Novel approaches to theoretical and numerical calculations of experimental observables in intense laser-atom interactions



Iustin Ouatu

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Acknowledgements

It has now been more than 14 years in which a great part of my endeavours, trials, attempts, expectations, sources of confusion or moments of serendipity, successes (or failures), and generally searches for understanding have been related to physics, mathematics or informatics. This journey has shaped my character and has culminated in the findings presented in this thesis. Overcoming the challenges associated with such a travel through time (and space) would have been much harder, and, to be frank, impossible to finish (and even start), without the persons who have been next to me, supporting me and providing unconditional help whenever I needed it. I am who I am due to them, the education they provided to me, and all the beautiful, insightful and unique interactions I had had with them over the years. I am grateful that I was lucky to have you in my life and that I was gifted with your company for periods ranging from my whole life to a short email-exchange.

First and the most, I would like to say thank you to my parents, who have always gone way and beyond to be on my side. Thank you for always believing in me, even when all the signs were showing that the light at the end of the dark tunnel might not appear at all. Thank you for motivating me to go the extra-mile and not resign, irrespective to what that has taken from you. Thank you for making this happen and being next to me when new Universes opened to me, be it in the sixth-grade when I learned about $G = mg$, or to when I firstly saw the rotation of the elliptical polarization and the lobe. We have come a long way since the beginnings ...

All these efforts finalize now, when I am part of Professor Peter Norreys' research team. Thank you, Peter, for accepting me as part of your group and allowing me to do my research under your supervision. All this thesis wouldn't have been possible without your never-ending trust in my abilities to carry out the work leading to the results presented in the chapters below. I would like to thank Ben from our group for helping me to get started at the beginning of my PhD, but also to Qingsong, Rusko, Sunny and Lily for all our insightful and fun interactions over the time.

I will forever keep in my soul and mind and cherish all the interactions I have had with Professor Dmitry A. Telnov. In today's ever-increasingly highly-transactional world, spending (even trivial) amounts of one's life with someone

whom you are fully aware that they cannot help you in any way with anything for your everyday endeavours is a very rare thing on its own. I thank you for everything. I will forever be indebted to you and I am grateful to have had the opportunity to be listened by such a human-being and researcher in quantum-mechanics (which one goes first, you choose). I hope that all the good things to come back to you as powerful and bright as the centre of a fully-converged PMD heatmap.

I would like to thank to (Dr.) Dmitry (Tumakov) for all our discussions and interactions. We have come a long way since I had the "idea" to try to calculate the (non-relativistic) scattering states with RADIAL ... I am grateful and lucky to have talked to you.

Professor Mitsuko Murakami-Korobkin has had a profound and dramatic positive impact on my evolution. I thank you for everything and that you went backwards in time to remember about past things for me. Imaginary time propagation over $-i\tau$ is non-unitary and in a multi-electron setting, the orbitals have to be re-orthonormalized at each timestep: that is costly and painful, we both know that. Thank you for having taken the time and investing the effort to perform the movement throughout the history for me.

Professor Erik Lötstedt, thank you for everything. In life one many times does not succeed immediately, but even with something like that, the experience gained lasts forever and helps them in the next steps undertaken.

Going back to Pembroke, the tutorials and discussions with (Dr.) Federico (Nova) and (Professor) Alfons (Weber) have had a profound impact on my scientific progress. Thank you for being such professionals, honestly representing the best teachers a student could have asked for. The students taught by you will have surely understood the scientific method and thinking, and even if they do not continue on pure physics research, they will apply the critical, researcher-like thinking throughout their whole lives. I guess that is, from your side, the best thing a teacher could have asked for. I will always remember with pleasure about your insights on how to think (like a physicist), not (only) remember facts.

Staying in Oxford, the researchers from the Atomic & Laser Physics sub-department have very much influenced me in a highly positive manner. Choosing the project aiming to analyse X-ray Thomson Scattering (XRTS) data through statistical techniques was one of the best decision of my life and I am grateful I got accepted to work on it. Although at the beginning it was a very steep learning curve for me, I got from not knowing much to being able to start tackling research ideas and questions on my own after the internship and the Masters' project. The period spent in the Simon room made me fully adhere to the idea that I want to pursue a PhD. Thank you to (Professor) Sam (Vinko), (Dr.) Muhammad (Kasim),

Professor Gianluca Gregori and Professor Justin Wark for accommodating me around you for that year and having the patience to answer all the questions of a (clueless at that time, but motivated to evolve) physics student. It was that period when I learned a lot about how many beautiful things can be done with computers in physics.

I have met many persons in Oxford who made me realise (again and again) that it is a place populated with many ultra-high-performers with whom you can also have a good laugh about life. I remember with pleasure about all the interactions I had at the Romanian Cultural Society. Apart from the good atmosphere I always found in there, I have come across people one only sees listed on international rankings, in Hall of Fames, people who make you amazed by how smart and intelligent one can be. It was nice to get to know you and learn from you in all kinds of ways. Thank you, Grigore, Nico, Emil, George, Ioan, Alex, Sebi, Florea, Diana, Ioana, Filip, Vlad, Cristi, Horia, Titus, for all the good memories we made over the years. A special thank you goes to Costel for always encouraging me to study more about informatics, offering insights and opinions, and never telling me: "Aaa, this is unknown and hard to you, you do not know anything, you cannot do it!", but always seeing in people the potential to succeed given that they work hard. Similarly, to Tamio, thank you for all our very insightful discussions and interactions over the time. I am sure you will be (or actually are) great leaders and that the Romanian IOI team is lucky to have you on their side.

The Busteni physics-summer-camps opened up a whole Universe for me. After 9 years, I still like quantum-mechanics. That is not to say I am good at it! I can only shut up and do the calculations - but even that is hard! Thank you, (Professor) Radu (Ionicioiu) and (Dr.) Stefan (Ataman) for what has meant to be the best introduction to the topic on which I focused most of my later years after the camps and I liked the most from physics. Is it even physics? We can measure it, sure, but I think it's more magic (or has a part of this into it), or at least this is how I still find it ... Your lectures were also magical and have positively influenced many of us, setting some very good initial conditions which dictated the temporal evolution due to TDSE applied to the moving wavepacket containing various energies we collectively call life. I thank to Professor Livius Trache, Professor Mihai Petrovici and (Professor) Alex(andru) (Bercuci) for offering me such amazing opportunities which shaped my thinking and understanding of life, helping me throughout all my scientific endeavours. I remember with great pleasure the amazing atmosphere from IFIN-HH, in the DFH and DFT sub-departments, all the people, (Dr.) Andrei (Herghelegiu), (Dr.) Madalina (Tarzila), (Dr.) Claudiu (Schiaua), (Dr.) Cristian (Andrei), the beautiful

way we were treated and the fact that we learned so much in our stays there. I would like to thank to (Dr.) Cris (Doloc) for showing us that if you learn physics well, it will help you throughout your whole life.

I would like to thank to ELI-NP (Magurele) for offering help and support, and in particular to (Professor) Catalin (Ticos) for his continuous support and good-will, kindness and openness, but also enthusiasm and availability to offer insights about (life and) physics. I can say that all what you said has proven to be true.

The high-school I graduated from, Colegiul National Iasi, has formed my thinking in a very healthy way. While offering the best possible general education one could aim for, I also had the chance, and I am grateful for the fact that I was part of the Olympiads system. The preparation I received in school enabled me to pursue a physics faculty. I thank to all the teachers from school, but also those from CEX Iasi, who spent so many hours working through physics problems with me in the weekends. Similarly, I would like to thank to the researchers from UAIC Iasi who offered us insightful discussions and preparations. The great memories from the physics Olympiads will stay in my mind as beautiful moments with great people: Jean Rotaru, Sorin Galer, Cristinel Miron, and the gang: Filip, Theo, Alex, Tavi, Marian, Cristiana, Camelia, Gabi, George, Bogdan, Tudor G. T., Georgiana, Dimi, Anca, Sebi, Roan, Rares, Cristi.

Finally, I thank you, the reader, for taking the time to go through a few phrases, a chapter, or all the thesis. I hope you enjoy it, find something interesting, useful, or just have a good time!

Abstract

The interaction of laser light with atoms has attracted a lot of attention over the years due to the questions it raises from the point of view of fundamental physics, but also due to the practical applications in which it is a core part. If the laser pulse is very intense, it can induce the excitation or ionisation of one or more electrons from the target. This opens up a plethora of physics phenomena to study and analyse.

In this thesis, the ionisation of noble-gas atoms is subject to numerical and theoretical investigations. Initially motivated by the potential of this interaction-scenario to be a method to place bounds on the peak-intensity of a laser beam-pulse, the ionisation of atomic xenon illuminated by extreme such pulses of light is firstly analysed by means of state-of-the-art field-ionisation-rates formulae, with care being taken that the simulation-setup respects the stringent experimental constraints which are found in a real target chamber. Then considerations about the time-dependent-Schrödinger-equation (TDSE) are presented, which showcase both the theory and the implementation of a fully-distributed numerical solver for very accurate studies on the laser-atom interaction. This has been used to carry out numerical simulations: specifically, the differences between the dipole-approximated and the beyond-dipole-approximation quantum-dynamics of ionisation are analysed for a hydrogenic-atom target. Then, the multi-electron nature of argon is taken into consideration from the point of view of a mean-field approach: the time-dependent-density-functional-theory (TDDFT). In the Kohn-Sham formalism, this reduces the problem to propagating in time nine coupled one-electron TDSE's, making it amendable to numerical studies performed in reasonable wall-clock times. Finally, conclusions are drawn and possible further work arising from the ideas discussed is proposed and analysed.

Contents

List of Figures **x**

List of Abbreviations **xx**

1 Introduction **1**
1.1 Historic background and introductory theory 1
1.2 Particle-in-Cell plasma-dynamics simulation codes 9
1.3 Radial Schrödinger Equation 15
1.4 The dipole-approximated laser-fields 17
1.5 Motivation 20
1.6 Conventional methods to diagnose the peak intensity 21
1.7 Structure of thesis 23
1.8 Contributions of the author 24

2 Inferring the peak intensity from the ionisation states of an ultra-rarefied noble-gas shot by an ultra-intense laser beam-pulse **25**
2.1 An atomic diagnostic to measure the laser peak intensity 27
2.1.1 1D PIC simulations results 30
2.1.2 Number of ions originating from the focus of the beam . . 32
2.1.3 Fit to the number of ions from the focus 35
2.1.4 2D PIC simulations and confirmation of the correctness of the reduced-dimensionality PIC simulation behaviour . . 44
2.1.5 Experimental considerations of the proposed diagnostic of the laser peak intensity 46
2.2 Conclusion and outlook 51

3 Atomic ionisation modelled through accurate numerical solutions of the time-dependent Schrödinger equation **52**
3.1 Introduction 53
3.2 Review 54
3.3 Theoretical considerations 57
3.3.1 Spectrum of the field-free Hamiltonian 58
3.3.2 Time-propagation 60

3.3.3 Photoelectron-momentum-distribution at the end of the simulation . . . 63
3.4 Implementation . . . 66
3.4.1 The "serial" workflow: parallelism exploited only inside one NUMA-region . . . 66
3.4.2 MPI - OpenMP paralellisation strategy . . . 74
3.5 Benchmark calculations . . . 86
3.5.1 Three-dimensional problem-setting . . . 86
3.6 Results and discussion: the parallel-scaling and performance of the solver . . . 91
3.6.1 The improvement in the timings per one timestep as more computational resources are used . . . 91
3.6.2 Serial-performance and insights . . . 98
3.6.3 Workload-distribution across the CPU-cores . . . 104
3.7 Conclusion and outlook . . . 106

4 Beyond-dipole-approximation electron quantum-dynamics investigations as performed using a length-gauge Hamiltonian correct up to first order in $1/c$ 111
4.1 Introduction . . . 112
4.2 The length gauge Hamiltonian, including non-dipole corrections up to first order in $1/c$ . . . 115
4.3 Results and discussion . . . 120
4.3.1 Ionisation probabilities at the end of the pulse . . . 122
4.3.2 Angle-integrated dipole versus non-dipole photoelectron-momentum-distributions . . . 126
4.3.3 Angle and momentum-resolved two-dimensional $p_z - p_y$ heatmaps . . . 129
4.4 Conclusion and outlook . . . 136

5 Mean-field approaches to atomic ionisation 138
5.1 Introduction . . . 138
5.2 Time-dependent density-functional-theory (TDDFT) and its application to strong-field atomic ionisation . . . 141
5.2.1 Ground-state density-functional-theory . . . 141
5.2.2 Time-dependence . . . 154
5.3 Hartree-Fock formalism . . . 173
5.3.1 Results, discussion and further work . . . 176
5.4 Conclusions and outlook . . . 178

6 Conclusions and outlook **180**
6.1 Conclusions 180
6.2 Possible further work 182

References **187**

List of Figures

1.1 Schematic representation of the laser-field induced ionisation regimes. Left panel: multi-photon-ionisation. Middle panel: tunnel-ionisation. Right panel: over-the-barrier ionisation. In the left panel, the solid blue line represents the core's attraction and the dashed dotted line is the electron's initial energy in its bound state. In the middle and right panels, the solid blue line is the result of the combination of the laser field at a given instant of time (solid orange line) with the nucleus' attraction. The vertical arrows in the left panel represent the photons that the electron absorbs to get promoted to the continuum. The arrow in the middle panel represents the quantum-mechanical tunneling through the lowered potential barrier. The arrow in the right panel represents a classical flight over the barrier, which is seen to be lowered by the laser field below the initial energy of the electron. 4

2.1 Ratio of the focused peak intensity of a beam pulse of central wavelength 800 nm and bandwidth of 50 nm having PFT degradation to the focused peak intensity of the STC-free beam pulse (thus called reduction), as a function of the beam's diameter for fixed degradation or as a function of the degradation's magnitude for fixed diameter. 27

2.2 Xe charge states in a 1D PIC simulation. The gas density was low enough to allow neglecting collective plasma effects and relativistic self focusing. 31

2.3 Xe representative charge states' populations at the end of 15,000 1D PIC simulations. These $c_n(I)$ curves can be used to calculate the number of ions produced in the focus. 33

2.4 Number of xenon ions produced in the focus of the Gaussian beam from Eq. (2.5). One datapoint on each curve corresponds to the number of ions arising from one, fixed, peak intensity I_m value in Eq. (2.5). Results on y-axis are numerical calculations for $N\left(Xe^{n+}\right)$ from Eq. (2.4). 35

2.5 Model evidence as a function of the polynomial order at the start of the training procedure. 41
2.6 Test set mean squared error as a function of the polynomial order at the start of the training procedure. 42
2.7 Predicted response as a function of the input for the model for which I listed the values of $\vec{\alpha}^*$ and β^*. The two-sigma deviation is highlighted in blue and represents the 95 % uncertainty in the predictions, as from Eq. (2.16). 44
2.8 (a): Xenon representative charge states' populations at the end of 1,000 2D gaussian beam PIC simulations. (b): squared electric-field at one timestep, across the 2D PIC simulation box, for a simulation with $a_0 = 707.19$. (c): Results of the integration of the 1D charge states across the 2D PIC simulation's box. 45
2.9 A `Geant4` [76] simulation visualisation: protons originating from a stochastic source (their kinetic energy was drawn at random from a Gaussian distribution centred at some user-defined input energy) are incident upon regions of parallel static electric and magnetic fields and are analysed at a "test plane" (pictured in green). . . . 47
2.10 A `Geant4` [76] simulation visualisation: protons originating from a stochastic source (their kinetic energy was drawn at random from a Gaussian distribution centred at some user-defined input energy) are incident upon regions of parallel static electric and magnetic fields: their trajectories are curved and they are deflected from their initial straight paths, based on their incident kinetic energy. 48
2.11 Parabolas drawn by the protons after their deflection by the Thomson spectrometer. Left panel: results obtained using the community-standard `Geant4` simulation framework. Right panel: results obtained using in-house developed codes, which are seen to be benchmarked against the community standard (i.e. the code developed at CERN). The differences in the actual positions across the parabolas are due to the random nature of the incident particles' kinetic energies. 49

3.1 The workflow and pre-processing pipeline of the GPS-TDSE solver. The external libraries are pictured as rounded-edges boxes and the important locations inside the solver are pictured as diamonds. The separation into preliminary calculations, temporal dynamics and the end is done through coloured-shapes. 75

3.2 Schematic diagram showing three ranks holding the radial-slices of the time-dependent wavefunction. The zero-th rank (the top row) holds two (l, m) pairs, the first rank (the middle row) holds one (l, m) pair, and the second rank (the bottom row) holds one (l, m) pair. The memory locations increment to the right and the dots signify elements which are not drawn in the figure due to lack of space. In the figure, `total_#_of_LM_pairs` = 4. 77

3.3 Pictorial representation of the starting distribution before the first `MPI_Alltoallv(...)` (denoted by blue arrows and braces) and the end distribution (denoted by orange arrows and braces). This AIM-distribution is to be obtained after the call and a necessary reshuffling have been performed. Three processes are shown ($size = 3$), $N_r = 8$ and `total_#_of_LM_pairs` = 4. 80

3.4 The memory buffers to be sent via the first `MPI_Alltoallv(...)` communication at each timestep of the solver. Top three rows: the buffers to be sent from process p_0, middle three rows: the buffers to be sent from process p_1, bottom three rows: the buffers to be sent from process p_2. For each process, the top line represents the data to be sent to p_0, the middle line represents the data to be sent to p_1 and the bottom line represents the data to be sent to p_2. It is observed that each process does not send the same amount of data to the other processes, including itself, thus the variable-communication patterns emerge. 81

3.5 Data re-distribution after the first `MPI_Alltoallv(...)` call. Top section: the movement inside p_0, middle section: the movement inside p_1, bottom section: the movement inside p_2. 82

3.6 Pictorial representation of the starting distribution before the second `MPI_Alltoallv(...)` (denoted by blue arrows and braces) and the end distribution (denoted by orange arrows and braces). This AIM-distribution is to be obtained after the call and a necessary reshuffling have been performed. Three processes are shown ($size = 3$), $N_r = 8$ and `total_#_of_LM_pairs` = 4. 83

3.7 The memory buffers to be sent via the second `MPI_Alltoallv(...)` communication at each timestep of the solver. Top three rows: the buffers to be sent from process p_0, middle three rows: the buffers to be sent from process p_1, bottom three rows: the buffers to be sent from process p_2. For each process, the top line represents the data to be sent to p_0, the middle line represents the data to be sent to p_1 and the bottom line represents the data to be sent to p_2. It is observed that each process does not send the same amount of data to the other processes, including itself, thus the variable-communication patterns emerge. 84
3.8 Data re-distribution after the second `MPI_Alltoallv(...)` call. Top section: the movement inside p_0, middle section: the movement inside p_1, bottom section: the movement inside p_2. 85
3.9 Two-dimensional heatmap showing the momenta in the laser polarisation plane p_z-p_y for an ellipticity of $\epsilon = 0.5$, at the wavelength $\lambda = 800$ nm, and peak intensity of 10^{14} W/cm^2, for a laser pulse of six optical-cycles in duration. 88
3.10 Two-dimensional heatmap showing the momenta in the laser polarisation plane p_z-p_y for circular polarisation (ellipticity $\epsilon = 1$), at the wavelength $\lambda = 400$ nm and at a peak intensity of 10^{14} W/cm^2 for a laser pulse of 2 optical-cycles in duration and a radial box-size of 240 a.u.. The well-known "attoclock" PMD-pattern is visible and the simulation results are in excellent quantitative agreement with panel a) of Fig. 4 from [98] and with panel a) of Fig. 5 from [96], all three of which match (both as the general shape and characteristics but also as the actual heatmap values) with each other. 90
3.11 Strong-scaling (`OpenMP` enabled): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The spatial discretisation was fixed at $N_r = 1255$, $nlat = 511$, $mmax = 10$. While generally following a linear trend up to a large number of processors, the start of the development of a plateau can be noticed towards a very large number of processors, when the communication is too expensive for the amount of computation performed by each processor. One `MPI`-rank was composed of 16 cores, each core launching one `OpenMP` thread, and there has been one such rank per one NUMA region of ARCHER2. The datapoints present their associated wall-clock time value, for ease of reading the results. 93

