requirements.

The simulation of particle systems has a history which dates back to work on diffusion [128]. Classical treatment of systems of interacting particles in more modern cases is typically done using one of two most

popular methods. The first such method, *molecular dynamics* (MD), is well-suited to studying the dynamics in systems with dimensions on the order of a few nanometres, where the electrostatic approximation is justified. This involves  $\mathcal{O}(N^2)$  operations ( $N$  being the number of simulation particles), whereby each particle is subjected to the binary Coulomb force from all the others. While this method has been shown to work well for cluster nanoplasmas [85, 129], chemical physics [130, 131], and materials science [132, 133], it does not account for electromagnetic radiation. The computational cost is also prohibitive on systems consisting of more than  $\sim 10^6$  particles. The hierarchical tree [134] and particle-particle/particle-mesh (P<sup>3</sup>M) [70] methods were developed to overcome this limitation but they also rely on the electrostatic approximation. MD, tree, and P<sup>3</sup>M methods are thus limited to simulate light-matter interaction in spatial domains whose size is only a small fraction of the laser wavelength, where radiation propagation effects like source depletion, retardation, and non-dipolar excitation can be neglected.

The second method, *particle-in-cell* (PIC), is a veteran in the field of computational plasma physics. It is formally treated in the works of Eastwood and Hockney [70] and Birdsall and Langdon [101] (see also [135]). PIC improves upon MD by being  $\mathcal{O}(N)$  and by including electromagnetic wave propagation explicitly. Charged particles or ensembles thereof (superparticles) are imposed upon a grid through the construction of weighted charge and current densities (Section 2.3), used as sources in Maxwell equations. The electromagnetic field is then time-evolved using the *finite-difference time-domain* (FDTD, Section 2.2) method [108]. Finally, the motion of the particles is computed with the Lorentz force (Section 2.5) by interpolating the electromagnetic field at the location of each particle (Section 2.3). In PIC, particles do not interact directly, but through the meshed electromagnetic field that contains all the field contributions from the particles. The advantage of this is the resulting  $\mathcal{O}(N)$  scaling. However, a single PIC particle averages over tens of thousands of physical particles and its size is much larger than the atomic length scale. As a result, short range interactions, such as two- and many-body collisions, and plasma micro fields cannot be resolved adequately. Advanced PIC schemes include statistical, Monte Carlo particle collision models based on binary-collisions theories to account for the underestimated short-range particle interactions [102, 103, 72, 136, 71, 73]. Still, the applicability of these methods is limited to simulating the dynamics of low-density, near-collisionless plasmas.

The newer tool MicPIC [106, 104, 105] effectively combines the MD and PIC approaches in a two-stage scheme in order to circumvent the electrostatic approximation in the former and the lack of fine-grained details in the latter. Large scale effects are treated with the traditional PIC approach, where each particle corresponds to an atom, ion, or electron. For computational efficiency the size of the PIC particles is chosen much larger than the atomic length scale. When particles come close to each other, the local electric field of each PIC particle is corrected with a short-range electrostatic contribution in the fashion of MD to account for the underestimated part in PIC. The result is a fully electromagnetic simulation method capable of accounting for collisions at the microscopic scale for individual charged particles. Initially, this approach could handle a maximum of  $\sim 10^6$  particles in serial runs [105, 106, 104]. In this chapter, the use of MicPIC for massively parallel simulations with up to  $\sim 10^{11}$  particles is reported, opening new avenues for large-scale modelling in nanophotonics, diffractive x-ray imaging, and strong-field science.

The rest of this chapter is organized as follows. First, in Section 3.2, details are presented regarding the individual numerical methods introduced in Chapter 2 and how they are extended into parallel form. A theoretical analysis of the expected scaling of MicPIC with respect to the different simulation parameters follows this in Section 3.3. Section 3.4 assesses parallel MicPIC's strong and weak parallel scaling characteristics by testing its performance in different scenarios on mid and large scale supercomputer systems, respectively. Final conclusions are given in Section 3.5.

## 3.2 The Method of Parallel MicPIC

The numerical methods of MicPIC are discussed in Chapter 2, the serial (single processor) implementation of which has been discussed extensively in [105]. Despite advances in the speed of processors and capacity of computer memory, the applications of a serial implementation are limited to small problem sizes, therefore motivating a desire to have access to an implementation of the MicPIC method capable of running on large numbers of computer processors simultaneously. The current parallel version of MicPIC presented here builds upon the serial version with the use of the *Message Passing Interface* (MPI) standard to implement communication of particle, field, current density, long-range force weighting, and short range data between processes, supplemented by specific particle and geometry data required for ionization and photoemission. Each simulation volume is decomposed into smaller sub-volumes, one of which is managed by a single process (see 3.2.1 below for more details). This therefore necessitates the exchange of certain field and particle data between processes in order to complete a time step. The need to correct for short range interactions is novel; in addition to PIC there is the need to develop parallelization schemes for short range corrections. Thus, each MicPIC step is interwoven with communications at key points (see below: Fig. 3.1). How MicPIC specifically incorporates distributed memory parallelism with short range force corrections for simulations with large numbers of processes is discussed below.

In practice, MicPIC is written in ISO standard C++11 with recent development of the code focusing heavily on an object oriented approach. It is worth noting that many of the large datasets are stored as contiguous arrays in memory. This is especially true in the case of the particle list, where random accesses, insertions, and deletions are necessary, while keeping the memory footprint of the most memory-demanding list in MicPIC as small as possible.

### 3.2.1 Parallel domain decomposition

The strength of MicPIC comes not only from its accuracy but also its parallelizability. The PIC method itself has been parallelized for use on shared memory [107, 104], dedicated (distributed) memory [137, 115], and even GPU clusters [138, 139, 140, 141, 142] with excellent scaling capabilities. These approaches can be ported directly to the PIC part of MicPIC. Moreover, it has been demonstrated that the microscopic corrections (the *Mic* part) in MicPIC are essentially local and naturally subdivided into a set of similar, MD-like operations. The use of a shared-memory version of MicPIC was effectively demonstrated in [107]. Below, the discussion focuses on a distributed-memory approach to simulate large-scale systems on massively-parallel high-performance-computing infrastructures.

For the benefit of the uninitiated, a brief introduction to a simplified distributed memory parallel processing paradigms is warranted. In MPI, processes are arranged into *groups* and every process in a group is assigned a *rank*, which is an integer signifying its identity in that particular context. Many different groups can exist simultaneously and a particular process can be a member of any number of them (or none). Communications of messages (data structures) between MPI processes is achieved by either *blocking* or *non-blocking* functions. The former type require a function to return (complete) before execution can proceed and accessing memory pertaining to data the communication is permitted, whereas the latter do not. Non-blocking communications must be completed at some later time by some collective function, permitting several communications to overlap in time and possibly yielding improved performance. *Collective* operations are simply those that must take place on all processes in a given group of processes. If a collective function requires the members of a group to wait until everyone in the group has returned, the function is said to be synchronizing.

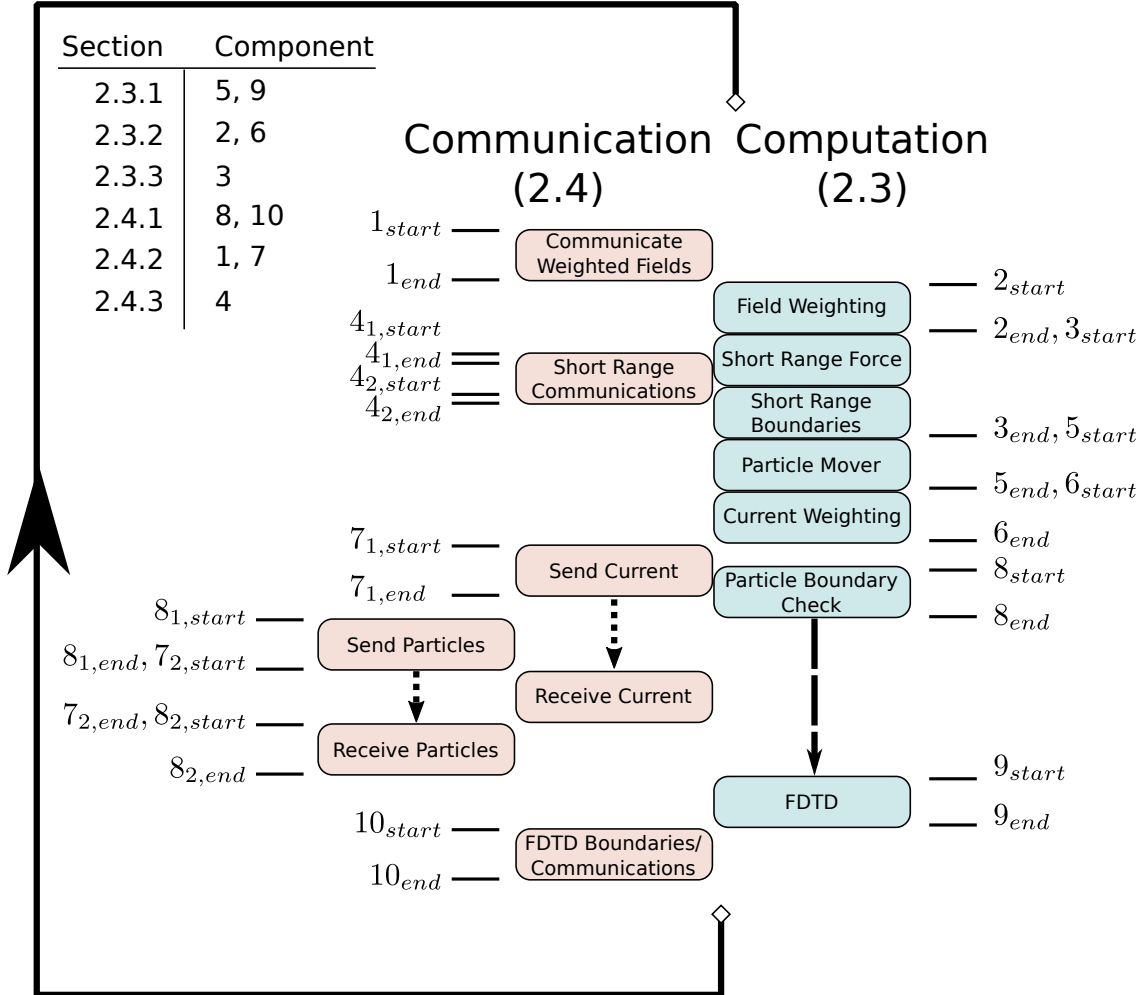


Figure 3.1: A schematic illustration of the MicPIC algorithm showing the location of functions related to communication between processes. Execution proceeds from top to bottom. Pink boxes (left side) indicate communication operations and blue boxes represent functions related to computation. Numbered lines indicate when time measurements are taken during a step (see text for details). Short range communications overlap with short range force and boundary corrections (see text below, 3.2.1). Currents (3.2.1) and particles (3.2.1) are sent and received between processes in a similar manner.

One of the goals in parallel processing is to minimize the number of synchronizing function calls and, when unavoidable, to minimize the time where processes are waiting for their counterparts.

Distributed memory parallelism is achieved in MicPIC *via* geometrical domain decomposition, where a large simulation volume is composed of small *subdomains* tiled in space to form a 3D mosaic. Each process is tasked with time evolving its own subdomain of the simulation’s fields and particles. The computation locality allows for primarily non-blocking, point-to-point exchange of only specific information and data located on the boundary regions between neighbouring subdomains, making optimal usage of the communication network bandwidth. Collective communications are only employed during the main time-stepping when processes are required to interact with more than just their sub-domain neighbours, such as (collective) filesystem accesses (non-blocking *MPI\_File\_iread\_at()* or blocking *MPI\_File\_read\_at()*), calculating byte offsets in globally accessed files via (collective) *MPI\_Exscan()*, or for calculating global quantities in post-step diagnostics via collective *MPI\_Reduce()* or *MPI\_Gather()*. Additionally, non-blocking send and receive operations allow one to overlap computation with communication. While this is difficult practically due to the sequential nature of the MicPIC algorithm, it is utilized wherever possible and is a focus of future optimization.

The present implementation of MicPIC is illustrated schematically in Fig. 3.1, which is a more detailed view of what is shown in Fig. 2.2. Execution of a step proceeds from top to bottom and is divided into computation (blue) and communication (pink) segments. At the beginning of a step, data from neighbouring subdomains are required to correctly interpolate the EM field experienced by a particle whose charge density extends across a boundary. Thus, communication of some field data (“Communicate weighted fields”, 1) takes place before the actual weighting operations for the long range force (“Field weighting”, 2). Once the long range forces are calculated the short range force corrections follow. This is subdivided into two segments (“Short range force” and “Short range boundaries”, 3). Particles belonging to a process undergo corrections while information about any which are near subdomain interfaces are exchanged between neighbours with MPI (“Short range communications”, 4). Their effects are included in the latter short range computation segment. Having calculated the forces on each particle, MicPIC integrates the particles’ equations of motion (“Particle mover”, 5) and assigns weighted current data to the FDTD grid (“Current weighting”, 6). Current density generated by particles extending across boundaries is exchanged between processes such that each process can accurately update its EM field. The current data are sent (“Send current”, 7) and received (“Receive current”, 7) but are done in a way that allows the communications to proceed while each process determines which of its particles are moving into a neighbour’s subdomain (“Particle boundary check”, 8). The particle communications overlap with those of the current density (“Send particles” and “Receive particles”, 8). Once the complete current density field is available the process will propagate the EM field with the FDTD scheme (“FDTD”, 9), which requires each process to both send (receive) some field data to (from) its neighbours (“FDTD boundaries/communications”, 10). The step concludes and the execution returns to the top of Fig. 3.1.

Currently, MicPIC decomposes the simulation volume according to a 1- or 2-D Cartesian processes topology, resulting in an array of rectangular slabs. This decomposition method works best for planar medium geometries and scales as long as the volume-to-surface ratio of the subdomains is relatively large, i.e., when the time needed to perform the neighbour-boundary computations and communications is negligible compared with the bulk volume computation (see Sec. 3.3 for more details). 2D topologies take on XY, XZ, or YZ configurations, depending on which two of the three dimensions are parallelized. An example of a 2D topology is shown in Fig. 3.2 and is discussed in more detail below. Future implementation of parallel MicPIC should consider full 3D decomposition, potentially with dynamic load balancing capabilities.

Parallel MicPIC implements *boundary conditions* as either *periodic* or *absorbing* (UPML) types of exterior

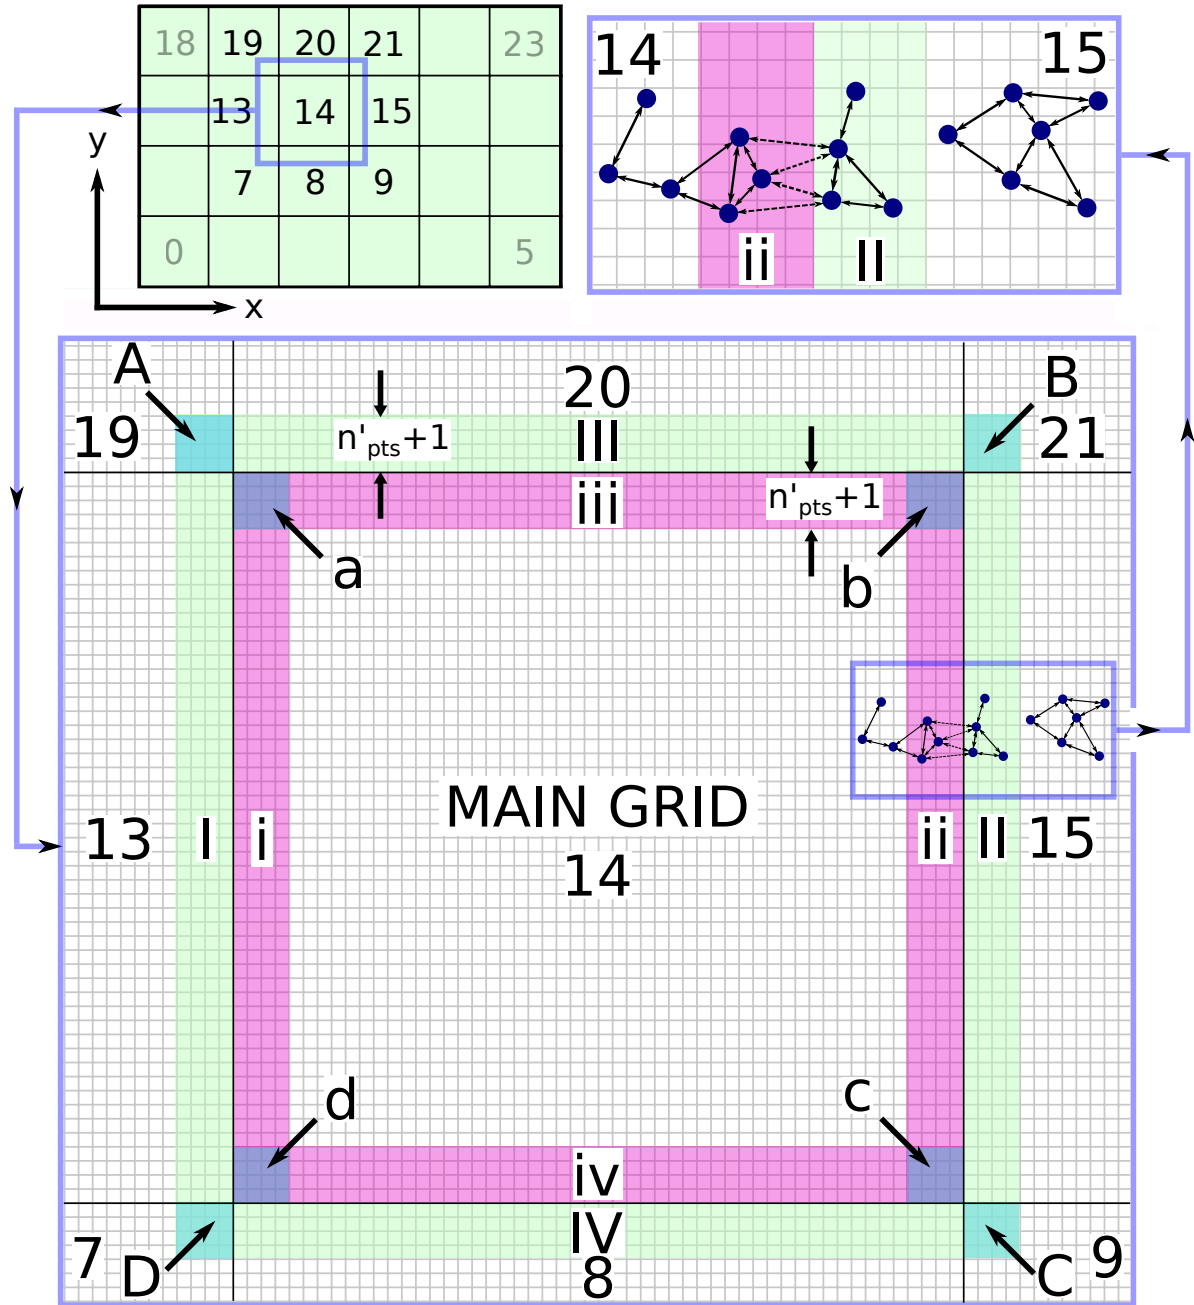


Figure 3.2: Schematic representation of a typical 2D [XY] domain decomposition, viewed down the  $z$ -axis, showing the sub-grids involved in communication of data between neighbouring processes. The main grid of process 14 is bounded by magenta boxes ( $i$ - $iv$ ), which represent current data from neighbours' particles extending across boundaries. Light green boxes ( $I$ - $IV$ ) are related to cells from which process 14 requires EM field data in order to properly update the weighted local field of particles near one or more boundaries. Particles whose short range correction radius intersects a boundary exchange their position and charge data across those boundaries in order to correct the PIC force of particles in the vicinity of a boundary. Edge buffers  $a$ - $d$  and  $A$ - $D$  (light blue, dashed) operate on the same principles.

boundaries at the global domain limits. Interfaces between subdomains interior to the global bounding surface are what can be referred to as parallel boundaries, since they utilize MPI to exchange information. When a subdomain (local) boundary is touching an exterior one, it inherits the associated global boundary conditions. If such a boundary condition is periodic, then it is treated the same as a parallel one and information is exchanged with the process on the opposite side of the global simulation domain.

Communications in parallel MicPIC depend on whether a subdomain’s boundary conditions inherited from the global level necessitate communication and what are referred to as the *geometric boundary type*. These stem from the rectilinear geometry of a subdomain and exist as *face*, *edge*, and *vertex* types, which are respectively associated with the neighbours with which they share faces, edges, or vertices. Additionally, these types dictate the subsets of grid and particle information to be communicated. For example, in Fig. 3.2 process 14 shares faces with processes 8, 13, 15, and 20, while sharing edges with 7, 9, 19, and 21. In general, a process will have six neighbouring processes along its face boundaries, four or twelve edge neighbours, depending on which particular type of communication is taking place, and eight vertex neighbours, for a total of up to 26 in total with which to exchange information a few times per time step (see Fig. 3.1). In the next subsections, the strategies employed for efficient boundary information communications are outlined.

### Particle and FDTD communications

Particle and electromagnetic field data are the simplest cases for which there is communication in parallel MicPIC. Each step, the positions of particles are checked against the limits of their sub-domain. If a particle is found to be crossing one of the 26 boundaries, it is handled in one of three possible ways. The primary case of interest is that of a parallel boundary. Particles crossing such a boundary are marked according to which of the 26 it is, corresponding to the *Particle Boundary Check* segment in Fig. 3.1. Once it is known which particles need to be sent and to where, they are all communicated simultaneously with calls to *MPI\_Isend()* and *MPI\_Irecv()* (*Send Particles* and *Receive Particles* in Fig. 3.1, respectively). Non-blocking send and receive operations allow the particle communications to proceed while the current density is received and added to the main grid (see below).

If a subdomain boundary is periodic in a dimension which is not parallelized, no communication is required and the particle re-enters from the opposite side. If the boundary is a UPML, meaning that a particle is leaving the simulation domain, then the particle is simply switched off. This is a gradual process, due to the low speed at which a particle can traverse a cell under normal conditions. As the particle’s charge density crosses the physical boundary, parts of it cease to have grid elements for interpolated current density. Thus, the charge of the particle is seen to ramp down until the centre of the charge density crosses the boundary, at which point the switch off occurs. This only takes place inside the UPML region, which attenuates any associate burst of radiation resulting from the remainder of the charge going to zero. In our experience to date, the remaining low-frequency components and static field are small enough to be ignored. Effectively, the particle is tracked by the last process to own it but no longer feels any force and does not contribute any current density.

Recall that in MicPIC the electromagnetic field is evolved numerically via the FDTD method. Updating the electromagnetic field in a distributed computation requires knowledge of certain field components one cell beyond the limits of the subdomains. Communication of the relevant field data across the appropriate parallel face boundaries is therefore necessary. This is facilitated with the use of MPI derived datatypes, which are user-defined data masks for selecting which field data to send and receive. In practice, create indexed datatypes are created (*MPI\_Type\_create\_indexed()*) during the initialization phase, based upon the

dimensions of the one cell thick slice of field mesh needing to be exchanged. This tells MPI exactly which field data to send, from where in the memory, and how to receive it.

### Current and long-range force communications

Using wide charge distributions on the PIC level implies that particle interpolation operations (current density and long-range force) require to read and write data beyond the subdomains. In other words, some current density values of a given process need to be added to the FDTD grid of another, while the same process needs parts of its neighbours' E and H fields to properly compute the electromagnetic long-range force on its particles. For example, in Fig. 3.2, current density data from process 14 is sent to 15 and added to cells contained in *II* (light green box) in the main grid of 15. On the other hand, the local field weighting of particles on process 14, which extend into the subdomain of 15, require EM field data from those neighbouring cells (again light green box *II*), that are sent to 14.

To implement these two-way communications, *weighting boundary grids* are employed. These sub-grids have a structure similar to that of the main FDTD mesh and are used to store field and current density data corresponding to regions belonging to neighbouring subdomains. In the case of E- and H-fields they represent a  $n'_{pts}$  cells thick slab of a neighbour's main FDTD mesh (sub-grids *i-iv* in Fig. 3.2) and a slab of equal thickness in the case of current density (sub-grids *I-IV* in Fig. 3.2).

The communications of field data required for the long-range force update must complete before weighting of the fields for each particle can occur. We employ MPI indexed datatypes in a similar manner to that used in the FDTD communications outlined above and non-blocking send and receive operations are used to exchange the data. Due to the aforementioned necessity of having data available before proceeding with the weighting, communications are only able to overlap with themselves. Specific grids are not written to more than once by any set of receive operations, since each boundary is distinct.

The exchange of current density information is very similar to the scheme outlined above for the long-range force, with two main differences. One is that the current density is not needed until it is time to perform the update of the EM field. The second is that the contributions to the local current density from neighbouring subdomains must be added to the main grid of a process. This provides the opportunity to overlap the sending and receiving of particles and currents. We initiate the sending of currents and then move on to the particle boundary checking and communications (see Sec. 3.2.1). After the particle send operations are posted, the current is ready to be received. MicPIC dynamically receives current density communications, meaning that the data is received on a first-come, first-served basis and the size of the message is determined upon receipt. Once a message containing current density information is received its components are added to the receiver's main grid while waiting for the remaining messages.

### Short-range corrections communications

Short range corrections in parallel MicPIC use similar custom-buffer communication strategies as the weighting routines (Sec. 3.2.1) in order to store the received data. The information contained in each buffer element is simply four *double* types: three for position and one for charge.

Using the top-right of Fig. 3.2 as an example, particles whose correction range intersects one or more boundaries have their position and charge information exchanged across those boundaries. These are represented by the particles joined by dotted lines. For periodic boundaries, the position of the particle is simply shifted in the periodic direction after it is buffered. In other words, a particle at  $z = 299$  seeing a particle at  $z = -1$  across a periodic boundary at  $z = 300$  treats it as being at  $z = 301$  instead.

Particles whose charge and position information needs to be communicated are determined in a manner akin to that which is used for the particle communications (Sec. 3.2.1). Any particle within a distance of  $r_{cut}$  from a physical subdomain boundary has its offset with respect to the beginning of the particle list recorded during the sorting of the main (local) cell list. In practice, the charge and position information are not contiguous, which motivates the use of *MPI\_Type\_create\_struct()*. The data are sent as an indexed type composed of these struct type elements, precluding the need for dedicated buffers for the data being sent (two smaller buffers are required to create the indexed MPI datatype).

With the use of nonblocking communication, computation of the short range force is allowed to proceed. Furthermore, if the non-parallelized direction has periodic boundary conditions, it is advantageous to apply them as well, while communication is ongoing. MicPIC employs dynamic non-blocking receiving of incoming short range force data. The receive buffers feature contiguous charge and position data, which allows the incoming data to be received as simple *MPI\_DOUBLE* primitive types.

After successful communication, the processes perform a second round of microscopic corrections for the particles near the subdomain boundaries that interact with particles in neighbouring subdomains, whose relevant information is now stored locally. A list of cells with particles within  $r_{cut}$  of geometric subdomain boundaries with parallel or periodic boundary conditions (eg. a face boundary in the case of Fig. 3.2, upper right) is maintained for each boundary for the duration of the short range force portion of a timestep. This facilitates rapid sorting of particles by cell for any within such a boundary region. From there, it is straightforward to find the neighbouring cells of each particle near one or more geometric subdomain boundaries, whose short range force is subsequently updated in the same manner used in the local updates (Section 2.4).

## Ionization

When an electron passes close to an ion, it will be conditionally involved in an impact ionization event, based on where it is in relation to the nearest ion. If this occurs near the boundary of a subdomain, is the nearest ion what is immediately available in the local particle list or is it just on the other side of the boundary? This is the motivating question behind the parallelization of electron impact ionization in MicPIC.

Recall that the parallel short range corrections proceed in a manner similar to that which is done for the local particle list once the communications complete. Consequently, MicPIC is able to update the result of the local nearest neighbour determined for a particle, during a subsequent check of the short range communication buffers, with the rank and list index of any particle in a neighbouring subdomain that is closer. During the following ionization phase, any electron matched with a nearest neighbour on a different process is flagged and information about its location, identification (rank, index), and potential energy is packaged in an *Impact\_Event* object. These objects only exist for the duration of the ionization phase of the simulation and are sent with MPI to the processes hosting the relevant ions, which undergo the same testing procedure used to identify successful impact events locally. If an event does occur, the host process of the impacting electron must be notified so that an impact electron can be generated. Once this second communication and the following spawning of the new electron complete, the parallel impact event ends.

In spatial terms, the generation of new electrons via the process of tunnel ionization only depends on the instantaneous position of the ion from which it is spawned. The weighted electromagnetic field which the ion experiences is known via the local and parallel long range force calculations (see above and Section 2.3). However, once an electron is generated by a successful tunnel ionization event, it must remain linked to its parent ion during the ramping phase (see Section 2.6), even across subdomain boundaries. This linking cannot

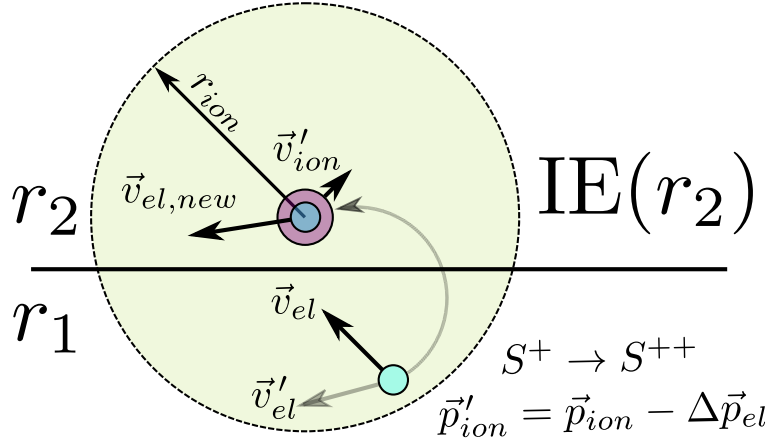


Figure 3.3: Graphical depiction of a parallel impact ionization event. An electron (blue) on rank  $r_1$  is made aware of the presence of the ion (pink) with charge state  $S^+$  on rank  $r_2$ , which is the nearest ion. A successful event caused by the electron entering the ionic radius  $r_{ion}$  (beige circle) with sufficient kinetic energy triggers the instantiation of `Impact_Event` object  $IE(r_2)$  on the host process of the ion, which facilitates the exchange of information between the two interacting particles. A new electron is created at the position of the ion with a random velocity  $\vec{v}_{el,new}$  and kinetic energy equal to the shifted ionization potential discussed in Section 1.1.3. Due to the inelastic nature of the process, total kinetic energy is not conserved, though the impacting electron loses the ion-relative kinetic energy to the release of the new electron. Momentum conservation yields the final velocities of the two original particles.

be accomplished with simple pointers, which are invalid when dealing with memory on different processes and must therefore be accomplished in the form of simple identifying integers.