3.12 Weak-scaling (no `OpenMP`): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The initial discretisation was $N_r = 64$, $nlat = 32$, $mmax = 31$. The increasingly finer discretisation have been calculated as: $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks and one such `MPI`-rank was composed of one core. The datapoints present their associated wall-clock time value, for ease of reading the results. 96

3.13 Weak-scaling (`OpenMP` enabled): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The initial discretisation was $N_r = 64$, $nlat = 32$, $mmax = 31$. The increasingly finer discretisations were calculated as $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks. There were 16 cores per one `MPI` rank, thus 16 `OpenMP` threads per process. 99

3.14 Results from applying the roofline-model to the TDSE solver by means of `likwid`. Across the vertical axis the number of floating point operations per second is pictured, while the degree of data locality, i.e. the number of floating point operations per byte moved, is shown on the horizontal axis. It is seen that some of the orchestrated (measured) parts of the solver are sitting to the left of the ridge-point, thus are memory-bound, while the others are compute-bound. 103

3.15 Workload distribution of the field-free propagation as shown by `Score-P`. The colourmap on the right shows across the horizontal axis the processes and along the vertical axis the threads spawn by each process. The colourbar for the data of the colourmap is presented on the bottom of the figure. As any external function which uses threads in its implementation is not currently seen as multi-threaded in `Score-P`, there are no threads visible, even if the `cblas_zgemv(...)` calls ran on multiple `OpenMP` threads. 108

3.16 The parallelisation scores given by `Score-P` for the field-free propagation. As this is a fully and equally distributed part of the solver, the obtained numbers are close to perfection. The maximum score, 1.0, is not achieved as the spatial-discretisation for these examples was especially chosen as not to be evenly divisible by the number of processes. 108

3.17 Spectral to grid transformation parallelisation scores as obtained using `Score-P`. As for each radial grid-point the same operations are performed, these grid-points can all be transformed in parallel. Thus the resulting scores for this operation are rather high, even for a grid-discretisation which is not evenly divisible by the number of the processors. 109

3.18 Spectral to grid transformation workload distribution as obtained using `Score-P`. Each `MPI`-process occupies one column from the heatmap on the right. The `OpenMP`-threads spawned by each process are visible across the vertical axis from the heatmap. The heatmap shows very similar colours across all its locations, indicating a rather well-balanced workload-distribution for the spectral to grid transformation. The colourbar for the data of the heatmap is presented on the bottom of the figure. 109

3.19 The second all-to-all (variable) communication parallelisation-scores obtained using `Score-P`. While the load-balance seems to approach a good one, the communication efficiency is only a fair one. This is because the costs associated with such a massive exchange of data across the processors are not trivial. 110

3.20 The workload distribution for the second all-to-all (variable) communication obtained using `Score-P`. As this communication is (potentially highly) non-uniform as each processor sends and receives (potentially highly) different amounts of data, the colours shown on the heatmap on the right are not the same across the processors, showcasing that this operation is rather costly and not perfectly (i.e. linearly) scalable across cores. The colourbar for the data of the heatmap is presented on the bottom of the figure. . . . 110

4.1 Total ionisation probabilities at the end of dipole-approximation (DA) and beyond dipole-approximation (BDA) TDSE-simulations with a laser-pulse of frequency $\omega = 3.5$ a.u. and various peak E-field strengths, for a $\sin^2$ temporal profile of 10 optical-cycles (upper panel), 15 optical-cycles (middle panel) and 5 optical-cycles (lower-panel). 125

4.2 Photoelectron energy spectrum (PES) plots as obtained at the end of TDSE-simulations. Top panels: the pulse was of 10 optical-cycles in duration. Bottom panels: the pulse was of 15 optical-cycles in duration. Left column: the pulse had a peak field strength of 30.0 a.u. Right column: the pulse had a peak field strength of 45.0 a.u. In all panels, the dashed red curve represents the dipole-approximated (DA) Hamiltonian's results and the solid, blue curve represents the beyond-dipole-approximation (BDA) Hamiltonian's results. The first ATI peak is seen at 3.0 a.u. of energy and the oscillations inside it (due to recolissions with the core) are diminished when going BDA. 127

4.3 Two-dimensional heatmaps showing the momenta along the laser polarisation direction ($\hat{z}$) and along the laser propagation direction ($\hat{x}$). The peak laser field was 30 a.u. Top row: DA-simulations, bottom row: BDA-simulations. Left panels: 10 optical-cycles laser field duration, right panels: 15 optical-cycles laser field duration. The emergence of the characteristic lobe against the propagation direction is clearly seen in the BDA-simulations. 131

4.4 Two-dimensional heatmaps showing the momenta along the laser polarisation direction ($\hat{z}$) and along the laser propagation direction ($\hat{x}$). The peak laser-field was 45 a.u. Top row: DA-simulations, bottom row: BDA-simulations. Left panels: 10 optical cycles laser field duration, right panels: 15 optical cycles laser field duration. The emergence of the characteristic lobe against the propagation direction is clearly seen in the BDA-simulations. 132

4.5 Photoelectron energy spectrum (PES) plot as obtained at the end of a TDSE-simulations. The laser field had its central frequency $\omega = 3.5$ a.u., it was 15 optical-cycles in duration with a sin-squared temporal envelope, and the laser peak-field strength was $E_0 = 45.0$ a.u., in a box extending up to $R_b = 300.0$ a.u.. 135

5.1 Triply ionised and four-times ionised argon atom probabilities at the end of TDDFT simulations, as a function of the peak intensity of the laser field for the simulation, at a fixed central frequency ω = 3.85 a.u., for a sin-squared temporal envelope of 40 fs duration. The box size was $R_b = 50$ a.u. and there were 1000 timesteps per one optical-cycle for all simulations. 162

5.2 Triply ionised and four-times ionised argon atom probabilities at the end of TDDFT simulations, as a function of the peak intensity of the laser field for the simulation, at a fixed central frequency ω = 3.85 a.u., for a sin-squared temporal envelope of 30 fs duration. The box size was $R_b = 50$ a.u. and there were 1000 timesteps per one optical-cycle for all simulations. 163

5.3 Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number Z = 18) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 10 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. In particular, the simulation domain extended up to 100 a.u. in the radial direction and the $\cos^{0.25}$ absorber started at 90 a.u. of radial extent from the origin. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The numerical discretisation for the governing equations has been performed with the aid of the generalized-pseudospectral-grids (GPS) formalism. For all the simulations, the number of radial-points was $N_r = 150$ and the number of angular-points was $N_\theta = 63$.166

5.4 Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number Z = 18) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 20 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here. 167

5.5 Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 30 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here. 168

5.6 Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 50 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here. 171

5.7 Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 100 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here. 172

List of Abbreviations

1D, 2D, 3D . . . One-, two- or three-dimensional.

RR . . . . . Radiation Reaction.

PIC . . . . . Particle-in-Cell

QED . . . . . Quantum-Electrodynamics

BW . . . . . Breit-Wheeler

STC . . . . . Spatio-Temporal coupling in a laser beam-pulse

FROG . . . . . Frequency-resolved optical gating

SPIDER . . . . Spectral phase interferometry for direct electric-field reconstruction

OAP . . . . . Off-axis parabola

ATI . . . . . Above-threshold ionisation

MPI . . . . . Depending on the context: either multi-photon ionisation, or message-passing-interface

FWHM . . . . . Full-Width-Half-Maximum

PFT . . . . . Pulse-Front-Tilt

ADK . . . . . Ammosov, Delone and Krainov laser-field ionisation rate

PPT . . . . . Perelomov, Popov and Terent'Ev laser-field ionisation rate

BM . . . . . Bauer and Mulser ionisation rate

LDA . . . . . Local density approximation

ALDA . . . . . Adiabatic local density approximation

XC . . . . . Exchange-correlation

DFT . . . . . Density-functional-theory

TDDFT . . . . . Time-dependent density-functional-theory

KS . . . . . Kohn-Sham

TDHF . . . . . Time-dependent Hartree-Fock

BSI . . . . . Barrier-Supression-Ionisation

TDSE . . . . . Time-dependent Schrödinger equation

TISE . . . . . Time-independent Schrödinger equation

BLAS . . . . . Basic Linear Algebra Subprograms

LAPACK . . . . Linear Algebra PACKage

FFT . . . . . Fast Fourier Transform

ppc . . . . . The particles-per-cell parameter in a PIC computer simulation

RVM . . . . . The Relevance-Vector-Machine Machine-Learning model

CMA-ES . . . . The covariance matrix adaptation evolution strategy for numerical optimization

MSE . . . . . Mean-Squared-Error metric

i.i.d. . . . . . Independent and identically distributed

RCF . . . . . Radio-Chromic Film(s)

IP . . . . . Image Plate(s)

1 Introduction

Contents

1.1 Historic background and introductory theory 1
1.2 Particle-in-Cell plasma-dynamics simulation codes 9
1.3 Radial Schrödinger Equation 15
1.4 The dipole-approximated laser-fields 17
1.5 Motivation 20
1.6 Conventional methods to diagnose the peak intensity 21
1.7 Structure of thesis 23
1.8 Contributions of the author 24

1.1 Historic background and introductory theory

A lot of the current technological applications which shape the lives of many people around the world can have their origin traced back to a solid understanding of the light-matter interaction from a fundamental point of view. Historically, Albert Einstein was among the firsts to predict quantum-mechanics effects in electromagnetic radiation-atom interaction: his pioneering work on the photoelectric effect could be considered to be the first such indication. There, one photon is shown to be able to liberate an initially bound electron, which consequently

flies into the continuum. However, it is now known that as the intensity of the light is increased, multiple photons can be absorbed and this process can still ionise an electron. The later process has been named, very representatively, above-threshold-ionisation (ATI) [1]. Generally, the minimum number of photons which needs to be absorbed such that the electron is ionised is equal to (atomic units are used throughout almost all the thesis, unless otherwise stated, that is, $e = m_e = \hbar = 1/(4\pi\epsilon_0) = 1$ and $c \approx 137.036$ a.u.):

$$M = ceil\left(\frac{IP_i}{\omega_L}\right) = \left\lceil\frac{IP_i}{\omega_L}\right\rceil, \tag{1.1}$$

where IP_i is the ionisation potential of the target, ω_L is the frequency of the light-field and the ceiling function rounds up its input argument to the nearest integer number. For very highly-energetic photons, only one of them is sufficient to ionize the target, $\frac{IP_i}{\omega_L} < 1$, and there might potentially still exist energy remaining: this excess is transferred to the electron in the form of kinetic-energy. Still, even if $M = 1$, multiple photons can be absorbed and the kinetic-energy spectrum of the freed electron will show very pronounced peaks spaced one photon-energy apart. On the other extreme, $M > 1$ and, perhaps more related to many experimental campaigns, $M \gg 1$ is achieved by using very intense near-infra-red (NIR) lasers, for example (or other long-wavelength radiation-sources).

Another quantity of interest which is cited many times in the literature is the Keldysh parameter, γ_K, defined as:

$$\gamma_K = \frac{\omega_L\sqrt{2IP_i}}{E_L}. \tag{1.2}$$

Above, the new quantity, E_L, is the laser-field's peak-strength. The physical significance of γ_K is that it is an indication of the degree of adiabacity of the ionisation process. This is because it is often written as:

$$\gamma_K = 2\pi \cdot \frac{t_K}{t_L}, \tag{1.3}$$

where $t_K = \frac{\sqrt{2\,IP_i}}{E_L}$ and $t_L = \frac{2\pi}{\omega_L}$. The quantity t_K is the Keldysh time and t_L is one laser-oscillation duration. If t_K is small compared to t_L, then $\gamma_K < 1$ and

tunnel-ionisation (TI) prevails. This means that the combined effect of the laser electric-field and the ion's Coulomb attraction leads to a potential barrier through which the electron can quantum-mechanically tunnel. If the tunneling process is rather fast, it can be considered that the electron escapes from its initially-bound state out to the continuum through a more or less constant-in-time barrier and this can be considered adiabatic ionisation. A sensible requirement for an ionisation process to be considered as such is that the tunneling time is way smaller than t_L. In tunnel-ionisation, the energy of the electron does not change, but its position modifies, up until it reaches the tunnel exit, from where it starts its flight in the continuum.

If γ_K is indeed way smaller than 1.0, there might be the possibility that the barrier is lowered that much that the electron can classically fly over-the-barrier in an interaction scenario called barrier-suppression-ionisation (BSI) [2]. This happens for atomic hydrogen for a peak field-strength of $1/16 = 0.0625$ a.u., and this value increases with an increase of the parameter $\frac{I_p^2}{4Z}$, where Z is the atomic number of the target under consideration ($Z = 1$ for atomic hydrogen). Once γ_K increases towards the value of 1.0, the barrier becomes more and more time-dependent and the degree of adiabacity of the process decreases.

For $\gamma_K > 1$, the ionisation-regime has been historically called multi-photon-ionisation (MPI). Here, before being born into the continuum, the position of the electron does not change, but its energy does, in such a way that after the absorption of several photons, it starts moving into the continuum: there is no tunneling involved, and the ionisation rate Γ_{MPI} is usually accurately described by the formula $\Gamma_{MPI} \propto I^n$, where n is the number of absorbed photons and I is the luminous intensity, which is considered to be $I < 10^{14}$ W/cm^2 for a perturbative treatment to correctly describe the process and consequently the quoted rate. In MPI, the electron is not born in the continuum at the exit point of the potential barrier, but really close to the nucleus. This is rather different than what happens in the tunneling regime, where the exit point can be quite far away from the nucleus (even several a.u.).

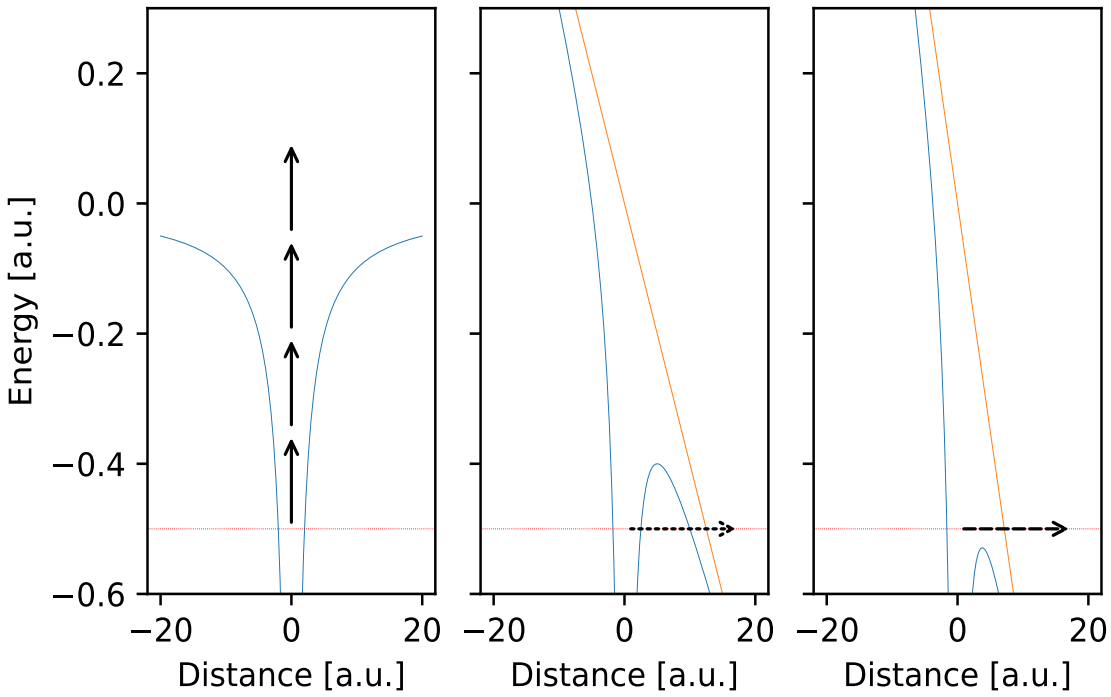


Figure 1.1: Schematic representation of the laser-field induced ionisation regimes. Left panel: multi-photon-ionisation. Middle panel: tunnel-ionisation. Right panel: over-the-barrier ionisation. In the left panel, the solid blue line represents the core's attraction and the dashed dotted line is the electron's initial energy in its bound state. In the middle and right panels, the solid blue line is the result of the combination of the laser field at a given instant of time (solid orange line) with the nucleus' attraction. The vertical arrows in the left panel represent the photons that the electron absorbs to get promoted to the continuum. The arrow in the middle panel represents the quantum-mechanical tunneling through the lowered potential barrier. The arrow in the right panel represents a classical flight over the barrier, which is seen to be lowered by the laser field below the initial energy of the electron.

A particularly interesting region from the values that γ_K can take is $\gamma_K \approx 1$. Here, the interpretation of the physics involved in the ionisation process is not very straightforward. A slight degree of non-adiabacity appears, but tunneling is also present to some extent.

Schematically, the regimes of MPI, tunnel ionisation and BSI are graphically depicted in Fig. 1.1.

On the left panel of this diagram MPI is shown: the electron is promoted through intermediate states towards the continuum. In the middle panel the diagram shows the electron tunneling through the instantaneous potential barrier created by the atomic core and the laser field and reaching the exit point, which

represents the liberation location, from which it flies into the continuum. On the right panel BSI is pictured: the laser field is more intense than in the middle panel so in combination with the core's attraction, it lowers the barrier that much that the electron escapes classically from its bound state towards the continuum. The same target, atomic hydrogen, was considered in all the panels, and the initially bound state had -0.5 a.u. as its eigenenergy (i.e. the ground state) for all of them.

After half a cycle of the laser-pulse, its electric field changes direction and the electron starts being accelerated in the opposite direction to the one in which it was flying. When it returns to the parent atomic-core, there are a few possibilities.

Firstly, it can recombine with it and emit the excess kinetic-energy as radiation, and this is called high-order harmonic generation (HHG) [3]. HHG cannot be interpreted perturbatively: this approach would predict that the harmonics spectrum decreases exponentially in a monotonically manner as their frequency increases and this is not the case, as experiments have shown. Instead, the strength of the harmonics has been found to be quite constant over an extended range of frequencies (the plateau). Then, at frequencies higher than the so-called cutoff frequency $\omega_{cutoff} = I_p + 3.17\,U_p$, where U_p is the ponderomotive potential of the electron, a drastic decrease in the harmonics spectrum appeared. The ponderomotive potential is equal to $U_p = \dfrac{I}{4\,\omega_L^2}$ [4].

Secondly, it can inelastically scatter with the core and produce what is called non-sequential double-ionisation (NSDI): a second electron is liberated to the continuum due to this.

Thirdly, if it elastically scatters (its velocity changes direction, but not magnitude), it then just flies to the continuum with a very high energy (because of a second acceleration step). In the case of backscattering, one obtains the maximum kinetic energy the electron can acquire in this third scenario and this is equal to $10.01\ U_p$ [5].

Multiple scattering events can occur: this means that more than one collision with the core happens. This is particularly true for whenever the electron is

liberated close to when the laser pulse's electric-field is in the near proximity of its peak-value.

Of course, the electron does not necessarily need to return towards the core: it can ionise directly if it has been liberated in the ascending parts of the pulse, or if the light is circularly-polarised.

HHG has been found to be useful to convert an infra-red (IR) pulse to a (soft-) X-ray pulse [6] or an attosecond pulse [7]. Generally, ATI has been the core part of imaging techniques for atomic/molecular targets, through laser-induced electron-diffraction [8].

The picture of the liberation of the electron and its recollision with the core originates from the three-step-model (or the simple man's model), due to P. Corkum [9]. This greatly simplifies the interpretation of an otherwise rather complex phenomenon, without a dramatic loss of accuracy.

Historically, many of the experimental studies related to laser-field induced ionisation could be performed due to the invention and successful realisation in practice of various techniques to generate short and ultra-short laser pulses. For example, attosecond pulses from a gaseous target can now be created due to the reduction of the driver-pulse duration down to O(10) fs and less. The invention of mode-locking in 1964 opened up the avenue for a steady reduction of the driving pulse duration over the following years. The laser-scientists made real progress back then, and it was mainly between 1980 and 1990 when light pulses of O(10) fs duration started to being routinely obtained in a laboratory setting. While nowadays triply-differential photoelectron-momentum-distributions (PMD's) can be measured using the cold target recoil-ion momentum spectroscopy (COLTRIMS) technique, many experiments from the past mainly focused on ion yields (counting the numbers of particles from each charge state), or singly differential spectra (electron energy distributions).

Starting with spectra, the first direct evidence (practical realisation) of ATI dates back to the 1979 experiment of Pierre Agostini (Nobel Prize in physics winner in 2023, distinction shared with Anne L'Huillier and Ferenc Krausz)

and implemented together with his collaborators from the Paris-Saclay research laboratories, in Gif-sur-Yvette, France [1]. There, the team used xenon as the target and a central-frequency of 1.17 eV laser-pulse, of 12 ns duration and amplified to 2 J of energy, which was frequency-doubled to investigate the six-photon ionisation of this noble-gas. The gas-pressure was very low (5.0×10^{-5} Torr) in order to be able to access the single-atom laser-interaction regime. The electrons were collected with the help of an electron-multiplier-tube (EMT), which stored integrated energy-spectra, thus a numerical differentiation of the obtained results was needed to be performed to obtain the probability to observe an electron at a given kinetic-energy.

The HHG signal, with its particular characteristics, has been clearly shown in practice in 1987 due to the experiment led by M. Ferray and A. L'Huillier, performed again in Gif-sur-Yvette [3]. They shot three noble-gas targets: xenon, krypton and argon. The laser used was an infra-red one having its central frequency equal to 1064 nm and was of 30 ps duration. It was amplified to 1 GW and the repetition rate was 10 Hz, this allowing collection of data from multiple shots in a small interval of time. The focusing optic had a focal length of 200 mm. The pressure inside the vacuum chamber was 10^{-8} Torr. With the help of different radiation detectors (each suitable for a particular wavelength-range), i.e. two (of different characteristics) gratings, a photo-multiplier and a windowless electron-multiplier, the harmonics have been recorded up to the 33^{rd} order for argon. For this particular one, the wavelength was 32.2 nm, thus in the XUV-regime. There was an initial, sudden drop in the intensity of the harmonics from the first ones to the fifth one, but then a clearly visible plateau appeared. The intensity of the harmonics fell very slowly as their order increased, up to the 27^{th} one. Then the cutoff appeared and this drastically reduced the brightness of the higher-in-order harmonics still observed in the experiment, all this being in close agreement with the theory and models known today.

Moving on to ion yields, the experimental discovery of the "knee"-structure in the probability of ionisation of a two-electron (or generally, a many-electron)

system as a function of the peak-intensity of the laser-pulse used in the interaction dates back to the experiments from 1992 and 1994 of Fittinghoff and Walker, respectively [10, 11]. While usually the single-active-electron-approximation (SAE) correctly describes ionisation of multi-electron targets illuminated by IR laser-light, data collected in the laboratory showed doubly-ionised atoms' yields orders of magnitudes larger than the theoretical predictions obtained using a purely sequential picture. While both the papers discuss and agree that the shake-off process is responsible for this apparent inconsistency, it is now established that the re-scattering arguments from the three-step model of Corkum [9] (theoretically supported by the complicated Feynman diagrams calculated by Becker and Faisal [12]) correctly explain this phenomenon.

Shake-off, one the other hand, is more related to the double ionisation produced by a single, highly-energetic photon (not the case for NIR and IR pulses). What it is about is that after the very fast release of the first electron, the resulting change in the ionic potential forces the remaining electron to relax into the eigenstates of this new potential. Some of these stationary states are in the continuum, and if the second electron "accesses" these positive energy ones, it also flies into the continuum, hand-in-hand, i.e. together, with the first one: they are more or less born at the same time. In contrast, in the recollision picture the electrons are born at radically different times.

The experiment of Walker and his collaborators [11] used helium as the target and laser pulses of O(100) fs duration, at a central frequency of 780 nm. Again, the interaction-chamber pressure was very low: 10^{-10} Torr. The pulse energy was 1.3 mJ and the repetition rate of the laser was 1 kHz. This allowed the experimentally-collected datapoints to actually be averages of as many as 60,000 laser-shots. A major direction in nowadays laser-engineering efforts and research is to increase the repetition rate of lasers, apart from making them more intense and more energetic.

Other (more recent) experiments have only looked at ionisation yields. E. A. Chowdhury continued the previous experiments of Walker and together, they

measured neon's charge states up to when it only has one electron remaining (Ne^{8+}) [13]. Interestingly (and to be expected), re-scatterings have been suppressed for the intensities used (10^{17} - 10^{18} W/cm^2) and a sequential-ionisation-view correctly fits the narrative of these efforts. Before this, even 10^{19} W/cm^2 has been used as the intensity of the driving light pulse, having a central frequency of 800 nm and of 25 fs duration. It was focused down to 3 μm and then illuminated argon and ionised it up to the Ar^{16+} charge state [14]. It was shown that the tunnel-ionisation rates accurately describe the obtained yields in a sequential fashion. A similar study was published in [15] and similar conclusions have been drawn for argon at roughly the same laser peak-intensities.

Analysing such experimental efforts mainly consists of comparing theoretical and numerical predictions with the data collected in the laboratory. Quantitatively, the interaction of a laser field with the atom can be described in terms of ionisation rates (usually the Perelomov, Popov and Terent'ev (PPT) rate-equation, discussed later) which produce the probabilities of the electron to be detached. When considering many such atoms illuminated by the beam-pulse and placed at its focus, the collective dynamics of all the resulting particles has to be taken into account if the gas density is not very low (otherwise the atoms are initially located extremely far away from each other and the interaction scenario is basically a single atom regime). This is also important if the laser field is ultra-intense and the liberated electrons start moving at large velocities, emanating light through the radiation-reaction (RR) mechanism.

1.2 Particle-in-Cell plasma-dynamics simulation codes

The fully-relativistic, self-consistent dynamics of a collisionless plasma is often solved by means of Particle-in-Cell (PIC) simulation codes. These are composed from the the close interplay between two numerical solvers and have attracted considerable attention over many years.

The first solver evolves the particles' coordinates and momenta along continuous trajectories under the influence of any electric and magnetic fields existent

throughout the simulation domain, also computing the currents arising from this relativistic movement of all the species considered in the interaction scenario.

The second solver discretises the Maxwell's equations on a real-space grid and calculates the change of the electric and magnetic fields due to the influence of the moving charges.

The equations solved are the Maxwell-Vlasov formalism (Eq. (1.4)) for the time-dependent phase-space distribution $f_s(\vec{x}, \vec{p}, t)$, depending on the spatial coordinates $\vec{x}$ and on the momentum $\vec{p} = m_s \gamma \vec{v}$, with $\vec{v}$ being the velocity and γ being the relativistic Lorentz factor and where subscript s denotes the species, together with the four Maxwell's equations (Eq. (1.5) - Eq. (1.8)):

$$\frac{\partial f_s(\vec{x}, \vec{p}, t)}{\partial t} + \frac{\vec{p}}{m_s \gamma} \cdot \nabla_{\vec{x}} f_s(\vec{x}, \vec{p}, t) + \vec{F}_L(\vec{x}, \vec{v}, t) \cdot \nabla_{\vec{p}} f_s(\vec{x}, \vec{p}, t) = 0, \tag{1.4}$$

where the Lorentz force $\vec{F}_L(\vec{x}, \vec{v}, t)$ is $\vec{F}_L(\vec{x}, \vec{v}, t) = q_s \left(\vec{E}(\vec{x}, t) + \vec{v} \times \vec{B}(\vec{x}, t) \right)$ and

$$\nabla \cdot \vec{B}(\vec{x}, t) = 0, \tag{1.5}$$

$$\nabla \cdot \vec{E}(\vec{x}, t) = \frac{\rho(\vec{x}, t)}{\epsilon_0}, \tag{1.6}$$

$$\nabla \times \vec{B}(\vec{x}, t) = \mu_0 \vec{J}(\vec{x}, t) + \mu_0 \epsilon_0 \frac{\partial \vec{E}(\vec{x}, t)}{\partial t}, \tag{1.7}$$

$$\nabla \times \vec{E}(\vec{x}, t) = -\frac{\partial \vec{B}(\vec{x}, t)}{\partial t}, \tag{1.8}$$

where the charge density $\rho(\vec{x}, t)$ and the current density $\vec{J}(\vec{x}, t)$ are obtained by taking moments with respect to the velocity (momentum) of the distribution function $f_s(\vec{x}, \vec{p}, t)$, i.e. $\rho(\vec{x}, t) = \sum_s q_s \iiint d^3p \, f_s(\vec{x}, \vec{p}, t)$ and
$\vec{J}(\vec{x}, t) = \sum_s q_s \iiint d^3p \, \vec{v} \, f_s(\vec{x}, \vec{p}, t)$.

The distribution function $f_s(\vec{x}, \vec{p}, t)$ is expanded in a sum of macro-particles, much like the splitting of a time-dependent wavefunction into a set of basis-functions. The number of macro-particles is directly related to the PIC-simulation-metric called particles-per-cell (**ppc**), which can be interpreted as a measure of how well the (7-dimensional) distribution function is sampled by the macro-particles. These could be understood from a statistical point of view - less

sampling, more noise and more inaccuracies on the computed particles' dynamics and fields' evolutions. With a higher number of ppc, the distribution-function f_s is better represented by the expansion onto the macro-particles, and more accurate results are obtained - still, at (many times significantly and prohibitively) higher computational costs. It is usually the case that some form of convergence is achieved with a given number of ppc, and this is checked by running the same simulation, but with a higher number of ppc, analysing the change of the quantities of interest between the two simulations. While formally this sits in the realm of the theory of converging sequences and the rate of change of any quantity has to decrease in magnitude as the number of ppc increases, a certain tolerance level is manually set and simulations show quantitatively-correct results once that tolerance level is achieved.

These basis functions are dependent on both the spatial coordinate and the momentum (velocity). The dependence on $\vec{x}$ is through position-localised shape-functions $S(\vec{x} - \vec{x}_i(t))$, where i denotes the macro-particle number, and the S function is the same for all the macro-particles $i \in [1, 2, ..., N_s]$ and is a **B-spline** [16]. It is to be mentioned that the family of triangles (a type of B-splines) can easily be used in practice in PIC-simulations. The dependence on the momentum $\vec{p}$ is through momentum-Dirac-delta functions $\delta(\vec{p} - \vec{p}_i(t))$:

$$f_s(\vec{x}, \vec{p}, t) = \sum_{i=1}^{N_{macro}} \frac{w_i}{V_c} S(\vec{x} - \vec{x}_i(t)) \, \delta(\vec{p} - \vec{p}_i(t)), \tag{1.9}$$

where w_i is the weight carried by this macro-particle and V_c is a real-space-cell's volume. Inserting this discretised(-onto-macro-particles) distribution function onto the Maxwell-Vlasov formalism from Eq. (1.4), one obtains that what is to be (numerically-) solved becomes a set of two coupled relativistic equations-of-motions for the macro-particle i of species s, along continuous trajectories ($\vec{x}_i$ for particle i):

$$\frac{d\vec{x}_i}{dt} = \frac{\vec{u}_i}{\gamma_i}, \tag{1.10}$$

$$\frac{d\vec{u}_i}{dt} = \frac{q_s}{m_s} \left(\vec{E}_i + \frac{\vec{u}_i}{\gamma_i} \times \vec{B}_i \right), \tag{1.11}$$

where $\vec{u}_i = \vec{p}_i / m_s$ and the subscripts on the E and B-fields denote their values at the i-th particle's location, obtained via an interpolation approach (fully described by the shape-function S):

$$\vec{E}_i(t) = \iiint d\vec{x}\ S(\vec{x} - \vec{x}_i)\ \vec{E}(\vec{x}, t), \tag{1.12}$$

$$\vec{B}_i(t) = \iiint d\vec{x}\ S(\vec{x} - \vec{x}_i)\ \vec{B}(\vec{x}, t). \tag{1.13}$$

The core approach for numerically solving the Maxwell's equations is the finite-differences in the time-domain method (FDTD) on a staggered, Yee grid [17]. That is equivalent to saying that the spatial derivatives (along the real-space) of the fields appearing in the Maxwell's equations (the discretised curl operator) are approximated by centred, second-order accurate, difference formulae. Thus, in this way the concept of a cell is introduced, and in one dimension it represents the linear-space between two consecutive grid points used in the discretisation of the fields, this concept being easily extendable to multiple dimensions. In this Yee setting, in a three-dimensional space, the three components of $\vec{E}(\vec{x}, t)$ and the three components of $\vec{B}(\vec{x}, t)$ are pinned to grid points in such a way that any E-field component is surrounded by four "circulating" B-field components The same centring of derivatives is sought-after also for the temporal changes: the E-fields are pinned to integer timesteps $l \in \mathbb{N}$ and the B-fields are pinned to half-integer timesteps $l + 0.5$.

The coupling with the particles' evolution is done through a leapfrog scheme, in which the particles' positions are pinned to integer timesteps l and their velocities to half-integer timesteps $l - 0.5$, where the same l has been used for the time-centering of the time-derivatives in Maxwell's equations solutions.

The spatial and temporal resolutions which are chosen always have to satisfy the Courant–Friedrichs–Lewy (CFL) condition for a correct time-evolution of the quantities of interest. This requirement can be phrased as the necessity that any information has to travel, in each timestep, strictly less than the distance between

two consecutive, adjacent, cells. Mathematically, the timestep of the simulation Δt has to be chosen (assuming a fixed spatial discretisation) as:

$$\Delta t < \left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right)^{-1/2}, \tag{1.14}$$

where Δx is the cell size along the x-coordinate, and similarly for other two Cartesian axes. For geometries which are not fully three-dimensional, the above inequality also has to hold, but only with two (for a two-dimensional Cartesian geometry) or one term (for a one-dimensional Cartesian geometry) inside the parentheses.

In a time-dependent PIC-simulation, the same iterative procedure composed of four actions happens at each timestep: **i)** interpolate the fields from the grid points to the particles' positions; **ii)** push the particles based on the forces they are acted upon by; **iii)** project the current and charge densities resulting from the particles' movement; **iv)** solve Maxwell's equations for the new values of the E and B-fields.

Diagnostics-data can be written to disk during any of the timesteps of a given simulation. Common examples are the values of the E and B-fields, the particle momenta, or the charge density across the simulation domain.

As any computer has finite memory, these physics-interaction scenarios are analysed by means of numerical experiments conducted in a simulation domain (many times called the box) which is not infinite. This real domain can have one dimension (x-coordinate), two dimensions (x and y coordinates) and can obviously spawn the whole real volume (x, y and z coordinates). Boundary conditions have to be in place to accommodate for the non-infinite slices from the real space in which the interaction takes place on the computer. Both the macro-particles and the fields have these. At the edges of the box the macro-particles can either hit the wall and reflect (having their momenta changed accordingly), loop-around to the other side of the box (for example, from $y = y_{max}$ to $y = y_{min}$), hit the wall and be stopped (momenta become identically 0), or be thermalised (so they are re-injected into the interior of the box with momenta sampled from

a Maxwellian distribution). The electromagnetic fields are either absorbed or injected at the walls, but they can also be reflected (just like the particles) or cycle-through to the other extreme of the coordinate's values.

Ionisation in a PIC-code is treated ad-hoc. This is to be seen from the fact that the core solvers and the four actions described above do not include it in any way. However, ionised-electrons are charged and, when accelerated by the fields, contribute to the plasma dynamics. It is usually the case that in PIC-codes, at each timestep, a time-resolved ionisation rate is calculated using the Ammosov, Delone, and Krainov (ADK) and the Perelomov, Popov, and Terent'ev (PPT) formulae. Then, based on a Monte-Carlo procedure, it is seen whether the atom (or ion) is ionised. The general scheme is that the single-ionisation probability P_1 is calculated as $1.0 - \exp(-\Gamma_{ADK}(t) \cdot \Delta t)$, then compared with a random real number (in between 0.0 and 1.0, both ends included) generated by a pseudo-random generator at each timestep ($\Gamma_{ADK}(t)$ is the ionisation rate at time t). Depending on the result of this comparison between two numbers (the random one and the one representing the probability), the atom (or ion) has one electron ejected (assuming this is possible, i.e. there is at least one bound electron to be ejected), or not.

The SMILEI PIC-code [18] in particular extends this scheme in a non-trivial way. The end result is that it allows for multiple (sequential) ionisations to happen during one timestep of the simulation [19]. This has huge advantages for when the laser-field is ultra-intense, as the user is not forced to artificially reduce the timestep in order to fully resolve the ionisation mechanisms.

The velocity of the newly-created electron is set artificially, by the simulation-code, to be equal to the parent ion's velocity, and this violation of the conservation of momentum is not corrected as it very small and does not change the dynamics in any observable way.

1.3 Radial Schrödinger Equation

The most accurate quantitative treatment of excitation and ionisation is gauging about the phenomena at their most fundamental level: solving the time-dependent-Schrödinger-equation (TDSE) for the laser-atom interaction. As initially the electron is in a stationary state of the atomic potential, the time-independent-Schrödinger-equation (TISE) is usually firstly considered in spherical coordinates to fully exploit and benefit from the symmetry of the target.

The TISE is firstly cast in the following form:

$$\mathcal{H}_0 \Psi(\vec{r}) = E\Psi(\vec{r}), \tag{1.15}$$

with

$$\mathcal{H}_0 = -\frac{\nabla^2}{2} - \frac{Z}{r}. \tag{1.16}$$

As mentioned, this field-free Hamiltonian has spherical symmetry. Thus, by expanding the Laplacian in spherical coordinates and expressing the total wavefunction $\Psi(r, \theta, \phi)$ as a product of three functions, each dependent only on one coordinate, it is obtained that the reduced radial wavefunction $y(r) = rR(r)$ respects a second order differential equation in the radial-coordiante r, and that the two functions depending on the angular variables (θ and φ), combined together, are the spherical-harmonics $Y_l^m(\theta, \varphi)$:

$$\nabla^2 \leftrightarrow \frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial}{\partial r}\right) + \frac{1}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial}{\partial\theta}\right) + \frac{1}{r^2\sin^2\theta}\frac{\partial^2}{\partial\varphi^2}, \tag{1.17}$$

$$\Psi(r, \theta, \phi) = R(r)\,P(\theta)\,Q(\varphi) = R(r)\,Y(\theta, \phi) = \frac{y(r)}{r}Y(\theta, \varphi), \tag{1.18}$$

$$-\frac{1}{2}\frac{Y(\theta,\varphi)}{r^2}\frac{d}{dr}\left(r^2\frac{dR(r)}{dr}\right) - \frac{1}{2}\frac{R(r)}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial Y(\theta,\varphi)}{\partial\theta}\right) - \tag{1.19}$$

$$\frac{1}{2}\frac{R(r)}{r^2\sin^2\theta}\frac{\partial^2 Y(\theta,\varphi)}{\partial\varphi^2} - \frac{Z}{r}R(r)Y(\theta,\varphi) = ER(r)Y(\theta,\varphi). \tag{1.20}$$

By the classical multiplication of the whole of Eq. (1.19) by $\frac{r^2}{R(r)Y(\theta,\phi)}$, one obtains two sides, one depending on the radial coordinate only, and the other one depending on the angles only:

$$-\frac{1}{2}\frac{1}{R(r)}\frac{d}{dr}\left(r^2\frac{dR(r)}{dr}\right)-Zr-Er^2=\frac{1}{2}\frac{1}{\sin\theta}\frac{1}{Y(\theta,\varphi)}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial Y(\theta,\varphi)}{\partial\theta}\right)+\frac{1}{2}\frac{1}{\sin^2\theta}\frac{1}{Y(\theta,\varphi)}\frac{\partial^2 Y(\theta,\varphi)}{\partial\varphi^2}.$$

The two sides have to be both equal to the same constant factor, which, by knowing that the right hand side resembles the angular momentum operator, is set to be $-l(l+1)/2$, where $l(l+1)$ for a given l is an eigenvalue of the angular momentum operator.