In MicPIC, tunnel electrons which cross parallel boundaries must track the host rank and list index of their parent ion and, if necessary, apply updates to the `Tunnel_Event` object with its new identification information. Similarly, the ion is linked to the ion in order to provide the new electron with updated identification information and to inform the new host process of the electron when the ramping is complete. Once this last step occurs, an `Unlock_Call` is communicated to each process for each parallel tunnel electron so that it can propagate normally.

### 3.3 Theoretical scaling analysis

Parallel MicPIC's self-consistent integration of particle dynamics and light propagation is composed of several computation and communication steps that are outlined in Section Section 2.1.2 (Fig. 3.1). Below, a simplified theoretical analysis of MicPIC's scaling behaviour is presented in terms of the particle number  $N$ , simulation volume  $V$ , FDTD grid resolution  $\Delta x$ , and number of processes  $N_p$ . A numerical confirmation of the theoretical scaling analysis is provided next in Sec. 3.4 with test simulations.

MicPIC scaling for situations where the computation is done by a single process was covered in previous publications [105, 106, 104]. Here, a slightly different derivation is provided to emphasize the role of the spatial discretization parameter  $\Delta x$ . First, it is assumed a cubic computational domain with side length  $L$  and volume  $V = L^3$  filled with a homogeneous plasma composed of  $N/2$  singly-charged ions and  $N/2$  electrons, for a total count of  $N$  particles. In a MicPIC calculation, the volume  $V$  is discretized in small

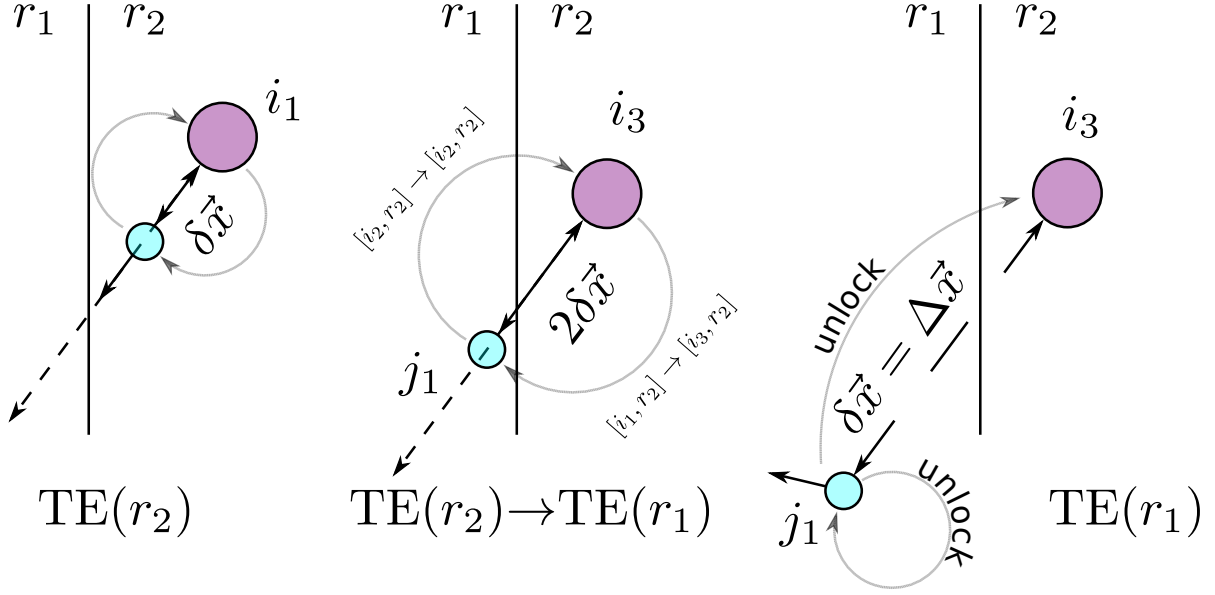


Figure 3.4: Graphical depiction of parallel tunnel event. A successful tunnel event at ion  $i_1$  (left, pink) on rank  $r_2$  triggers the generation of a new electron (blue). The ramping begins and moves the electron to its release position (dashed arrow) by spatial step size  $\delta\vec{x}$  while the two remain linked by `Tunnel_Event` object  $TE(r_2)$  (curved arrows). During the ramping, the electron crosses into the subdomain of rank  $r_1$  (middle); meanwhile the ion’s memory location can change due to list operations. The two exchange updated identifications via the `Tunnel_Event`, which accompanies the electron to  $r_1$ . Finally, the electron is unlocked (released) at the final ramp position  $\Delta\vec{x}$  and an `Unlock_Call` is sent to the parent ion so that it can participate in future tunnel ionization events.

cube elements of side length  $\Delta x$ . For stable computation of the electromagnetic field propagation with the FDTD technique, the time-step  $\Delta t$  must obey the following inequality:  $\Delta t \leq \Delta x/(\sqrt{3}c)$ , where  $c$  is the vacuum speed of light. For the analysis, simply fix  $\Delta t/\Delta x = \text{const.}$  For one time step  $\Delta t$ , the evaluation of the microscopic corrections scales as the number of particles  $N$  times the number of particles within the correction sphere  $N_c = (4\pi/3)(N/V)r_{cut}^3$ , which gives  $O_{step}^{mic} = \alpha N r_{cut}^3$  when  $N_c \ll N$  (see also discussion in Section 2.4), with  $\alpha$  being a constant. On the other hand, one step of the PIC part requires calculation of the current densities and forces, which is an  $\mathcal{O}(N)$  operation, and field update, which is determined by the number of cells to compute  $V/\Delta x^3$ . As a result, the PIC part scales as  $O_{step}^{pic} = \beta N + \gamma/\Delta x^3$ , where both  $\beta$  and  $\gamma$  are constants. Using the fact that  $r_{cut} \propto \Delta x$  (see, e.g., [105, 106, 104]) and that  $\Delta t \propto \Delta x$  [see Eq. (2.15)], the total MicPIC execution time for a fixed physical time interval is:

$$O^{\text{MicPIC}}(N, \Delta x) = \underbrace{\left( \alpha' \Delta x^2 + \frac{\beta'}{\Delta x} \right) N}_{\text{particles}} + \underbrace{\frac{\gamma'}{\Delta x^4}}_{\text{field}} \quad (3.1)$$

where the parameters  $\alpha'$ ,  $\beta'$ , and  $\gamma'$  are constant prefactors for the short-range, long-range, and field propagation parts of MicPIC, respectively. Alternatively, using the fact that  $N/V = \text{const.}$ , Eq. (3.1) can be written in terms of the simulation volume  $V$  as:

$$O^{\text{MicPIC}}(V, \Delta x) = \left( \underbrace{\alpha'' \Delta x^2}_{\text{Mic}} + \underbrace{\frac{\beta''}{\Delta x} + \frac{\gamma''}{\Delta x^4}}_{\text{PIC}} \right) V, \quad (3.2)$$




Figure 3.5: Support sketches for theoretical scaling analysis and parallel decomposition scenarios. In (a), a 2D cut of a 3D subdomain showing the geometry of the boundary layer used for inter-process communications. Also shown in (b) and (c), four-process examples of 2D parallelization approaches corresponding to (b) strong and (c) weak scaling situations. In (b), the calculation of a certain physical problem contained in a volume  $V = L^3$  with  $N$  particles is split in 4 subdomains, each handled by different processes. As a result of the domain decomposition, each process has a fourth of the total workload. In (c), an initial simulation of a volume  $V$  containing  $N$  particles is duplicated to simulate a total volume four times as big. The two parallel strategies are typically used: (b) to reduce the global computation time and (c) to benefit from the large memory capacities of distributed simulation clusters to study big material volumes.

where  $\alpha''$ ,  $\beta''$ , and  $\gamma''$  are constants. It is readily observed that a MicPIC computation scales linearly with the volume of matter that is modelled. It is seen also that for coarse grid resolution ( $\Delta x \rightarrow \infty$ ), the short-range (Mic) part is the dominant contribution to the global computational workload. In this limit, MicPIC is equivalent to MD. On the other hand, for fine grid resolution ( $\Delta x \rightarrow 0$ ), the computation is essentially PIC (long-range and field). Effectively,  $\Delta x$  can be optimized to balance the load between the Mic and PIC parts to minimize the total computation time. This typically leads to values of  $\Delta x$  that are smaller than needed to resolve classical solid-density plasma processes. This topic is further developed in Sec. 3.4.1.

As emphasized in previous sections, inter-process communications in MicPIC require a buffer zone (ghost layer) beyond the subdomains' limits for the exchange of boundary information needed to perform all local computations at each time step. Effectively, the overhead per time step associated with the copy and communication of particle and field data scales respectively with the number of particles  $N_b = (N/V)V_b$  and the number of FDTD cells  $V_b/\Delta x^3$  in the buffer layer volume  $V_b$ . If the buffer region has a thickness  $l$ , then  $V_b = 4(lL + l^2)L$  [see Fig. 3.5(a)]. Furthermore, there is the overhead from the corrections to local particles that have short-range particle interactions across parallel boundaries (see Fig. 3.2). This extra step is equivalent to doing a ‘‘Mic’’ computation over the  $N_b$  particles in the buffer boundary volume  $V_b$  and scales as  $\Delta x^2 N_b \propto \Delta x^2 V_b$  [see Eqs. (3.1) and (3.2)]. Then, assuming an homogeneous load over all the processes and perfect synchronization between them, the parallel MicPIC overhead for a fixed physical time interval and constant particle density  $\rho = N/V = N_b/V_b = \text{const.}$  is:

$$O_{\text{overhead}}^{\text{MicPIC}}(V_b, \Delta x) = \left( \delta \Delta x^2 + \frac{\epsilon}{\Delta x} + \frac{\zeta}{\Delta x^4} \right) V_b, \quad (3.3)$$

where  $\delta$ ,  $\epsilon$  and  $\zeta$  are constants.

The particular dependence on  $\Delta x$  in Eq. (3.3) should be interpreted with care. In facts, the thickness  $l$  of the buffer layer has to be chosen large enough to cover both the penetration of the particles' Gaussian shape function and short-range correction sphere into the neighbouring subdomains. The relevant parameters are the Gaussian function sampling radius  $r_G = (n'_{pts} + 1)\Delta x$  and the cutoff radius  $r_{cut} \propto w_{pic} \propto \Delta x$  (see Sections 2.3 and 2.4). It follows that  $l \propto \Delta x$  such that  $V_b = \mu\Delta x L^2 + \nu\Delta x^2 L$ , where  $\mu$  and  $\nu$  are constants. This extended definition of  $V_b$  reveals a more complex dependence on the resolution parameter  $\Delta x$ :

$$O_{\text{overhead}}^{\text{MicPIC}}(L, \Delta x) = \left( \delta' \Delta x^4 + \epsilon' \Delta x + \frac{\zeta'}{\Delta x^2} \right) L + \left( \delta'' \Delta x^3 + \epsilon'' + \frac{\zeta''}{\Delta x^3} \right) L^2, \quad (3.4)$$

where  $\delta'$ ,  $\epsilon'$ ,  $\zeta'$ ,  $\delta''$ ,  $\epsilon''$ , and  $\zeta''$  are constants. The numerical results presented later in Fig. 3.6 are consistent with this fairly complex scaling equation.

Eqs. (3.3) and (3.4) indicate that MicPIC's parallel overhead can be minimized with an appropriate choice of  $\Delta x$ , just like the computation part [see Eq. (3.2)]. However, parallel MicPIC is designed to use a minimal number of communications to transfer a restricted set of data between neighbour processes with point-to-point communications. We show later that with typical subdomain process workloads, inter-process communications in MicPIC are effectively faster than the computation part by an order of magnitude (see, e.g., Fig. 3.6), even in the case of massive-scale simulations (see, in particular, Fig. 3.12). In this context, optimization of the parallel overhead with respect to  $\Delta x$  has no benefit.

Ultimately, combining Eqs. (3.2) and (3.3) gives:

$$O_{\text{parallel}}^{\text{MicPIC}}(V', V_b, \Delta x) = \left( \alpha'' \Delta x^2 + \frac{\beta''}{\Delta x} + \frac{\gamma''}{\Delta x^4} \right) V' + \left( \delta \Delta x^2 + \frac{\epsilon}{\Delta x} + \frac{\zeta}{\Delta x^4} \right) V_b, \quad (3.5)$$

where  $V'$  is the subdomains' volume. We recall that for ideal scaling, one must ensure that the time spent in communications is negligible compared with the total computation time. We see from the equation above that it naturally happens in MicPIC when all the subdomains' computational volumes  $V'$  are much larger than the volume  $V_b$  associated with the boundary buffer layer, i.e., when  $V'/V_b \gg 1$ . This condition implies that the boundary layer must be thin compared with the size of the domain, i.e., when  $L \gg l$ , in each three spatial dimensions. In particular for a cubic domain fulfilling this condition,  $V_b \simeq 4LL^2$ , which is effectively proportional to the cube surface ( $A = 6L^2$ ). Optimal MicPIC parallel scaling is thus expected for subdomains with large volume-to-surface ratios. This is best achieved with 3D geometrical decomposition in large cubes. Nevertheless, it is shown later that with a 2D decomposition topology it is possible to achieve parallel efficiency in the 80% range and above (see, in particular, Fig. 3.9).

## 3.4 Tested performance

The approximate scaling equations given above in Sec. 3.3 do not account for the full complexity of parallel code execution and collective process synchronization in large-scale, multi-user simulation cluster environments. In the upcoming sections, full-scale test simulations are presented to validate the theoretical analysis of Sec. 3.3 and assess the strong and weak scaling characteristics of parallel MicPIC in a real production context. It should be noted that all test runs of parallel MicPIC were performed on systems where jobs have exclusive access to the nodes being utilized. To identify the optimal choice of the resolution parameter  $\Delta x$  (see Sec. 3.4.1), a series of test simulations were conducted with different material atomic densities. The execution times for the most intensive components of parallel MicPIC are outlined and analytic fits using the equations from theoretical analysis of Sec. 3.3 show excellent agreement.

For strong scaling (see Sec. 3.4.2), a wide and flat initial simulation volume is split into  $N_p$  smaller pieces (subdomains) using a 2D decomposition as in Fig. 3.2. The initial workload is then distributed over  $N_p$  computer processes, with each  $(1/N_p)$ th of the problem to compute. As an example, consider a domain split in the  $x$  and  $y$  directions in  $N_x$  and  $N_y$  parts, respectively, such that  $L_x = L/N_x$  and  $L_y = L/N_y$  ( $L_z \ll L$ ). The individual process subdomains volumes are then  $V' = L_x L_y L_z = L^2 L_z / N_p$ , where  $N_p = N_x N_y$  is the total number of subdomains, each containing  $N' = (N/V)V' = N/N_p$  particles [see also Fig. 3.5(b)]. On the other hand, the buffer volume associated with the communications to neighbours is  $V_b = (L_x + 2l)(L_y + 2l)L_z - V' = 2[lL(N_x + N_y)/N_p + 2l^2]L_z$ . Then  $V'/V_b = L^2/[2lL(N_x + N_y) + 4l^2 N_p]$ , which indicates that the subdomains' volume-to-surface ratio drops with increasing number of processes. We see later that this leads also to a drop in the parallel efficiency (see, in particular, Fig. 3.9).

For weak scaling (see Sec. 3.4.3), a near-cubic initial simulation domain is duplicated in one or two spatial dimensions to create a bigger one [see Fig. 3.5(c)]. If the initial, “small” simulation has  $N'$  particles in a cube of side length  $L'$  and volume  $V' = L'^3$ , then the total  $N_p$  process simulation will be composed of  $N = N_p N'$  particles contained in a volume  $V = N_p V'$ . In this case, the subdomains' volume-to-surface ratio is defined by the initial domain (the first “tile” of the 2D mosaic) and remains constant with increasing  $N_p$  and, by extension, the ratio of communication time to computation time. Thus, the parallel efficiency is expected to remain constant up to the capacity of the computer system (see, in particular, Fig. 3.12). We show that this scaling behaviour allows to simulate large particle ensembles, up to  $\sim 10^{11}$  on a 4-rack IBM Blue Gene/Q supercomputer with a total of 65 536 processes.

### 3.4.1 Optimum Cell Size

We have seen in the previous sections that MicPIC effectively solves a full  $N$ -particle problem on a global scale with an electromagnetic PIC simulation, along with  $N$  local MD simulations performed on small ensembles to correct for short range interactions. As was emphasized in Section 2.4, in optimal conditions MicPIC scaling is linear with the total number of particles  $N$  (see also [105]). An important optimization parameter is the FDTD cell size  $\Delta x$  that defines the respective loads on the MD (Mic) and PIC parts of MicPIC. For a given physical problem, the load can be split in a way that the execution time is the shortest.

The choice of ideal cell size in MicPIC is first dictated by the necessity to resolve microscopic density fluctuations and atomic-scale collisions. Except for x-rays that may require a better resolution, this is usually enough to resolve the wavelength of the incident radiation and its principal harmonics (typically 2nd and 3rd), each of which require at least 20 cells per optical cycle. Provided the size of the subdomains is sufficiently large compared to their boundary layers (see Sec. 3.2.1 and Fig. 3.2), the communication time will not significantly affect the optimum cell size.

In MicPIC, the time step  $\Delta t$  is linearly proportional to the cell size  $\Delta x$  [see Eq. (2.15)]. Instinctively, for larger cell sizes one would predict that the fewer time steps required to simulate a given physical time would lead to a significant decrease of the total computation time. However, it is not always the case. This results from varying overhead on competing processes. For example, doubling cell size will decrease the load on PIC by a factor of  $\sim 8$ , but increases the load on the short range corrections because there are now more particles per cell (see Section 2.4). Conversely, cells which are too small will lead to particles having a very low number of short-range partners, if any. Such a simulation will therefore have a higher percentage of the execution time dedicated to the PIC components of MicPIC. Effectively, for a given material density, there is an optimal cell size that minimizes the global execution time.

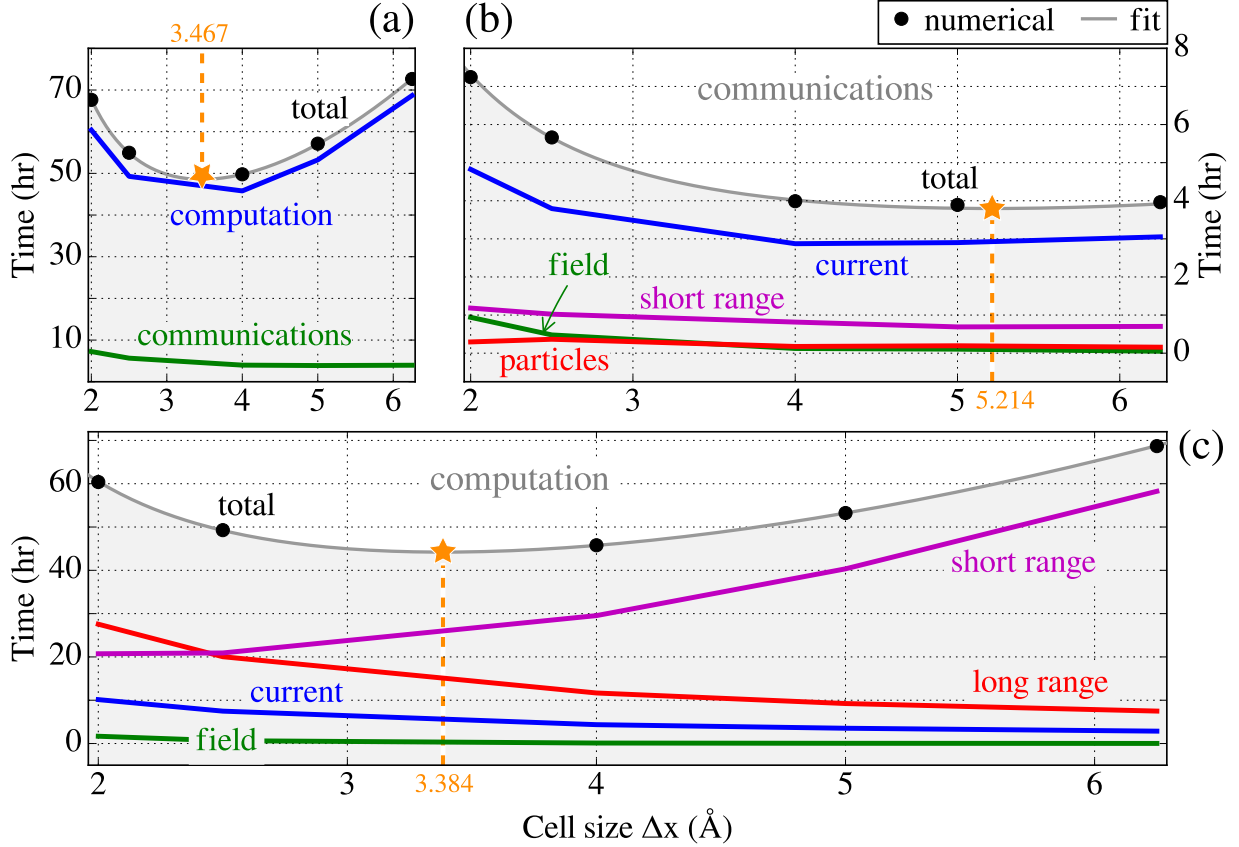


Figure 3.6: Execution times for the most intensive components of parallel MicPIC for 48 process test series over a range of cell sizes and atomic density  $\rho = 22\text{nm}^{-3}$  (see text for more details). All times are averaged over all processes. The top subplots show (a) the total execution times and (b) the time spent on communication functions, including synchronization and overhead related to counting, buffering, etc. Long-range force, short range corrections, current density, and field propagation computation times appear in (c). Totals are shown as large black dots, that are fit using the scaling equations found in Sec. 3.3 [see, in particular, Eqs. (3.2) and (3.4)]. The shaded grey areas outlined by a darker grey curve effectively correspond to (c)  $f_{\text{comp}} = 1.38\Delta x^2 + 92.5/\Delta x + 140.0/\Delta x^4$ , (b)  $f_{\text{comm}} = 2.4 + 0.09\Delta x + 4 \times 10^{-4}\Delta x^3 + 2.5 \times 10^{-4}\Delta x^4 + 18.5/\Delta x^2 + 0.8/\Delta x^3$ , and (a)  $f_{\text{tot}} = f_{\text{comp}} + f_{\text{comm}}$ . Optimal values of  $\Delta x$  for fastest execution were found with downhill simplex minimization and are shown as orange stars. In (a), including both computation and communications, the optimum is found to be  $\Delta x \simeq 0.35$  nm.




Figure 3.7: Performance of MicPIC for 48-process test runs with varying cell size for different materials, represented as plasmas with a single free electron per ion. The simulation volume was held constant and particle number fixed according to the specific atomic density associated with gold (Au,  $58.85 \text{ nm}^{-3}$ ), silicon (Si,  $49.88 \text{ nm}^{-3}$ ), quartz ( $26.38 \text{ nm}^{-3}$ ), and gaseous nitrogen ( $N_2$ ,  $0.05 \text{ nm}^{-3}$ ). a) Arithmetic mean execution time in milliseconds for each step. Violin plots show the distribution of times for a step, each of which is the arithmetic mean over all 48 processes. b) Total execution time required to simulate 1.117fs in hours for each material. Note the distinction between a) and b): the former is per step while the latter is the total time required to complete the number of steps needed to simulate the fixed time window.

## Experiment

A series of tests was performed with 48 processes on four nodes on the Briarée cluster maintained by Calcul Québec, each with two Intel Westmere EP X5650 six-core processors running at 2.667GHz, each with a 12MB cache. Briarée’s nodes feature a minimum of 24GB of memory per node and are interconnected via an Infiniband QDR network. These tests of MicPIC were done using the GNU compiler `gcc 4.9.2` with the compiler flags `-std-c++11 -O3` and the `openmpi 1.8.8` library. Execution times were collected with the `LibSciBench 0.2.2` library at the points indicated in Fig. 3.1.

Tests were on fixed-volume systems with dimensions  $20\text{nm} \times 138\text{nm} \times 200\text{nm}$  and the cell size was varied from 0.2 nm to 0.625 nm. Model materials were represented by a plasma with one free electron per atom. The spatial distribution of the singly-ionized atoms was defined by the physical, crystal structures of the different model materials.

## Performance

The results presented in Fig. 3.6, where the most important contributions to the MicPIC workload are analyzed in detail, show excellent agreement with the theoretical analysis presented in Sec. 3.3. Measured times for computation and communication, as well as the total execution time, as functions of cell size for a fixed time window (5.411fs) are shown in Fig. 3.6a. There is a clear minimum in the computation time where the effects of grid operations balance with the demands of particle-based operations. On the other hand, communication time tends to be largest for small cells. While overall communication time decreases with increasing cell size (Fig. 3.6a), the individual communications do show differences (Fig. 3.6b). For example, in these tests current density communications suffer significantly increased times for small cells, as do those of the field (FDTD and long range force). Increasing the size of a cell does result in more rapid communication of short range data as well, although this effect is less pronounced. Lastly, particle communications are not affected by cell size, which is reflected in Fig. 3.6b.

Interestingly, when the cell size is small, computations involving the FDTD grid are affected (Fig. 3.6c), including the current density and long range force calculations, even though they do not involve iteration over cells. This suggests that the larger memory buffers themselves, with more cells, play a role in the increased computation times. The field propagation (FDTD), however, does strongly depend on the cell size but its contribution to the overall execution time is so small it is effectively negligible. Short range force calculation time sees an opposite effect. Smaller cells mean fewer neighbouring particles within the correction radius and therefore shorter execution.

Conversely, the effect of particle density on the performance of MicPIC is emphasized in Fig. 3.7. Fig. 3.7a shows the distribution of arithmetic mean execution times per step for all 48 processes in the test run as violin plots [143, 144] for each of the four simulated substances. Violin plots display a kernel density estimation plot of the probability density function of a distribution of values. The width of the “violin”-shaped figure corresponds to the likelihood of a measurement. This provides a concise way to illustrate detailed statistical information. Coloured bars indicate the maximum and minimum values of the distributions, including outliers. The middle bars represents the medians of the distribution. In the second subplot, Fig. 3.7b, the total execution times are plotted for a fixed time window of 1.117 fs.

For the lowest density ( $\text{N}_2$ , gas phase), the total particle count is so low that MicPIC behaves essentially like FDTD alone, scaling as  $(\Delta x)^{-4}$ . The next most dense medium, quartz, is modelled as a collection of  $\text{SiO}_2^+ + e^-$  units positioned at the centres-of-mass of the silicon and two of their neighbouring oxygen atoms in a trigonal crystal (space group  $\text{P3}_121$ ). The minimum execution time for this case occurs for  $\Delta x \simeq 0.4$  nm.

The final two media (Au, face-centred-cubic and Si, diamond cubic,  $Fd\bar{3}m$ ) each have an optimum cell size of  $\Delta x \simeq 0.25$  nm. It is worth mentioning that for performance tests of parallel MicPIC, the physical arrangement of the atoms or molecules is of little consequence; only the density matters. We can therefore conclude that for most solid density materials 0.25 nm is a good choice of cell size to minimize the total execution time. Assuming 20 sampling points per wavelength, it appears that this is enough to resolve the propagation of extreme ultraviolet and soft x-ray radiation.

### 3.4.2 Parallel Speedup and Efficiency

We characterized the performance of parallel MicPIC in terms of its parallel speedup and efficiency, often referred to as *strong scaling*. The speedup is a measure of the increase in speed at which a set of instructions are executed relative to some reference time. If the time required for one process to execute a job is defined as  $T(1)$  and the time for  $n$  processes to complete the same job in parallel is  $T(n)$ , the speedup attained by using  $n$  processes is [145]

$$S_1(n) = \frac{T(1)}{T(n)}. \quad (3.6)$$

If an  $n$ -process job is  $n$  times faster, the speedup is  $S_1(n) = n$  and the scaling is said to be linear. While it is often true that the speedup of a code is sub-linear, due to bandwidth limitations and synchronization, it is possible for the speedup to be super-linear. This simply implies that a set of instructions is more efficiently executed with more processes, often by reducing local memory access operations. This naturally suggests the introduction of a second metric, the parallel efficiency, defined as

$$E_1(n) = \frac{S_1(n)}{n} \equiv \frac{T(1)}{nT(n)}. \quad (3.7)$$

A parallel efficiency of 1 (100%) is equivalent to a linear speedup.

It can also be useful to determine the relationship between number of processes  $n$  and the fraction of time spent on the computation ( $f_p$ ) and communication ( $f_c = 1 - f_p$ ) during execution. This provides an alternative method to estimate where communications begin to saturate execution time and the value of adding more processes to a job. The two quantities  $n$  and  $f_p$  are linked via the speedup of execution for  $n$  processors in the form of Amdahl's Law [146]:

$$S_b(n) = \frac{1}{[1 - f_p(b)] + \frac{b}{n}f_p(b)} \quad (3.8)$$

where  $f_p(b)$  is for a base case of  $b$  processes. Corollary to this, increasing the number of processes in a job does not affect the communication time.

In terms of measured times  $T_p(n)$  and  $T(n)$ , the computation and total (computation+communication) times, respectively, the fraction of execution time spent on computation by  $n$  processes is

$$f_p(n) = \frac{T_p(n)}{T(n)} = \frac{T_p(n)}{T_p(b)} \frac{T_p(b)}{T(b)} \frac{T(b)}{T(n)} \quad (3.9)$$

Ideally the parallelizable portions of the code are perfectly so: that is, the execution time for  $n$  processes is  $n/b$  times faster than the  $b$ -process base case. Thus, the first term in the latter part of Eq. (3.9) is  $\frac{b}{n}$ . The second term is the expression for computation time fraction for the base case,  $f_p(b)$ , and the third is simply the speedup from Eq. (3.6). Combining these with Eq. (3.8) yields

$$f_p(n) = \frac{b}{n} \frac{f_p(b)}{[1 - f_p(b)] + \frac{b}{n}f_p(b)} \quad (3.10)$$

This simple model serves as a guide to predict the fraction of time a certain number of processes will spend on the parallelizable portions of MicPIC. It should be noted that deviations from the expected behaviour will arise due to the *fraction* of time for the communication portion of the code not being strictly constant and also due to subdomain geometry.

To get the speedup and efficiency as a function of the number of processes, the execution time was measured that parallel MicPIC took to complete 100 time steps of some test simulations described below. The global physical dimensions were chosen in such a way as to attempt to keep a reasonable subdomain volume-to-surface ratio for each process over a wide range of number of processes.

## Experiment

A first set of simulations considered  $\sim 10^8$  particles distributed over an FCC lattice structure to mimic a model silica crystal ( $\rho = 22\text{nm}^{-3}$ ), henceforth referred to as *Series A*. The cell size was set to 0.4 nm, and the physical size was  $200\text{nm} \times 152\text{nm} \times 80\text{nm}$ . Domain decomposition was performed in the first two dimensions ( $x$  and  $y$ ) and done such that the desired total number of processes was divisible by an integer number of processes per node. Times were measured at the points indicated in Fig. 3.1 and were used to calculate the parallel speedup and efficiency for this scenario using a series of tests with process counts ranging from 4 to 120 on the Briarée cluster described in 3.4.1.

With the aim of measuring the parallel performance of MicPIC for larger parallel jobs, another set of test simulations, (*Series B*), were conducted with a  $52\text{nm} \times 343\text{nm} \times 2.4\mu\text{m}$  domain containing  $5.02 \times 10^9$  particles ( $\text{Au}^+ + \text{e}^-$ ) in a FCC structure. The test series involved process counts ranging from 168 to 2040 with the simulation domain being split in the  $y$  and  $z$  directions and a cell size of 0.25 nm, corresponding to the minimum of the curve for gold (Au) density in Fig. 3.7.

We emphasize that modern, distributed-memory supercomputer architectures are characterized by interconnected computer nodes that typically have multi-core central processing units (CPUs). So far, “subdomain” and “process” referred interchangeably to MPI instances that run on different nodes and exchange information through the network, or to instances running on the same node—often even on the same CPU—and sharing information *via* on-chip communications.

## Performance

The results for the fixed problem size in *Series A* are illustrated in Figs. 3.8 and 3.9, which show the effect of increasing the amount of processes involved in the simulation on performance. In particular, Fig. 3.8 shows a breakdown by component for both computation and communication, while Fig. 3.9 shows the parallel performance characteristics discussed in Section 3.4.2 in addition to the largest memory allocations for each test.

The subset of computation functions illustrated in Fig. 3.8a-c indicates a common trend toward decreasing overall execution time with increasing number of processes. Indeed, each of these sees an order-of-magnitude decrease in execution time from, for example, 12 to 120 processes. While the second subset of functions – the communications (Fig. 3.8d-h) – does incorporate some operations dependent upon subdomain size, there is no noticeable trend akin to what is observed in the computation times. Instead, the communication times exhibit much more variation and no dependence on the number of processes.

In general, the particle and short range communications have the weakest dependence on the physical dimensions of the subdomain, roughly scaling with the number of particles in the subdomain and in the boundary region, respectively. The remaining components, while generally faster, depend on the precise

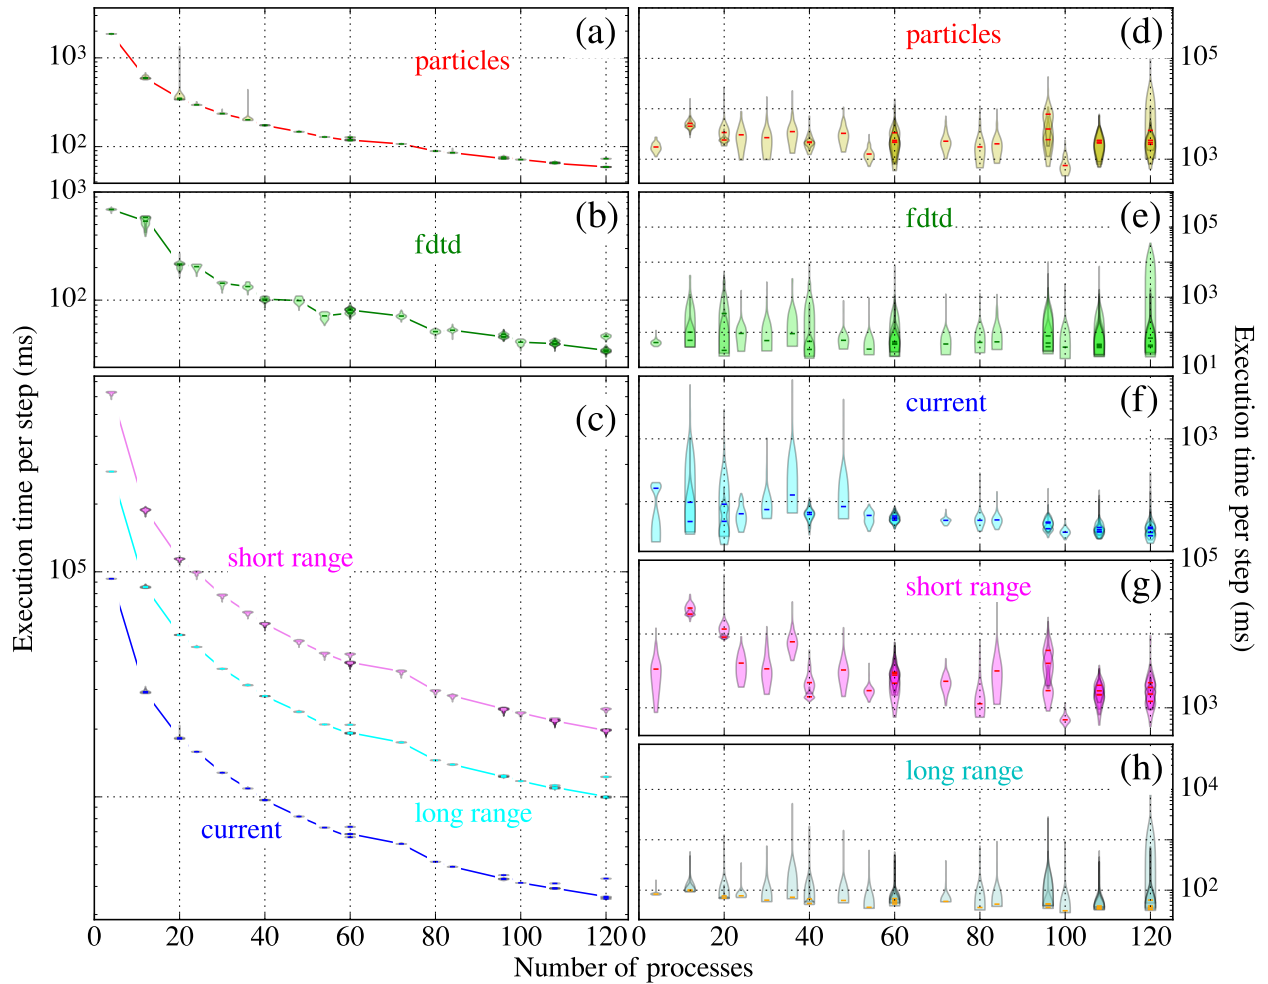


Figure 3.8: Execution times of the main components of the parallel MicPIC algorithm vs number of processes for a test system of  $1.04 \times 10^8$  particles. Each test is shown as a violin plot of the distribution of median execution time of a step across all processes. (a) Particle mover time. (b) FDTD field propagation. (c) Top to bottom: short range force, long range force, and current density. (d-h) Respectively, particle, FDTD, current, short range, and long range communication times. All times are in milliseconds.

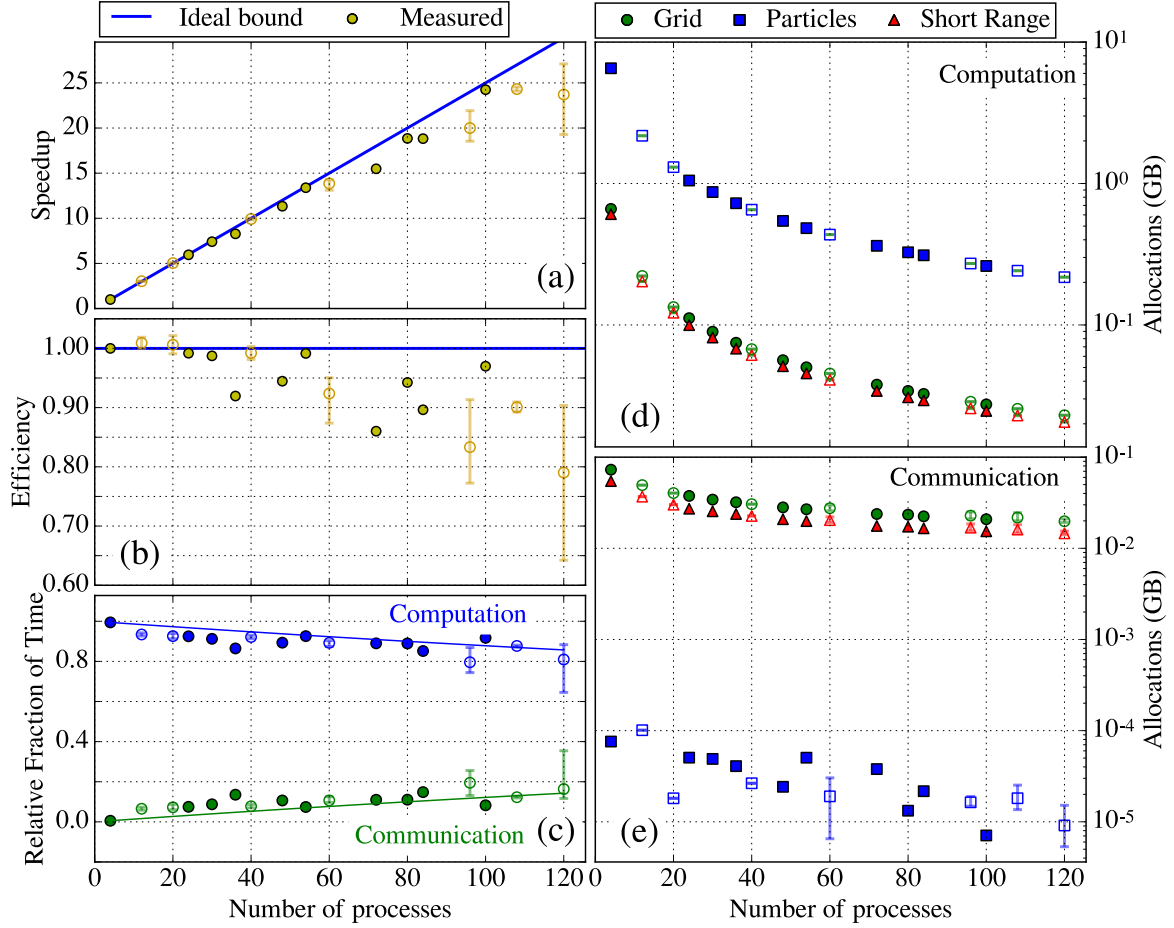


Figure 3.9: Parallel speedup (a) and efficiency (b) of Parallel MicPIC for a  $20\text{nm} \times 138\text{nm} \times 200\text{nm}$  system of  $\sim 10^8$  particles. In (b), as the size of the subdomains becomes smaller with increasing number of processes, there is a corresponding drop in the parallel efficiency caused by an increase in the relative fraction of the execution times of the parallel functions (c). Solid lines correspond to Eq. (3.10). (d) and (e) show the largest mean persistent allocations of memory for computation and communication. Some process counts were tested for multiple domain decompositions (max of 4), which are summarized as bars (extrema of the samples) and hollow markers. In a), b), d), and e) markers correspond to the arithmetic means of the speedup and efficiency values, whereas in c) they correspond to the geometric means of the fractions of computation and communication time. This mean is suggested as a way to summarize ratios [1].

dimensions and resolution of the FDTD grid and will therefore exhibit more variability with respect to subdomain geometry.

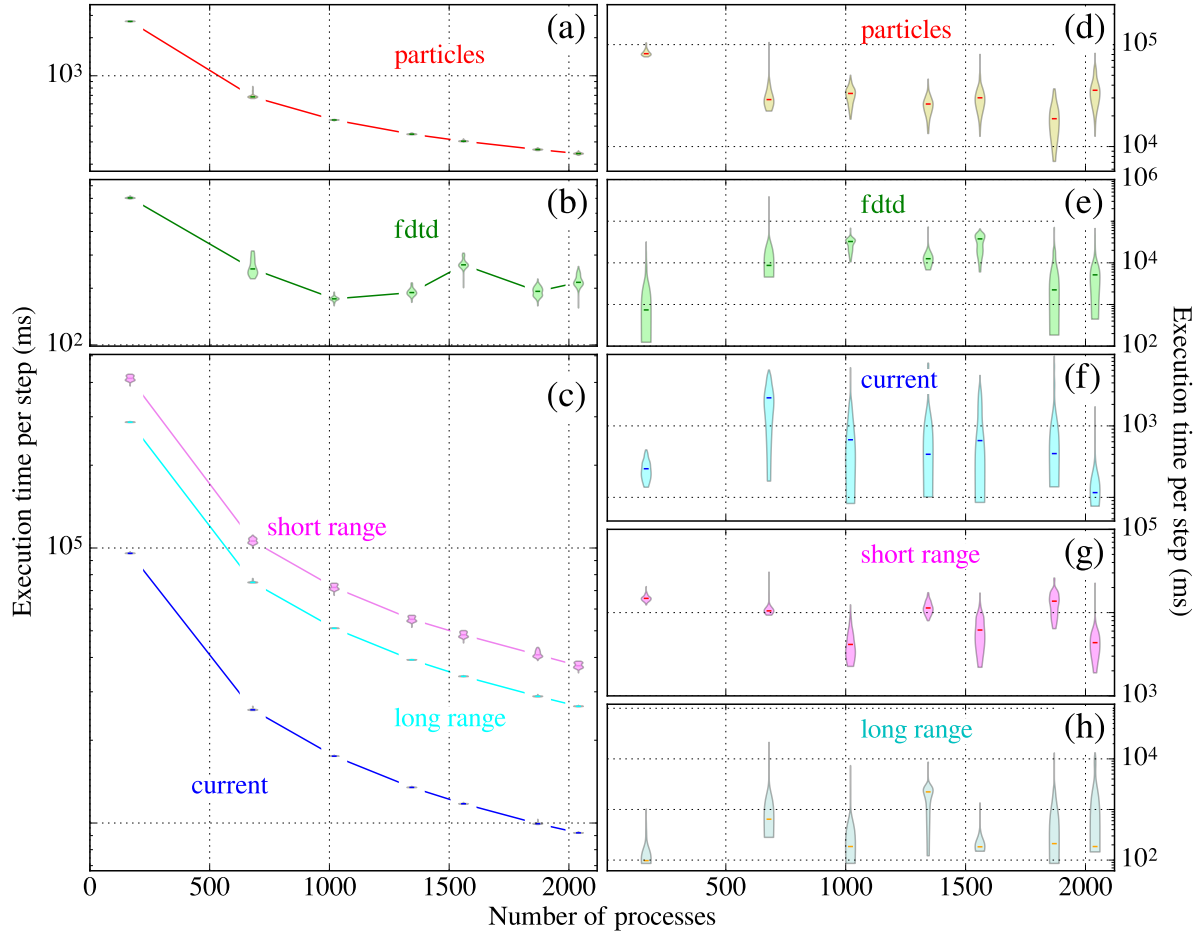


Figure 3.10: Execution times of the main components of the parallel MicPIC algorithm vs number of processes for a test system of  $5.02 \times 10^9$  particles. Each test is shown as a violin plot of the distribution of median execution time of a step across all processes. The medians of the distributions shown are visible as ticks within the violins. (a) Particle mover time. (b) FDTD field propagation. (c) Top to bottom: short range force, long range force, and current density. (d-g) Respectively, particle, current, short range, and field communication times. All times are in milliseconds.

Fig. 3.9 shows that for jobs with the fewest number of processes, the speedup is almost linear (Fig. 3.9a), resulting in an efficiency of 90% or greater for the majority of tests (Fig. 3.9b). We have shown in the introduction to Section 3.4 that with an increase in the number of processes, the subdomains' volume shrinks with respect to the parallel boundary surface layer. As a consequence, a larger portion of the execution time and memory is used for buffering, communications, and synchronization (Fig. 3.9c). This results in the saturation of the speedup and a decrease in efficiency seen in Fig. 3.9a-c, wherein the fraction of execution time spent on communication for a job (green dots) increases with more processes. Extrapolating suggests the communication times will become comparable to or greater than the computation fraction (blue dots). Solid lines in Fig. 3.9c are derived from Eq. (3.10) by substituting the measured times for the four process base case in this series and serve as a guide for when to expect communications to become significant.

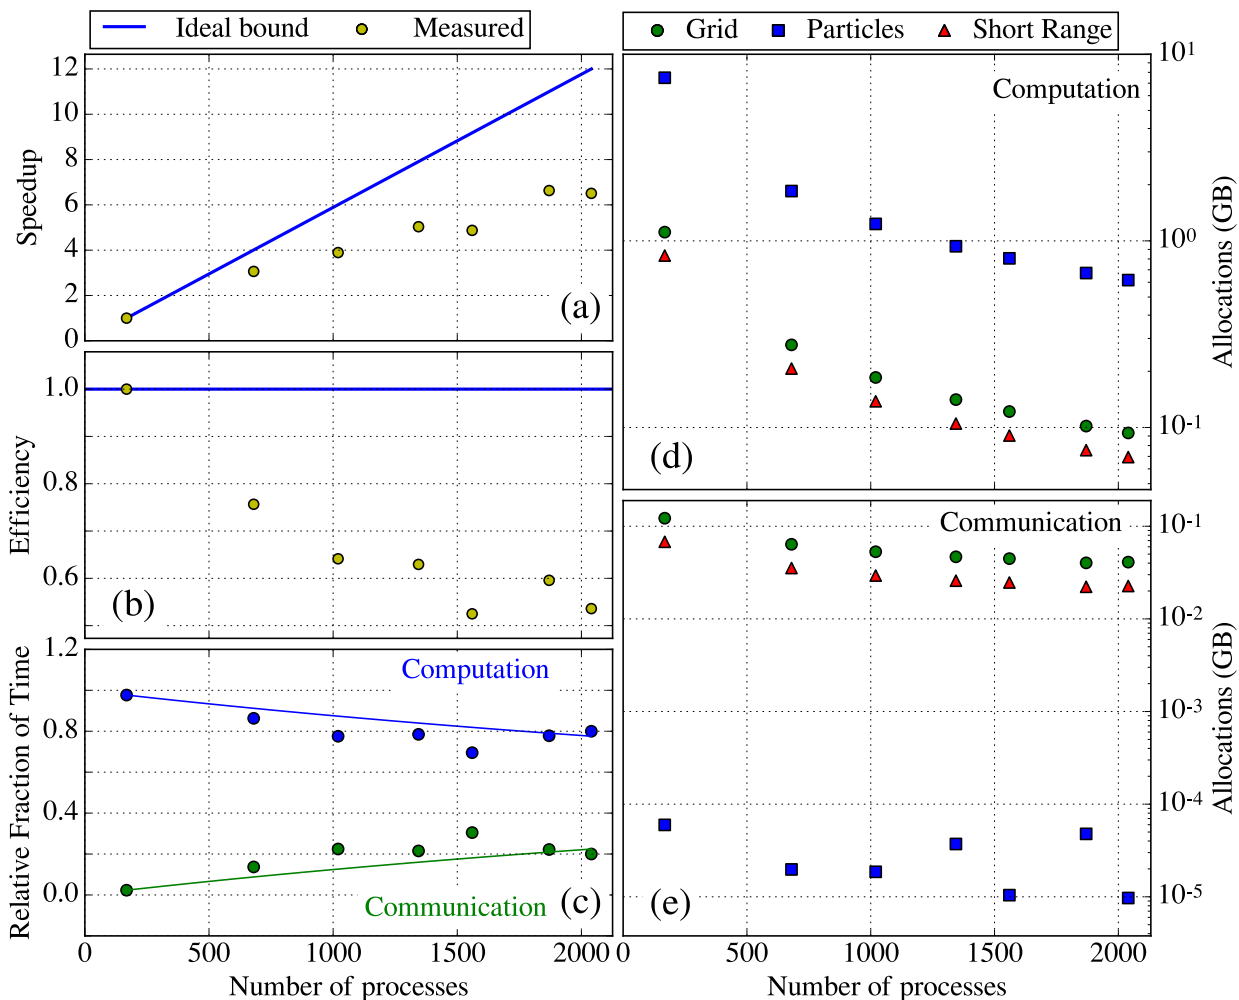


Figure 3.11: Parallel speedup (a) and efficiency (b) of Parallel MicPIC for a  $52\text{nm} \times 343\text{nm} \times 2.4\mu\text{m}$  system of  $5.02 \times 10^9$  particles. Similar to what is seen in Fig. 3.9, both quantities tend to decrease as the sizes of the subdomains become smaller with increasing number of processes. (c) and (d) display the mean persistent allocations of memory on a process for the grid, microscopic corrections, and particles.

Deviations from the behaviour predicted by Eq. (3.10) arise due to differences in parallel efficiency with respect to the base case and subdomain geometry. The effect of subdomain geometry on these deviations is illustrated in Fig. 3.9d-e, which shows the mean size of the main blocks of memory allocated for use in computations (Fig. 3.9d) and communications (Fig. 3.9e) by particle (squares), FDTD and weighting (circles), and short range (triangles) functions for a process; in essence, the size of a subdomain in memory. Functions involving data stored in these memory blocks themselves have various smaller, ancillary, temporary allocations. The data depicted correspond to persistent allocations, except in the case of particle communications, which vary from step to step and whose maximum values for each test are shown (Fig. 3.9e, squares). Note how the allocations related to local computations decrease nearly two orders of magnitude, whereas those related to communication decrease by less than one. Furthermore, increasing the number of processes leads to the main memory allocations required for computation to become comparable to what is needed for communication. This tendency to have comparable computation and communication related allocations for jobs

with increasing process counts correlates to the trends seen in Fig. 3.9c.

The performance results of tests on *Series B* are displayed in Figs. 3.10 and 3.11. The overall trend is the same seen in *Series A*. The computation components of the code scale quite well, while communications exhibit more variability as well as wider distributions of times. The observed variations in the individual component times are thought to be primarily the result of synchronization times resulting from less-than-optimal load balancing, fluctuating infiniband network utilization, highlighting the challenge of choosing an appropriate domain decomposition, and the number of processes per node. To emphasize this last point, note that some tests were decomposed in such a way as to necessitate fewer than the maximum number of processes per node.

In Fig. 3.11, it is observed that for this particular model system the parallel efficiency steadily decreased up to 2040 processes, suggesting that parallel MicPIC can realize high parallel efficiency for moderately sized jobs provided there is a sufficient number of particles and the average subdomain size is large. Higher process counts see further gains in speedup but less efficient parallel execution. It is evident in Fig. 3.11 that the efficiency suffers from beyond 1500 processes, after which point the communications begin to consume more of the execution time.

The cause of this drop is apparent in Fig. 3.10 and results from particle and field communications taking significantly longer than expected. We suspect that this is rooted, in addition to the above fact, in how these operations overlap with other communications and is a focus of further optimization. Overall, parallel efficiency is relatively high when the ratio of computation to communication memory utilization is also high (Fig. 3.11c,d).

For the tests in *Series A* and *Series B*, it was necessary to consider various  $\text{node} \times (\text{process}/\text{node})$  options and observed important variations in the parallel efficiency. This effect is most evident in Fig. 3.11(b). Unfortunately, it was not possible to perform a rigorous statistical sampling in the 200-2000 processes range with the available computational resources. Still, it is seen in Fig. 3.11(b) that the global trend is a decrease of the parallel efficiency with increasing process count, which is explained by the increase in relative fraction of execution time dedicated to parallel operations.

We emphasize that optimization of the  $\text{node} \times (\text{process}/\text{node})$  strongly depends on the computer architecture and the resource allocation policies of the computer system (e.g., user-dedicated versus shared node allocation). Ideally, MicPIC parallelization should be properly hybridized (MPI + multithreading) to fully benefit from the multiple CPU cores available on the nodes, and, potentially, from hyper-threading. Nevertheless, parallel efficiency of current parallel MicPIC remains serviceable for process counts up to  $\sim 1500$ . Suboptimal synchronization times and subdomain volume-to-surface ratios resulting in unnecessarily large ratios of communication time to computation time as well as physical resource utilization indicate that better can be expected with basic load optimization.

### 3.4.3 Scalability for massive-scale jobs

We emphasized in the introduction that the ultimate goal in modelling strong-field light-matter interactions is to resolve the microscopic, atomic-scale processes along with the associated electromagnetic radiation over distances comparable to the laser wavelength. Ultimately for infrared laser light and solid-density materials, it sums up to simulating volumes of material that are approximately  $1\mu\text{m}^3$  and tracking approximately  $10^{10} - 10^{11}$  particles. We show below that MicPIC can effectively reach this goal by handling such large systems of particles on a massively parallel high-performance computer architecture, here in particular, on a 4-rack IBM Blue Gene/Q supercomputer with a total of 65 536 processes.

The Blue Gene/Q computing cluster administered by the Southern Ontario Smart Computing Innovation Platform (SOSCIP) is Canada’s largest computer. It is effectively hosted at the University of Toronto’s SciNet HPC facility. The machine itself consists of 4,096 1.6GHz PowerPC CPUs, each with 16 cores. Each compute node has 16GB RAM, making up to 64TB RAM available. Tests of the scalability of parallel MicPIC on that particular supercomputer system involved measuring its performance with increasing simulation size. This was done by setting first a small physical domain with  $\sim 10^6$  particles, which can reasonably fit on single node, and by constantly increasing the total number of particles by adding more nodes (while keeping a constant load per node) to ultimately reach the  $10^{10} - 10^{11}$  particle range.

## Experiment

The basic test system for the smallest block on the machine (64 nodes; 1024 processes) consisted of a thin plasma film with atomic density of gold with the dimensions  $256\text{nm} \times 46.25\text{nm} \times 1024\text{nm}$  ( $\text{Au}^+ + \text{e}^-$ ), which was chosen to yield a reasonable average subdomain volume-to-surface ratio. Each subsequent larger test was simply double the size of the previous one, doubling first in one parallelized direction and then the other. Due to the lower amount of memory available per process, it was necessary to have smaller subdomains and fewer particles per subdomain than in the tests performed for the speedup/efficiency runs in Section 3.4.2. However, due to the larger number of nodes available, much larger systems of particles were realizable. Effectively, parallel MicPIC is shown to consistently perform with nearly constant execution time for simulations ranging in size from  $1.4 \times 10^9$  to  $4.5 \times 10^{10}$  particles using 1 024 to 32 768 processes (see Fig. 3.12).

## Performance

In the light of this MicPIC weak scaling test, it appears that using twice as many processes to run a simulation twice as big has a minimal impact on the total communication load. This is primarily due to the fact that MicPIC uses a point-point communication approach, where a specific subdomain needs to exchange a minimal set of boundary information with only a few neighbour processes, overall minimizing the inter-node communications and usage of the network bandwidth. As seen in Fig. 3.12, this leads to the global time taken by the communications to remain nearly constant even for these massively parallel simulations, typically two orders of magnitude less compared with the computation part. However, on this particular supercomputer it was observed that communications do reach a limit when all the 65 536 processes of the machine were involved, but this is not yet well understood. Most likely, MicPIC was here limited by this specific computer system architecture. It is safe to infer that scaling up to even bigger parallel jobs would be possible on larger systems.

## 3.5 Conclusion

In this chapter, it has been shown that the current distributed parallel implementation of the microscopic particle-in-cell method can simulate the dynamics of massive ensembles of particles at solid densities. Using the largest supercomputer in Canada (IBM Blue Gene/Q with 65 536 physical cores), it has been shown that it is effectively capable of dealing with up to  $10^{11}$  particles, which corresponds approximately to a  $1\mu\text{m}^3$ -volume of matter at solid gold density. In particular, the implementation of MicPIC’s current density and electromagnetic field weighting schemes were outlined, as well as the short range correction routine that implements the atomic-scale particle collisions. The distributed memory parallelization of these components was given special attention along with details of how the different boundary communications are handled.