The part dependent on the θ and φ angles has solutions which are denoted as $Y_l^m(\theta,\varphi)$ and called the spherical-harmonics, being a complete set with which one can expand a well-behaved function defined on the unit-sphere.

The radially-dependent part becomes:

$$\frac{-1}{2R(r)}\frac{d}{dr}\left(r^2\frac{dR(r)}{dr}\right)-Zr-Er^2=-\frac{l(l+1)}{2}. \tag{1.21}$$

Using the reduced radial wavefunction $y(r)=rR(r)$ the below is obtained:

$$\frac{d}{dr}\left(r^2\frac{d}{dr}\left(\frac{y(r)}{r}\right)\right)=\frac{d}{dr}\left(r\frac{dy(r)}{dr}-y(r)\right)=r\frac{d^2y(r)}{dr}, \tag{1.22}$$

and it is used to convert Eq. (1.21) to:

$$\left(-\frac{1}{2}\frac{r}{y(r)}\right)r\frac{d^2y(r)}{dr^2}-Zr-Er^2=\frac{-l(l+1)}{2}. \tag{1.23}$$

Multiplying by $\frac{y(r)}{r^2}$, the working-equation, i.e. the time-independent-Schrödinger equation, is casted in the eigenvalue-form below:

$$-\frac{1}{2}\frac{d^2y(r)}{dr^2}+\left(\frac{l(l+1)}{2r^2}-\frac{Z}{r}\right)y(r)=Ey(r). \tag{1.24}$$

This produces eigenvectors (y functions) and eigenvalues (E values) and has been discretised in this thesis using the generalized-pseudospectral-grids

(GPS) method. Of course, the above equation has solutions which are known analytically: they are orthogonal and normalised on the real, three-dimensional, infinite space. But these are not orthogonal on the numerical grid used to evolve the three-dimensional-time-dependent wavefunction (which does not have an analytically-known temporal evolution). Thus the propagation is not unitary if these analytical solutions are used in the field-free Hamiltonians' propagators from the split-operator in the energy-representation time-evolution scheme. All this will be discussed at length in the thesis.

Unlike the TISE, the TDSE involves the electrical field of the laser-pulse with all its temporal dependence, i.e. the (non-trivial) physical-perturbation on the quantum-system under consideration. I will now discuss how this is usually included in the mathematical formalism.

1.4 The dipole-approximated laser-fields

The two Maxwell's equations from Eq. (1.5) and Eq. (1.8) lead to the existence and possibility to formally define the A-vector potential $\vec{A}(\vec{x},t) = \vec{A}(\vec{r},t)$ with $\vec{r} = \vec{x}$ the position vector by the following equation:

$$\nabla \times \vec{A}(\vec{r},t) = \vec{B}(\vec{r},t). \tag{1.25}$$

This in turn means that:

$$\nabla \times \left(\vec{E}\left(\vec{r},t\right) + \frac{\partial \vec{A}\left(\vec{r},t\right)}{\partial t} \right) = 0, \tag{1.26}$$

so the E-field can be written as (because the curl of the vector field in the parentheses is identically 0):

$$\vec{E}\left(\vec{r},t\right) = -\frac{\partial \vec{A}\left(\vec{r},t\right)}{\partial t} - \nabla\phi(\vec{r},t), \tag{1.27}$$

where it has been introduced the scalar quantity $\phi\left(\vec{r},t\right)$.

The physical-fields are the $\vec{E}$ and $\vec{B}$-fields. The mathematical constructs $\vec{A}$ and ϕ do not move the particles in any way. It is thus the case that with any change in these objects which leave the physical fields the same, the physical

observables remain the same. This leads to the concept of a gauge, and switching between these means exploiting the freedom to re-cast the $\vec{A}$ and ϕ quantities in any way in which the physical-fields are the same. Different gauges have different advantages and disadvantages when numerically propagating in time the Schrödinger equation. In the limit of infinite time and space resolution in the numerical propagation, the results obtained with any of all the possible and mathematically-correct gauges are identically the same. The transformations of $\vec{A}$ and ϕ are:

$$\vec{A}_{new}(\vec{r},t) = \vec{A}_{old}(\vec{r},t) + \nabla K(\vec{r},t), \tag{1.28}$$

$$\phi_{new}(\vec{r},t) = \phi_{old}(\vec{r},t) - \frac{\partial K(\vec{r},t)}{\partial t}, \tag{1.29}$$

with freedom on the scalar function $K(\vec{r},t)$.

The Maxwell's equations in a vacuum without generating sources admit plane-wave solutions for the E and B-fields:

$$\vec{E}(\vec{r},t) = \hat{n}\,E_0 \cos\left(\vec{k}\cdot\vec{r} - \omega t + \delta\right), \tag{1.30}$$

$$\vec{B}(\vec{r},t) = \hat{k} \times \vec{E}(\vec{r},t)\,\frac{1}{c}, \tag{1.31}$$

where $\hat{n}$ is the polarisation direction and $\hat{k}$ is the propagation direction, with $\hat{n}\cdot\hat{k} = 0$. A temporally-short laser-pulse (so not of infinite temporal extent as the above plane-waves, which are also infinite along the three spatial directions x, y, z) is often described by an envelope function $f(t)$, which is non-zero only in a finite time-interval:

$$\vec{E}(\vec{r},t) = f(t)\,\vec{E}_{peak} \cos\left(\vec{k}_0\cdot\vec{r} - \omega_0 t + \delta_0\right), \tag{1.32}$$

where now the centre frequency is ω_0 with its corresponding wavevector $\vec{k}_0$ and where $\vec{E}_{peak}$ includes the polarisation direction. The continuous-wave is obtained for $f(t) = 1.0, \forall\, t \in [0,\infty)\, s$.

For a $\sin^2$ temporal envelope, which has been used in many of the numerical-simulations from this thesis, $f(t)$ is written as:

$$f(t)_{sin-sq} = \sin^2\left(\frac{\pi t}{T}\right), \tag{1.33}$$

where T is the duration of the laser pulse ($T = N \cdot 2\pi/\omega_0$, with N the number of optical-cycles in the laser pulse).

There are many ways to reach the time-dependent Schrödinger equation for one electron in the laser-field, so the starting point of the discussion below is its form:

$$i\frac{\partial \Psi(\vec{r},t)}{\partial t} = \left(\frac{\left(\underline{\hat{p}} + \vec{A}(\vec{r},t)\right)^2}{2} - \phi(\vec{r},t) \right) \Psi(\vec{r},t), \tag{1.34}$$

with the operator $\underline{\hat{p}} = -i\nabla$. The dipole approximation implies that $\vec{A}(\vec{r},t)$ is a function of time only $\vec{A}(t)$. This means that the products $\underline{\hat{p}} \cdot \vec{A}(t)$ and $\vec{A}(t) \cdot \underline{\hat{p}}$ are actually the same when acting on $\Psi(\vec{r},t)$ (however for any spatial dependence of the vector-potential this would not be true, unless the Coulomb-gauge condition is in place, i.e. $\nabla \cdot \vec{A}(\vec{r},t) = 0$). The TDSE then reads, for a general $\phi(\vec{r},t)$:

$$i\frac{\partial \Psi(\vec{r},t)}{\partial t} = \left(\frac{\underline{\hat{p}}^2}{2} + \vec{A}(t) \cdot \underline{\hat{p}} + \frac{\vec{A}(t)^2}{2} - \phi(\vec{r},t) \right) \Psi(\vec{r},t). \tag{1.35}$$

In the absence of any generating charges, $\phi_{old}(\vec{r},t) = 0$, where the subscript has been re-introduced to start referring below to the actual transformations to the length-gauge and the velocity-gauge.

The length-gauge (used throughout the thesis) is obtained by taking $K(\vec{r},t) = -\vec{A}(t) \cdot \vec{r}$. This way, $\vec{A}_{new}(t) = 0$ and $\phi_{new}(\vec{r},t) = +\frac{\partial \vec{A}(t)}{\partial t} \cdot \vec{r} = -\vec{E}(t) \cdot \vec{r}$, thus the total Hamiltonian, including the attraction with the core, reads:

$$\mathcal{H}_{Length} = \frac{\underline{\hat{p}}^2}{2} + \vec{r} \cdot \vec{E}(t) - \frac{Z}{r}, \tag{1.36}$$

where $Z = 1$ for atomic hydrogen.

The velocity-gauge is many times considered to be superior to the length-gauge in terms of convergence of the quantities of interest with respect to the number of terms included in the expansion in spherical harmonics, but it possesses the disadvantage that it can easily require significantly more grid-points along the radial coordinate in a spherical-geometry setting. It can also

become cumbersome to be implemented for numerical propagation. The choice to be made to obtain this gauge is $K(\vec{r},t) = -\int_{\infty}^{t} \vec{E}\,(t')\,dt' \cdot \vec{r}$, so $\vec{A}(t) = \vec{A}(t)$ and $\phi_{new} = (-\vec{r}\cdot\vec{E}(t)) + (\vec{r}\cdot\vec{E}(t)) = 0$, thus the Hamiltonian reads:

$$\mathcal{H}_{Velocity} = \frac{\hat{\underline{p}}^2}{2} + \vec{A}(t)\cdot\hat{\underline{p}} + \frac{\vec{A}(t)^2}{2}. \tag{1.37}$$

Any wavefunction which solves the TDSE with the Hamiltonian from Eq. (1.37) will also solve the TDSE in the so-called reduced-velocity-gauge:

$$\mathcal{H}_{Velocity-Reduced} = \frac{\hat{\underline{p}}^2}{2} + \vec{A}(t)\cdot\hat{\underline{p}}, \tag{1.38}$$

because the difference between the two wavefunctions is just the multiplication of one of them by a (global) phase-factor to obtain the other. That phase disappears when evaluating the observables of interest from the time-dependent wavefunction. It is usually the case that the reduced-velocity gauge is denoted as the velocity-gauge and worked with it in the literature.

1.5 Motivation

The recent progress in the creation and usage of ultra-short and ultra-intense laser pulses of light has led the community to almost being able to experimentally measure and verify different sorts of new, complex, physics phenomena, which were not accessible up until now because of the previous limits on the value of the focused peak intensity. One of the scientific aims of some ultra-intense laser-matter interaction experimental facilities from across the globe is to increase the achievable peak intensities beyond even 10^{23} W/cm^2. On another note, a complementary aim is to increase the laser repetition rate, this effectively allowing a higher rate of scientific discovery in the laboratory.

On the high-intensity side, in particular, the two Extreme-Light-Infrastructure (ELI) pillars which are devoted to such ultra-intense laser-matter studies, ELI-NP near Bucharest, in Magurele, Romania [20–23] and ELI-Beamlines near Prague, Czech-Republic [24], plan to be operational at O(10) PW of pulsed peak power.

ZEUS, in Michigan, USA [25] plans to become a multi-petawatt-class laser facility. The CoReLS facility in South-Korea currently holds the record for the highest focused peak intensity [26]. In Salamanca, Spain, the CLPU ultra-intense laser laboratory has the capability to perform experiments at more than 2×10^{20} W/cm^2 [27]. Apollon, in Paris, at LULI, has started its commissioning phase in March 2023 with the aim to become a petawatt-class user-accessible ultra-high-power laser facility [28]. These currently under commissioning or just commissioned research facilities are preparing to start trying to identify which theoretical model of the radiation-reaction (RR) phenomenon is applicable to the laser-plasma interactions. Prolific pair-production [29, 30] will be possible due to the Breit-Wheeler (BW) mechanism. Theoretical and numerical predictions for relativistic tunnelling will be assessed in the laser laboratories [31]. The onset of Quantum-Electrodynamics (QED) cascades will prevail [32]. Thus a change of paradigm in the plasma dynamics comes with the birth of the above mentioned new physics: the transport of energy throughout the irradiated targets has to be understood in this new context, and also how the laser-energy is absorbed and partitioned has to be adapted to the newly created species.

The condition to come close to achieving these experimental milestones is to have a large value of the focused peak intensity on the target.

1.6 Conventional methods to diagnose the peak intensity

Due to the complexity of the laser beamlines in any ultra-high power laser facility, spatio-temporal couplings (STC's) are inevitable, and they negatively affect the true (i.e. real) value of the peak intensity obtained in the interaction chamber. Broadly speaking, one way to understand this is that the pulse-beam wavefront is non-trivially distorted and the focusability of the light is moving away more and more from the ideal flat-wavefront scenario. Another way is to see that due to the various types of STC's, the pulse duration and focal volume are increased, this effectively lowering the focused intensity on target in the

interaction chamber. The effect is more pronounced when the beam diameter is larger and when the pulse duration is shorter, exactly the case required to reach ultra-high peak intensities. For example, the Apollon beam diameter is 400 mm [33] and the ELI-NP beam diameter is 550 mm [34]. Generally, from the mathematical point of view, for any STC, the beam-pulse electric-field cannot be simply split into a temporal part and a spatially-dependent part anymore, i.e. a coupling appears: $\vec{E}(\vec{r},t) \neq g(\vec{r}) \cdot f(t)$.

Conventional peak intensity detection techniques assume that the splitting above can be safely made and, subsequently, separately measure the temporal profile of the pulse, $f(t)$, and the spatially dependent part of the beam, each on their own, i.e. the measurements being independent of each other.

Autocorrelation-based measurements are usually the method of choice for the temporal shape: FROG [35], SPIDER [36], D-SCAN [37], or any other variants based on any of these three.

The spatial part is dealt with by putting a calorimeter in the beam path and measuring the energy in the beam.

Then, at low power, the focal spot area A (not to be understood as the vector potential $\vec{A}(\vec{r},t)$ of a laser beam-pulse) is measured via an imagining technique.

These three characteristics, i.e. the (usually Full-Width-Half-Maximum (FWHM) of the Gaussian/Lorentzian-shape fitted) temporal distribution, the energy and the spatial shape, define a measurement of the focused intensity. [38]

The problem with the above approaches are that they cannot be performed at ultra-high power, i.e. on the beam-pulse that is used for ultra-high intensity laser-matter experiments. They are only viable as extrapolation from low-power to high-power methods. Measurements are taken at low power, where the (optical and electrical) instruments can be actually placed in the beam path. They all have damage thresholds above which they are not usable anymore as they are destroyed by the extreme laser-energy density in the focal volume. As a result, it is then assumed that what is thought to be obtained (i.e. the

extrapolation) when the actually used ultra-high power laser is shot, based on the low-power measurements, is true.

Unfortunately, the assumption does not take into account the STC's presented earlier. They reduce the actual peak intensity on the target via a complex change in the pulse-beam characteristics in the far-field.

1.7 Structure of thesis

The structure of this thesis is as follows: I will initially present an atomic diagnostic aimed at measuring the peak intensity of an ultra-intense beam-pulse of light. This is based on simulations with state-of-the-art laser-field induced ionisation rate-equations inserted into a fully-relativistic PIC code. I also concentrated on the numerical modelling of the detection of the charge states originating from the focus of the beam. Overall, my work presents a computational workflow which produced results to be compared with experimental values from the laboratory.

I will then move to the TDSE. There, I show my design and implementation of a fully distributed computational workflow to solve the TDSE for one electron atom in full-dimensionality by means of the generalized-pseudospectral-grids numerical discretisation. I also measured the performance of my fully-parallelised algorithm and gained insights about its behaviour. I then used this solver for the following problems: using atomic hydrogen as the target, I propagated a Hamiltonian which includes non-dipole corrections up to first order in $1/c$, where c is the speed of light, in the length-gauge. Specifically, I looked at the ionisation probabilities at the end of the pulse, at the photo-electron energy spectra and finally, at fully-resolved PMD's. Striking differences between the dipole-approximated and beyond-dipole electron-dynamics were obtained and carefully analysed.

Finally, the multi-electron nature of noble-gases has been taken into account using mean-field theories. I numerically solved the coupled TDSE-like equations which produced the temporal evolutions of the Kohn-Sham (KS) orbitals to

gauge about the ionisation dynamics of argon using the time-dependent-density-functional-theory (TDDFT) formalism. This has been possible due to and was the natural generalisation of my previous work on one-electron-TDSE. Using a statistical-approach based on the survival probabilities of each orbital, I obtained the charge states' distribution as a function of the peak intensity of the laser-field used in the simulation. Then I performed a parameter scan across a range of intensity values to obtain and interpret the behaviour of atomic argon when the laser peak field strength is varied.

1.8 Contributions of the author

The project initially aimed to create and implement a peak-intensity atomic diagnostic, in the spirit of well known direction composed from the measurements of charge states' distributions against the peak field-strength. The task of extending this to higher intensities which have not been measured up until now was suggested by my supervisor, Professor Peter Norreys. I found all the new field ionisation rates in the recent literature, I implemented them as a new class in SMILEI and then I ran the numerical simulations.

All the other chapters and parts of the thesis are my ideas and C++/Python implementations, unless indicated otherwise, and I carried the work under the supervision of Professor Peter Norreys.

I have also applied for beamtime in two large-scale experimental campaigns, one at ELI-NP and one at the CLF. I discussed the drafts I completed with my supervisor and we both sat through the interview process held by the International-Scientific-Advisory-Board (ISAB) of ELI-NP.

2

Inferring the peak intensity from the ionisation states of an ultra-rarefied noble-gas shot by an ultra-intense laser beam-pulse

"If you cannot measure it, it is not physics ..."

Contents

2.1 An atomic diagnostic to measure the laser peak intensity . . . 27
2.1.1 1D PIC simulations results 30
2.1.2 Number of ions originating from the focus of the beam . 32
2.1.3 Fit to the number of ions from the focus 35
2.1.4 2D PIC simulations and confirmation of the correctness of the reduced-dimensionality PIC simulation behaviour 44
2.1.5 Experimental considerations of the proposed diagnostic of the laser peak intensity 46
2.2 Conclusion and outlook 51

The conventional methods to diagnose the peak intensity in a laboratory setting do not produce accurate results at ultra-high laser field intensities. As mentioned in the previous chapter, theoretical extrapolations from data collected at low peak power to values which are thought to be obtained at ultra-high peak power produce peak intensities dramatically different than what is actually

obtained in the interaction chamber.

As an illustration, I performed numerical simulations to show the reduction in the peak intensity in the far-field which is caused by one of the simplest possible STC: the pulse-front-tilt (PFT). The PFT appears, for example, when the diffraction gratings from the compressor are misaligned. Of course, in practice perfect alignment is impossible, and for what is generally considered to be just a slight misalignment of 40 μrad (PFT of 0.2 fs/mm) [33], the reduction in the focused peak intensity is dramatic.