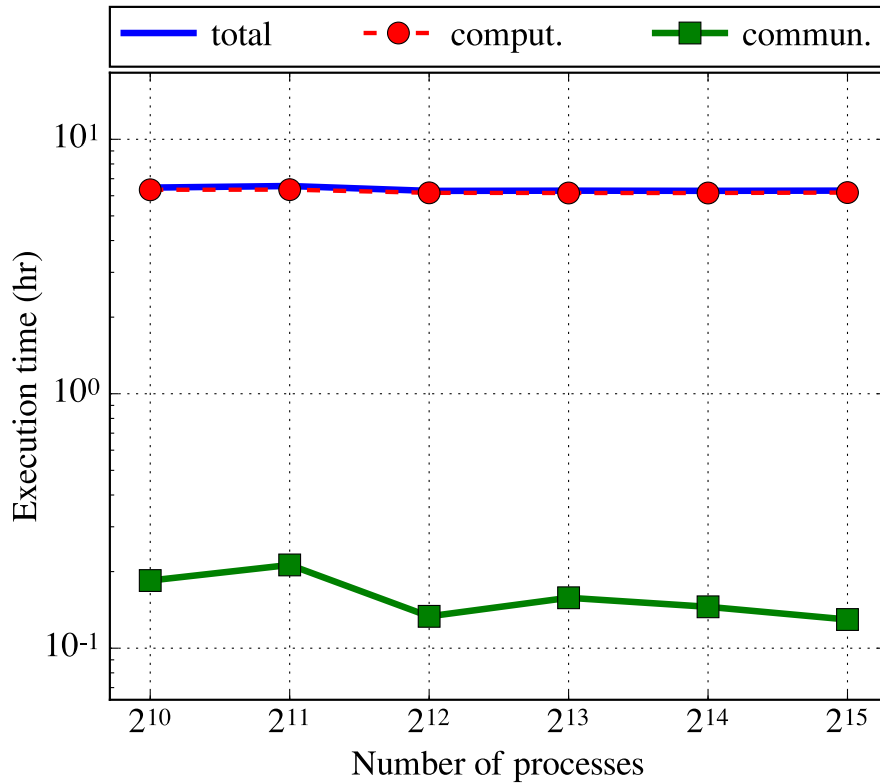


Figure 3.12: Results from the BlueGene/Q scalability tests. The base-10 logarithm of the execution time in hours is plotted against the base-2 logarithm of the number of processes. The test series begins with  $1.4 \times 10^9$  particles on 1024 processes and doubles in size up to  $4.5 \times 10^{10}$  particles on 32768 processes. It appears that using twice as many processes to run a simulation twice as big has a minimal impact on the global computation and communications execution time.

Parallel MicPIC’s performance was characterized by analyzing the execution time of its principal components with test simulations. On a mid-scale supercomputer, strong scaling results for  $5 \times 10^9$ -particle computations with up to 2040 processes indicate that parallel efficiency is reduced significantly beyond  $\sim 1500$  processes. This suggests certain avenues for future optimization. Furthermore, the weak scaling analysis performed with large-scale simulations on a 4-rack IBM Blue Gene/Q (up to 65 536 physical cores) suggests that parallel MicPIC’s capabilities can be extended even beyond  $10^{11}$  particles on larger supercomputers.

---

CHAPTER

FOUR

---

# COLLISION DYNAMICS IN STRONGLY COUPLED PLASMAS

## 4.1 Introduction

Numerical modelling of plasma dynamics has been restricted, until recently, to small systems in the electrostatic limit or large systems in the fluid limit, where approximations and statistical models are required. Simulating hot, dense, mesoscopic and macroscopic plasmas fully microscopically has not been possible until now. Presented here is the first such large-scale microscopic simulation of a *strongly coupled* plasma interacting with an intense, ultrashort laser pulse on length scales greater than the laser wavelength. As an example, the effects of collisional dynamics on the evolution of the system of particles and generation of high harmonics will be explored, although there are numerous possible applications.

Plasmas share a history with computational physics dating back to the middle of the last century [147]. Perhaps the most popular method, particle-in-cell (PIC), is a mainstay in the field of plasma simulation and is formally treated in the literature [70, 101]. Traditionally, PIC models particles or statistical ensembles of particles (superparticles) on a discrete mesh, while the particles' motion and charge is used to generate current densities to be used as sources of electromagnetic field. It can also fully account for radiation propagation from externally introduced sources and from the motions of particles and has been used to study HHG and other nonlinear phenomena in solids over the years [148, 37, 92, 149, 150]. Not without its disadvantages, PIC uses particles which are typically much larger than the atomic scale and therefore suffer from reduced accuracy when dealing with n-body collisions and microscopic fields. Moreover, representing atoms, ions, and electrons as individual PIC particles does not accurately reproduce the inter-particle forces needed to resolve collisional effects.

Microscopic Particle-in-Cell (MicPIC) is a more recent technique building upon PIC by incorporating corrections to the inter-particle force within the neighbourhood of each particle [106, 104]. MicPIC is a fully 3D simulation which allows interactions between closely packed particles and their trajectories to be computed with ångström resolution, without the aid of additional statistical models for collisions. Taking advantage of advances in computing power, MicPIC has also been shown to exhibit good strong and weak

parallel scaling and to be capable of handling systems of nearly  $10^{11}$  particles [151].

Plasmas have played an important role in the understanding of HHG, since it was first observed, as technology allows researchers to steadily progress towards probing higher densities, with greater intensities, on shorter timescales [15, 13]. More recently, there has been considerable interest in utilizing solid density plasmas as a means to characterize and control the emission of high-order harmonics over a broad range of energies [37, 92, 149, 39, 36, 152].

This section is presented as follows. Being well-suited to investigating the behaviour of a plasma mirror interacting with an intense femtosecond (fs) laser pulse, MicPIC is used to demonstrate particle collision effects with microscopic detail in a set of test simulations, highlighting the subsequent harmonic spectrum generated. In Section 4.2 the theory of harmonic generation in overdense plasmas is discussed. Following this, Section 4.3 outlines the simulations that were performed, including the numerical techniques and physical parameters used. Finally, we analyze the results of our simulations in Section 4.4, discuss how they compare to the theory outlined below, the role played by collisions between particles, and how they fit into HHG from plasmas as a whole.

## 4.2 Theoretical Background

An intense, ultrashort laser pulse incident upon a solid density material rapidly ionizes the atoms or molecules near the vacuum interface. The resultant plasma exists at similar densities and is highly reflective to radiation at frequencies less than the plasma frequency,  $\omega_p = \sqrt{n_e e^2 / \epsilon_0 m_e}$ . At solid density,  $w_p$  has values corresponding to the near-infrared (NIR) to extreme ultraviolet (XUV) range of the electromagnetic spectrum. This plasma mirror, as it has come to be known [39], exhibits a sharp discontinuity in the density at the vacuum interface. Due to the femtosecond-scale duration of the interaction with the incident pulse, the plasma is only able to expand a fraction of the laser wavelength. This leads to strong electric fields ( $> 10^9$  V/m), which rapidly accelerate electrons in the vicinity of the plasma-vacuum interface to high energies on attosecond timescales.

The non-uniform plasma formed by the ionization in the presence of the laser at angular frequency  $\omega_L$  typically exhibits a density profile extending from zero up to some maximum, which is many multiples of the critical plasma frequency  $\omega_c = \omega_L$ , corresponding to a critical density of  $n_c = \epsilon_0 m_e \omega_L^2 / e^2$ . Resonant absorption allows plasmons excited by the incident pulse to efficiently couple to electromagnetic field and, through a process called linear mode conversion, to excite higher harmonics of the laser [35, 150]. When an electromagnetic wave at frequency  $\omega = \omega_c$  propagates deeper into the plasma, it can excite a second harmonic, in combination with a plasma wave at the same frequency. This process can continue in a cascading fashion up to the maximum density available in the material.

The applicability of resonance absorption and linear mode conversion breaks down in those cases for which the characteristic electron oscillation amplitude  $x_{osc} = eE_L / m_e \omega_L^2$  becomes comparable to or exceeds the plasma density scale length  $L = (|\partial n / \partial x|_{n=n_c})^{-1}$ , resulting either from relativistically intense laser radiation ( $\gtrsim 10^{18}$  W/cm<sup>2</sup>) or very steep plasma densities [37, 35]. Two of the most promising theories put forth to deal with these cases theoretically are the *Sliding* or *Relativistic Oscillating Mirror* (ROM) model [37, 153, 40] and *Coherent Wake Emission* (CWE) [92, 149, 39]. The former takes place at moderate intensities. Brunel demonstrated an alternative energy transfer mechanism for these cases, whereby bunches of electrons absorb energy through vacuum heating and carry it back into the interior of the plasma [38].

HHG in plasmas has been demonstrated in both theory and experiment primarily via these two mecha-

nisms. ROM is the more mature of the two [37, 153] and describes radiation from plasma density surface oscillations in cases where the electron quiver velocity  $v_{osc} = eE/\gamma m_e \omega$  is relativistic, where  $\gamma = \sqrt{1 + a^2}$  is the Lorentz factor and  $a = eE_L/m_e \omega_L c = \gamma v_{osc}/c$  is the normalized field amplitude. Such conditions occur when  $a \gtrsim 1$  ( $I\lambda^2 \gtrsim 1.37 \times 10^{18} \text{ W cm}^{-2} \mu\text{m}^2$ ). Due to the inherent relativistic motion of the particles in this model, it is currently beyond the reach of MicPIC.

The second regime, CWE, occurs for well below non-relativistic ( $a < 1$ ) intensities [92, 149, 39]. This is characterized by a roughly linear dependence of the conversion efficiency on laser intensity and a harmonic cutoff around  $\omega = \omega_{p,max}$ , where  $\omega_{p,max}$  is the maximum local plasma frequency. For intense fs-pulses, the plasma formed by the laser-matter interaction expands on time scales much longer than the pulse duration, which contrasts with early experiments with HHG using nanosecond pulses on solid targets [23, 24]. In these early cases, the highest generated harmonic was thought to arise from the interplay between the expanding plasma density profile and the radiation pressure applied by the laser. Since this does not have time to occur with fs-pulses, there must exist some other mechanism at work.

Brunel showed that electrons emitted at different times from the surface of an overdense plasma with a density gradient are subsequently accelerated by the laser electric field and absorb many multiples of the laser photon energy before a fraction of them return to the plasma simultaneously [38]. The speed of these electrons' return is dependent upon when during the laser cycle they leave the overdense plasma region, with electrons returning later having significantly higher speeds (up to  $\sim 0.3c$  possible) than those beginning their excursion earlier in the cycle. This behaviour results in bunches of hot electrons crossing the critical density surface of the plasma and carrying the energy absorbed from vacuum heating into the deeper portions of the plasma where the electric field is unable to penetrate [39].

### 4.3 Numerical Methods

PIC codes are well-suited to simulate many aspects of the ultrafast dynamics described above. This numerical approach is able to model large many-particle systems in a relativistic and fully electrodynamic fashion. Particles are typically grouped into larger ensembles (superparticles), whose weighted charge densities act as source terms in Maxwell's equations. The interactions between these particles are mediated by the long-range fields, which, coupled with the spatially extended nature of the particles themselves, provides a course-grained picture of the dynamics of these systems. Over the past several decades PIC simulations have played an important part in providing a more detailed view of plasma dynamics. Indeed, there have been several different incarnations of PIC used to support theoretical models as well as experimental findings [153, 37, 92, 39]

At solid densities and high temperatures, the plasma becomes strongly coupled and the number of particles  $N$  in the Debye sphere of radius  $\lambda_D = \sqrt{\epsilon_0 k_B T / e^2 n}$  is on the order of unity [10]. In these conditions, short range interactions between particles become important and collisions play a role in the dynamics of the plasma.

It is currently not known precisely what the nature of the differences will be when using MicPIC to investigate the effects of strong fields interacting with solid density materials on length scales spanning many wavelengths. This is particularly interesting in the case of coherent wake emission processes, where the intense electric fields produce fast moving of electrons, which propagate through and along the surface of a sample. In this section, the first preliminary results of simulations are presented of an intense, few-cycle pulse incident upon a thin sample of solid density plasma on the micron scale. No other simulation is known on

such a scale, which provides similar fully microscopic and electrodynamic detail. In order to place the results in the proper context, a second simulation was performed on exactly the same physical conditions using only the PIC method, simply by deactivating the short range force corrections in MicPIC. It is important to note here, however, that this is PIC in the “MicPIC limit”, where the particles are PIC-like in that they receive no microscopic corrections but are still treated as individual particles. In this sense, PIC is taken to a scale beyond which it can grow no smaller and MicPIC is the next logical step. If any differences arise between these two cases, it is reasonable to assume that such differences are minimized with respect to the granularity the PIC simulation.

The results presented here are derived from data from the aforementioned set of PIC and MicPIC runs. The sample is a thin film of amorphous silica (number density  $n = 22 \text{ nm}^{-3}$ ) whose surface is coplanar with the  $yz$ -plane and has  $(x, y, z)$  dimensions equal to  $3.00 \mu\text{m} \times 72 \text{ nm} \times 200 \text{ nm}$ . The simulation domain spans a region with spatial dimensions  $3356 \text{ nm} \times 436 \text{ nm} \times 20 \text{ nm}$ , centred on the sample, and is periodic in  $z$ , yielding a semi-infinite sample in that direction, while the  $x$  and  $y$  directions are bounded by 8 cell thick PMLs. The cell size itself was chosen to be  $0.4 \text{ nm}$ , which is close to the optimum value for this density (see Section 3.4.1), yielding a time step with a duration of  $7.696 \times 10^{-19} \text{ s}$ . The depth of the sample in the  $y$ -direction was chosen to provide a reasonably but not unrealistically thin sample that would execute in a reasonable amount of time, while still modelling surface effects and nonlinear plasma wave phenomena.

The incident laser pulse was a p-polarized cosine plane wave of wavelength  $800 \text{ nm}$  with a Gaussian temporal profile of  $8 \text{ fs}$   $e^{-1}$ -width and zero carrier-envelope phase (CEP), propagating in the direction  $\hat{k} = (1/\sqrt{2}, 1/\sqrt{2}, 0)$ , corresponding to a  $45^\circ$  angle to the plane of the surface, with an intensity of  $5 \times 10^{15} \text{ W/cm}^2$  near the upper end of the CWE regime. With these parameters, the pulse requires approximately  $22 \text{ fs}$  to transit the sample from the leading edge of the pulse at  $e^{-1}$  intensity to the counterpart trailing edge of the pulse. This was introduced via the TFSF method (see Section 2.2.2), the origin of which was located at cell coordinates  $(x, y, z) = (447, 24, 0)$ . The peak of the pulse arrived at this origin after a delay of  $10 \text{ fs}$  or roughly  $1.3 \times 10^4$  time steps.

Particles were initialized as ion-electron pairs, where the mass of an individual ion was set to  $60.0843 \text{ amu}$ , corresponding to the atomic mass of  $\text{SiO}_2$ . Microscopic corrections to the PIC force were applied within a three cell radius of each particle in the domain in the case of MicPIC.

Each simulation was performed on the Briarée cluster using the same compiler options and hardware parameters previously provided in Section 3.4.1. The simulation domains were decomposed into 408 equally sized subdomains in the  $x$ -direction, each split further in two in the  $y$ -direction. The PIC and MicPIC simulations required a total time of approximately 87 and 217 hours, respectively, although it was possible on the available resources to execute both simultaneously. The simulated time in both was  $31.4 \text{ fs}$ .

## 4.4 Simulation Results

The two cases are compared below by analyzing the electromagnetic field and particle velocity data generated by the code. Four instants in time have been selected to demonstrate the interaction between laser and medium, and to illustrate the subtle effects attributable to the microscopic interactions. To this end, selected regions are shown in greater detail due to the similar appearance of the electron speeds on the scale of the whole domain. Comparisons are made between the MicPIC and PIC electric fields as well as the speed distributions of the electrons. Finally, preliminary harmonic spectra are presented, although at the time of this writing, the simulations have not yet fully completed the propagation of the incident pulse, yielding a

lower signal-to-noise ration than is desired.

#### 4.4.1 Particle Velocities

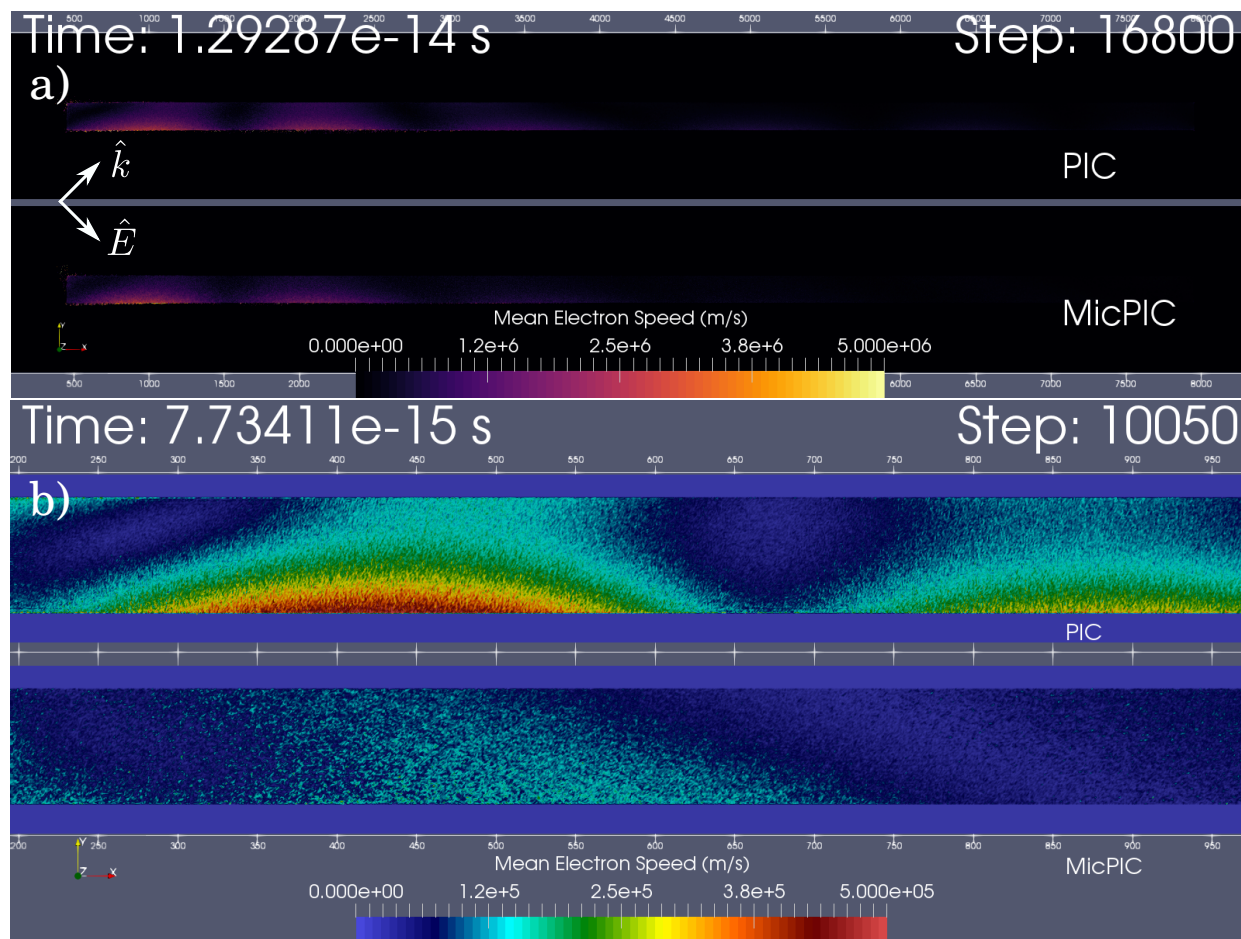


Figure 4.1: The laser is propagating from the lower left corner in all figures at a  $45^\circ$  angle to the upper right. Polarization is in the same plane ( $x$  increases to the right,  $y$  increases upward). a) Global view looking down the  $z$ -axis (into the page) of the PIC and MicPIC simulations at step 10050 (indicated), at a time when the laser begins exciting surface plasmons. At this step, the pulse has not yet arrived. Electrons have otherwise not travelled more than a few nanometres from the surface and still exhibit mostly collective behaviour. b) Closer view of the sample in the same direction as a), focused on the far left end of the sample. The weaker particle interactions in PIC permit surface plasmons, whereas microscopic effects in MicPIC prevent electrons from reaching high speeds.. See text for discussion.

The figures discussed in this section are generated from raw particle data output by MicPIC during its execution. Velocities are represented as heat maps and are calculated by postprocessing by sampling the average electron on a grid with half the resolution of the FDTD in each direction. Thus, the resolution shown is a factor of 8 lower than what is ultimately available. Visualized in each case is the middle of the sample in the periodic direction,  $z = 24$  to  $z = 25$ .

After 10050 time steps ( $\sim 7.73$  fs), the laser pulse has begun accelerating electrons on the incident side of

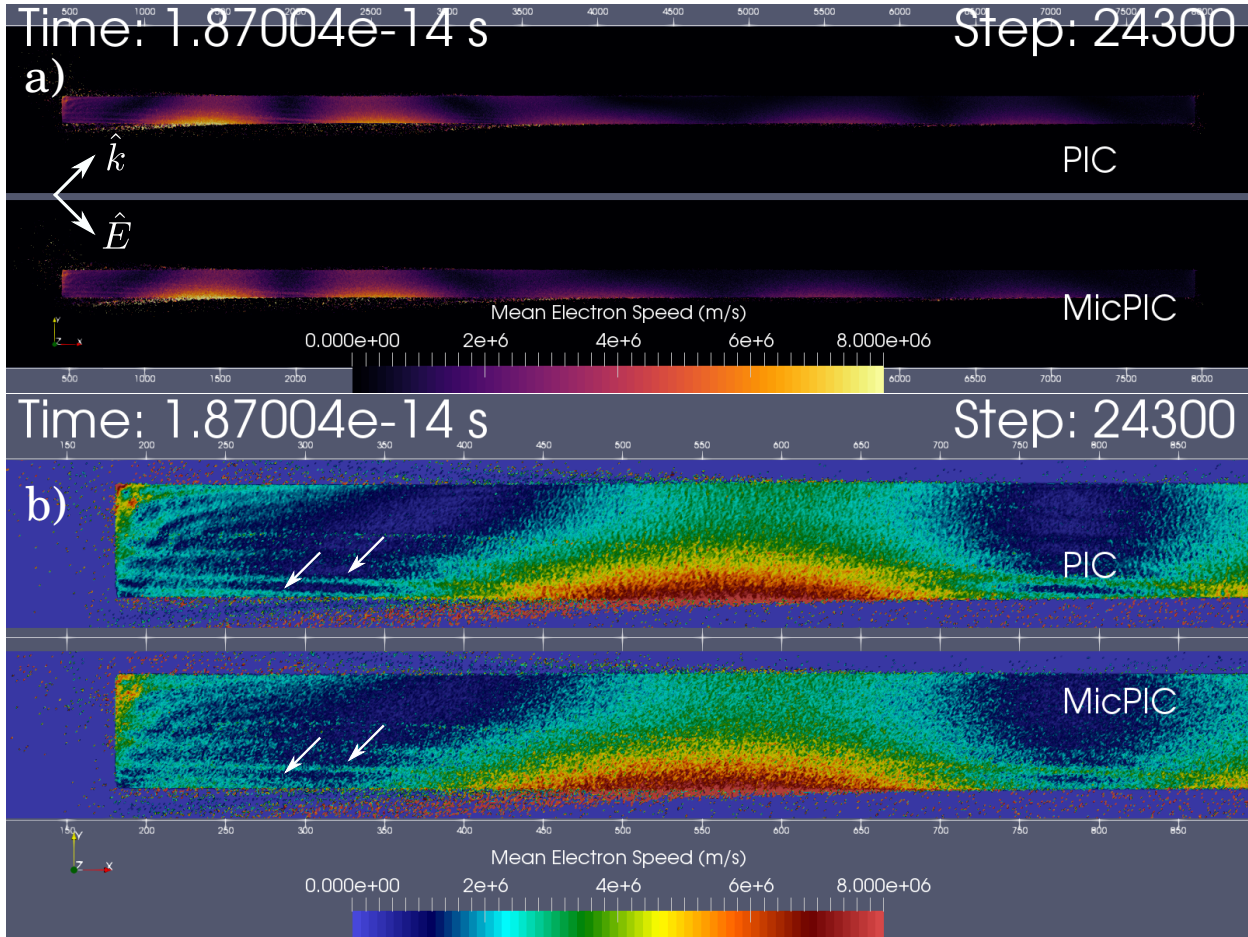


Figure 4.2: The laser is propagating from the lower left corner in all figures at a  $45^\circ$  angle to the upper right. Polarization is in the same plane ( $x$  increases to the right,  $y$  increases upward). a) Global view looking down the  $z$ -axis (into the page) of the PIC and MicPIC simulations at step 24300 (indicated), shortly before the arrival of the peak of the laser pulse. Waves of electrons are visible propagating through the sample and a wake of electrons begins to leave the medium. b) Closer view of the sample in the same direction as a), focused on the far end of the sample. Surface plasmons exhibit similar features in both, while the small amplitude waves propagating within tend to be more distinct in the case of PIC. See text for discussion.

the left end of the sample (see Fig. 4.1a). This effect is more pronounced in PIC, where weaker coupling of the plasma allows electrons to be more mobile, reacting more readily to the applied field. The behaviour is quite visible from the perspective of the entire simulation domain, where PIC-only electrons exhibit higher speeds. The microscopic corrections in MicPIC, at this early stage, allow electrons to collisionally distribute the energy absorbed from the laser amongst themselves and also with the ions (not shown). Fig. 4.1b provides a closer view of this behaviour, indicating that a surface plasmon is more readily excited in the PIC simulation, with electron speeds already reaching  $10^5$  m/s. This contrasts with the MicPIC case, where speeds even at the surface are only approximately  $1/4$  of those in PIC but are more uniform with depth.

At later times, shortly before the arrival of the peak of the pulse, the differences between MicPIC and PIC tend to become overwhelmed by the intense fields and the subsequent response of the electrons. An example of this is shown for the full extent of the simulation domain in Fig. 4.2a, wherein both simulations now exhibit

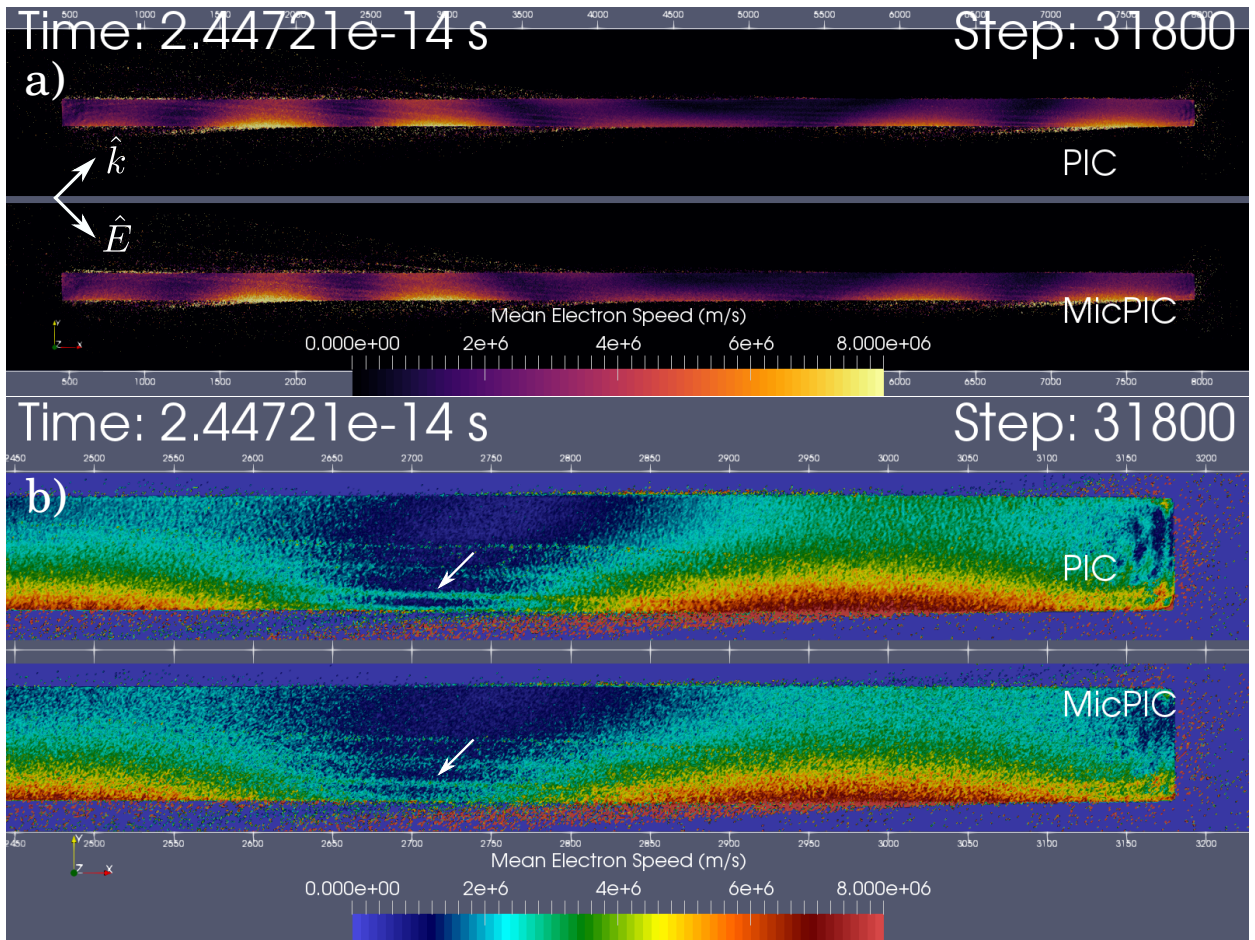


Figure 4.3: The laser is propagating from the lower left corner in all figures at a  $45^\circ$  angle to the upper right. Polarization is in the same plane ( $x$  increases to the right,  $y$  increases upward). a) Global view looking down the  $z$ -axis of the PIC and MicPIC simulations at step 31800 (indicated), at a time when the laser has traversed the majority of the length of the sample ( $x$  increases to the right,  $y$  increases upward). Many electrons have left the medium, forming a wake, and those that are returned have excited numerous short period plasma waves inside the material. b) Closer view of the sample in the same direction as a), focused on the far right end of the sample. Surface plasmons reflect off the discontinuity at the boundary in  $x$ . Small amplitude waves are more visible in the case of PIC, where the attenuation of short period plasma waves is less pronounced. See text for discussion.

surface plasmons propagating to the right with electrons attaining similar peak speeds. Electrons also leave the lower surface at high speed in the form of a wake, more clearly visible in a detailed view of the left end of the sample Fig. 4.2b. Additionally, short period waves (indicated by arrows in Fig. 4.2b and Fig. 4.3b) of relatively fast-moving electrons now become visible and are thought to be the result of Brunel electrons returning to the medium after absorbing energy via vacuum heating. These features are more defined in the PIC simulation, where the electrons are able to propagate more coherently over a longer distance inside the sample, whereas in MicPIC, collisions tend to more effectively exchange energy with the slower moving electrons of the bulk.