This workflow is inspired by [39], where the authors analysed a realistic set-up generating STC's by carefully considering the details of four diffraction gratings. The complex interplay of these four optical elements created the analysed wavefront distortions. In contrast to their approach, I considered PFT from the start and analytically included it in the phase profile. I considered a Gaussian temporal (frequency) profile for the pulse and a flat-top E-field distribution (i.e. a top-hat function which takes non-zero values only inside a given, fixed value, beam diameter). The central wavelength was $\lambda_0 = 800$ nm and the bandwidth was $\Delta\lambda = 50$ nm on both sides of the central frequency. The focal length of the focusing optics was f = 2 m. The beam diameter took values in the range $[1, ..., 600]$ mm. The PFT degradation took values in the range $[0.01, ..., 1.0]$ sec/m. I took two routes to show the degradation of the peak intensity due to the STC's. For one set of simulations, I fixed the beam diameter to 0.5 m and varied the PFT. For the other set of simulations, I fixed the PFT value to 0.2 fs/mm and varied the beam diameter. I defined the reduction as the ratio between the STC-corrupted beam-pulse's peak intensity and the STC-free beam-pulse's peak intensity. The peak intensity was the maximum value of the far-field electrical-field, and the far-field electrical field was obtained using a Fraunhoffer diffraction integral performed using highly-optimized and shared-memory-parallelised Fast-Fourier-Transform (FFT) routines. I wrote the computer code to perform this analysis from scratch using Python with its wrapper to the FFT library FFTW3 [40].

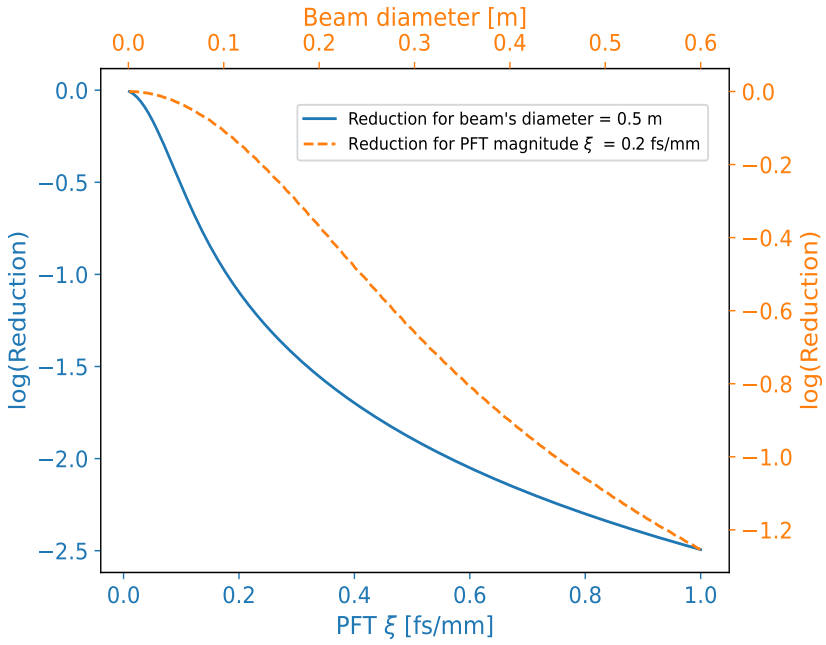


Figure 2.1: Ratio of the focused peak intensity of a beam pulse of central wavelength 800 nm and bandwidth of 50 nm having PFT degradation to the focused peak intensity of the STC-free beam pulse (thus called reduction), as a function of the beam's diameter for fixed degradation or as a function of the degradation's magnitude for fixed diameter.

The results I obtained are presented in Fig. 2.1. As expected, the beam diameter increase worsens the degradation of the peak intensity. Similarly, increasing the PFT-value makes the reduction in the peak intensity larger and larger. It is thus clear that for common sets of physical parameters, the effect of even more complex STC's is significant.

2.1 An atomic diagnostic to measure the laser peak intensity

To circumvent the difficulties posed by the presence of STC's and their accumulation along the beamline, indirect methods to measure the peak intensity are needed. Work done towards this indirect direction can roughly be categorised in three categories: Thomson scattering [41–47], ponderomotive scattering of charged particles from the laser focus [48–52] and laser-field ionisation of high-Z, ultra-rarefied, noble-gas atoms [53–56].

My work has focused on the laser-field ionisation technique. I have extended the theory behind it, made it accessible (feasible) to computational study, and showed that it can be implemented as an *in-situ* diagnostic. This diagnostic is easy to operate in the laboratory and is suitable for the future plans of increasing the repetition rate of the ultra-high power laser facilities from across the globe.

The physics phenomenon underlying this peak intensity retrieval is the laser-field induced ionisation of atoms. It is commonly described in terms of the Keldysh parameter [57]:

$$\gamma_K = \frac{E_K}{E_l} = \omega_L \frac{\sqrt{2 m_e \, IP_i}}{e \, E_l}, \tag{2.1}$$

where m_e is the electron's rest mass, IP_i the ionisation potential (IP) of the atom (or ion), ω_L the light's frequency, e the fundamental unit of charge and E_l the peak electric field of the laser pulse. γ_K dictates in which regime the ionisation takes place: multi-photon (MPI) if $\gamma_K > 1$ or tunnel (TI) if $\gamma_K < 1$. For current state-of-the-art laser-matter interaction experiments, where $\gamma_K \ll 1$, the Perelomov, Popov and Terent'Ev (PPT) theory for field ionisation rates [58, 59] is often used. Following [60], the PPT static electric field ionisation rate is:

$$\begin{aligned} w_{lm} &= \omega_a \, \kappa^2 \, C^2_{\kappa l^*} \, (2l+1) \left(\frac{2}{F}\right)^{2n^* - |m| - 1} \\ &\quad \times \frac{(l+|m|)!}{2^{|m|}(|m|)!(l-|m|)!} \, exp\left(\frac{-2}{3F}\right), \\ C^2_{\kappa l^*} &= \frac{2^{2n^*}}{n^* \Gamma(n^* + l^* + 1)\Gamma(n^* - l^*)}, \end{aligned} \tag{2.2}$$

where l and m are the orbital and magnetic quantum numbers of the electron, ω_a the atomic unit of frequency $\omega_a = 4.1341 \times 10^{16}$ s^{-1}, $\kappa^2 = IP_i / IP_H$ with $IP_H = m_e e^4 / (2\hbar^2) \approx 13.6$ eV, $F = E/(\kappa^3 E_a)$ with E the local value of the electric field and E_a the atomic unit of electric field $E_a = 5.1422 \times 10^{11}$ V/m, $n^* = Z/\kappa$, $l^* = n^* - 1$, with Z the ion charge number, and $\Gamma(x)$ the gamma function, defined as $\Gamma(x) = (x-1)!$ for a natural number x and $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt$ for a complex number z with a strictly positive real part.

However, quantum tunnelling only occurs when the potential barrier caused by the combined atomic and laser fields is higher than the initially unperturbed energy level of the electron. For the case of higher laser-pulse electric fields, the electron classically escapes above the potential barrier [2]: BSI happens for $E \gtrsim E_{cr} = E_a \kappa^4/(16Z)$ and then the quantum-tunnelling phenomenon is not present anymore. For $E > E_{cr}$ the tunnel rates strongly (dramatically) deviate from numerically calculated rates [61], as expected. The transition region between TI and BSI is correctly described by the Bauer-Mulser (BM) rate $w_{BM} = 2.4\omega_a \left(E/E_a\right)^2 \left(IP_H/IP_i\right)^2$, proven to work for $E \sim E_{cr}$ [62]. Furthermore, w_{BM} deviates from numerically calculated rates as $E \gg E_{cr}$ [60, 62]. Derived both classically and quantum mechanically (in the motionless and free electron approximation), the rate dependence on E for $E \gg E_{cr}$ is linear [60]: $w_{BSI} = p\,\omega_a (E/E_a)\sqrt{(IP_H/IP_i)}$. The coefficient p (a purely real number) ranges from 0.62 to 0.87, depending on the potential used in the derivation: for the three-dimensional Coulomb potential $p \approx 0.8$. Numerically integrating the time-dependent Schrödinger equation (TDSE) shows that both approximations used in the theoretical (analytical) derivation are indeed applicable [60]. w_{lm}, w_{BM} and w_{BSI} are grouped into a piecewise-defined rate [63] which describes the whole range of possible E values at the location of the atom:

$$w(E) \approx \begin{cases} w_{lm}, & E \leq E_1 \\ w_{BM}, & E_1 < E \leq E_2 \\ w_{BSI}, & E > E_2 \end{cases}, \tag{2.3}$$

where E_1 and E_2 are found by imposing continuity on w at the transition points.

To model the interaction between the laser pulse and the atom, I made use of the open-source, massively-parallel, particle-in-cell (PIC) `C++` code SMILEI [18]. It has already had the PPT rate implemented when I started the work presented in this chapter. I modified the source code and included both the BM and BSI rates presented above, but also the mechanism to transition between the three rates, in a general, unified, ionisation model.

The diagnostic is based on theoretical predictions of the ionisation charge-state dynamics, which are compared with the experimental observables: the peak intensity is inferred from the number of ions of different charge states which are collected in the laboratory using Thomson-parabola detectors. The experimental set-up is rather simple and is the starting point for many laser-matter experiments: an off-axis parabola (OAP) with a short focal length focuses the beam as close as possible to a diffraction-limited spot, where the interaction with an ultra-rarefied gas takes place. The laser-field ionisation develops and the atom is stripped of its electrons. Ions of different charge states are created due the passage of the laser pulse and are then collected by means of Thomson parabolas. No extrapolation from low-power to high-power is needed as the obtained results fully and correctly include the real, high-power laser beam-pulse.

2.1.1 1D PIC simulations results

I ran one-dimensional (1D) PIC simulations using the modified SMILEI code and analysed the charge states' populations after the passage of the laser pulse.

It is very important that the gas is ultra-rarefied as I must avoid all plasma-effects which can potentially arise. The diagnostic has to measure the pulse itself, not spurious electrical fields which can potentially appear due to how particles from the focal volume move and interact with the pulse. Basically, the diagnostic is sensitive to the local-value of the electric-field and this has to be only the laser-pulse electric field itself with no additions or modifications to it. Xenon (Xe), atomic number Z = 54, was used in my SMILEI simulations, and the initially neutral gas density n_0 was chosen such that the electron number-density for the fully ionised gas, n_e, is low enough to make both collective effects and relativistic self-focusing negligible. The first condition led me to setting the plasma frequency $\omega_p = \sqrt{n_e e^2 / (m_e \epsilon_0)}$ to be such that the plasma oscillation period is significantly longer than the laser pulse (of O (10) fs at ultra-high power laser facilities). The second condition led me to setting n_e to be such that $P_L = 10\ \text{PW} < P_{cr} = 17 \left(\omega_L/\omega_p\right)^2$ GW. The right-hand-side of the previous equality

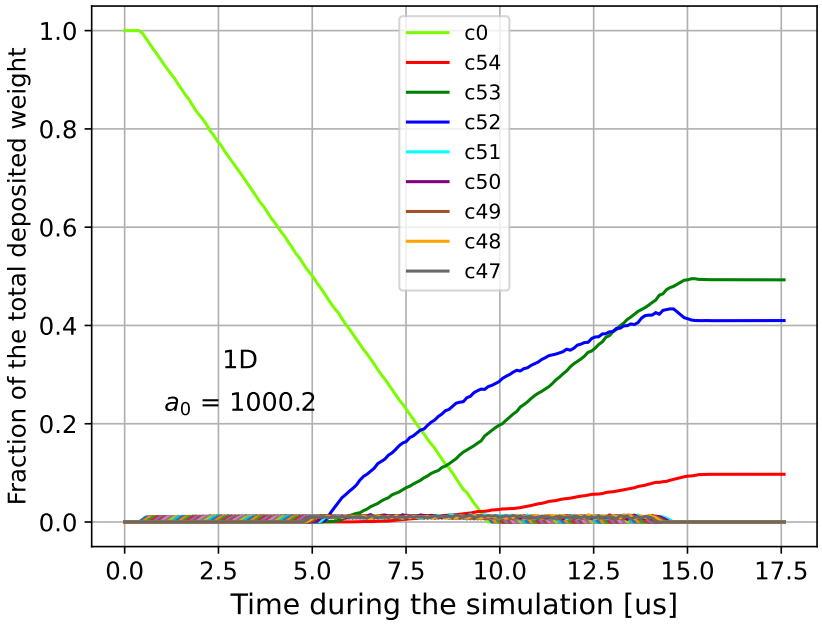


Figure 2.2: Xe charge states in a 1D PIC simulation. The gas density was low enough to allow neglecting collective plasma effects and relativistic self focusing.

is the critical power for relativistic self-focusing. To meet both constraints, the neutral Xe number density $n_0 = 1.98 \times 10^{12}$ cm^{-3} was chosen.

For a fixed value of the laser's normalised vector potential $a_0 = eE_l/\left(m_e\omega_L c\right)$, a 1D PIC simulation shows how the proportion of the different charge states evolves during the passage of the laser pulse. Typical results of a 1D simulation are shown in Fig. 2.2.

The charge states' dynamics are displayed as a function of time during the simulation. The $\hat{y}$-axis represents the deposited weights from a particle binning diagnostic for different ion charge states, as a fraction of the total Xe species population's deposited weight. Generally speaking, one value of any of the c_n curves represents the fraction from the total number of the ions that particular charge state represents, at that particular time during the simulation. The electrons resulting from ionisation were injected into the simulation as a separate species which had its radiation-reaction module turned on. The laser pulse had a wavelength $\lambda_L = 800$ nm and a Gaussian temporal profile with a full-width at half-maximum (FWHM) of six optical cycles (six T_L). The extent of

the laser pulse's temporal profile was limited to 10 T_L. The simulation box was 12 λ_L long and the spatial resolution was set to 8. On both sides, Silver-Müller boundary conditions were chosen for the fields. The simulations lasted for 22 T_L and the number of particles per cell (ppc) for Xe was 32. The same simulations, but with 128 ppc, gave the same results as for those with 32 ppc. This shows that the results are converged with respect to this simulation parameter.

I collected the results obtained at the end of each such simulation. Running 1D simulations to cover the range $a_0 \in [6.8, ..., 2162.8]$ ($\sim I \in [10^{20}, ..., 10^{25}]$ W/cm^2 for a plane wave), the ionisation behaviour as I, the plane-wave intensity, is varied, can be obtained. Fifteen-thousand 1D simulations were run to finely sample this intensity range. The results for the whole range of intensities appear in Fig. 2.3. Each data point from one c_n curve corresponds to the fraction of the deposited weight of the charge state Xe^{n+} relative to the total deposited weights of all states at the end of one simulation at a particular input intensity. The charge states values at the end of each simulation are plotted against the corresponding intensities.

2.1.2 Number of ions originating from the focus of the beam

To diagnose I, the numerically-estimated functions $c_n(I)$'s shown in Fig. 2.3 are used to infer the number of Xe ions expected to originate from the focus of the laser pulse. A suitable model for $I(\vec{r})$ is chosen and the number of Xe ions of charge state $n+$ is obtained from an integral over the focal volume:

$$N\left(Xe^{n+}\right) = n_0 \iiint c_n\left(I\left(\vec{r}\right)\right) d^3r, \tag{2.4}$$

where

$$I(\vec{r}) = \frac{I_m}{1 + \dfrac{z^2}{z_R^2}} \exp\left(-\frac{2\rho^2}{w_z^2}\right) \tag{2.5}$$

and $\rho = \sqrt{x^2 + y^2}$, I_m is the peak intensity, $z_R = \pi w_0^2/\lambda_L$ is the Rayleigh range, $w_0 = 3$ μm is the beam waist and $w_z = w_0\sqrt{1 + (z^2/z_R^2)}$. Other descriptions of the focus can be employed [56].

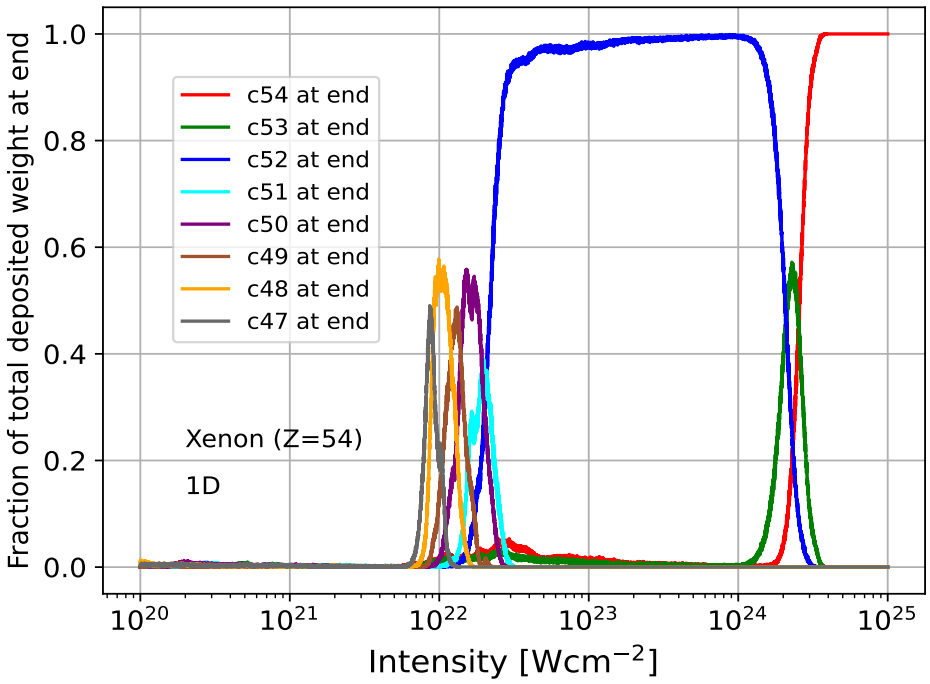


Figure 2.3: Xe representative charge states' populations at the end of 15,000 1D PIC simulations. These $c_n(I)$ curves can be used to calculate the number of ions produced in the focus.

Even if the focus' geometrical shape is not known analytically, one can think of an optimisation campaign to infer a complete description of the beam. Usually, the real beam shape in the physics-laboratory is not a smooth, fully well-behaved function as the textbook example of the Gaussian beam: it can have regions of high intensity distributed across the focal volume, often at random. The computational workflow from this chapter can be employed in a guided search through the parameter space in which all possible beam shapes reside and the most probable conditions which generated the collected experimental observables are inferred by comparing to the PIC simulations plus the focal-volume integration.

For the Gaussian beam presented above, to numerically calculate the integral from Eq. (2.4), linear 1D interpolators are used to obtain the values of c_n at I-values not simulated using the PIC code but required by the integration routine. For I-values outside the PIC-simulated range of $[10^{20}, ..., 10^{25}]$ W/cm^2, manual extrapolations for the highly charged states (i.e. c_{47} to c_{54} ones) are done. From

physical considerations (confirmed by Fig. 2.3), any highly charged state is not excited for $I < 10^{20}$ W/cm^2. Additionally, only the fully ionised state remains for $I > 10^{25}$ W/cm^2.

A midpoint rule in two dimensions (z and ρ) is employed. A two-dimensional integration scheme is used because Gaussian beams are symmetric with respect to the azimuthal dimension (ϕ). I performed a 2D cylindrical integral, this being then multiplied by 2π to calculate the full 3D integral.

The integration ranges are obtained by considering that below a threshold intensity the c_{47} to c_{54} charge states are not excited. From Fig. 2.3 $I_{thr} = 10^{20}$ W/cm^2.

The integral in z is performed across $[-z_{max}, z_{max}]$, with $z_{max} = z_R(I_m/I_{thr} - 1)^{0.5}$.

At each step along z, the integration along ρ is performed across $[0, \rho_{max}]$, where ρ_{max} is obtained by imposing that the intensity at that particular z and ρ_{max} equals I_{thr}, thus

$$\rho_{max} = \sqrt{\frac{-w_z^2}{2} \ln\left(\left(1 + \frac{z^2}{z_R^2}\right)\frac{I_{thr}}{I_m}\right)}. \tag{2.6}$$

The number of steps in z was $N_z = 2^{12}$ and the number of steps in ρ, for each z value, was $N_\rho = 2^{12}$. I have written the integration script from scratch, using the Python's wrapper to BLAS, i.e. the `Numpy` scientific-computation library. I have parallelised the script using the Python's Message-Passing-Interface (`MPI`) implementation from `mpi4py`. The parallelisation strategy was to partition the $\hat{z}$-axis across the `MPI` processes (domain decomposition). Each such process computed the chunk of the 1D radial integrals at fixed z-values it received from the total set of 1D radial integrals at fixed z-values, set which spawned the whole $\hat{z}$-axis range.

The numerical errors of the computational method have to be estimated and analysed in order to ensure the veracity of the results obtained. The $N(Xe^{n+})$'s from Eq. 2.4, obtained using the 2D midpoint rule, are presented in Fig. 2.4.

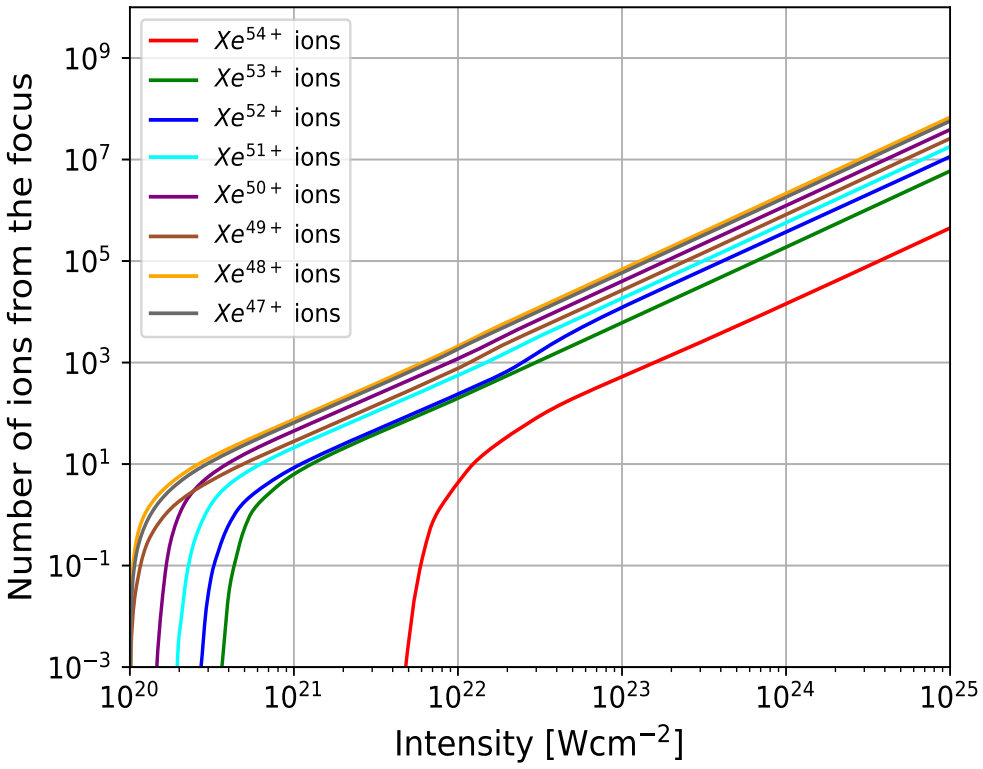


Figure 2.4: Number of xenon ions produced in the focus of the Gaussian beam from Eq. (2.5). One datapoint on each curve corresponds to the number of ions arising from one, fixed, peak intensity I_m value in Eq. (2.5). Results on y-axis are numerical calculations for $N\left(Xe^{n+}\right)$ from Eq. (2.4).

There, integral estimates in the theory of converging sequences of numerical approximations [64] are shown (after being multiplied by n_0). The result estimate for one integral is denoted as I_M, where M is the highest number of points my laptop could accommodate, $M = 2^{12}$ for both z and ρ axes. The same integral, with half the number of points along both z and ρ, was performed and stored as $I_{M/2}$. The numerical error on a result I_M was determined to be $|I_M - I_{M/2}|$. After being multiplied by n_0, these approximations to the numerical errors are too small to be visible on the plot.

2.1.3 Fit to the number of ions from the focus

It is possible to train a regressor on the curves from the Fig. 2.4 and then query it an unseen examples. This is a tool I was informed that physicists working at high-power laser facilities would benefit from. Curve-fitting when the underlying

data-generating distribution is unknown is more and more prevalent nowadays in the form of machine-learning (ML). I fitted the curve $N\left(Xe^{n+}\right)$ as a function of I using a Relevance Vector Machine (RVM) model [65]. I have implemented it from scratch using `Python-Numpy`.

Bayesian thinking was adopted: this framework naturally penalises heavily complex models which overfit drastically when tested on the hold-out set. The resulting fitting equation is envisioned to be practical, thus Bayesian linear regression with the basis functions being the standard polynomials was performed. The regressor acquires a Mean-Squared-Error (MSE) on the test set of around 2×10^{-3} when the targets have been standardised (mean made to be equal to 0 and variance made to be equal to 1). This is comparable to results which I obtained using state-of-the-art ML algorithms, such as the XGBoost [66] version of the ensemble of weak-learners (i.e. boosted trees). The optimisation algorithms implemented in Optuna [67] were used for setting the hyperparameters of the XGBoost (in particular an algorithm based on the evolutionary optimisation strategy called CMA-ES and an algorithm based on trees called TPE). These algorithms were let run to try to find suitable values for the hyperparameters for 15,000 trials and the objective of their minimisation was the MSE on the validation set.

As means of example, the dataset is the curve with the fewest datapoints from Fig. 2.4, i.e. the fully ionised ions curve. It is denoted as $\mathcal{D} = \{(x_i, t_i) : i = 1, 2, ..., N\}$, where $N = 681$, x_i represents the intensity and t_i represents the number of ions from the focus. The usual over-top arrow for vectors is used and a vector will be expressed through a lower-case letter, while a matrix will be expressed as an upper-case letter. Anything written in normal font and lower-case letters is a real-valued scalar. The noise is modelled as Gaussian, i.e. it comes from a normal distribution with mean 0 and variance $1/\beta^2$. The polynomial features are expressed, for one input example x_n, as: $\vec{\phi}(x_n) = \left[1, x_n, x_n^2, ..., x_n^M\right]$. The 2-dimensional design-matrix Φ contains across its rows the $\vec{\phi}(x_n)$'s, for

$n \in \{1, 2, ..., N\}$. Although non-linear in the input x, the model is linear in the fit-parameters, i.e:

$$p(t|x, \vec{\theta}) = \mathcal{N}\left(t \mid \vec{w}^T \vec{\phi}(x), \frac{1}{\beta^2}\right), \tag{2.7}$$

where $\vec{\theta}$ is a generic notation for the parameters of the model and $\vec{w}$ is a column vector with M elements. The likelihood of the data $\vec{t} = [t_1, t_2, ..., t_N]^T$ and $X_i = [x_i]$ for $i \in \{1, 2, ..., N\}$ (thus X is shape $N \times 1$), i.e. the joint-likelihood, given the assumption that the datapoints are drawn in an i.i.d. fashion (independently and identically distributed) is:

$$p(\vec{t}|X, \vec{w}, \beta) = \prod_{i=1}^{N} p(t_n|x_n, \vec{w}, \beta), \tag{2.8}$$

thus a Gaussian probability distribution function. A prior is placed on the weights $\vec{w}$:

$$p(\vec{w}|\vec{\alpha}) = \prod_{i=1}^{M} \mathcal{N}\left(w_i \mid 0, \frac{1}{\alpha_i}\right), \tag{2.9}$$

where $\vec{\alpha} = [\alpha_0, \alpha_1, \alpha_2..., \alpha_M]^T$. The posterior on the weights, from the Bayes rule, is proportional to the product of the joint-likelihood and the prior on the weights. As both of the terms are Gaussian, the resulting product is again Gaussian:

$$p(\vec{w}|\vec{t}, X, \vec{\alpha}, \beta) = \mathcal{N}(\vec{w} \mid \vec{m}, \Sigma), \tag{2.10}$$

where:

$$\begin{aligned} \vec{m} &= \beta \Sigma \Phi^T \vec{t} \\ \Sigma &= (A + \beta \Phi^T \Phi)^{-1}, \end{aligned} \tag{2.11}$$

where $A = diag(\vec{\alpha})$ (thus shape $M \times M$). For an infinitely broad prior, $\alpha_i \to 0$ for $i \in \{1, 2, ..., M\}$, the mode of the posterior $\vec{w}_{MAP} = \arg\max_{\vec{w}} \; p(\vec{w}|\vec{t}, X, \vec{\alpha}, \beta)$ tends to the Maximum Likelihood Estimate of Type I for the weights, $\vec{w}_{MLE}$, commonly used in simple fits of data (which can easily overfit).

The priors on $\vec{\alpha}$ and β are obtained through Empirical Bayes (Maximum Likelihood Type II). They are those values which maximise the marginal (log)

likelihood, i.e. the integral of the likelihood against the prior on the weights across all the parameters space $\vec{w}$:

$$p(\vec{t}|X,\vec{\alpha},\beta) = \int_{\vec{w}} p(\vec{t}|X,\vec{w},\beta)p(\vec{w}|\vec{\alpha})d\vec{w}. \tag{2.12}$$

This integral is the correlation of two Gaussians and using the "Bayesian-Gaussian identity": $\int_{\vec{z}} \mathcal{N}(\vec{y}\,|\,W\vec{z}+\vec{b},\,\Sigma_y)\,\mathcal{N}(\vec{z}\,|\,\vec{\mu}_z,\,\Sigma_z)\,d\vec{z} = \mathcal{N}(\vec{y}\,|\,W\vec{\mu}_z+\vec{b},\Sigma_y+W\Sigma_y W^T)$, I obtain that:

$$p(\vec{t}|X,\vec{\alpha},\beta) = \mathcal{N}(\vec{t}|\vec{0},C), \tag{2.13}$$

where

$$C = \beta^{-1}\mathbb{1} + \Phi A^{-1}\Phi^T, \tag{2.14}$$

with $\mathbb{1}$ being the identity matrix. This marginal likelihood has to be maximised w.r.t. $\vec{\alpha}$ and β and this is a non-convex optimisation problem. One could perform the Expectation–Maximisation (EM) algorithm [68] on it, but I have followed [65, 69]: $\vec{\alpha}$ and β have been iteratively updated as below, up until a convergence criterion was achieved:

$$\begin{aligned} \alpha_i^{new} &= \frac{\gamma_i}{m_i^2} \\ (\beta^{new})^{-1} &= \frac{||\vec{t}-\Phi\vec{m}||^2}{N-\sum_i \gamma_i}, \end{aligned} \tag{2.15}$$

where $\gamma_i = 1-\Sigma_{ii}$, m_i is the i-th component of $\vec{m}$ and $||\vec{a}||^2 = \sum_i a_i^2$ for a real-valued $\vec{a}$ and where the sum runs over all elements of $\vec{a}$. In practice, the EM algorithm tends to provide a slower convergence than what I used [65]. The hierarchical Bayesian inference model has been stopped here: the optimisation algorithm was simply started from very small values of $\vec{\alpha}$ and β. I mention that one could put a hyper-prior on $\vec{\alpha}$ and β as well in order to be even more Bayesian (Gamma distributions, which in turn have parameters, but set by hand to fixed values), an idea known in the literature [65].

After the optimisation routine was run, I found that some of the elements of $\vec{\alpha}$ increased dramatically, up to infinity. This translates into the fact the the posterior on those weights w_i for which the $\alpha_i \to \infty$ is centred at 0 and has a virtually-0

variance [65, 69]. The explanation for this is that the model prunes the basis functions corresponding to these w_i's: they have no predictive power (or have a very low one) for the problem which is to be solved, given the dataset. This way the model naturally balances complexity with the avoidance of overfitting, beautifully incorporating the Occam's razor idea. Instead of using regularisers and expensive k-fold-cross-validation (CV) runs, this Bayesian approach to learning achieves roughly what *l1* and *l2* regularizations (for example) achieve. Details can be found in [69].

During the run of the optimisation routine I was confronted with numerical instabilities because of those α_i which shoot up to ∞, especially when the maximum order of the polynomial was really high (> 30, for example). I followed the original idea in Reference [65] and, as explained in Appendix B.1 from the mentioned paper, the columns of Φ corresponding to those values of α_i which became large during the optimisation routine were removed: in particular, when $\gamma_i < \epsilon$, where ϵ is the machine-precision, that element i from the vector $\vec{\alpha}$ was removed, similarly with column i from Φ. These caused no harm for the remaining steps of the optimisation routine running to maximise the marginal likelihood (details in the same Appendix).

Having found the optimal $\vec{\alpha}$ and β values, denoted as $\vec{\alpha}^*$ and β^*, the predictive distribution of the model is again a Gaussian and reads:

$$p(t|x, X, \vec{t}, \vec{\alpha}^*, \beta^*) = \int_{\vec{w}} p(t|x, \vec{w}, \beta^*) p(\vec{w}|X, \vec{t}, \vec{\alpha}^*, \beta^*) d\vec{w} = \mathcal{N}(t|\vec{m}^T \vec{\phi}(x), \sigma^2(x)), \tag{2.16}$$

where:

$$\sigma^2(x) = (\beta^*)^{-1} + \vec{\phi}(x)^T \Sigma \vec{\phi}(x), \tag{2.17}$$

with Σ from Eq. (2.11) with β set to β^* and A arising from $\vec{\alpha}^*$.

Because it is envisioned that the fitted equation is used in real-time laboratory settings, the maximum order of the polynomial cannot be high. In order to choose which model to use, the Bayesian view was again adopted and the model evidence has been used to rank the models. That is, models with $M \in \{1, 2, 3, ...\}$

were tried: empirical Bayes was applied to each of them and then the model evidence was calculated for each of them. The model evidence is the normalizing constant (the denominator) from the Bayes rule for the parameters $\vec{w}$:

$$p(\vec{w}|D,\mathcal{H}) = \frac{p(D|\vec{w},\mathcal{H})p(\vec{w}|\mathcal{H})}{p(D|\mathcal{H})}. \tag{2.18}$$

For no a-priori preference for a particular model, given a set of models $\{\mathcal{H}_j : j = 1,2,...,J\}$, $p(\mathcal{H}_j) = 1/J$, so there is no preference for any of the models. Bayes rule then reads:

$$p(\mathcal{H}_j|\mathcal{D}) \propto p(\mathcal{D}|\mathcal{H}_j)p(\mathcal{H}_j), \tag{2.19}$$

and the ranking of two competing models is performed based on the Bayes factor, which is equal to the ratio of posterior probabilities of the models in light of the data $\mathcal{D}$, in turn equal to the models' evidences for no a-priori preferences over them:

$$B_{ab} = \frac{p(\mathcal{H}_a|\mathcal{D})}{p(\mathcal{H}_b|\mathcal{D})} = \frac{p(\mathcal{D}|\mathcal{H}_a)}{p(\mathcal{D}|\mathcal{H}_b)} \tag{2.20}$$

Calculating a generic model evidence is a complicated task and is well known in the literature [70]. In the model presented here, because everything is analytical, the model evidence is in closed form: usage of Markov-Chain-Monte-Carlo (MCMC) [71] or simplifying approximations, such as the Laplace one, or any of the information-criteria appearing in the literature (Akaike [72], Bayesian or Schwarz [73], etc.) which are used for model comparison, has not been made. The log of the marginal-likelihood, that is, the log of the model evidence, from Eq. (2.13), reads, for one model:

$$\ln p(\vec{t}\,|\,X,\vec{\alpha}^*,\beta^*) = -\frac{1}{2}\left[\, N\ln(2\pi) + \ln|C| + \vec{t}^{\,T} C^{-1}\vec{t}\,\right], \tag{2.21}$$

where C is evaluated from Eq. (2.14) with $\beta \leftarrow \beta^*$ and $\vec{\alpha} \leftarrow \vec{\alpha}^*$ and $|C|$ represents the determinant of the matrix C.

The obtained log-model-evidences (from Eq. (2.21)) as a function of the polynomial order of the model are presented in Fig. 2.5.

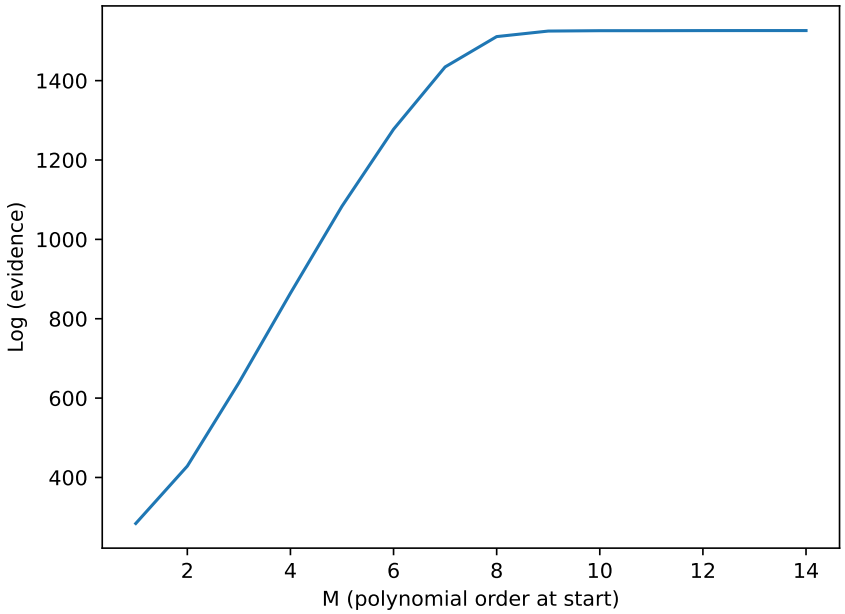


Figure 2.5: Model evidence as a function of the polynomial order at the start of the training procedure.

Because of how the RVM algorithm is designed, it removes and does not count on the basis functions which make the model infer a complexity penalty greater than the benefit obtained from a better fit to the train data. This is what has been explained above when $\alpha_i \to \infty$.

Because the fitting equation has to be practical, the maximum polynomial order, that is M, which was chosen is one for which the number of remaining (non-pruned) weights is less than a reasonable number, which translates in the ability to list all the non-zero elements from the resulting $\vec{\alpha}^*$. I chose that reasonable number to be seven. For this, $M = 14$ in my case: this polynomial order remains, after pruning, with seven non-zero weights. There might be the case that as M increases way more than what was chosen, there will be so many pruned basis functions in those high-M models that the number of the remaining ones would be less than my reasonable number. This issue is highlighted here because for those large M's, the model evidence might be higher than the model evidence for what was chosen ($M = 14$). I studied this and found, in the current work, that

the remaining number of basis functions greatly exceeds the reasonable number, making it clear that any of those higher-M models are not to be reported.

The choice of the polynomial order can be analysed also from the point of view of the Mean-Squared-Error (MSE) on the test-set. The test set contained 103 examples. The MSE values as a function of the polynomial order is presented in Fig. 2.6. Although a stochastic quantity suffering from noise (especially in the low-amounts-of-data regime), this MSE on the test-set provides an insight on the performance of the model and is widely used in the frequentist view of model selection (albeit with expensive cross-validation runs).