Fig. 4.3 demonstrates the response of the electrons as the peak of the pulse moves along the surface. By

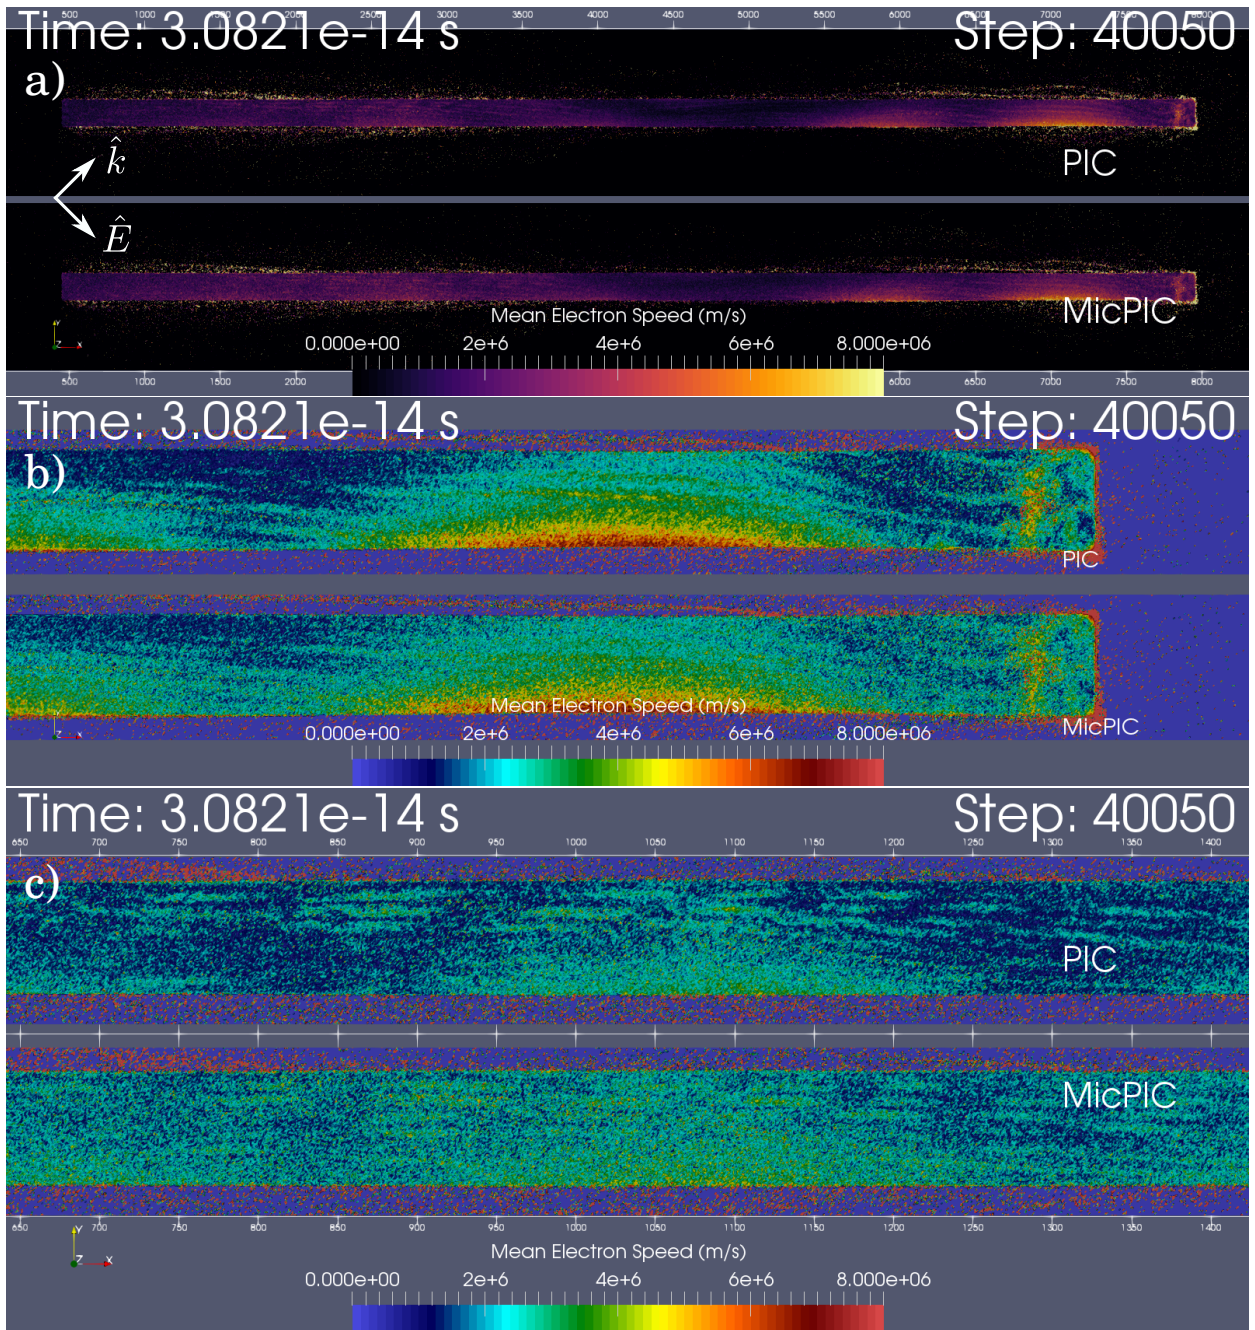


Figure 4.4: The laser is propagating from the lower left corner in all figures at a  $45^\circ$  angle to the upper right. Polarization is in the same plane ( $x$  increases to the right,  $y$  increases upward). a) Global view looking down the  $z$ -axis of the PIC and MicPIC simulations at step 40050 (indicated), just before the peak of the laser pulse passes beyond the end of the sample ( $x$  increases to the right,  $y$  increases upward). In both simulations, many fast electrons escape the medium, either due to field enhancement at the end of the sample or by accelerating in the wake of the laser. b) Closer view of the sample in the same direction as a), focused on the far right end of the sample. Interfering short period plasma waves are more clear in the PIC case, another example of which is shown to the left of the middle in c). See text for discussion.

step 31800, the pulse has propagated approximately 950 nm along the surface, roughly corresponding the second plasmonic peak from the left in Fig. 4.3a. Even in a full view of the samples, a wake of fast moving electrons is visible travelling outward from the surface. Still, a fraction of these electrons are returned to the material when the polarization reverses direction in the latter half of the cycle. Finer details are emphasized in Fig. 4.3b, where electron speeds appear to be very similar overall.

As the pulse approaches the end of the sample, surface plasmons reflect off the discontinuity there and interfere with those that follow. Complex interference patterns of plasma waves inside the medium appear as a result of the constituent electrons reflecting off the opposite surface (upper side). Additional electron wakes are visible propagating along this opposite surface and are also a source of electrons returning to the sample. This behaviour is shown in Fig. 4.4a, while finer details near the location of the pulse peak are visible in a closer view of the right end of the sample in Fig. 4.4b. Here, the fast electrons behave very similarly in the MicPIC and PIC results and very strong fields occurring at the vacuum interface at the end of the sample distort its corners, where the electrons are ejected. The subtle differences are more visible behind the pulse, where the short period waves inside the material interfere with one another (see Fig. 4.4c). Again, the more accurate microscopic interactions available with MicPIC lead to the breaking of these waves sooner, distributing their energy to their surroundings.

The most striking comparison of MicPIC and PIC in these simulations comes in the form of the distributions of electron speeds. Four examples of these differences are shown in Fig. 4.5, where data corresponding to what is shown in Figs. 4.1–4.4 are represented as coloured curves. In the times where the strongest parts of the laser pulse are interacting with the medium, it is very apparent that, in general, electrons in PIC are not as hot as their counterparts in MicPIC. The likely cause of this is collisional heating brought about via the stronger coupling between particles in MicPIC, which is supported by the heatmap images of the mean speeds presented above. Thus, electrons in MicPIC are more reluctant to move initially but more effectively transfer their kinetic energy amongst themselves and with the host material. Another consequence of this can be seen in the bottom plot in Fig. 4.5, which shows a zoomed in view of the distributions in MicPIC and PIC at step 10050. Here, the distribution peaks are reversed, with MicPIC having the higher lower peak velocity. This ties back into the what is seen in Fig. 4.1b, where the PIC electrons are initially able to be accelerated to higher velocities. It is worth emphasizing that in all cases, the maps of speeds in Figs. 4.1–4.4 are cell-averaged mean speeds, and provide a more detailed look at the spatial distribution of electrons and how fast they are moving, as opposed to the statistical nature of their speeds (Fig. 4.5).

## 4.4.2 Variations in the Electric Field

To illustrate the behaviour of the electric field, Fig. 4.6 shows a representative sample in the middle of the domain in the periodic direction ( $z = 25$ ) as the peak of the pulse is propagating along the surface. This particular time step was chosen in order to correspond with Fig. 4.3, although it takes place 50 time steps ( $\sim 3.18 \times 10^{-17}$  s) later. Short period waves within the sample analogous to those seen in the velocity data are apparent in the field as well, primarily in the  $y$ -direction, which is to be expected, since the motion of the electrons causing these waves is largely in the  $y$ -direction. Plasmonic fields are visible as intense perturbations to the  $y$  component of the field in both cases (Fig. 4.6a and b). The  $z$  component of the field is composed of primarily of noise, since the polarization of the laser is in the  $xy$ -plane, and fluctuations within are driven by collisional processes. Fig. 4.6c illustrates the numerical difference between the PIC and MicPIC field components. Subtle variations are due to a slight delay in the initial excitation of the medium in MicPIC, most likely cause by the stronger coupling, similar to what is seen in Fig. 4.1b, where electrons

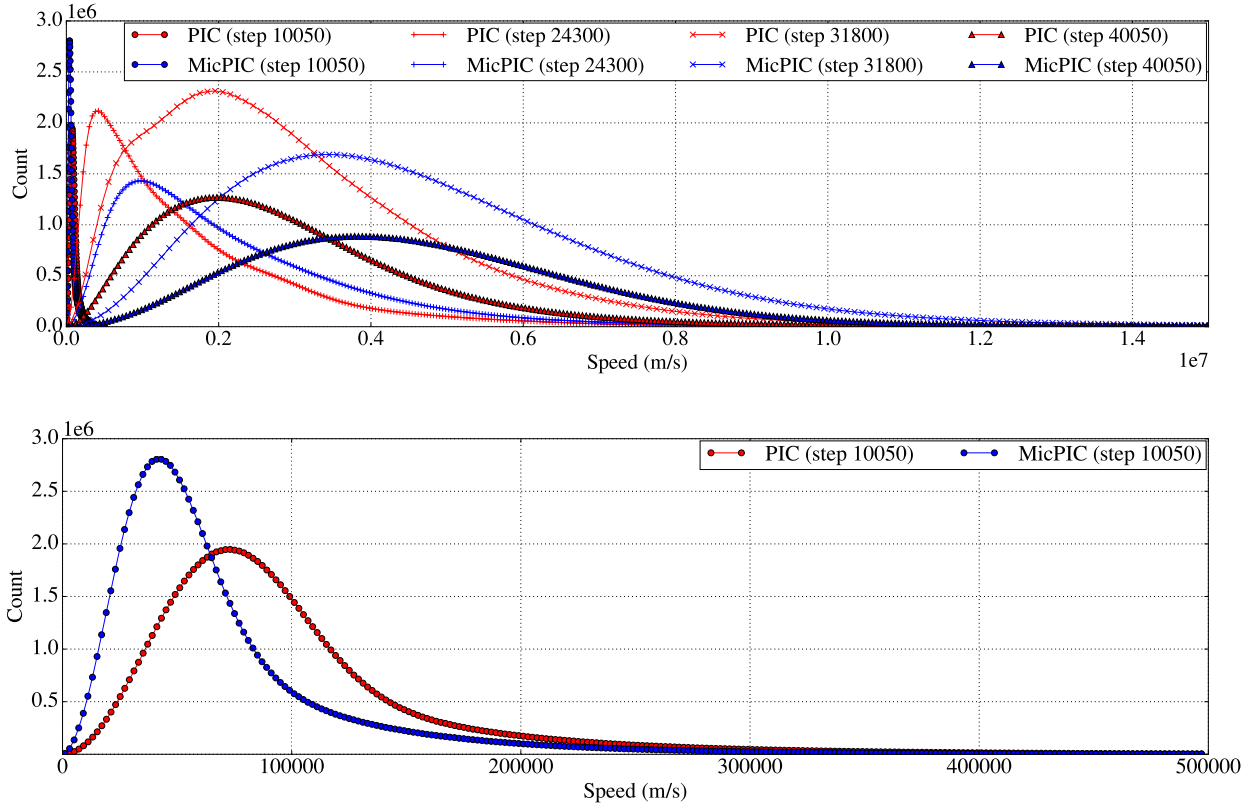


Figure 4.5: Comparison of the distributions of electron speeds in PIC and MicPIC. Red curves and markers indicate PIC data; MicPIC data are represented in blue. Top: In both cases, magnitudes of electron velocities were sampled into 250 speed bins, shown as markers on the curves, spanning ranges of velocities specific to the data from a given timestep. The data are shown up to  $10^7$  m/s, beyond which the particle counts become negligible compared to the peaks and interfere with visibility. Bottom: Plot of only the results from step 10050.

in PIC accelerate more readily.

### 4.4.3 Harmonic Spectrum

The harmonic spectrum generated by strong field processes is of significant interest and can provide important information about the dynamics atomic processes on attosecond timescales. Furthermore, due to the differences in the behaviours of the MicPIC and PIC simulations, corresponding differences in the resulting harmonic spectra would be expected. As Figs. 4.7 and 4.8 demonstrate the harmonic spectra near either end of the medium, given the data available at the time of writing, sampled at points located within two planes of observation. These planes are each located 182 nm from the surface of the sample in the  $y$ -direction, with the incident and transmitted sides corresponding to  $y = 18$  and  $y = 1088$ , respectively. Due to time and storage constraints, the laser pulse is still propagating out of the simulation domain at the final timestep, which unfortunately results in significant amounts of noise in the harmonic signal. Nevertheless, it is still possible to observe up to the  $\sim 5$ -th harmonic in both the PIC and MicPIC signals on the transmitted side. The left

end on the incident side shows a strong second harmonic generation in Fig. 4.7, while the transmitted side at the same end suggests possible higher harmonics. At the far end, higher harmonics are more visible on the transmitted side of the material, with both case showing similar spectra overall, although MicPIC results indicate more efficient third harmonic generation while the converse appears to be true for the PIC case and the fourth harmonic. Continuation of the simulations to conclude the propagation of the incident pulse is required, in order to drastically increase the signal-to-noise ratio and resolve the higher harmonics currently hidden within.

## 4.5 Conclusions

In this chapter, the discussion focused on the comparison of two simulations beginning from identical initial conditions, where the only difference between them was the inclusion of short range collisional effects available with MicPIC Section 2.4. Overall, similarities between the two cases are very subtle and most striking difference is the distribution of electron speeds. This is actually an important point and is illustrative of the effects that collisional processes play in strongly coupled plasmas. The stronger interactions of the electrons with their parent ions at first seems to inhibit the absorption of energy from the laser compared to PIC until later in the pulse, when they have been exposed to stronger fields for a longer period of time. In the end, collisions distribute the energy absorbed from the laser more effectively via collisional heating. Therefore, accurate modelling of strong field processes in solid density, strongly coupled plasmas should include microscopic effects in order to achieve greater accuracy.

However, comparing the electric field and the harmonic spectra of these PIC and MicPIC results presently does not indicate much difference in the physics that can be extracted from the more time-consuming latter case. This is misleading and disingenuous to a degree, however, for two reasons. The electron speed distributions indicate clearly that the plasma is more cohesive in MicPIC, and it is expected that this will affect the generation of higher harmonics through the modified response time of the plasma as a whole as well as the interaction between the bulk plasma and short period excitations induced by surface plasmons and Brunel electrons. Additionally, and perhaps most importantly, the PIC results shown here are PIC in the limit of least size for such a simulation. Particles are modelled individually and on nanometre scales, which implicitly includes some amount of microscopic many-body effects. Indeed, this “small size” PIC is in effect a softened MicPIC. Thus, traditional PIC codes employing superparticles will not have any of the same microscopic effects, by virtue of their excluding numerically what amounts to the higher order terms in the BBGKY hierarchy. Near-term, future investigations will investigate the differences between MicPIC and a more traditional PIC simulation involving superparticles. Lastly, PIC simulations of solid-density plasmas in the literature (eg. [92, 39]) begin with a presupposed plasma density gradient resulting from ionization from the laser. This is required for processes like CWE to take place but are absent in the simulations presented here since the MicPIC code does not presently account for the ionization of silica molecules. It is therefore the subject of an expanded investigation that will be conducted in the near-term, as it is expected that ionization dynamics will play an important role in the plasma formation and subsequent CWE process and MicPIC is at the forefront of providing a look “under the hood” of this and other ultrafast optical processes from tart to finish.

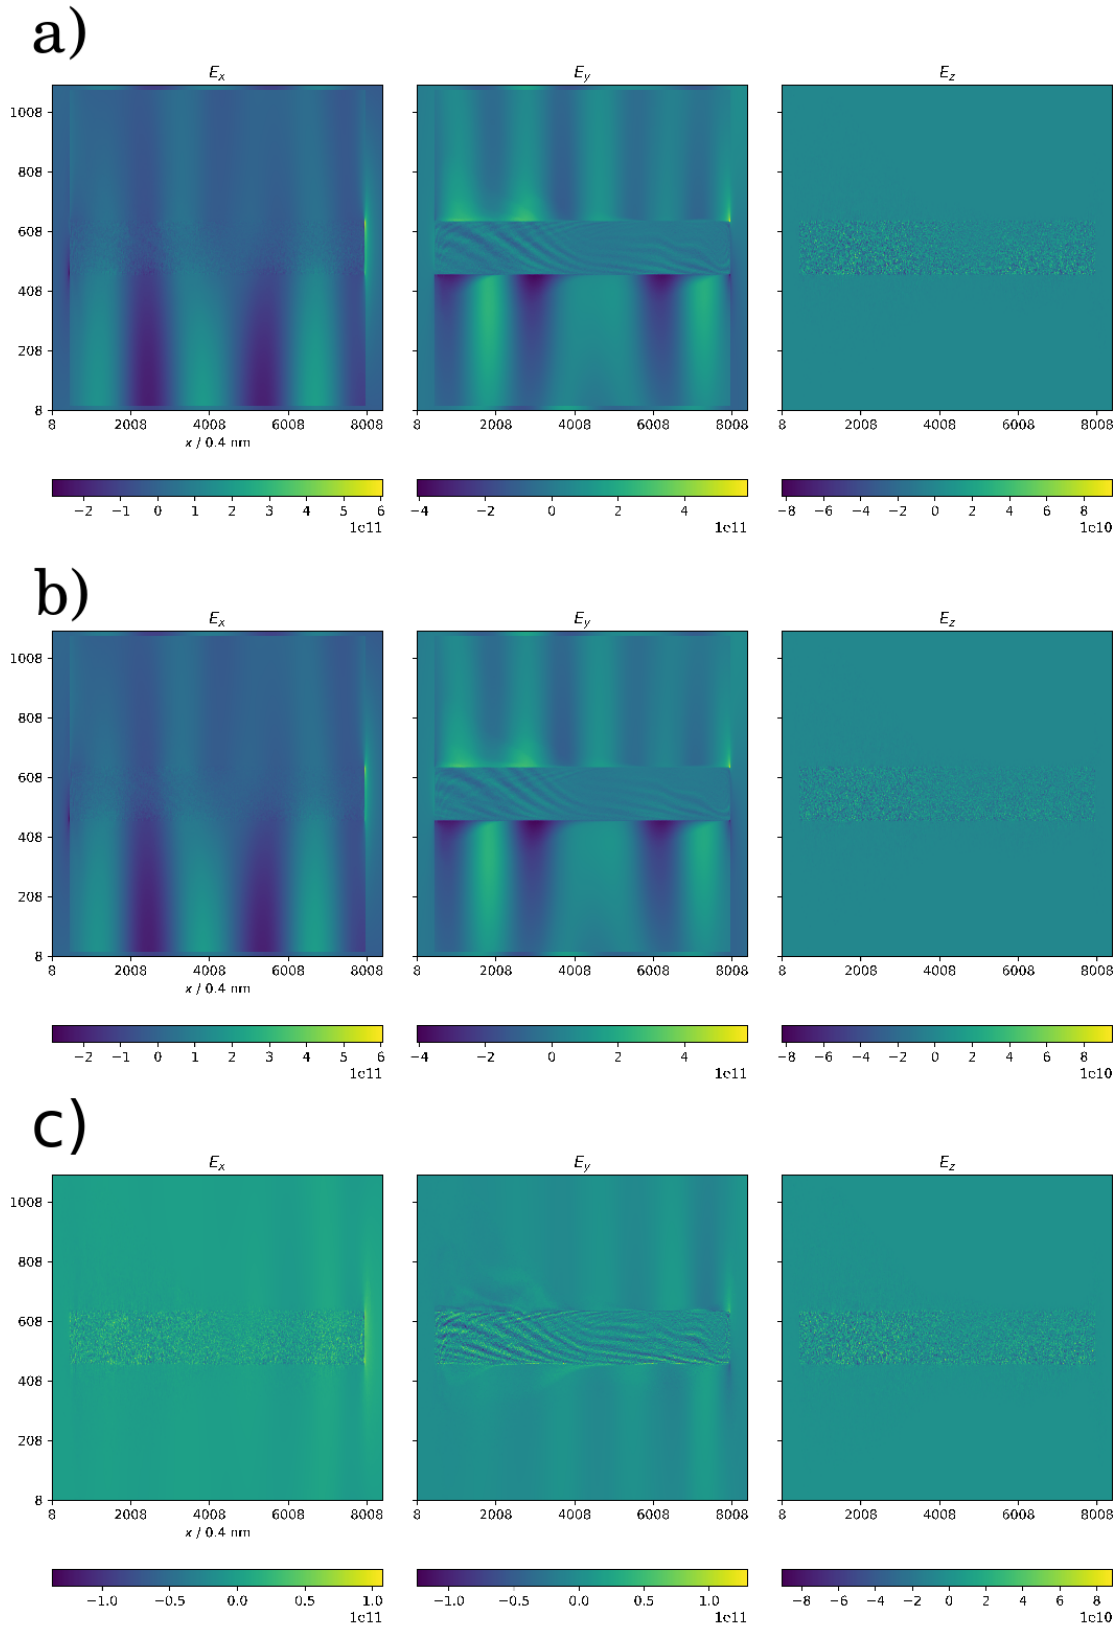


Figure 4.6: Comparison of components of electric fields in PIC (a) and MicPIC (b) simulations and the difference between the two (c). Sample is visible in the middle of the subdomain, primarily in plots of  $E_y$ . TFSS interfaces are visible as discontinuities in field components near the peripheries.

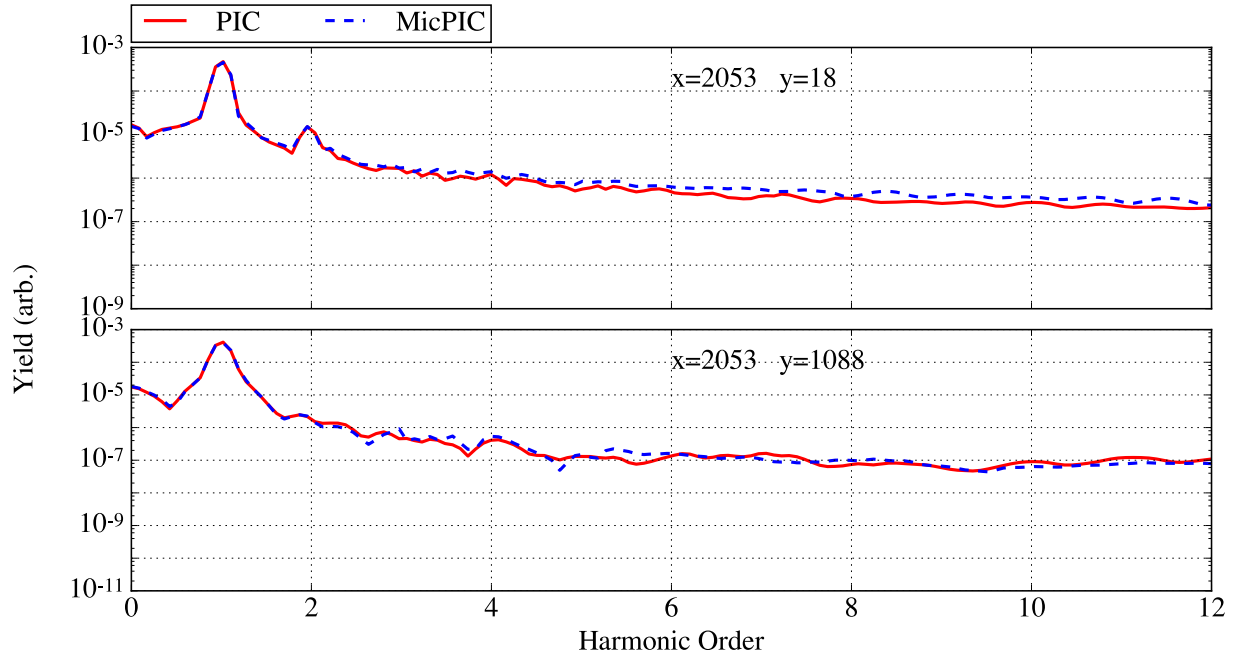


Figure 4.7: The harmonic spectrum for the electric field taken at observation points on the incident and transmitted sides of the sample. In this case, the probe point is taken at a location 373 nm from the left end of the sample and approximately 182 nm from either side in the  $y$ -direction.

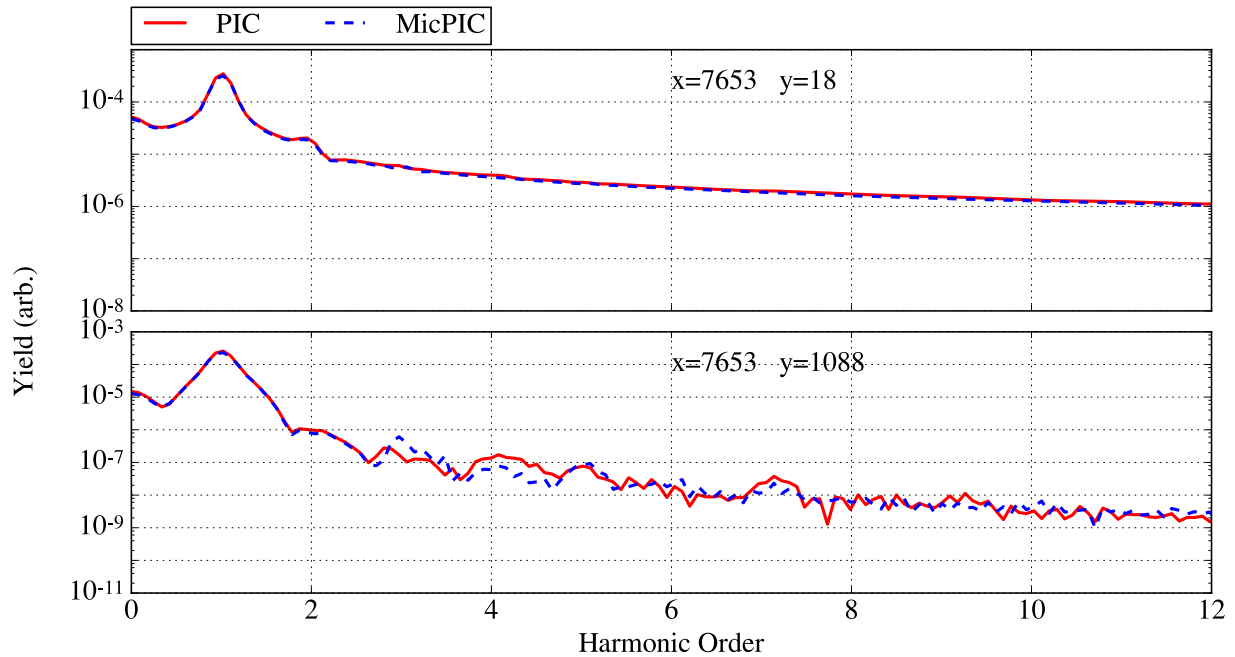


Figure 4.8: The harmonic spectrum for the electric field taken at observation points on the incident and transmitted sides of the sample. In this case, the probe point is taken at a location 120 nm from the left end of the sample and approximately 182 nm from either side in the  $y$ -direction.

---

CHAPTER

FIVE

---

# PHOTOEMISSION OF ELECTRONS FROM METAL NANOTIPS

## 5.1 Introduction

Tunnelling of electrons from metal surfaces into free space is enabled by the presence of strong electric fields. Ultrashort laser pulses offer researchers the ability to deliver high intensity radiation over femtosecond timescales, opening up the realm of strong field physics and the atom to deeper exploration [12, 11, 13]. Using few-cycle pulses of visible and NIR combined with atoms, molecules, and nanometre scale objects allows for the production of new sources of coherent sub-fs XUV, x-ray [154], and electron pulses [65, 155, 156].

Attosecond electrons in the gas phase have been shown to be useful for probing the ionization dynamics [157] and molecular imaging [158], where the atomic or molecular species are often subjected to high intensity radiation. Experiments with solid surfaces indicated that electron emission processes were taking place, analogous to what occurs in gaseous samples, such as tunnel ionization and ATI, except at much lower intensities. Spatially extended sources are limited by the damage threshold of the bulk material [159] but are prone to emit electrons at sites of random surface roughness where significant field enhancement can occur [160, 161].

Sharp metallic nanotips have been shown to be good sources of temporally and spatially coherent electrons, as the geometry of the tip provides significant field enhancement without damage to the material, while keeping the electron emission sites localized to nanometre sized regions [67, 162, 163]. Moreover, these photoelectrons, once free of the metal, can be described classically by the three-step or simple man's model [33, 34] (Section 1.1.3) and are able to be controlled on sub-cycle time scales [66]. This attosecond control has enabled tip-enhanced electron microscopy [162, 164] probing of the carrier-envelope phase of the laser pulse [98], and high harmonic generation [13]. More recently, ultrafast nanoscale devices based on nanotips have been proposed, including vacuum-tube diodes [165], which suggest the possibility of constructing electronics operating at optical frequencies.

## 5.2 Theoretical Background

Nonlinear photoemission from metal surfaces can be treated much in the same way as ionization of atomic systems. Field enhancement near the end of a metal nanotip during the passage of a femtosecond pulse allows local electric field values to exceed the damage threshold of the material without any lasting effect [166]. This allows the electron to undergo three possible distinct processes, depending on laser intensity, frequency, and material properties: *above-threshold photoemission* (ATP), *multi-photon photoemission* (MPP, sometimes called MPI), and optical field emission (tunnelling). This last case corresponds to the solid state analog of an electron tunneling through a potential barrier which has been lowered by a strong electric field.

The characteristic Keldysh parameter  $\gamma = \sqrt{I_p/2U_p}$ , where  $I_p$  is the ionization potential (work function) of the atom (surface) and  $U_p = q^2 E^2/4m_e \omega^2$  is the ponderomotive energy, defines a continuum between the two limiting cases of the multiphoton and tunnelling regimes [80]. It should be noted that this only serves as a guide and breaks down if the intensity of the laser enters into the relativistic regime and magnetic field effects must be taken into account [167]. However, the intensities used in experiments involving photoemission from metal nanotips or nanoparticles (typically  $\sim 10^{11} - 10^{13} \text{ W cm}^{-2}$ ) are well below this range [166, 168, 169, 170]. Typically numerical approaches involve integration of the time-dependent Schrödinger equation [13, 66] or approach the problem through Fowler-Nordheim theory [162], hence ranging from the purely quantum mechanical, to the semiclassical, and often in the case of free photoelectrons, classical Newtonian mechanics.

Hot electrons in a metal can be modelled classically in cases where the degeneracy parameter (Eq. (1.2)) is greater than unity. Furthermore, it is possible to model tunnelling out of metal nanostructures in just such a classical plasma simulation (MicPIC). In the remainder of this chapter, the implementation of such a method is provided in detail and its application to a test scenario is presented.

Photoemission from metal nanotips involves particles interacting with potential barriers in the neighbourhood of the metal surface and therefore motivates an analytic description of the transmission coefficient of such particles. The WKB approximation (Section 1.1.4) is one such approach to the problem and assumes that the wavelength  $\lambda$  of the particle varies slowly in the region containing the potential  $U(x)$ . In other words, the kinetic energy of the particle must be sufficiently large in order for the approximation to be valid. MicPIC represents a numerical solution of the classical Liouville equation for a hot dense plasma, yielding classical electron orbits. Such hot electrons impinging on a potential barrier that varies smoothly by a few electron volts per nanometre, should have small de Broglie wavelengths ( $\lesssim 10^{-10} \text{ m}$ ) and adequately satisfy the conditions imposed by the WKB approximation. A solution to the integral for the WKB exponent in Eq. (1.50) requires accurate knowledge of the energy landscape in the vicinity of the electron as well as along its direction of propagation.

## 5.3 Implementation

With the mathematical formalism for particles tunnelling through steep potential barriers, it becomes possible to simulate the process, even within classical or semi-classical many-particle simulations. MicPIC provides the means to incorporate quantum mechanical tunnelling in a plasma physics setting. Specifically, the long and short range aspects of MicPIC are used to determine the mechanical energy of an electron, how likely it is to tunnel, and where it would classically “reappear”. To do this, MicPIC takes advantage of several tools and resources already in place in the core of the software but also new development as well, which are discussed below. The tunnelling probability calculation, in brief, requires computation of the electric

potential along the velocity vector of the electron, the configuration of its neighbours, and minimizing the effect of the particle on itself. Additionally, it is necessary to incorporate parallelization into the method, since simulations are generally large and many computation processes share the load.

### 5.3.1 Determining the turning points

The primary step in the procedure is to determine which particles should undergo the photoemission process, when, and from where their tunnelling paths should begin. To do this, MicPIC employs a short chain of successive binary logic tests, which all must be satisfied in order to receive a tunnelling probability. Only those particles which are electrons are considered, since ions move too slowly to have any appreciable chance of tunnelling. Furthermore, the electron must be moving outward - that is, a velocity with a component directed perpendicular to the tangent plane to the surface in the neighbourhood of the electron's position (shaded region delineated by "object boundary" in Fig. 5.1). If this criterion is satisfied, the electric field at the electron's position, determined by the core MicPIC functions, is used to predict the electron's classical turning point within the metal according to Newton's second law. If the motion of the electron is arrested in the direction perpendicular to the material surface within the next time step,

$$\Delta t = \frac{2m_e v_{\perp}}{eE_{\perp}} \leq dt \quad (5.1)$$

then the electron will be considered for the primary tunnelling probability calculation, whose classical turning point is approximately

$$\vec{x}_{tp} = \vec{x}(t) + \vec{v}(t)\Delta t - \frac{eE}{2m_e}(\Delta t)^2 \quad (5.2)$$

### 5.3.2 Trilinear Interpolation and Calculation of the Electric Potential

The most extensive component of the tunnelling probability calculation is the electric potential. Once the classical turning point of the electron is determined, the simulation will sample the electric field from the grid along the velocity vector of the electron. In order to do this sampling MicPIC uses trilinear interpolation, the three-dimensional extension of linear interpolation. This requires the values of the electric field stored on the vertices of the grid cell occupied by the electron during the timestep it reaches its turning point. Specifically, to sample the individual components of  $\vec{E}$ , one must make note of their arrangement in the Yee cell: eg.  $E_x$  values lie on 2D sub-grids of  $\vec{E}$  with integer  $j$  and  $k$  values and half-integer  $i$  values. Explicitly, the sampling is performed according to

$$E_x(\vec{x}) = (1 - z_d)c_{0x} + z_dc_{1x} \quad (5.3a)$$

$$E_y(\vec{x}) = (1 - z_d)c_{0y} + z_dc_{1y} \quad (5.3b)$$

$$E_z(\vec{x}) = (1 - y_d)c_{0z} + y_dc_{1z} \quad (5.3c)$$

The interpolated values  $c_{0x}$ ,  $c_{1x}$ , etc. are given by

$$c_{0x} = (1 - y_d)c_{00x} + y_dc_{10x} \quad (5.4a)$$

$$c_{1x} = (1 - y_d)c_{01x} + y_dc_{11x} \quad (5.4b)$$

$$c_{0y} = (1 - x_d)c_{00y} + x_dc_{10y} \quad (5.4c)$$

$$c_{1y} = (1 - x_d)c_{01y} + x_dc_{11y} \quad (5.4d)$$

$$c_{0y} = (1 - x_d)c_{00y} + x_dc_{10y} \quad (5.4e)$$

$$c_{1y} = (1 - x_d)c_{01y} + x_dc_{11y} \quad (5.4f)$$

with the further interpolated values being determined directly from the field according to

$$c_{00x} = (1 - x_{d,1/2})f(\vec{x}_{i+1/2,j,k}) + x_{d,1/2}f(\vec{x}_{i-1/2,j,k}) \quad (5.5a)$$

$$c_{10x} = (1 - x_{d,1/2})f(\vec{x}_{i+1/2,j+1,k}) + x_{d,1/2}f(\vec{x}_{i+1/2,j+1,k}) \quad (5.5b)$$

$$c_{01x} = (1 - x_{d,1/2})f(\vec{x}_{i+1/2,j,k+1}) + x_{d,1/2}f(\vec{x}_{i+1/2,j,k+1}) \quad (5.5c)$$

$$c_{11x} = (1 - x_{d,1/2})f(\vec{x}_{i+1/2,j+1,k+1}) + x_{d,1/2}f(\vec{x}_{i+1/2,j+1,k+1}) \quad (5.5d)$$

$$c_{00y} = (1 - y_{d,1/2})f(\vec{x}_{i,j+1/2,k}) + y_{d,1/2}f(\vec{x}_{i,j-1/2,k}) \quad (5.6a)$$

$$c_{10y} = (1 - y_{d,1/2})f(\vec{x}_{i+1,j+1/2,k}) + y_{d,1/2}f(\vec{x}_{i+1,j-1/2,k}) \quad (5.6b)$$

$$c_{01y} = (1 - y_{d,1/2})f(\vec{x}_{i,j+1/2,k+1}) + y_{d,1/2}f(\vec{x}_{i,j-1/2,k+1}) \quad (5.6c)$$

$$c_{11y} = (1 - y_{d,1/2})f(\vec{x}_{i+1,j+1/2,k+1}) + y_{d,1/2}f(\vec{x}_{i+1,j-1/2,k+1}) \quad (5.6d)$$

$$c_{00z} = (1 - z_{d,1/2})f(\vec{x}_{i,j,k+1/2}) + z_{d,1/2}f(\vec{x}_{i,j,k-1/2}) \quad (5.7a)$$

$$c_{10z} = (1 - z_{d,1/2})f(\vec{x}_{i,j+1,k+1/2}) + z_{d,1/2}f(\vec{x}_{i,j+1,k-1/2}) \quad (5.7b)$$

$$c_{01z} = (1 - z_{d,1/2})f(\vec{x}_{i+1,j,k+1/2}) + z_{d,1/2}f(\vec{x}_{i+1,j,k-1/2}) \quad (5.7c)$$

$$c_{11z} = (1 - z_{d,1/2})f(\vec{x}_{i+1,j+1,k+1/2}) + z_{d,1/2}f(\vec{x}_{i+1,j+1,k-1/2}) \quad (5.7d)$$

Finally, the integer cell index weighting coefficients are defined as

$$x_d = \frac{x - x_0}{x_1 - x_0} \quad (5.8a)$$

$$y_d = \frac{y - y_0}{y_1 - y_0} \quad (5.8b)$$

$$z_d = \frac{z - z_0}{z_1 - z_0} \quad (5.8c)$$

and the half-integer coefficients as

$$x_{d,1/2} = \frac{x - x_0 + 1/2}{x_1 - x_0} \quad (5.9a)$$

$$y_{d,1/2} = \frac{y - y_0 + 1/2}{y_1 - y_0} \quad (5.9b)$$

$$z_{d,1/2} = \frac{z - z_0 + 1/2}{z_1 - z_0} \quad (5.9c)$$

The denominators in Eqs. 5.8-5.9 reduce to unity when measuring positions in cell-based coordinates, whereby the sampling cells are congruent with those of the FDTD grid.

With the above equations 5.3-5.9 it becomes possible to compute the potential difference along a path from  $\vec{x}_a$  to  $\vec{x}_b$  according to

$$\phi(\vec{x}_b) - \phi(\vec{x}_a) = \int_{\vec{x}_a}^{\vec{x}_b} \vec{E} \cdot d\vec{x} \quad (5.10)$$

where  $\vec{x}_a = \vec{x}_{tp}$  is the turning point of the electron given by Eq. 5.2 and  $\vec{x}_b$  is that point in space where the potential difference undergoes a change of sign from negative to positive. That is, it is the point where the electron should “reappear” on the other side of the barrier, having zero kinetic energy in the direction parallel to the tunnelling path.

Numerically, the integral becomes a discrete sequence of sampled values at  $n$  points, separated equally by  $\vec{\Delta}s = s\hat{n}$ , where  $s$  is the number of samples per grid cell and  $\hat{n}$  is the unit vector along the tunnelling path. Then with the condition that  $\phi(\vec{x}_{tp}) = 0$ ,

$$\phi(\vec{x}_k) = \phi(\vec{x}_{k-1}) + \frac{\vec{\Delta}s}{2} \cdot [\vec{E}(\vec{x}_{k-1}) + \vec{E}(\vec{x}_k)] \quad (5.11)$$

and thus, the potential difference with respect to the turning point  $\vec{x}_{tp}$  at the  $n$ -th sample point  $\vec{x}_n = \vec{x}_{tp} + n\vec{\Delta}s$  is

$$\phi(\vec{x}_n) = \frac{\vec{\Delta}s}{2} \cdot \left\{ \vec{E}(\vec{x}_{tp}) + \vec{E}(\vec{x}_n) + 2 \sum_{k=1}^{n-1} [\vec{E}(\vec{x}_{k-1}) - \vec{E}(\vec{x}_k)] \right\} \quad (5.12)$$

In practice, the point at  $\vec{x}_b$  does not coincide with a field sample point. It is instead necessary to integrate to the first point where  $\phi > 0$ . A further interpolation is required to predict the approximate position of the zero crossing  $\vec{x}_b$ . To a linear approximation, and provided that  $\phi(\vec{x}_n) > 0$  and  $\phi(\vec{x}_{n-1}) < 0$ , this point is situated at

$$\vec{x}_b = \vec{x}_{tp} + (n-1)\vec{\Delta}s + \frac{\vec{\Delta}s}{\phi(\vec{x}_n) - \phi(\vec{x}_{n-1})} \quad (5.13)$$

### 5.3.3 Short Range Field Effects

The second significant effect is that of particles near to the particle which is being considered for tunnelling. In the core MicPIC algorithm, electric forces are adjusted to increase the accuracy of inter-particle dynamics at short range, but this is restricted to the particle objects as they exist within the simulation and not every point in space. Therefore, while computing electric potential values at a series of sample points, it is necessary to apply a similar short range correction procedure to the one described in Section 3.2.1 and [105] in order to update the values of the electric field at these points.

The photoemission component of MicPIC follows immediately after the short range force calculations, and is therefore ideally situated within the execution to take advantage of the already sorted and partitioned cell lists (see Section 2.4). This provides a convenient method for quickly locating those particles within the correction radius  $r_{cut}$  of a sample point, represented as yellow circles surrounding each particle in Fig. 5.1. The procedure is similar to the main short range force calculation, with the exception that only corrections to the electric field are calculated and only at the sample points.

The position of each sample point, expressed as a C++ type `double` and converted into units of grid cells, can be used to quickly determine the appropriate element in the partitioned and sorted cell list. From there, a cell with indices  $(i, j, k)$  must include any necessary corrections to the electric field from particles in neighbouring cells with indices in the range  $(i - n_{cell}, j - n_{cell}, k - n_{cell}), \dots, (i, j, k), \dots, (i + n_{cell}, j + n_{cell}, k + n_{cell})$ .

Any particle within  $n_{cell}$  cells of a sample point  $\vec{x}_s$  contributes a correction to the field at the point according to Eq. (2.10b), reproduced here for convenience and in the context of this discussion:

$$\vec{E}^{(mic)}(\vec{x}_s) = \frac{q_i}{4\pi\epsilon_0} \left\{ \frac{1}{r_{is}^2} \left[ \operatorname{erf}\left(\frac{r_{is}}{w_0}\right) - \operatorname{erf}\left(\frac{r_{is}}{w_{PIC}}\right) \right] - \frac{1}{r_{is}} \frac{2}{\sqrt{\pi}} \left[ \frac{e^{-\frac{r_{is}^2}{w_0^2}}}{w_0} - \frac{e^{-\frac{r_{is}^2}{w_{PIC}^2}}}{w_{PIC}} \right] \right\} \hat{r}_{is} \quad (5.14)$$

where  $\hat{r}_{is} = \frac{\vec{r}_s - \vec{r}_i}{\|\vec{r}_s - \vec{r}_i\|}$  is the unit vector in the direction of the sample point  $\vec{r}_s$  with respect to the  $i$ -th particle with charge  $q_i$ .

In the event that a sample point is within a distance of  $r_{cut}$  cells from a subdomain boundary, the contribution of any particles in the neighbouring subdomain adjacent to the boundary must be included. This is visible in Fig. 5.1 as the set of points on the line segment from particle 1 to the point  $\vec{r}_f$  (red star) which intersect the union of the correction volumes of particle 2 and 3 (yellow circles); eg.  $\vec{r}_s$  (red cross). This is facilitated by the short range correction buffers (Sec 3.2.1) left over from the core short range force routines immediately preceding photoemission, which contain partitioned and sorted lists of the necessary particles. A procedure much like what is used for the local short range corrections for interpolation yields the required boundary cells and any particles they may contain. In both local and boundary cases, simple pointer comparisons prevent the electron in question from contributing to the electric field at any nearby sample points, thereby avoiding contamination of its own tunnelling probability.

### 5.3.4 Removing the Self-Action

It is worth emphasizing that the electric field data used in the interpolation includes the contribution of all particles, *including* the particle in question, so it is necessary to remove its contribution to its surroundings. In Section 2.3, the procedure is outlined for the weighting of current densities to and electric and magnetic fields from the grid within a radius of  $n_{pts} + 1$  cells from each particle, represented by the green circle in particle 1 in Fig. 5.1. This is reflected here in the removal of the self-force, for which values of  $\vec{E}^{(PIC)}$  are computed within the same radius. Recall that the charge densities of particles in MicPIC are modelled as Gaussian distributions of width  $w_{PIC}$  according to Eqs. (2.1a) and (2.1b), whose potential energy on the (coarse) PIC level is

$$\vec{E}^{PIC}(\vec{r}_i) = -\frac{e}{4\pi\epsilon_0} \left[ \frac{1}{r_{is}^2} \operatorname{erf}\left(\frac{r_{is}}{w_{PIC}}\right) - \frac{1}{r_{is}} \frac{2}{\sqrt{\pi}} \frac{e^{-\frac{r_{is}^2}{w_{PIC}^2}}}{w_{PIC}} \right] \hat{r}_{is} \quad (5.15)$$

This approach is used for two reasons. Presently, the photoemission code is new and un-optimized and completely removing the PIC field due to a particle would be much less efficient, essentially requiring a miniature version of the PIC algorithm to be implemented. Eq. 5.15 is a compromise between speed and accuracy for the tunnelling calculation. Furthermore, keeping a look-up table available, similar to what is used in the short range force corrections, makes computing multiple field values very fast. In the event of a particle's charge density extends across a boundary, the path segment calculation in the neighbouring subdomain can set the correct field value at the relevant sample points by calculating the distance from the sample point in question to the particle from the index in the list of samples and sample step size.

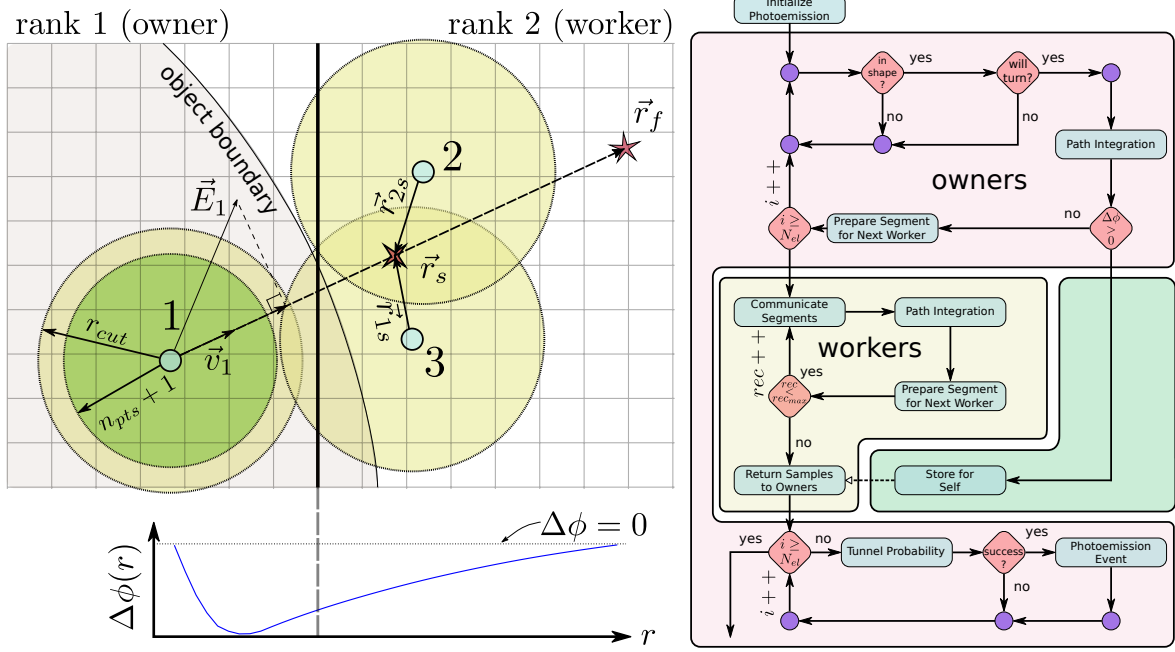


Figure 5.1: Schematic depiction of a parallel photoemission event. Left: electron 1 approaches the bounding surface of some material (shaded region, left). Components of velocity and electric field in the direction normal to the surface determine whether potential difference relative to the classical turning point is calculated (bottom left, actual MicPIC result) until final point  $\vec{r}_f$  (red star). Green shaded region corresponds to Eq. (5.15); any point within  $r_{cut}$  of other particles (red cross) is corrected via Eq. (5.14). Right: schematic representation of photoemission procedure. Checks are made for location and turning points; path integrations are performed conditionally. Every path starts on its owner (pink region, right). Any path reaching a boundary is passed to workers for further integration, recursively (yellow region). Some paths may be short; these are stored locally (green region) until full paths are reassembled after communication of segments from workers to owners. Tunnel probability calculations are made on owners via Eqs. (1.49) and (1.50). More detail in text.

### 5.3.5 Parallelization

In Chapter 3 the nature of the parallelization of MicPIC was discussed and analyzed in detail. In particular, the domain decomposition was a central feature and how information is exchanged between the processes involved in the parallel execution of a simulation plays an important role in its performance. The decomposed nature of the simulation domain presents a challenge for the potential difference calculation via the path integration through the electric field, the inclusion of short range effects, and the exclusion of self-action.

The approach taken towards a parallel implementation of photoemission in MicPIC follows the same principle of domain decomposition as its core components (see Chapter 3). The path integration algorithm in MicPIC can, in general, proceed for an arbitrary distance though an equally arbitrary electric field until the potential difference becomes positive ( $\Delta\phi \geq 0$  in Fig. 5.1) or a global domain boundary is reached (ie. PML). Parallelization of such a situation would be quite inefficient involving an unspecified number of point-to-point communications, where segments of the integration path would have to be passed back and forth for every electron whose tunnelling probability was being calculated. Instead, the procedure is truncated and single collective function is called recursively up to some maximum level of recursion.

The integration paths originate in the subdomain occupied by the particle whose tunnelling probability is to be determined. The process managing this subdomain is the *owner* for any such path originating here and the level of recursion is its lowest (0). Subdomain-local path integration is then performed according to the procedure outlined in Sections 5.3.1–5.3.4 until either the potential difference relative to the turning point becomes positive or a boundary condition must be applied (pink block in Fig. 5.1). The former case is the simpler of the two, in which the sampled potential difference data are stored until the remainder of the path integration phase is completed. In the latter case, relevant information about the integration path must be communicated to the process managing the subdomain in which the path is continued. Following the inter-process exchange of path integral information, the level of recursion is increased and the same procedure repeats, excepting the initial operations involving determining which electrons are to be tested for tunnelling probabilities. Each process is now a *worker*, which will *continue* path integration begun by any owners (tan block, Fig. 5.1). This continues recursively up to some user-defined maximum level of recursion, typically 2-4, and if a positive potential difference for a given electron is found, the operation is stopped. When the maximum level of recursion is reached, any paths which do not encounter zero potential difference are assumed to yield negligibly small tunnelling probabilities due to their lengths, and are discarded. Each worker is aware of the owner of a path, regardless of the level of recursion, allowing the remaining *path integral segments* to be returned to the owners in a single collective communication.

### 5.3.6 Tunnelling

The tunnelling probability of an electron reflecting off the inner surface of a metal is determined via Eq. (1.50). Potential differences are used to ultimately calculate the energy landscape ahead of the electron in question along its velocity vector normal to the surface, simulating the evanescent part of the electronic wave function. At the classical turning point, the total energy of the electron in the surface-normal direction is entirely bound up in potential form. Thus, the energy dependent integrand in Eq. (1.50) depends on the path integration result  $U(x) = -e\Delta\phi(x)$ , where  $x$  is the position along the path and  $e$  is the elementary charge, and is calculated with simple trapezoid integration. Successful tunnel events are determined by comparing the probability calculated via Eq. (1.49) to a pseudo-random number.

A photoemission even triggers the instantiation of two objects to mediate the tunnelling process. Due to the discrete nature of the simulation grid and the demands of causality, electrons cannot be arbitrarily switched on and off or move superluminally without creating unphysically intense fields. Instead, in MicPIC an electron which tunnels via the photoemission routines experiences a ramping of its charge to zero with what is referred to as an **Annihilate\_Event**. Simultaneously, a new electron is ramped on at the same rate at the “exit point” of the tunnel path ( $\vec{r}_f$  in Fig. 5.1) with a **Create\_Event**. These objects ensure the correct electron is ramped and provide the means to satisfy the condition of causality and conserve charge at every time step during the ramp process while allowing the two particles to behave independently, the latter being essential in cases where the tunnel path ends in another subdomain. During the ramp process, unlike the case of atomic and molecular tunnel ionization (see Section 2.6), photoelectrons are free to propagate. Tests during development of the photoemission code using a five-step ramp indicate that this is a smooth process.

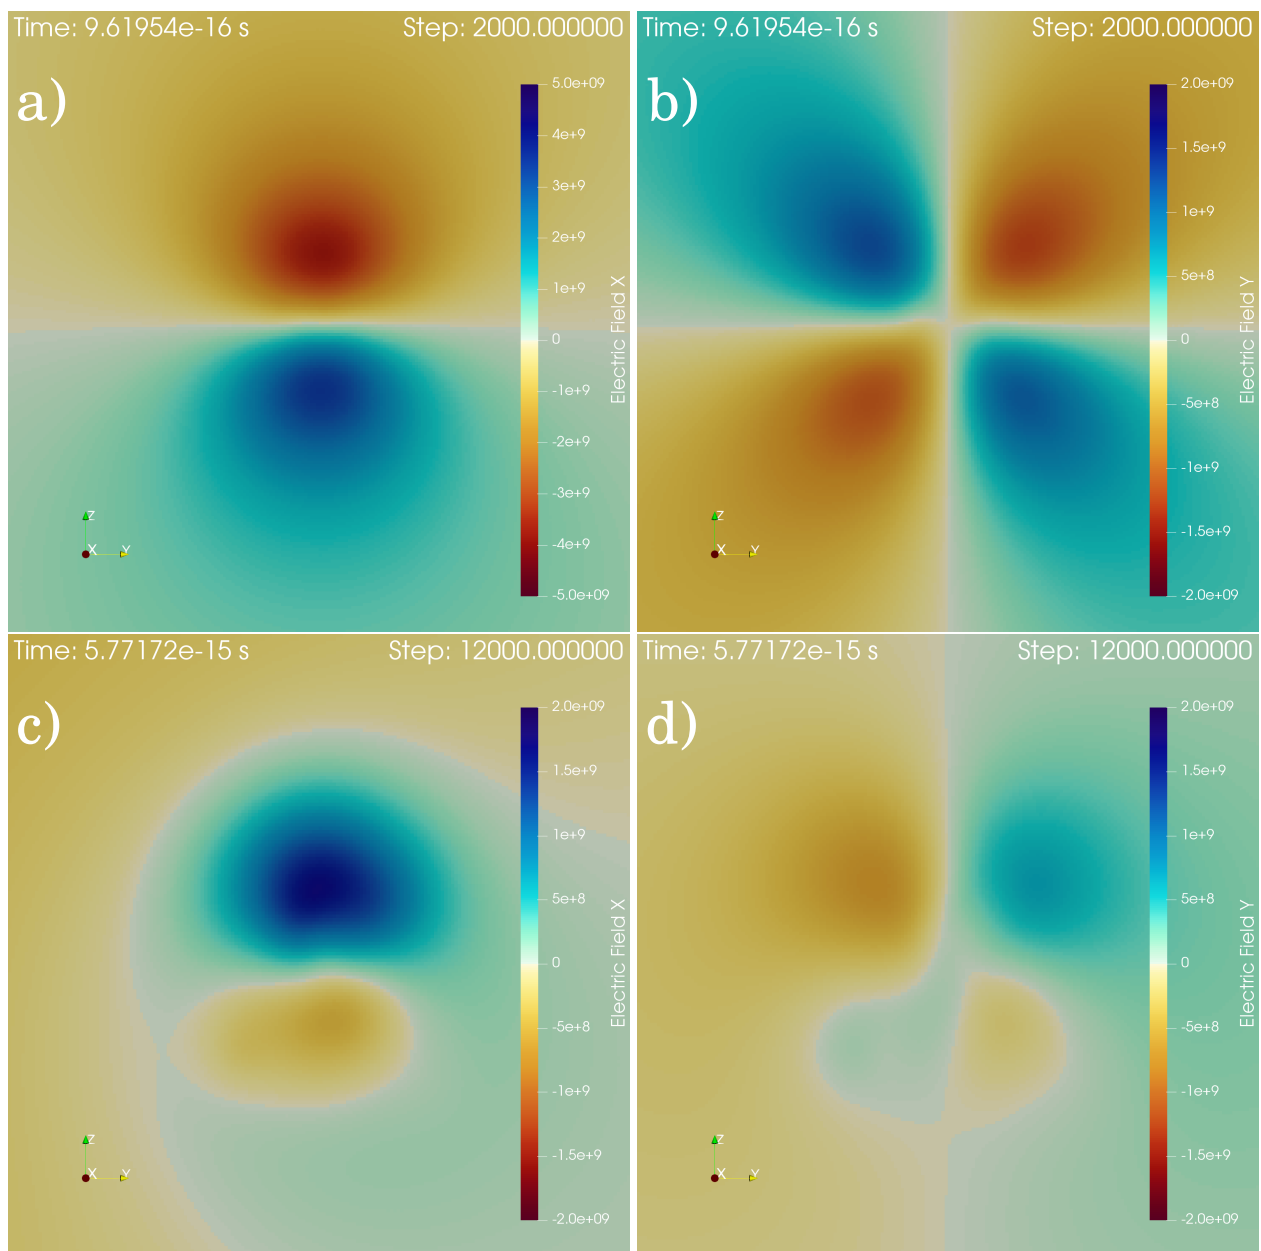


Figure 5.2: Electric field components  $E_x$  (a and c) and  $E_y$  (b and d) in a slice at  $x = 620$  or  $5$  nm beyond the end of the tip. a) and c) show the deviation of the field in the polarization direction as the laser intensity increases. b) and d) show similar behaviour in the propagation direction.  $y$  and  $z$  are increasing to the right and top of the figures, respectively. See text for more detail.