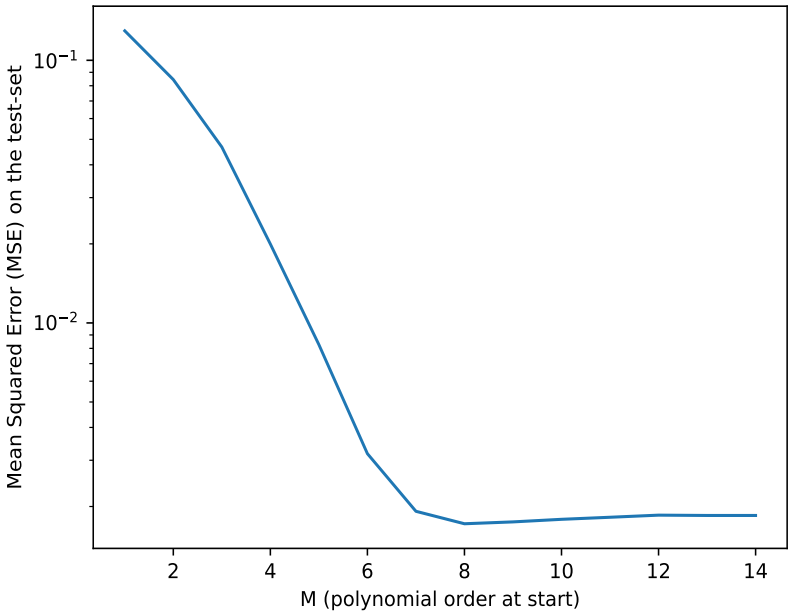


Figure 2.6: Test set mean squared error as a function of the polynomial order at the start of the training procedure.

One ought not to discard the models with approximately equal or slightly lower evidence than the evidence of the "best performing" model. Bayesian Model Averaging (BMA) needs to be undertaken (i.e. predictions done using all the ranked models, each of the contributions being weighted by the corresponding evidence), but in this simple fit I performed a point-wise approximation to BMA and the model with the highest evidence was chosen, understanding the

limitations this approach poses, but also the benefits it offers. The relationship between Bayesian Model Selection and frequentist approaches is way beyond the scope of the work from this chapter and will not be further discussed here, but details can be found in [74].

I have encountered large numerical instabilities when I tried the algorithm on the original, real dataset. As mentioned earlier, I eventually standardised it to have mean 0 and variance 1. This is a technique to make the algorithms which are deployed on the dataset more robust to numerical instabilities and proved itself useful for my task. Before subtracting the mean and dividing by the standard deviation, I also took the logarithm in base 10 of both the response but also of the predictors, again of the whole dataset. This was because the dataset spans many orders of magnitude (the intensity between 10^{20} W/cm^2 and 10^{25} W/cm^2 and the number of ions between 10^{-3} and 10^6, or even 10^8, depending on the considered charge state). Doing this, the numerical instabilities disappeared. Thus when predicting the response at new, unseen examples, the input to the regressor is scaled accordingly, similarly for the obtained response.

The obtained non-zero elements of $\vec{\alpha}^*$ and the scalar β^* are now listed to two significant digits:

$$\begin{aligned} &\alpha_0^* = 4.7 \times 10^1,\ \alpha_1^* = 2.2,\ \alpha_3^* = 1.6 \times 10^2, \\ &\alpha_4^* = 1.1 \times 10^3,\ \alpha_5^* = 3.3 \times 10^3, \\ &\alpha_{11}^* = 9.4 \times 10^4,\ \alpha_{12}^* = 3.7 \times 10^5, \\ &\beta^* = 5.9 \times 10^3 \end{aligned} \tag{2.22}$$

The plot of the predictions on the test set is presented in Fig. 2.7, where the region in blue highlights the two-sigma uncertainty of the predicted quantity (the posterior predictive is a Gaussian, see Eq. (2.16)).

Having this tool to predict the number of ions of any charge state at any input intensity value which was included in the simulation runs' targeted intensities it is helpful as it enables one to utilise the theoretical c_n curves to infer the peak intensity in the laboratory.

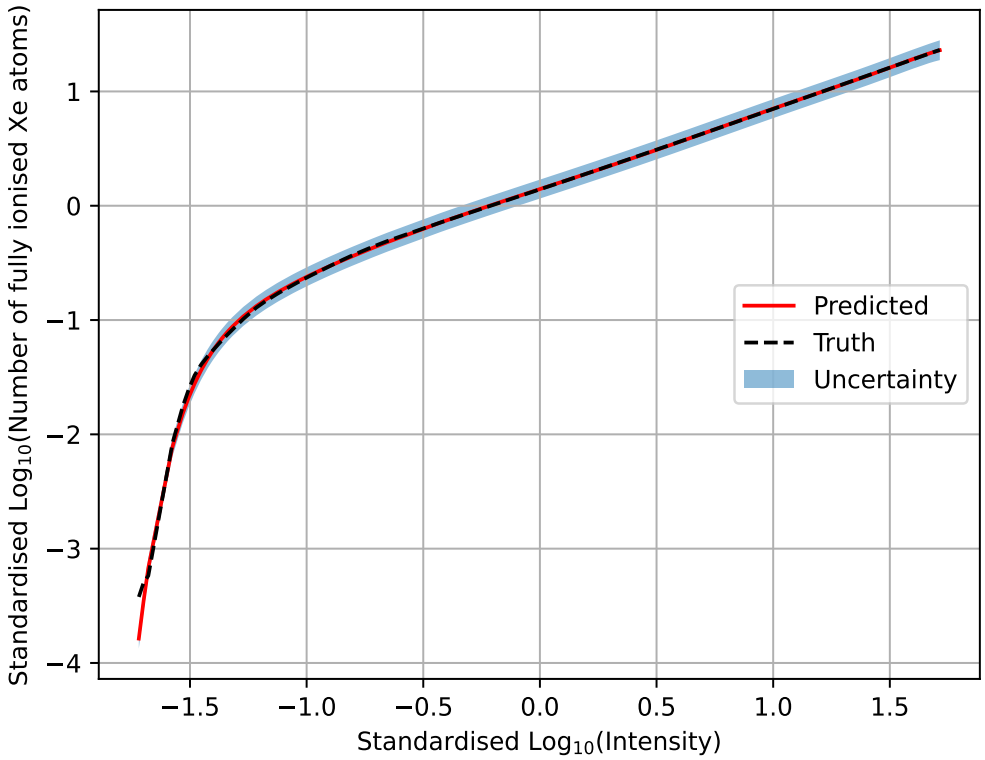


Figure 2.7: Predicted response as a function of the input for the model for which I listed the values of $\vec{\alpha}^*$ and β^*. The two-sigma deviation is highlighted in blue and represents the 95 % uncertainty in the predictions, as from Eq. (2.16).

2.1.4 2D PIC simulations and confirmation of the correctness of the reduced-dimensionality PIC simulation behaviour

To verify that the reduced dimensionality (i.e. 1D) of the PIC simulations was not impacting the results, I also investigated the relationship between I and charge state distribution in 2D PIC simulations. For a beam with a Gaussian transverse mode, a 3 μm radius at focus and the same λ_L and the same temporal profile as above, the charge states at the end of the simulation for 1,000 different a_0 values are presented in the left panel of Fig. 2.8. The middle panel of Fig. 2.8 shows the spatial distribution of a representative beam's squared electric field from one such 2D simulation. The 2D PIC simulation box was 12 λ_L along x and 20 λ_L along y, with Silver-Müller boundary conditions along x and periodic boundary conditions along y. I have found that it is essential to have computational box sizes large enough that the laser field does not touch the boundaries perpendicular

to the laser propagation direction. Reflection of the electromagnetic wave from those box boundaries, even if set to be absorbing, often poses computational stability issues. While much work has been done by the scientific community to try to alleviate this problem, the classical example being the invention of perfectly-matched layers [75] and the body of literature devoted to improving and rephrasing it in more performant terms afterwards, a simpler solution I used was to increase the box size up until there is almost no touching of the box-wall and the laser electric field in the directions perpendicular to the laser propagation direction. This is seen in the middle panel of Fig. 2.8: the box extends vertically upwards and downwards considerably.

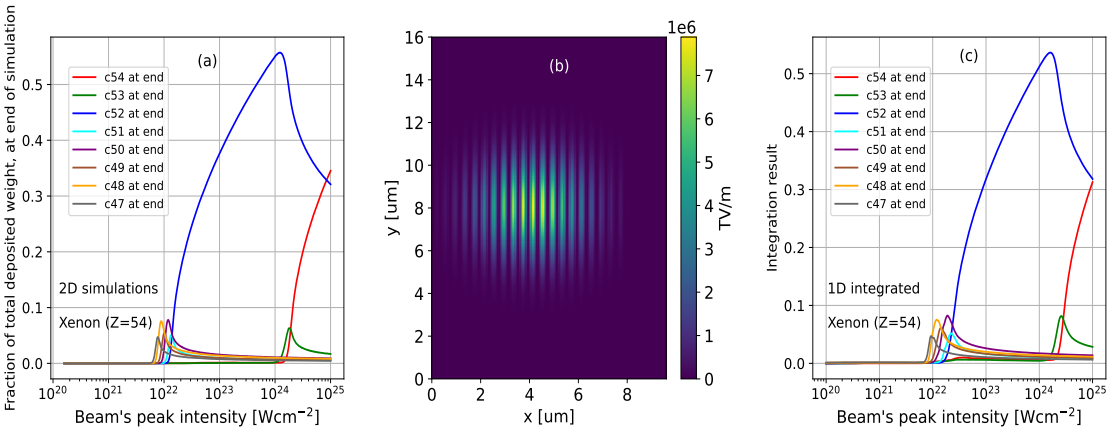


Figure 2.8: (a): Xenon representative charge states' populations at the end of 1,000 2D gaussian beam PIC simulations. (b): squared electric-field at one timestep, across the 2D PIC simulation box, for a simulation with $a_0 = 707.19$. (c): Results of the integration of the 1D charge states across the 2D PIC simulation's box.

To check the consistency of the results between the 1D and 2D simulations, a 2D version of the 3D integral from Eq. (2.4) of the c_n was performed, while keeping the grid size the same in the simulation boxes. The right panel of Fig. 2.8 confirms that the charge states at the end of a 2D simulation reproduce the 1D results.

2.1.5 Experimental considerations of the proposed diagnostic of the laser peak intensity

To count the different charge states, filtering of the resulting ions based on their charge-to-mass ratio is needed and this is normally done using a Thomson spectrometer. The spectrometer is composed of a metallic box in which static and parallel electrical and magnetical fields act upon the particles incident into the detector's volume. At the end of the particle-flight region, a (many times passive) planar ion detector is placed and collects the deflected charged particles. Such a passive ion detector can be any of the following: Image Plates (IP's), CR-39 nuclear track detectors, Radio-Chromic Film(s) in a stack (RCF's), etc. For a fixed charge-to-mass ratio, the spectrometer disperses the ions based on their incident energy and the ions will "draw a parabola" on the planar detector. For a fixed incident energy, the spectrometer deflects the different charge-to-mass ratios incident upon its volume and these will lie at the end of the flight-region on different parabolas, each of them corresponding to one such ratio. If the deflection is not strong enough (i.e. the fields are not high enough, or if the geometry of the apparatus is not carefully thought of), the parabolas will start to overlap and the differentiation between species is no longer possible. Careful consideration is needed and numerical simulations are usually performed to gain insights into the details of the deflection mechanism and how/where on the screen one expects to see the collected ions. Then the comparison between the data collected in the laboratory and the numerical predictions allows one to differentiate between the curves, see the incident kinetic energies and count the ions.

I have written a computer code in `Python` (`Numpy`) which simulates the flight of a charged particle through regions of static parallel electric and magnetic fields by means of an adaptive step-size Runge-Kutta 45 (RK45) Fehlberg integration method. It solves the relativistic equations of motion (EOM) of the charged particle given its initial conditions and the integration stops once the particle hits the detector at the end of the box or goes out of the box's volume by hitting the walls (i.e. when "clipping" occurs). It makes efficient use of all the CPU cores

by running the same solver, in parallel, on multiple input particles, as many as the number of cores which the target machine on which the code runs has. I then benchmarked it against the community-standard, widely used, highly performant `C++` code developed at CERN, `Geant4` [76].

Having the object-oriented philosophy at its core, `Geant4` allows simulating arbitrarily complex geometries by creating volumes of any shape and size. It also offers state-of-the-art particles-tracking precision and accuracy. The visualisation of the detector is done using `Qt` (although other methods exist and have been shown to work with `Geant4`) and one can clearly observe the particles' trajectories and the box geometry. The downside of this complex infrastructure built over an extended period of time is that it is difficult to run even simple, textbook, simulations using it. For example, I needed to devise a strategy to incorporate parallel E and B fields in the same part of the "world volume", i.e. in the same sub-volume. While it was a purely computational (Application-Programming-Interface's (API's) inner workings understanding) effort, it shows that even if `Geant4` is highly efficient and performant, simple numerical simulations can be done using in-house developed codes and that they are an appealing alternative to using it. Two snapshots of what I managed to achieve and which fully recreates the physics inside a Thomson spectrometer are in Fig. 2.9 and in Fig. 2.10.

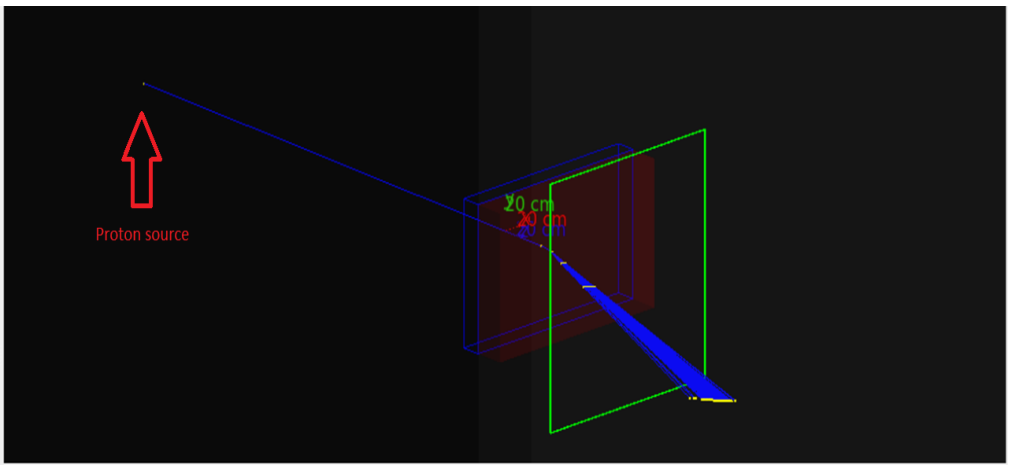


Figure 2.9: A `Geant4` [76] simulation visualisation: protons originating from a stochastic source (their kinetic energy was drawn at random from a Gaussian distribution centred at some user-defined input energy) are incident upon regions of parallel static electric and magnetic fields and are analysed at a "test plane" (pictured in green).

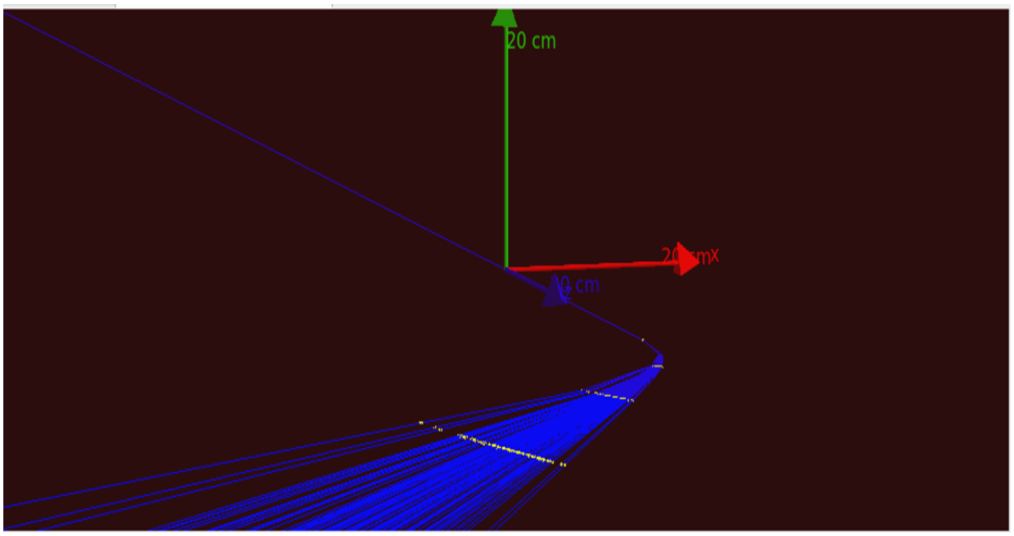


Figure 2.10: A `Geant4` [76] simulation visualisation: protons originating from a stochastic source (their kinetic energy was drawn at random from a Gaussian distribution centred at some user-defined input energy) are incident upon regions of parallel static electric and magnetic fields: their trajectories are curved and they are deflected from their initial straight paths, based on their incident kinetic energy.

The comparison between results obtained using my EOM solver and `Geant4` appear in Fig. 2.11. The differences in the positions across the parabola are due to the random nature of the incident particles' kinetic energies (these were drawn at random from a Gaussian distribution centred at a fixed kinetic energy inputted by myself), which mimic how the particles enter the detector in a real world experiment.

A reasonable estimate for the minimum number of ions of a given charge state that needs to be collected for a clear differentiation of that charge state is $N_{min} = 50$. One therefore needs to calculate the minimum initial gas density to allow N_{min} ions to be collected, taking into account realistic distances from the spectrometer's pinhole to the interaction point.

The ions are assumed to expand isotropically from the focus (confirmed by the 2D PIC simulations). A circular aperture of diameter $D = 500$ μm gives a solid angle of $\Omega = (\pi D^2/4)/R^2$ (the planar area of the aperture is used because its dimensions need to be at least one order of magnitude smaller than the source-aperture distance R based on experimental considerations). The total number of xenon ions produced in the focus is read from Fig. 2.4 for any charge state

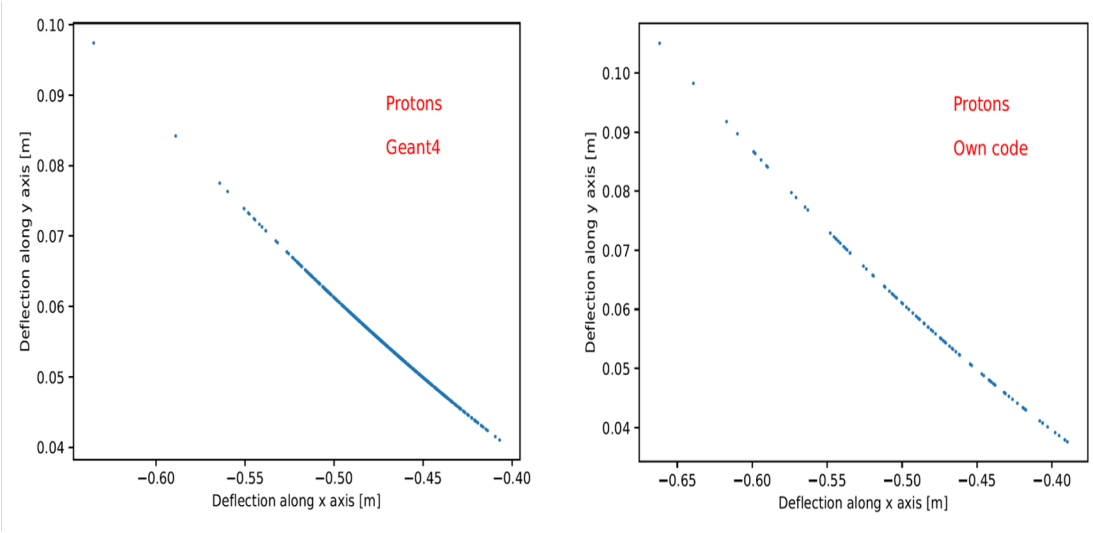


Figure 2.11: Parabolas drawn by the protons after their deflection by the Thomson spectrometer. Left panel: results obtained using the community-standard `Geant4` simulation framework. Right panel: results obtained using in-house developed codes, which are seen to be benchmarked against the community standard (i.e. the code developed at CERN). The differences in the actual positions across the parabolas are due to the random nature of the incident particles' kinetic energies.

and is denoted as N_{total}^{n+} for charge state $n+$. N_{total}^{n+} corresponds to the full 4π sr for an isotropic ejection of the ions.