## 5.4 Simulation

In the remainder of this chapter, preliminary results are presented from a simulation of a gold nanotip using MicPIC and its photoemission routines. The density of gold typically results in more lengthy computation times compared to dielectric materials and noble gas clusters, which motivated the use of a relatively small nanotip. Moreover, it was desired to use a laser intensity that reflected those which are encountered in the laboratory in investigations of photoemission from metal nanotips.

With these considerations, a plane wave Gaussian-profile laser pulse with a 4 fs  $e^{-1}$ -width with an intensity of  $1.1 \times 10^{12}$  W/cm<sup>2</sup> and carrier-envelope phase of  $-\pi/2$ ,  $x$ -polarized and propagation in the  $y$ -direction illuminated a gold nanotip oriented in the polarization direction. The peak intensity of the pulse arrived at the apex of the nanotip after 10.15 fs or roughly  $2.11 \times 10^4$  timesteps. The corresponding time step duration is  $\Delta t \sim 4.81 \times 10^{-19}$  s, given the choice of FDTD cell size of  $\Delta x = 0.25$  nm. The nanotip itself is composed of a crystal lattice of pairs of electrons and singly-charged gold ions arranged in an FCC configuration congruent with the FDTD grid. The widest portion of the nanotip is a cylinder of radius 40 nm and length 47.5 nm, which tapers via a conical volume whose opening angle is  $18^\circ$  to a rounded tip 12.5 nm long having a radius of 10 nm, ultimately yielding  $2.59 \times 10^7$  ion-electron pairs.

The simulation was performed on the Briarée cluster described in Section 3.4.1. The simulation domain was decomposed unevenly to improve load balancing between processors, where it was attempted to maintain a roughly constant number of particles per subdomain. In total, 96 processors were used to simulate approximately  $2.9 \times 10^4$  time steps, requiring roughly 300 hours of computation time.

## 5.5 Results

At present, the most interesting features in the interaction of an laser pulse at moderate intensity with a small gold nanotip appear in the electromagnetic field; the electric field in particular displays some details not observable in a bulk model. Early in the simulation, the laser pulse is still relatively weak and the material composing the nanotip is thermalizing. Within the following few femtoseconds, perturbations to the electric field from the laser induce surface plasmons, which tend to propagate toward the end of the tip, interfering with discontinuities in the surface, due in part to the way the tip is “constructed” out of more primitive geometric solids. An example of these deviations are visible in Fig. 5.2a-d, which show the change over  $10^4$  timesteps from a collective response to a more nonuniform one, indicating a slight push from the incident laser on the electrons. This behaviour is also visible in Fig. 5.3a and b, which shows the  $x$  and  $z$  components of the electric field in a plane bisecting the nanotip in the direction of propagation of the laser. The surface plasmons initially tend to suppress any surface enhancement at the end of the tip until later times. An example of this later situation is shown in Fig. 5.4, where strong field enhancement is visible in both the  $x$  and  $z$  components, resulting in larger amplitude electron oscillations along the polarization direction. Also visible are a small number of electrons which have left the nanotip, mainly by classically propagating along the steep gradient present in the region at the tip’s end during the current half-cycle of the laser. In both cases, and those that follow, the fluctuations visible in the interior of the nanotip are due to the fields of individual particles in the crystal lattice.

These figures both hint at a curious effect present throughout the simulation. This is shown in an enhanced view of the end of the nanotip in Fig. 5.5. The effect begins as a variation in the field at the vacuum interface during the thermalization phase, which subsequently takes on a quasi-standing wave pattern as it slowly proceeds deeper into the material. It does not oscillate with the carrier frequency and is perhaps linked to the energy absorbed by the material. This could be a numerical artifact due to the proximity of the TFSF interface to the rear end of the tip (low  $x$ ) although it is also possible that it is due to structural changes taking place inside the ionic lattice. More simulations are required to identify a definitive cause of the phenomenon.

MicPIC has the capacity to perform the necessary calculations for determining the tunneling probability according to Eqs. (1.49) and (1.50) and Section 5.3.2. Presented here are some of the first ever results of the

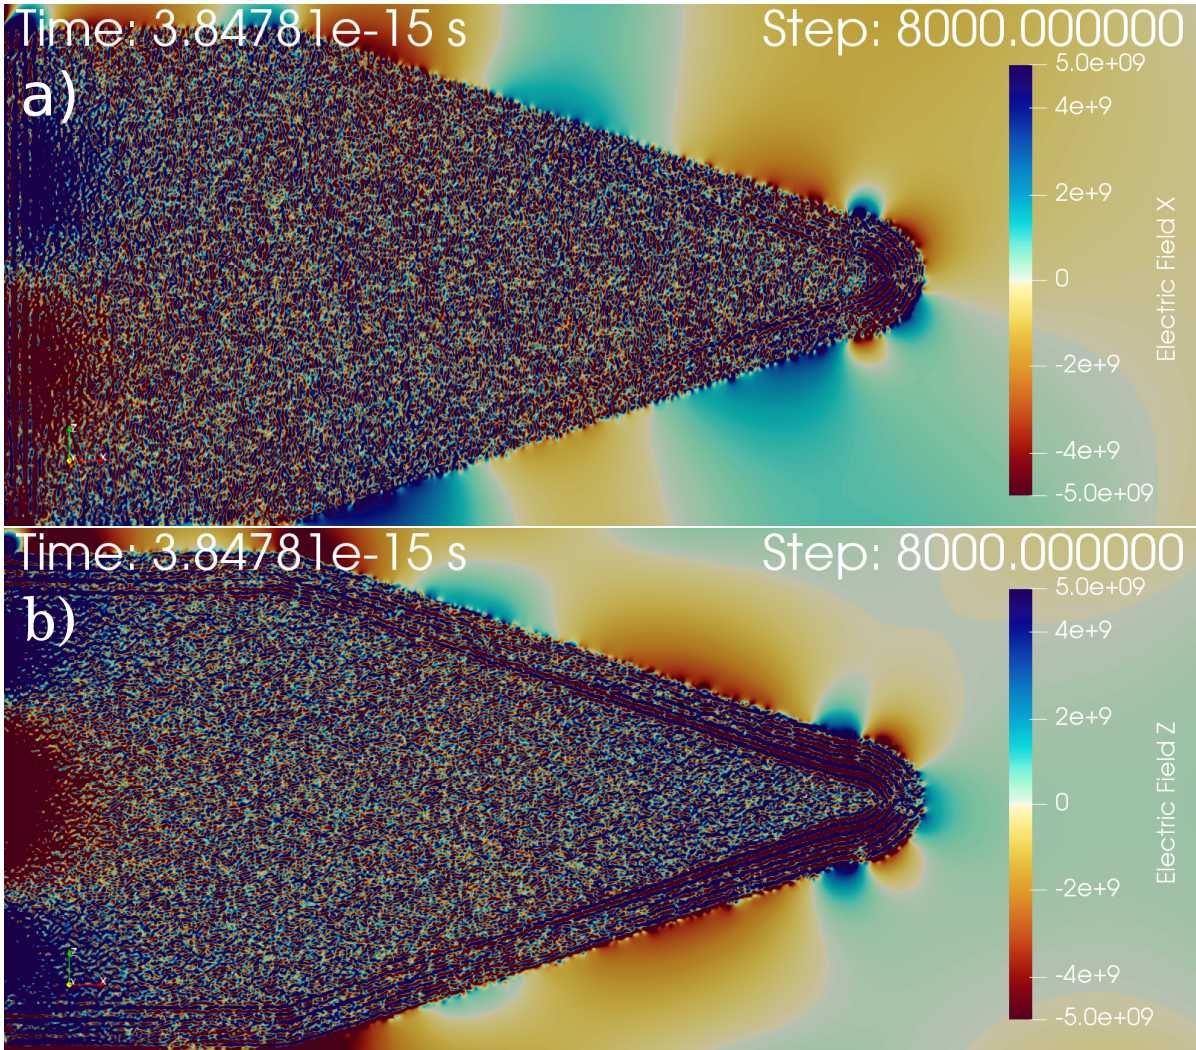


Figure 5.3: A slice of the  $x$  (a) and  $z$  (b) components of the electric field bisecting the tip in the propagation direction at step 8000.  $x$  and  $z$  are increasing to the right and top of the figures, respectively.

photoemission component of MicPIC. The reader is reminded of that these results are preliminary and are even being continuously generated/updated at the time of writing; currently available results are visualized in Fig. 5.6a and b. Fig. 5.6 shows an angled view of the length of each tunneling path from steps 10000 to 26000 exaggerated by a factor of 10 for visibility and coloured by the time step at which the event occurs. Surprisingly, the photoemission of electrons appears to take place over a wide area of the surface of the nanotip, with the majority on the lower  $z$  portion.

These photoemission times are shown in Fig. 5.6 compared to the laser pulse magnitude at the TFSF origin  $(x, y, z) = (4, 4, 4)$  cells. Since the photoemission events are asymmetrically distributed over the surface of the sample, there would appear to be more complex dynamics at play than simply the laser polarization and field enhancement due to collective electron motions. The photoemission method used in MicPIC (Section 5.3) is sensitive to microfield fluctuations, which appear at all visible points on the surface of the metal. For this reason, occurrences of photoemission events may be enhanced due to strong local microscopic variations in the electric field. A complete set of results will allow a more detailed look at the photoelectron energies as

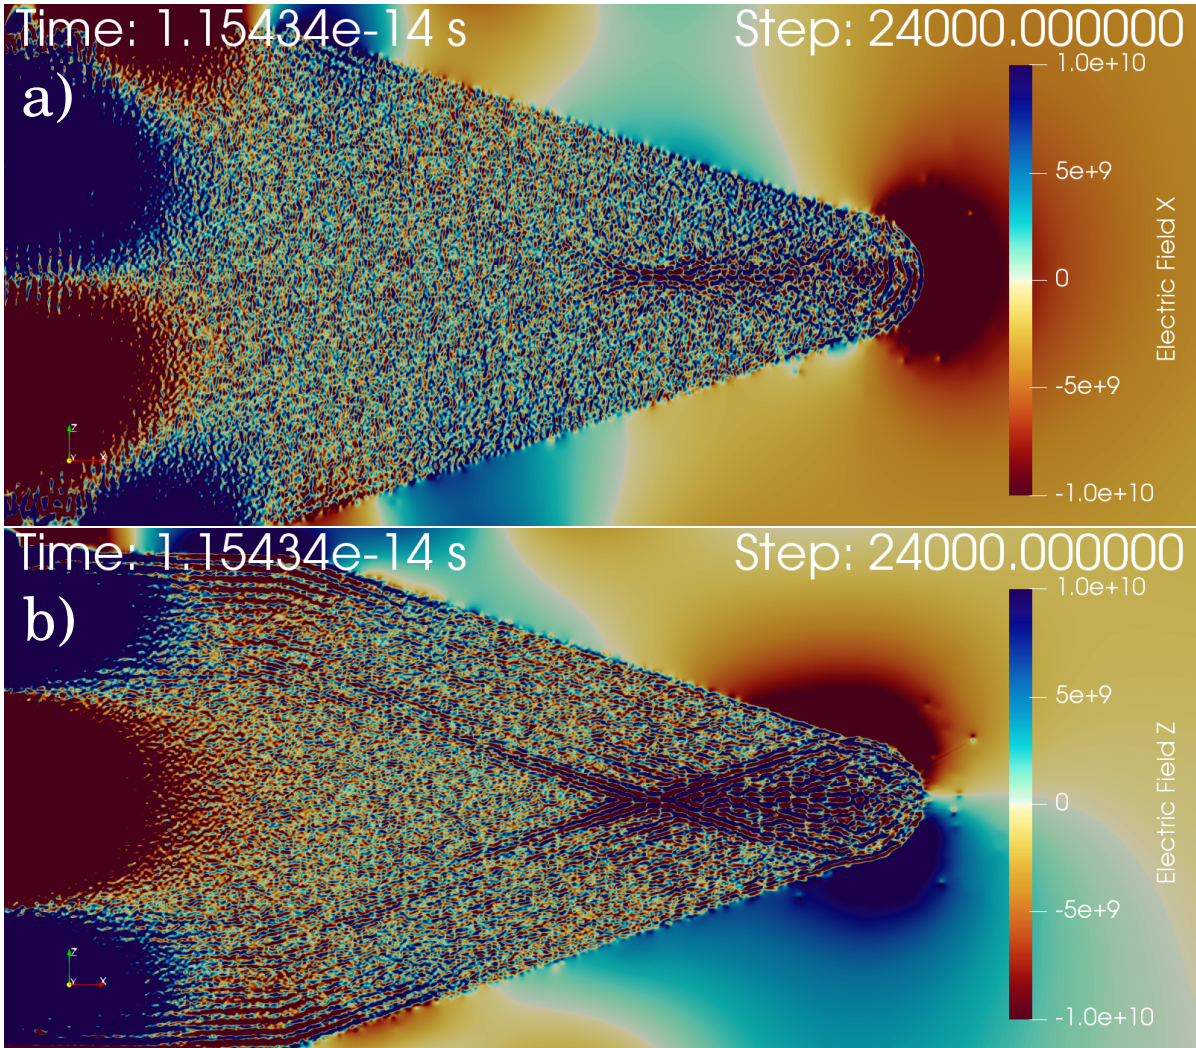


Figure 5.4: A slice of the  $x$  (a) and  $z$  (b) components of the electric field bisecting the tip in the propagation direction at step 24000.  $x$  and  $z$  are increasing to the right and top of the figures, respectively.

well as the electron-electron correlation.

## 5.6 Conclusion

This chapter introduced in detail the implementation of photoemission of electrons from surfaces in MicPIC, demonstrating the tantalizing possibility of modelling a wide range of interesting intense field phenomena in nanostructures. The electric field is shown to deviate from a bulk plasmon response, which can be attributed to a slight off-axis push of the electrons until later times, when the electronic response to the laser polarization becomes stronger than the plasmonic fluctuations visible early in the simulation. Additionally, the appearance of an intriguing phenomenon inside the nanotip is presented. Presently, there are insufficient results to attempt a rigorous explanation of this effect. Lastly, the photoemission of electrons is shown in terms of the photoemission paths in three dimensions and in terms of event times compared to the laser pulse amplitude. These are still very preliminary results and future efforts will be directed at simulations of larger

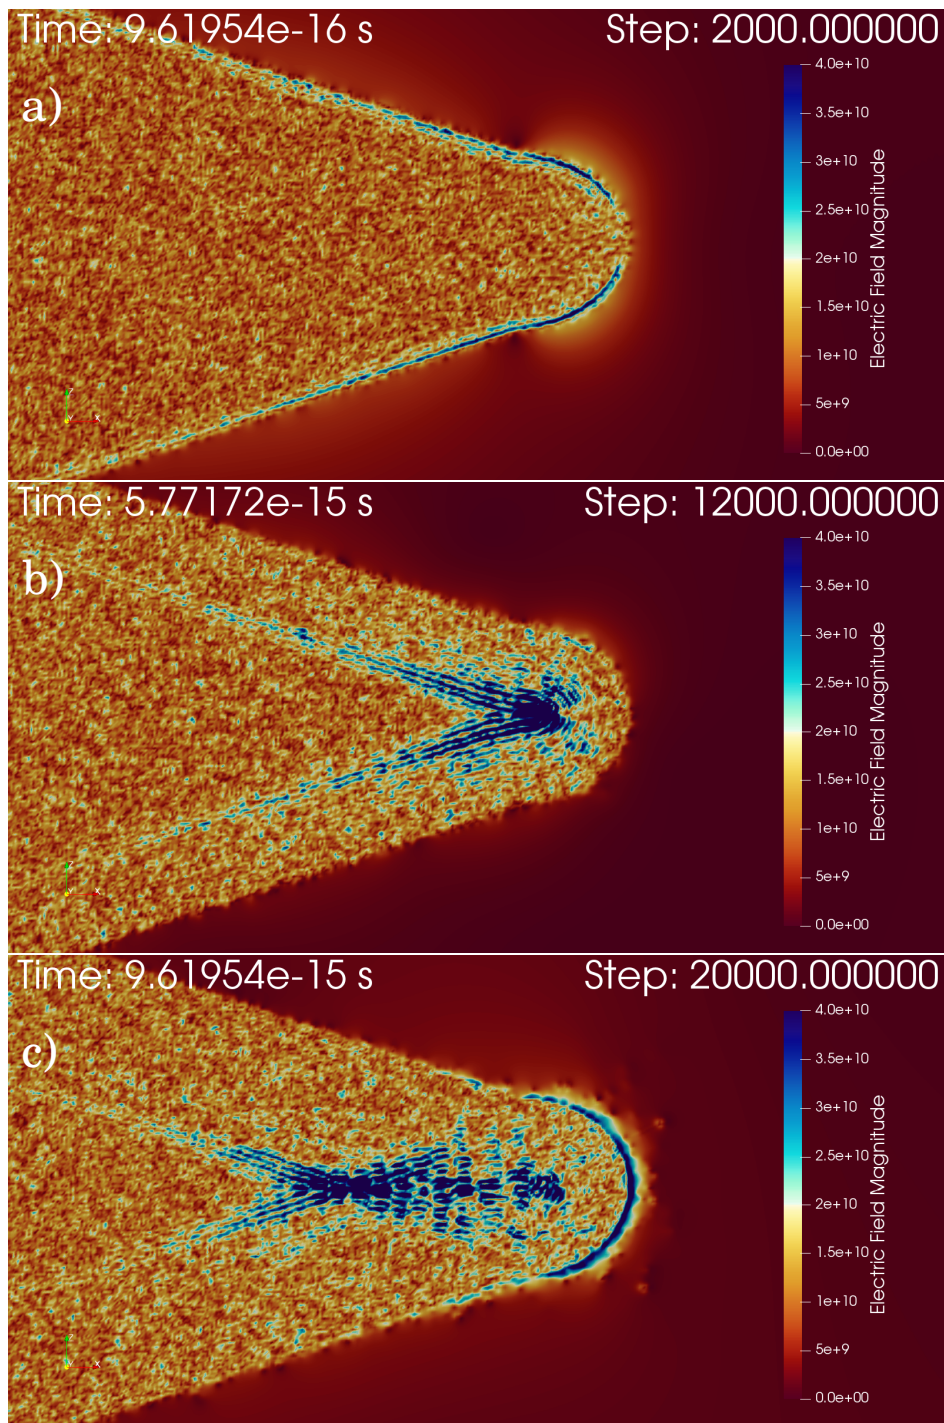


Figure 5.5: Slices of the magnitude of the electric field bisecting the tip in the propagation direction at steps 8000 (a), 12000 (b), and 20000 (c).  $x$  and  $z$  are increasing to the right and top of the figures, respectively. The pattern in the field slowly moves into the interior of the sample and is does not appear to be linked to the carrier frequency. The small distortions near the end of the tip are individual electrons.

nanotips and other experimental parameters, such as DC bias and carrier-envelope phase. Unfortunately, harmonic spectra of the nanotip fields are not currently available. Nevertheless, the utility and level of detail


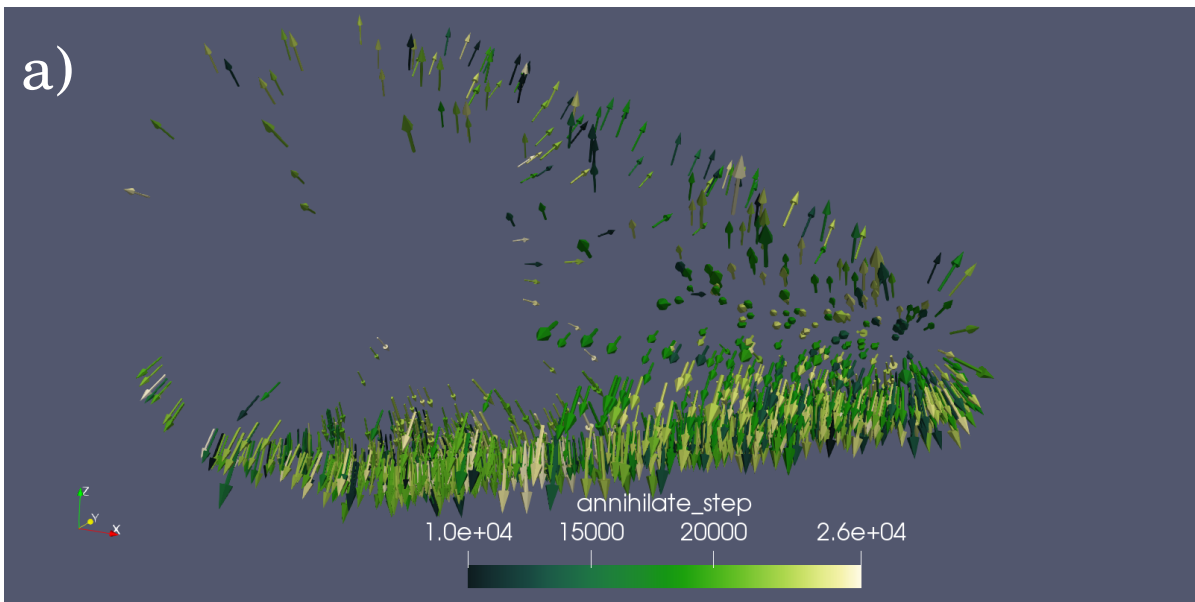


Figure 5.6: a) Photoemission paths shown exaggerated by a factor of 10 for visibility. Darker green arrows indicate earlier events; conversely, lighter arrows correspond to events at later times. b) Comparison of photoemission times with the incident laser pulse amplitude. The peaks of the photoemission tend to correspond with the peak amplitudes of the corresponding cycle.

available with MicPIC opens a new avenue for investigating photoemission from metal surfaces.

---

## CHAPTER

# SIX

---

## CONCLUSIONS

This thesis presented a numerical tool, MicPIC, for simulating the interaction of ultrafast optical processes in strongly coupled plasmas spanning a range of sizes from the atomic to the micron scale, with attosecond resolution. The motivation for such a tool was discussed in the context of current numerical methods in plasma physics. Special attention was paid to the theory and implementation of MicPIC, including the novel corrections to the short range force action on a particle, as well as how it incorporates tunnel and impact ionization processes. The utility of MicPIC has been greatly enhanced by its parallelization, allowing very large, distributed memory simulations of strongly coupled plasmas to be performed on timescales of weeks instead of years. The code has been shown to perform well, possessing both strong and weak linear scaling on some of Canada's largest computing clusters. Two applications of MicPIC were shown via preliminary results. In the first case, comparison of MicPIC to PIC showed subtle differences, the most notable of which was the effectiveness of MicPIC at distributing energy absorbed by the laser to the surrounding environment. Both case showed similar behaviour of plasma waves and surface plasmons. In the data presently available, it was not possible to resolve more than the 5th harmonic of the applied laser field, likely being attributable to incomplete data in the time domain. In the second case, photoemission of electrons from a small metal nanotip was demonstrated, with intriguing results. The theory behind this inherently quantum mechanical process and how it is implemented in a massively parallel classical plasma simulation has been described in detail.

In future efforts with MicPIC, ionization of more complex molecules will be included in the code, so that the entire process of plasma formation leading to high harmonic generation can be modelled, with the aim of enhancing the dynamics of the calculations presented in Chapter 4. The resulting plasma will move unequivocally into the strong-coupling regime, where MicPIC is able to provide the most accuracy, and will ultimately yield more plasma density variations, steeper electric field gradients, and stronger HHG. Additionally, near term simulations will compare MicPIC with much more PIC-like results, using super-particles and larger cell sizes to more clearly emphasize the difference, as the results shown here involve the small size limit of PIC. More results are forthcoming with nanotip simulations, as post-processing tools for analyzing the results are still being designed. Only a small sample of the available data for the two tests have been presented; there is much still to extract from both.

In short, MicPIC is now a powerful, massively parallel tool for investigating strong field laser plasma

physics phenomena in solid density materials.

# BIBLIOGRAPHY

- [1] T. Hoefler, R. Belli, Scientific benchmarking of parallel computing systems: twelve ways to tell the masses when reporting performance results, International Conference for High Performance Computing, Networking, Storage and Analysis (2015) 1–12.
- [2] L. Tonks, The birth of “plasma”, American Journal of Physics 35 (1967) 857.
- [3] I. Langmuir, Oscillations in ionized gases, Proc. Natl. Acad. Sci. 14 (1928) 627–637.
- [4] H. M. van Horn, Dense astrophysical plasmas, Science 252 (1991) 384–389.
- [5] M. D. A., M. Rosenberg, Cosmic dusty plasma, Annu. Rev. Astron. Astrophys. 32 (1994) 419–463.
- [6] P. J. Coleman Jr., Turbulence, viscosity, and dissipation in the solar-wind plasma, The Astrophysical Journal 153 (1968) 371–388.
- [7] A. Y. Potekhin, D. A. Baiko, P. Hansel, D. G. Yakovlev, Transport properties of degenerate electrons in neutron star envelopes and white dwarf cores, Astron. Astrophys. 346 (1999) 345–353.
- [8] W. L. Boeck, O. H. Vaughn, R. J. Blakeslee, B. Vonnegut, M. Brook, The role of the space shuttle videotapes in the discovery of sprites, jets, and elves, J. Atmospheric Sol.-Terr. Phys. 60 (1998) 669–677.
- [9] T. Neubert, On sprites and their exotic kin, Science 300 (2003) 747–749.
- [10] S. Ichimaru, Statistical Plasma Physics, Vol. 1: Basic Principles, Addison-Wesley, New York, 1991.
- [11] F. Krausz, M. Ivanov, Attosecond physics, Rev. Mod. Phys. 81 (2009) 163–234.
- [12] P. B. Corkum, F. Krausz, Attosecond science, Nature Physics 3 (2007) 381–387.
- [13] M. F. Ciappina, J. A. Pérez-Hernández, A. S. Landsman, W. A. Okell, S. Zherebstov, B. Förg, L. Schötz, J. Seiffert, T. Fennel, T. Shaaran, T. Zimmermann, A. Chacón, R. Guichard, A. Zaïr, J. W. G. Tisch, J. P. Marangos, T. Witting, A. Braun, S. A. Maier, L. Roso, M. Krüger, P. Hommelhoff, M. F. Kling, F. Krausz, M. Lewenstein, Attosecond physics at the nanoscale, Rep. Prog. Phys. 80 (2017) 054401.
- [14] T. Brabec, F. Krausz, Intense few-cycle laser fields: frontiers of nonlinear optics, Rev. Mod. Phys. 72 (2000) 545–591.
- [15] F. Krausz, The birth of attosecond physics and its coming of age, Physica Scripta 91 (2016) 063011.

- [16] A. J. de Maria, W. H. Glenn, M. J. Brienza, M. E. Mack, Picosecond laser pulses, *Proc. IEEE* 57 (1969) 2–25.
- [17] R. L. Fork, C. H. Brito Cruz, P. C. Becker, C. V. Shank, Compression of optical pulses to six femtoseconds by using cubic phase compensation, *Optics Letters* 12 (1987) 483–485.
- [18] P. M. W. French, The generation of ultrashort laser pulses, *Rep. Prog. Phys.* 58 (1995) 169–262.
- [19] M. Dantus, M. J. Rosker, A. H. Zewail, Real-time femtosecond probing of "transition states" in chemical reaction, *J. Chem. Phys.* 87 (1987) 2395–2397.
- [20] A. H. Zewail, Femtochemistry: Atomic-scale dynamics of the chemical bond, *J. Phys. Chem. A* 104 (2000) 5660–5694.
- [21] M. Hentschel, R. Kienberger, C. Spielmann, G. A. Reider, N. Milosevic, T. Brabec, P. Corkum, U. Heinzmann, M. Drescher, F. Krausz, Attosecond metrology, *Nature* 414 (2001) 509–513.
- [22] N. H. Burnett, H. A. Baldis, M. C. Richardson, G. D. Enright, Harmonic generation in CO<sub>2</sub> laser target interaction, *Appl. Phys. Lett.* 31 (1977) 172–174.
- [23] R. L. Carman, D. W. Forslund, K. M. Kindel, Visible harmonic emission as a way of measuring profile steepening, *Phys. Rev. Lett.* 46 (1996) 29–32.
- [24] R. L. Carman, C. K. Rhodes, R. F. Benjamin, Observation of harmonics in the visible and ultraviolet created in co<sub>2</sub>-laser-produced plasmas, *Phys. Rev. A* 24 (1996) 2649–2653.
- [25] A. L’Huillier, P. Balcou, High-order harmonic generation in rare gases with a 1-ps 1053-nm laser, *Phys. Rev. Lett.* 70 (1993) 774–777.
- [26] J. J. Macklin, J. D. Kmetec, C. L. Gordon III, High-order harmonic generation using intense femtosecond pulses, *Phys. Rev. Lett.* 70 (1993) 766–769.
- [27] D. von der Linde, T. Engers, G. Jenke, P. Agostini, G. Grillon, E. Nibbering, A. Mysyrowicz, A. Antonetti, Generation of high-order harmonics from solid surfaces by intense femtosecond laser pulses, *Phys. Rev. A* 52 (1995) R25.
- [28] T. T. Luu, M. Garg, S. Y. Kruchinin, A. Moulet, M. T. Hassan, E. Goulielmakis, Extreme ultraviolet high-harmonic spectroscopy of solids, *Nature* 521 (2015) 498–502.
- [29] S. Ghimire, A. D. DiChiara, E. Sistrunk, P. Agostini, L. F. DiMauro, D. A. Reis, Observation of high-order harmonic generation in a bulk crystal, *Nature Physics* 7 (2011) 138–141.
- [30] O. Schubert, M. Hohenleutner, F. Langer, B. Urbanek, C. Lange, U. Huttner, D. Golde, T. Meier, M. Kira, S. W. Koch, R. Huber, Sub-cycle control of terahertz high-harmonic generation by dynamical Bloch oscillations, *Nature Photonics* 8 (2014) 119–123.
- [31] G. Vampa, B. G. Ghamsari, S. Siadat Mousavi, T. J. Hammond, A. Olivieri, E. Lisicka-Skrek, A. Yu Naumov, D. M. Villeneuve, A. Staudte, P. Berini, P. B. Corkum, Plasmon-enhanced high-harmonic generation from silicon, *Nature Physics* 13 (2017) 659–662.
- [32] I. McKinnie, H. Kapteyn, High-harmonic generation: Ultrafast lasers yield x-rays, *Nature Photonics* 4 (2010) 149–151.

- [33] P. Corkum, Plasma perspective on strong-field multiphoton ionization, *Phys. Rev. Lett.* 71 (1993) 1994–1997.
- [34] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L’Huillier, P. B. Corkum, Theory of high-harmonic generation by low-frequency laser fields, *Phys. Rev. A.* 49 (1994) 2117–2132.
- [35] P. Gibbon, High-order harmonic generation in plasmas, *IEEE J. Quantum Electron.* 33 (1997) 1915–1924.
- [36] A. Borot, A. Malvache, X. Chen, A. Jullien, J.-P. Geindre, P. Audebert, G. Mourou, F. Quéré, R. Lopez-Martens, Attosecond control of collective electron motion in plasmas, *Nature Physics* 8 (2012) 416–421.
- [37] S. V. Bulanov, N. M. Naumova, F. Pegoraro, Interaction of an ultrashort, relativistically strong laser pulse with an overdense plasma, *Phys. Plasmas* 1 (1994) 745–757.
- [38] F. Brunel, Not-so-resonant, resonant absorption, *Phys. Rev. Lett.* 59 (1987) 52–55.
- [39] C. Thaury, F. Quéré, High-order harmonic and attosecond pulse generation on plasma mirrors: basic mechanisms, *J. Phys. B* 43 (2010) 213001.
- [40] A. S. Pirozhkov, S. V. Bulanov, T. Z. Esirkepov, M. Mori, A. Sagisaka, H. Daido, Attosecond pulse generation in the relativistic regime of the laser-foil interaction: the sliding mirror model, *Physics of Plasmas* 13 (2006) 013107.
- [41] U. Huttner, M. Hohenleutner, F. Langer, O. Schubert, M. Knorr, R. Huber, S. W. Koch, M. Kira, High-harmonic generation in solids, *Proc. of SPIE* 9746.
- [42] N. Tancogne-Dejean, O. D. Mücke, F. X. Kärtner, A. Rubio, Impact of the electronic band structure in high-harmonic generation in solids, *Phys. Rev. Lett.* 118 (2017) 087403.
- [43] G. Vampa, C. R. McDonald, G. Orlando, D. D. Klug, P. B. Corkum, T. Brabec, Theoretical analysis of high-harmonic generation in solids, *Phys. Rev. Lett.* 113 (2014) 073901.
- [44] A. F. Kemper, B. Moritz, J. K. Freericks, T. P. Devereaux, Theoretical description of high-order harmonic generation in solids, *New J. Phys.* 15 (2013) 1–15.
- [45] J. D. Cox, A. Marini, F. J. García de Abajo, Plasmon-assisted high-harmonic generation in graphene, *Nature Communications* 8 (2017) 1–7.
- [46] T. Tajima, J. M. Dawson, Laser electron accelerator, *Phys. Rev. Lett.* 43 (1979) 267–270.
- [47] S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori, L. O. Silva, Exploring laser-wakefield-accelerator regimes for near-term lasers using particle-in-cell simulation in lorentz-boosted frames, *Nature Physics* 6 (2010) 311–316.
- [48] J. P. Couperus, R. Pausch, A. Köhler, O. Zarini, J. M. Krämer, M. Garten, A. Heubl, R. Gebhardt, U. Helbig, S. Bock, K. Zeil, A. Debus, M. Bussmann, U. Schramm, A. Irman, Demonstration of a beam loaded nanocoulomb-class laser wake field accelerator, *Nature Communications* 8 (2017) 487.
- [49] E. G. Gamaly, A. V. Rode, B. Luther-Davies, V. T. Tikhonchuk, Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics, *Physics of Plasmas* 9 (2002) 949–957.

- [50] K. Wædegaard, M. Frislev, P. Balling, Femtosecond laser excitation of dielectric materials: experiments and modeling of optical properties and ablation depths, *Appl. Phys. A* 110 (2013) 601–605.
- [51] P. Dittrich, R. Bartolme, G. Montemezzani, P. Günter, Femtosecond laser ablation of dast, *Appl. Surf. Sci.* 220 (2003) 88–95.
- [52] H. Kumagai, K. Midorikawa, K. Toyoda, S. Nakamura, T. Okamoto, M. Obara, Ablation of polymer films by a femtosecond high-peak-power ti:sapphire laser at 798nm, *Appl. Phys. Lett.* 65 (1994) 1850–1852.
- [53] M. D. Shirk, P. A. Molian, A review of ultrashort pulsed laser ablation of materials, *Journal of Laser Applications* 10 (1998) 18–28.
- [54] B. Chimier, O. Utéza, N. Sanner, M. Sentis, T. Itina, P. Lassonde, F. Légaré, F. Vidal, J. C. Kieffer, Damage and ablation thresholds of fused-silica in femtosecond regime, *Phys. Rev. B* 84 (2011) 094104.
- [55] R. R. Gattas, E. Mazur, Femtosecond micromachining in transparent materials, *Nature Photonics* 2 (2008) 219–225.
- [56] L. Cerami, E. Mazur, S. Nolte, C. B. Schaffer, *Femtosecond Laser Micromachining*, Springer International Publishing, Heidelberg, 2013, pp. 287–321.
- [57] P. Balling, J. Schou, Femtosecond-laser ablation dynamics of dielectrics: basics and applications for thin films, *Rep. Prog. Phys* 76 (2013) 036502.
- [58] V. R. Bhardwaj, E. Simova, P. P. Rajeev, C. Hnatovsky, R. S. Taylor, D. M. Rayner, P. B. Corkum, Optically produced arrays of planar nanostructures inside fused silica, *Phys. Rev. Lett.* 96 (2006) 057404.
- [59] L. Englert, M. Wollenhaupt, C. Sarpe, D. Otto, T. Baumert, Morphology of nanoscale structures on fused silica surfaces from interaction with temporally tailored femtosecond pulses, *Journal of Laser Applications* 24 (2012) 042002.
- [60] F. Chen, J. R. Vásquez de Aldana, Optical waveguides in crystalline dielectric materials produced by femtosecond-laser micromachining, *Laser Photonics Rev.* 8 (2014) 251–275.
- [61] K. W., S. Mitterer, S. Krüger, W. Husinsky, G. Grabner, Femtosecond-pulse laser ablation of human corneas, *Appl. Phys. A* 58 (1994) 513–518.
- [62] I. Ratkay-Traub, T. Juhasz, H. C, C. Suarez, K. Kiss, I. Ferincz, R. Kurtz, Ultrashort pulse (femtosecond) laser surgery: initial use in lasik flap creation, *Ophthalmolog. Clin. North Am.* 14 (2001) 347–355.
- [63] Z. Nagy, A. Takacs, T. Filkorn, M. Sarayba, Initial clinical evaluation of an intraocular femtosecond laser in cataract surgery, *Journal of Refractive Surgery* 25 (2009) 1053.
- [64] C. D. Lin, A.-T. Le, Z. Chen, T. Morishita, R. Lucchese, Strong-field rescattering physics–self-imaging of a molecule by its own electrons, *J. Phys. B.* 43 (2010) 122001.
- [65] M. T. Hassan, Attomicroscopy: from femtosecond to attosecond electron microscopy, *J. Phys. B* 51 (2018) 032005.

- [66] M. Krüger, M. Schenk, P. Hommelhoff, Attosecond control of electrons emitted from a nanoscale metal tip, *Nature* 475 (2011) 78–81.
- [67] P. Hommelhoff, Y. Sortais, A. Aghajani-Talesh, M. A. Kasevich, Field emission tip as a nanometer source of free electron femtosecond pulses, *Phys. Rev. Lett.* 96 (2006) 077401.
- [68] T. Fennel, K.-H. Meiwes-Broer, J. Tiggesbäumker, P.-G. Reinhard, P. M. Dinh, E. Suraud, Laser-driven nonlinear cluster dynamics, *Rev. Mod. Phys.* 82 (2010) 1793–1842.
- [69] S. C. Wilks, W. L. Kruer, W. B. Mori, Odd harmonic generation of ultra-intense laser pulses reflected from an overdense plasma, *IEEE Trans. on Plasma Science* 38 (2010) 1793–1842.
- [70] R. W. Hockney, J. W. Eastwood, *Compute Simulation Using Particles*, McGraw-Hill, New York, 1981.
- [71] V. Vahedi, M. Surrendra, A monte-carlo collision model for the particle-in-cell method: applications to argon and oxygen discharges, *Comput. Phys. Commun.* 87 (1995) 179–198.
- [72] K. Nanbu, Theory of cumulative small-angle collisions in plasmas, *Phys. Rev. E* 45 (1997) 4642–4652.
- [73] C. K. Birdsall, Particle-in-cell charged-particle simulations, plus monte carlo collisions with neutral atoms, *pic-mcc*, *IEEE Trans. on Plasma Science* 19 (1991) 65–85.
- [74] F. F. Chen, *Plasma Physics and Controlled Fusion*, 2nd Edition, Vol. 1: Plasma Physics, Springer, New York, 2006.
- [75] E. E. Salpeter, Electron screening and thermonuclear reaction, *Aust. J. Phys.* 7 (1954) 373–388.
- [76] S. Ichimaru, Strongly coupled plasmas: high-density classical plasmas and degenerate electron liquids, *Rev. Mod. Phys.* 54 (1982) 1017–1059.
- [77] M. Kardar, *Statistical Physics of Particles*, Cambridge University Press, New York, 2007.
- [78] I. L. Klimontovich, On the method of "second quantization" in phase space, *Soviet Physics JETP* 6 (1958) 753–760.
- [79] J. R. Oppenheimer, Three notes on the quantum theory of aperiodic effects, *Phys. Rev.* 13 (1928) 66–81.
- [80] L. V. Keldysh, Ionization in the field of a strong electromagnetic wave, *Soviet Physics JETP* 20 (1965) 1307–1314.
- [81] A. M. Perelomov, V. S. Popov, M. V. Terent'ev, Ionization of atoms in an alternating electric field, *J. Exptl. Theoret. Phys.* 23 (1966) 924–934.
- [82] H. R. Reiss, Effect of an intense electromagnetic field on a weakly bound system, *Phys. Rev. A* 22 (1980) 1786–1813.
- [83] M. V. Ammosov, N. B. Delone, V. P. Krainov, Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field, *Soviet Physics JETP* 64 (1986) 1191–1194.
- [84] W. Lotz, An empirical formula for the electron-impact ionization cross-section, *Z. Phys.* 206 (1967) 205–211.

- [85] T. Fennel, L. Ramunno, T. Brabec, Highly charged ions from laser-cluster interactions: Local-field-enhanced impact ionization and frustrated electron-ion recombination, *Phys. Rev. Lett.* 99 (2007) 233401.
- [86] Y. Gontier, M. Poirier, M. Trahin, Multiphoton absorptions above the ionization threshold, *J. Phys. B* 13 (1980) 1381–1387.
- [87] E. A. Martin, L. Mandel, Electron energy spectrum in laser-induced multiphoton ionization of atoms, *Appl. Opt.* 15 (1976) 2378–2380.
- [88] P. Agostini, F. Fabre, G. Mainfray, G. Petite, N. K. Rahman, Free-free transitions following six-photon ionization of xenon atoms, *Phys. Rev. Lett.* 42 (1979) 1127–1130.
- [89] L. A. Lompré, G. Mainfray, C. Manus, G. Farkas, Multiphoton absorption by electrons, *Phys. Rev. Lett.* 43 (1979) 1243–1245.
- [90] M. Protopapas, C. H. Keitel, P. L. Knight, Atomic physics with super-high intensity lasers, *Rep. Prog. Phys.* 60 (1997) 389–486.
- [91] T. F. Gallagher, Above-threshold ionization in low-frequency limit, *Phys. Rev. Lett.* 61 (1988) 2304–2307.
- [92] F. Quéré, C. Thauray, P. Monot, S. Dobosz, P. Martin, Coherent wake emission of high-order harmonics from overdense plasmas, *Phys. Rev. Lett.* 96 (2006) 125004.
- [93] R. H. Fowler, L. Nordheim, Electron emission in intense electric field, *Proc. R. Soc. A* 119 (1928) 173–181.
- [94] F. V. Bunkin, M. V. Fedorov, Cold emission of electrons from the surface of a metal in a strong radiation field, *Soviet Physics JETP* 21 (1965) 896–899.
- [95] R. G. Forbes, J. H. B. Deane, Reformulation of the standard theory of fowler-nordheim tunnelling and cold field electron emission, *Proc. R. Soc. A* 463 (2007) 2907–2927.
- [96] D. Bohm, *Quantum Theory*, Dover Publications Inc., New York, 1989.
- [97] T. Pfeiffer, C. Spielmann, G. Gerber, Femtosecond x-ray science, *Rep. Prog. Phys.* 69 (2006) 443–505.
- [98] P. Hommelhoff, M. A. Kealhofer, C. Kasevich, Ultrafast electron pulses from a tungsten tip triggered by low-power femtosecond laser pulses, *Phys. Rev. Lett.* 97 (2006) 247402.
- [99] T. Ditmire, Simulation of exploding clusters ionized by high-intensity femtosecond laser pulses, *Phys. Rev. A* 57 (1998) R4094–R4097.
- [100] U. Saalman, J. M. Rost, Electron dynamics in strong laser pulse illumination of large rare gas clusters, *Eur. Phys. J. D.* 36 (2005) 159–164.
- [101] C. K. Birdsall, A. B. Langdon, *Plasma Physics via Computer Simulation*, McGraw-Hill, New York, 1985.
- [102] D. Tskhakaya, R. Schneider, Optimization of pic codes by improved memory management, *J. Comput. Phys.* 225 (2007) 829–839.

- [103] K. Matyash, R. Schneider, F. Taccogna, A. Hatayama, S. Longo, M. Capitelli, D. Tskhakaya, F. X. Bronold, Particle in cell simulation of low temperature laboratory plasmas, *Contrib. Plasma Phys.* 47 (2007) 595–634.
- [104] C. Peltz, C. Varin, T. Brabec, T. Fennel, Fully microscopic analysis of laser-driven finite plasmas using the example of clusters, *New J. Phys.* 14 (2012) 065011.
- [105] C. Varin, C. Peltz, T. Brabec, T. Fennel, Light wave driven electron dynamics in clusters, *Ann. der Phys.* 526 (2014) 135–156.
- [106] C. Varin, C. Peltz, T. Brabec, T. Fennel, Attosecond plasma wave dynamics in laser-driven cluster nanoplasmas, *Phys. Rev. Lett.* 108 (2012) 175007.
- [107] C. Peltz, C. Varin, T. Brabec, T. Fennel, Time-resolved x-ray imaging of anisotropic nanoplasma expansion, *Phys. Rev. Lett.* 133 (2014) 133401.
- [108] A. Taflove, S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, 2nd Edition, Artech House, Boston, 2000.
- [109] K. Yee, Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media, *IEEE Trans. Antennas Propagat.* 14 (1966) 302–307.
- [110] I. Çapoğlu, A. Taflove, V. Backman, Computation of tightly-focused laser beams in the fdtd method, *Optics Express* 21 (2013) 87–101.
- [111] J.-P. Berenger, A perfectly matched layer for the absorption of electromagnetic waves, *J. Comp. Phys.* 114 (1994) 185–200.
- [112] S. D. Gedney, An anisotropic perfectly matched layer-absorbing medium for the truncation of fdtd lattices, *IEEE Trans. Antennas Propagat.* 44 (1996) 1630–1639.
- [113] J. P. Verboncoeur, A. B. Langdon, N. T. Gladd, An object-oriented electromagnetic pic code, *Comput. Phys. Comm.* 85 (1995) 199–211.
- [114] D. Tskhakaya, K. Matyash, R. Schneider, F. Taccogna, The particle-in-cell method, *Contrib. Plasma Phys.* 47 (2007) 563–594.
- [115] S. Markidis, G. Lapenta, R.-U., Multi-scale simulations of plasma with ipic3d, *Mathematics and Computers in Simulation* 80 (2009) 1509–1519.
- [116] M. P. Allen, D. J. Tildesley, *Computer Simulation of Liquids*, Oxford University Press, New York, 1989.
- [117] J. P. Boris, Relativistic plasma simulation - optimization of a hybrid code, *Proc. 4th Conf. Numerical Simulation of Plasmas* (1970) 3–67.
- [118] H. Qin, S. Zhang, J. Xiao, J. Liu, Y. Sun, W. M. Tang, Bonding of glass with femtosecond laser pulses at high repetition rates, *Physics of Plasmas* 20 (2013) 084503.
- [119] L. V. Keldysh, Ionization in the field of a strong electromagnetic wave, *J. Exptl. Theoret. Phys.* 20 (1965) 1307–1314.

- [120] S. S. Mao, F. Quéré, S. Guizard, X. Mao, R. E. Russo, G. Petite, P. Martin, Dynamics of femtosecond laser interactions with dielectrics, *Appl. Phys. A* 79 (2004) 1695–1709.
- [121] E. G. Gamaly, A. V. Rode, B. Luther-Davies, V. T. Tikhonchuk, Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics, *Physics of Plasmas* 9 (2002) 949–957.
- [122] A. Vogel, J. Noack, G. Hüttman, G. Paltauf, Mechanisms of femtosecond laser nanosurgery of cells and tissues, *Appl. Phys. B* 81 (2005) 1015–1047.
- [123] T. Tamaki, W. Watanabe, J. Nishii, K. Itoh, Welding of transparent materials using femtosecond laser pulses, *Japanese Journal of Applied Physics* 44 (2005) 20–23.
- [124] S. Richter, S. Döring, A. Tünnermann, S. Nolte, Bonding of glass with femtosecond laser pulses at high repetition rates, *Appl. Phys. A* 103 (2011) 257–261.
- [125] R. Osellame, G. Cerullo, R. Ramponi (Eds.), *Femtosecond Laser Micromachining*, Springer, New York, 2012.
- [126] C. S. R. Natala, A. Ajami, A. A. Ionin, S. I. Kudryashov, S. V. Makarov, T. Ganz, A. Assion, W. Husinsky, Experimental study of fs-laser induced sub-100-nm periodic surface structures on titanium, *Optics Express* 23 (2015) 5915–5929.
- [127] R. R. Thomson, T. A. Birks, S. G. Leon-Saval, A. K. Kar, J. Bland-Hawthorn, Ultrafast laser inscription of an integrated photonic pattern, *Optics Express* 19 (2011) 5698–5705.
- [128] R. W. Hockney, Computer experiment of anomalous diffusion, *Phys. Fluids* 9 (1966) 1826–1835.
- [129] T. Döppner, J. Müller, a. Przystawik, S. Göde, J. Tiggesbäumker, K.-H. Meiwes-Broer, C. Varin, L. Ramunno, T. Brabec, T. Fennel, Steplike Intensity Threshold Behavior of Extreme Ionization in Laser-Driven Xenon Clusters, *Phy. Rev. Lett.* 105 (2010) 2–5.
- [130] M. Bishop, M. H. Kalos, H. L. Frisch, Molecular dynamics of polymeric systems, *J. Chem. Phys.* 70 (1979) 1299–1304.
- [131] S. Karaborni, J. P. O’Connell, Molecular dynamics of hydrocarbon chains, *J. Chem. Phys.* 92 (1990) 6190–6194.
- [132] Y. D. Yan, J. J. Zhang, T. Sun, W. D. Fei, Y. C. Lang, S. Dong, Nanobending of nanowires: a molecular dynamics study, *Appl. Phys. Lett.* 93 (2008) 241901.
- [133] Y. Takato, M. E. Benson, S. Sen, Rich collision dynamics of soft and sticky nanoparticles: numerical experiments, *Phys. Rev. E* 92 (2015) 032403.
- [134] J. Barnes, P. Hut, A hierarchical  $O(N \log N)$  force-calculation algorithm, *Nature* 324 (1986) 446–449.
- [135] J. M. Dawson, Particle simulation of plasmas, *Reviews of Modern Physics* 55 (1983) 403–447.
- [136] A. Bergmann, Parallelized pic code for kinetic modelling of a “classical” scrape-off layer, *Contrib. Plasma Phys.* 38 (1998) 231–235.

- [137] R. Fonseca, L. Silva, F. Tsung, V. Decyk, W. Lu, C. Ren, W. Mori, S. Deng, S. Lee, T. Katsouleas, J. Adam, Osiris: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators, in: Computational Science – ICCS 2002, Vol. 2331 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2002, pp. 342–351.
- [138] H. Bura, R. Widera, W. Hnig, G. Juckeland, A. Debus, T. Kluge, U. Schramm, T. E. Cowan, R. Sauerbrey, M. Bussman, Picongpu: A fully relativistic particle-in-cell code for a gpu cluster, *IEEE Transactions on Plasma Science* 38 (10) (2010) 2831–2839.
- [139] S. Bastrakov, R. Donchenko, A. Gonoskov, E. Efimenko, A. Malyshev, I. Meyerov, I. Surmin, Particle-in-cell plasma simulation on heterogeneous cluster systems, *Journal of Computational Science* 3 (2012) 474–479.
- [140] G. Stantchev, W. Dorland, N. Gumerov, Fast parallel particle-to-grid interpolation for plasma pic simulations, *J. Parallel Distrib. Comput.* 68 (2008) 1339–1349.
- [141] V. K. Decyk, T. V. Singh, Adaptable particle-in-cell algorithms for graphical processing units, *Computer Physics Communications* 182 (2011) 641–648.
- [142] F. Rossi, P. Londrillo, A. Sgattoni, S. Sinigardi, G. Turchetti, Towards robust algorithms for current deposition and dynamic load balancing in a gpu particle in cell code, *AIP Conference Proceedings* 1507 (2012) 184–192.
- [143] R. McGill, J. W. Tukey, W. A. Larsen, Variations of box plots, *The American Statistician* 32 (1) (1978) 12–16.
- [144] J. L. Hintze, R. D. Nelson, Violin plots: A box plot-density trace synergism, *The American Statistician* 52 (2) (1998) 181–184.
- [145] J. L. Henessey, D. A. Patterson, *Computer Architecture: A Quantitative Approach*, 5th Edition, Morgan Kaufman, Waltham, MA, USA, 2012.
- [146] G. M. Amdahl, Validity of the single processor approach to achieving large scale computing capabilities, in: *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference, AFIPS '67 (Spring)*, ACM, New York, NY, USA, 1967, pp. 483–485.
- [147] M. W. Evans, F. H. Harlow, The particle-in-cell method for hydrodynamic calculations, Tech. rep., Los Alamos Scientific Laboratory (1957).
- [148] G. Bonnaud, C. Reisse, Particle code study of the influence of non-monochromaticity of laser light on stimulated raman scattering in laser-irradiated plasmas, *Nuclear Fusion* 26 (1986) 633–646.
- [149] F. Quéré, C. Thauray, H. George, J. P. Geindre, E. Lefebvre, G. Bonnaud, P. Monot, P. Martin, High-order harmonic generation using plasma mirrors, *Plasma Phys. Control. Fusion* 50 (2008) 124007.
- [150] Z.-M. Sheng, K. Mima, J. Zhang, H. Sanuki, Emission of electromagnetic pulses from laser wakefields through linear mode conversion, *Phys. Rev. Lett.* 94 (2005) 095003.
- [151] G. Bart, C. Peltz, N. Bigaouette, T. Fennel, T. Brabec, C. Varin, Massively parallel microscopic particle-in-cell, *Comput. Phys. Commun.* 219 (2017) 269–285.

- [152] H. Vincenti, S. Monchocé, S. Kahaly, M. Bonnaud, P. Martin, F. Quéré, Optical properties of relativistic plasma mirrors, *Nature Communications* 5 (2014) 1–9.
- [153] R. Lichters, J. Meyer-ter Vehn, A. Pukhov, Short-pulse laser harmonics from oscillating plasma surfaces driven at relativistic intensity, *Phys. Plasmas* 3 (1996) 3425–3437.
- [154] N. Zhavoronkov, Y. Gritsai, M. Bargheer, M. Woerner, T. Elsaesser, F. Zamponi, I. Uschmann, E. Förster, Microfocus Cu  $K_\alpha$  source for femtosecond x-ray science, *Opt. Lett.* 30 (2005) 1737–1739.
- [155] B. J. Siwick, J. R. Dwyer, R. E. Jordan, M. R. J. D., Ultrafast electron optics: propagation dynamics of femtosecond electron packets, *J. Appl. Phys.* 92 (2002) 1643–1648.
- [156] B. J. Siwick, J. R. Dwyer, R. E. Jordan, M. R. J. D., An atomic-level view of melting using femtosecond electron diffraction, *Science* 302 (2003) 1382–1385.
- [157] H. Niikura, F. Légaré, R. Hasbani, A. D. Bandrauk, M. Y. Ivanov, D. M. Villeneuve, P. B. Corkum, Sub-laser-cycle electron pulses for probing molecular dynamic, *Nature* 417 (2005) 917–922.
- [158] C. D. Lin, A.-T. Le, Z. Chen, T. Morishita, R. Lucchese, Strong-field rescattering physics – self-imaging of a molecule by its own electrons, *J. Phys. B.* 43 (2010) 122001.
- [159] R. W. Boyd, *Nonlinear Optics*, Academic Press, New York, 2003.
- [160] M. Aeschlimann, C. A. Schmuttenmaer, H. E. Elsayed-Ali, R. J. D. Miller, J. Cao, Y. Gao, D. A. Mantell, Observation of surface enhanced multiphoton photoemission from metal surfaces in the short pulse limit, *J. Chem. Phys.* 102 (1995) 8606–8613.
- [161] R. Bormann, M. Gulde, A. Weismann, S. V. Yalunin, C. Ropers, Tip-enhanced strong-field photoemission, *Phys. Rev. Lett.* 105 (2010) 147601.
- [162] C. Ropers, D. R. Solli, C. P. Schulz, C. Lienau, T. Elsaesser, Localized multiphoton emission of femtosecond electron pulses from metal nanotips, *Phys. Rev. Lett.* 98 (2007) 043907.
- [163] S. A. Hilbert, A. Neukirsch, C. J. G. J. Uiterwaal, H. Batelaan, Exploring temporal and rate limits of laser-induced electron emission, *J. Phys. B.* 42 (2009) 141001.
- [164] H. S. Park, J. S. Baskin, O.-H. Kwon, A. H. Zewail, Atomic-scale imaging in real and energy space developed in ultrafast electron microscopy, *Nano Letters* 7 (2007) 2545–2551.
- [165] T. Higuchi, L. Maisenbacher, A. Liehl, P. Dombi, P. Hommelhoff, A nanoscale vacuum-tube diode triggered by few-cycle laser pulses, *Appl. Phys. Lett.* 106 (2015) 051109.
- [166] M. Krüger, M. Schenk, M. Förster, P. Hommelhoff, Attosecond physics in photoemission from a metal nanotip, *J. Phys. B* 45 (2012) 074006.
- [167] H. R. Reiss, Limits on tunneling theories of strong-field ionization, *Phys. Rev. Lett.* 101 (2008) 043002.
- [168] S. T. Purcell, V. T. Bin, N. Garcia, M. E. Lin, R. P. Andres, R. Reifenberger, Field emission from narrow bands above the fermi level of nanometer-scale objects, *Phys. Rev. B* 49 (1994) 17259–17263.
- [169] M. Krüger, M. Schenk, P. Hommelhoff, G. Wachter, C. Lemell, J. Burgdörfer, Interaction of ultrashort laser pulses with metal nanotips: a model system for strong-field phenomena, *New Journal of Physics* 14 (2012) 085019.

- [170] T. Shaaran, M. F. Ciappina, R. Guichard, J. A. Pérez-Hernández, L. Roso, M. Arnold, T. Siegel, A. Zair, M. Lewenstein, High-order-harmonic generation by enhanced plasmonic near-fields in metal nanoparticles, *Phys. Rev. A* 87 (2013) 041402(R).