Only $N_{Th} = N_{total}^{n+}\ \Omega/(4\pi)$ ions will be actually collected by the pinhole, for any $n+$. For 50 ions of each charge state to pass through the pinhole, the aperture has to be placed at around or less than $R = 0.1$ cm from the beam's focus if the focused intensity is around 10^{23} W/cm^2. This is not realisable in practice. One needs to increase the interaction point - pinhole distance. Doing so, multiple shots are needed to collect enough statistics for a good differentiation of the parabolas. Increasing D decreases the spectrometer's capability to discriminate between different charge states. While providing a higher ion flux at the detector, a very large aperture is not suitable for the envisioned experiments.

The solution for a single-shot experiment is to increase the neutral gas number density n_0. However, one needs to avoid relativistic self-focusing: it would provide an artificially high estimate of the peak intensity. The results presented so far are for PIC simulations with $n_0 = 1.98 \times 10^{12}$ cm^{-3} for Xe. I have repeated these simulations for n_0's up to three orders of magnitude higher. I observed no

self-focusing in the 2D simulations. Increasing n_0 leads to a higher number of ions produced in the focus: a single shot measurement of the peak intensity using Xe is more easily achieved. Depending on the experimental constraints, one can run additional simulations to check when self-focusing appears and choose the most suitable n_0 for the optimum source - pinhole distance.

The laser pre-pulse does not influence the results obtained via the diagnostic I presented. For a contrast ratio of $10^8 : 1$, common at state-of-the-art laser facilities, shooting a pulse at 10^{23} W/cm^2 (i.e around the highest imaginable intensity achievable nowadays), the pedestal is 10^{15} W/cm^2 and laser-field ionisation only starts to develop: only charge states from the lower end will be populated. The work from [54] demonstrated that any initial state of the target, as long as it contains non-zero populations only in the lower end of the charge states, produces the same charge state distribution after the passage of the laser pulse.

I performed the calculations using xenon (atomic number 54), but iodine (I, atomic number 53), for example, could be used as well in an actual experiment, if noble gases are not available. However, I emphasise that using noble gases is the easiest way to perform an experiment. With iodine, for example, the preliminaries are different: the main laser beam would be focused in a region of plasma located away from the actual iodine target, but originating from it. This is because the calculations are valid for very low electron number densities. Before performing the experiment with iodine as the target, the location where to focus the main laser pulse has to be pre-determined by performing numerical simulations by means of a radiation hydrodynamic simulation computer code. The code needs to dictate the location at which the number density in the iodine plasma is low enough to allow neglecting any plasma effects, in order to be sure the actual laser pulse focused there is measured, and not any spurious electric fields arising from plasma waves, for example. `FLASH4` [77] would be suitable to perform this type of hydrodynamic calculations. Then the collection of iodine ions would be performed by means of Thomson parabola detectors in a similar fashion as explained earlier.

2.2 Conclusion and outlook

This chapter presented a new computational workflow to be implemented before the practical realisation of the peak intensity diagnostic in the laboratory. Based firstly on highly efficient 1D PIC simulations, then on parallelised, computationally-efficient, numerical integration for the focal volume averaging and on very precise (Bayesian) fitting which results in an analytical prediction formula for the number of ions of each charge state, the computations required for the diagnostic can be done on a current-day personal computer.

The next chapters will delve deeper into the physics behind the laser-field induced ionisation of the noble gas atoms. A purely quantum-mechanical phenomenon, laser-field induced ionisation is most accurately studied by solving the underlying equation governing the ionisation dynamics in the non-relativistic regime: the time-dependent Schrödinger equation (TDSE). In Chapter 3, I will present how I developed from scratch a highly-performant, parallelised (distributed across CPU-cores), numerical solver for the full-dimensionality (3D) TDSE. I will then show in Chapter 4 results for how the non-dipole effects, which are deemed to appear both in the ultra-short wavelength regime (X-ray pulses of light from, for example, EuXFEL [78]) and in the long-wavelength regime (> 800 nm) at ultra-high intensities [79], are changing (increasing) the ionisation probabilities of the atom. I will also show how the TDSE solver has been used to obtain photo-electron momentum distributions (PMD's) in various types of laser fields.

In the end, Chapter 5 will be devoted to full-electron Time-Dependent Density Functional Theory (TDDFT) studies of the ionisation of an all-electron argon atom, a mean field-approach which makes the many-body problem feasible to be simulated numerically.

Conclusions are presented in Chapter 6, together with potential future work arising from this thesis.

3

Atomic ionisation modelled through accurate numerical solutions of the time-dependent Schrödinger equation

"In order to solve a physics problem, you need to spend some of your life in the company of it . . . "

— Rodica Luca in "We learn physics by solving problems / Invatam Fizica rezolvand probleme"

Contents

3.1 Introduction 53
3.2 Review 54
3.3 Theoretical considerations 57
3.3.1 Spectrum of the field-free Hamiltonian 58
3.3.2 Time-propagation 60
3.3.3 Photoelectron-momentum-distribution at the end of the simulation 63
3.4 Implementation 66
3.4.1 The "serial" workflow: parallelism exploited only inside one NUMA-region 66
3.4.2 MPI - OpenMP paralellisation strategy 74
3.5 Benchmark calculations 86
3.5.1 Three-dimensional problem-setting 86
3.6 Results and discussion: the parallel-scaling and performance of the solver 91
3.6.1 The improvement in the timings per one timestep as more computational resources are used 91

3.6.2 Serial-performance and insights 98
3.6.3 Workload-distribution across the CPU-cores 104
3.7 Conclusion and outlook 106

3.1 Introduction

Above-threshold-ionisation (ATI) [1] and the high-order-harmonic-generation (HHG) [4, 7, 80, 81] attracted a lot of interest over the last decades. Experimental and theoretical effort has been invested towards obtaining a better understanding of the quantum-behaviour of the electrons in these complex processes. All this has been made possible due to advances in the laser technology which enabled physicists to routinely generate short and intense pulses of light in a laboratory setting.

The data collected from such laser-matter experiments is usually compared with theoretical predictions. As a result of the steady improvement in the resolution of the instruments and detectors, the numerical results obtained from considering and solving the governing physical laws needed to keep up with this ascending trend.

At the most fundamental level, one has to solve the time-dependent-Schrödinger-equation (TDSE) in order to fully grasp and understand the observed phenomena. For any other system than the hydrogen atom, this is a computationally intractable problem to be solved in full-dimensionality: when considering all the degrees of freedom which exist in reality, 'the curse of dimensionality' prevails. As an example, to correctly obtain the full ground-state configuration (thus no laser-field present at all) of a single water molecule H_2O (a system with 10 electrons) and assuming each coordinate from the real three-dimensional (3D) space is sampled with a very modest number of 20 grid-points, one would require to store at least $20^{3\times10}$ values in memory: this is far from being achievable now and in the foreseeable future. Even fully-correctly solving the excitation and ionisation of the two-electron helium atom is still not a trivial task for modern

workstations or supercomputers. Thus, much of the recent computational effort is devoted to solving the TDSE in the single-electron-approximation (SAE).

Another route to obtain predictions of the quantum-mechanical observables is to make use of approximate theoretical methods such as the strong-field-approximation (SFA) (i.e., the Keldysh-Faisal-Reiss (KFR)-type theories) [57, 82, 83] or the semi-classical rescaterring model (commonly known as "the three-step model" [9] and detailed in Chapter 1). While being fast to compute, each of them neglects at least some parts of the physics of the interaction. Experiments can measure fully-differential photoelectron-momentum-distributions (PMD's), but a fully-quantitative agreement with this data is only possible via solving the TDSE. Thus there is a great incentive to make such TDSE numerical solutions faster and more accurate.

This chapter is devoted to presenting my efforts in this SAE-direction: I implemented, then fully-distributed across CPU-cores from multiple nodes the Generalized-Pseudospectral-Grids (GPS) 3D-TDSE solver. To best of my knowledge, designing an algorithm for its parallelisation has not appeared in the literature before, not to mention the advanced and very detailed performance measurements which I will present. Firstly, I checked its correctness by benck-marking its predictions against results (mainly fully differential PMD's) from the literature. Secondly, I analysed its computational performance and scalability. The results from this chapter show that it enables one to carry out studies on ATI processes in reasonable wall-clock-times.

3.2 Review

Over time, various numerical spatial discretisations and time-stepping algorithms for TDSE have been proposed and used in intense laser-atom or molecule interaction studies.

The split-operator approach to predict quantum dynamics by means of numerically solving TDSE [84–86] enabled numerical schemes based upon switching

between the two representations of the wavefunction (real space and the momentum space), each of which is suited for the application of specific quantum operators, to be used in theoretical and computational studies, e.g. [87, 88].

The radial coordinate can be discretised using a set of functions particularly suited to describe both bound and continuum states: the B-splines [16], which coupled with the, for example, Crank-Nicolson (Cayley form) time-stepping [89] together with the split-operator approach, makes it feasible to obtain accurate results of the evolution of the wavepacket.

Without using the split-operator, B-splines have found applications in highly-performant, parallelised solvers based on the Cayley propagator alone, at each timestep a preconditioner being used before solving the set of coupled-linear equations by means of the generalised minimal residual algorithm [90].

The matrix iteration method [91] again does not employ operator-splitting and intelligently avoids the direct inversion of the potentially too-large-to-be-efficient matrix involved in the Crank-Nicolson stepping, providing a way to accurately treat HHG and ATI in a full-dimensionality TDSE simulation.

An improved Lanczos propagator, together with the finite element (FEM) discrete-variable-representation (DVR) (FEM-DVR) for the radial coordinate, is employed in Ref. [92]: it proves that this method is competitive for long-wavelength (large angular momenta involved in the simulation) laser pulses interacting with matter.

Based on the pioneering work of Muller [93], the solver from [94] builds a simulation framework for which the cost of obtaining a full dimensionality TDSE numerical solution scales linearly with the problem size.

It is also to be mentioned that a highly mature solver which can also take into account the multi-electron effects present in atoms different than hydrogen is QPROP [95–97].

A very recently published (only intra-node) parallelised TDSE solver which has gained considerable traction is QPC-TDSE [98] due to its easily extendable nature, as it has been written under the philosophy of going full object-

oriented-programming (OOP), as per the canonical way to produce modern C++ programming style software.

The split-operator scheme in the energy representation [99, 100] proposes again an alternating approach, but in another, innovative way: the propagation with the laser-field-free part of the Hamiltonian is done in the spectral (coefficients) space and with the laser-field part in the grid (real space) representation. It uses the GPS discretisation [101, 102] and it has been widely used in studies of light-atom [103–110] and light-molecule [111–115] interaction, potentially being computationally more efficient than more traditional-schemes based, e.g., on finite differences [116]. Being a DVR-type discretisation, the GPS-discretisation shows exponential convergence, rather than polynomial convergence as seen in, e.g., the finite difference schemes.

I have thought of, and developed from scratch, a high-performance, fully-parallelised (`MPI + OpenMP`) CPU-only TDSE-solver which uses the GPS discretisation and is written in `C++`.

The rest of this chapter is organised as follows. I first present in Sec. 3.3 the theoretical treatment of the TDSE in the GPS formalism, then also show the theory behind the PMD-retrieval algorithm by means of direct projection of the time-dependent wavefunction at the end of the simulation onto the scattering states of the atom. Sec. 3.4 then details the choices made to implement the discretised partial differential equation's (PDE)'s matrix-vector operations which are at the core of the time-stepping algorithm. That is, in this section the parallelisation strategy is discussed, together with the communication patterns which happen between the `MPI`-ranks. After this, in Sec. 3.5 some benchmark, well known in the literature, numerical-calculations are presented. The results part of the chapter is in Sec. 3.6. Here, timings for the various parts of the iterative time-stepping algorithm are illustrated, in the weak-scaling regime but also in the strong-scaling regime. The roofline-model, as obtained using the `likwid` high-performance-computing (HPC) tool [117], is shown and discussed. This enables one to think about potential ways to speed up the serial execution of the solver

(where serial in this context means its execution within 1 rank only). Finally, this chapter ends up with diagnostic-plots obtained using `Score-P` [118] which target the workload on the processors over the duration of the execution of one simulation. Unless otherwise stated, atomic units are used throughout this chapter ($e = m_e = \hbar = 1/(4\pi\epsilon_0) = 1$). A summary is provided at the end of the chapter, in Sec. 3.7.

3.3 Theoretical considerations

The TDSE is solved in the length-gauge for the hydrogen atom in a spherical coordinates geometry. Generalisations for other atoms treated in the SAE approximation are straightforward.

$$i\frac{\partial \Psi(\vec{r},t)}{\partial t} = \mathcal{H}\,\Psi(\vec{r},t) \tag{3.1}$$

Full dimensionality (i.e. a 3D real-space) is assumed from the start: $\vec{r} = [r,\theta,\varphi]^T$, where ^{T} means the transpose operation on a vector (or matrix).

The time-dependent part of $\mathcal{H}$ is the coupling with the laser electrical-field (E-field) expressed here in the length gauge $V(\vec{r},t) = \vec{r}\cdot\vec{E}(t)$.

The time-independent part $\mathcal{H}_0$ represents the field-free (unperturbed) Hamiltonian and is composed from: 1) the attraction with the core ($V_C(r) = -1/r$ for atomic hydrogen, for other atoms pseudopotentials could be used) and 2) the centrifugal part and the radial, second order derivative (together in the $-\nabla^2/2$).

In the dipole-approximation (DA), the laser E-field is a function only of the time-parameter: $\vec{E} = \vec{E}(t)$. For such a laser pulse, the form of $\vec{E}(t)$ is either linearly-polarised (along $\hat{z}$-axis and the problem becomes two-dimensional (2D): there will be no φ-angle dependence because the system possesses rotational symmetry around the $\hat{z}$-axis) or elliptically-polarised with ellipticity $\epsilon \in (0,1]$ (and the problem is known to become fully 3D).

3.3.1 Spectrum of the field-free Hamiltonian

Before the start of each TDSE simulation, the electronic structure of the atom is firstly solved for by diagonalising the field-free Hamiltonian discretised using the GPS method. The reason for this is that the eigenvectors and eigenvalues of the field-free Hamiltonian are at the core of the iterative algorithm which relates two consecutive timesteps in the time-integration of the TDSE.

The time-independent-Schrödinger-equation (TISE) (no laser pulse electromagnetic field) is solved in spherical coordinates (considering the radial direction only as the atomic-core potential is spherically symmetric) for the atom's spectrum: the eigenvalues (eigenenergies) ϵ_{nl} and eigenvectors ψ_{nl}, enumerated by the index n, for each partial wave l, i.e. each of the $\mathcal{H}_0^l$, are obtained. This gives a complete basis-set within the simulation-box: both bound (eigenenergy < 0 a.u.) and (pseudo-) continuum (eigenenergy > 0 a.u.) states are calculated. It might be the case that the last (psuedo-) continuum states obtained have very high energies, depending on the laser parameters. Such parts of the spectrum are very unlikely to be populated in the time-dependent simulation, again this depending on the laser-field parameters. Discarding these highly-energetic states does not influence the results (the dynamics) in a significant way and thus it is safe to do so. This is helpful when the numerical stabilities of the simulations are low, a problem found especially in TDDFT-simulations, but not that much in TDSE-simulations. Not including these highly-energetic states in the radial propagators which will be detailed later can be seen to improve the stability of the TDDFT-simulations. I have not found many TDSE simulations in which neglecting these states has a significant effect.

Thus both the excitation and ionisation phenomena can be treated in the simulation by solving:

$$\mathcal{H}_0^l \, \psi_{nl}(r) = \epsilon_{nl} \, \psi_{nl}(r), \tag{3.2}$$

where

$$\mathcal{H}_0^l = -\frac{1}{2}\frac{\partial^2}{\partial r^2} + \frac{l\,(l+1)}{2r^2} - \frac{1}{r}. \tag{3.3}$$

The symmetrised and discretised matrix $\underline{\underline{\mathcal{H}_0^l}}$, having real entries (double-precision values), is diagonalisable using, for example, the DSYEV subroutine from LAPACK [119]). Hints on how to deduce the form of this matrix are given in, e.g. [99, 103, 116] and later in this chapter will be provided in full.

The GPS method allows for a non-uniform, optimal discretisation of the radial coordinate r. This is important: for a fixed memory size and compute time, one almost always wishes for a fine discretisation next to the core where its attraction is the strongest, and a way coarser discretisation a large distance away from the core, where its attraction has decreased as dictated by its dependence on the radial coordinate. This non-uniform discretisation of the radial-coordinate is achieved by means of an (e.g. algebraic) mapping [102]:

$$r\,(X) = L\,\frac{1+X}{1-X+\alpha}, \tag{3.4}$$

where L and α are mapping parameters, with $\alpha = \frac{2L}{R_\mathrm{b}}$, where R_b is the radial coordinate value at the outer end of the simulation domain. The choice of a capital letter X has been made to showcase the difference with respect to the x-coordinate from the real 3D-space. For the algebraic mapping, the X-grid points are related to the radial-coordinate $r = \sqrt{x^2+y^2+z^2}$ by means of Eq. (3.4) presented above. The X-grid points are the roots of the derivative of the Legendre polynomial of order equal to the number of grid-points used in the discretisation minus one (that is, if $N_r + 2$ points have been used, then there are N_r zeroes of $P'_{N_r+1}(X)$). For a fixed number of such grid-points and a fixed outer-boundary location, L helps modifying the distribution of these radial-grid points. This is important for large box sizes and extremely large number of grid-points, where the sampling of the radial coordinate is almost exclusively done near to the core and the very far continuum is not very finely discretised. This can pose problems with the absorber function (discussed below, in the temporal-propagation part of the solver), which will have non-trivial (non-zero) elements

only over very few grid points from the end of the simulation box. Such a mask will create artificial reflections if it is not smooth enough.

These preliminary calculations are time-independent and have to be performed only once before the beginning of the time-dependent simulation. They are usually extremely fast to compute inside 1 Non-Uniform-Memory-Access (NUMA) region only. If multiple time-dependent simulations with slightly different physical parameters (laser frequency, peak-strength and duration) are performed, the eigenvectors and eigenvalues can be safely reused between them.

It is usually the case that the outer boundary of the simulation domain is chosen to be significantly larger than the classical excursion length of the electron in the laser field (which is equal to $E_{\max}/\omega^2$, where $E_{\max}$ is the peak value of the laser E-field and ω is the frequency of the fast oscillations of the laser field). This way all the dynamics can be treated accurately because the artificial reflections from the boundary are kept small, or the wavefunction is slightly absorbed at the wall if a mask function [120] is used. The time-dependent wavefunction consists of contributions from various momenta $k = \sqrt{2E_{\mathrm{K}}}$, where E_{K} represents the kinetic energy of the electron. It is a simple calculation, for a fixed box size R_{b}, to see what is the contribution having the maximum momentum which is kept inside the box up until the end of the time-dependent simulation.

3.3.2 Time-propagation

At $t = 0$, the time dependent wavefunction $\Psi(\vec{r}, t = 0)$ is prepared as the ground state of the hydrogen atom (usually, for reasonably chosen discretisation parameters, the numerically-obtained ground state has its eigenenergy very close (almost to machine precision) to the ionisation potential of atomic hydrogen, $I_{\mathrm{p}} = -0.5$ a.u).

For a general 3D problem, the time-dependent wavefunction is expanded in spherical harmonics [121] (for a 2D problem, the φ-dependence is trivial and the expansion is done using the Legendre polynomials) at each timestep of

the simulation:

$$\Psi\left(\vec{r},t\right)=\sum_{l=0}^{l_{\max}}\sum_{m}\frac{R_{l}^{m}(r,t)}{r}\,Y_{l}^{m}\left(\theta,\varphi\right), \tag{3.5}$$

where $l_{\max}$ is the maximum angular momentum included in the simulation. How many angular momentum values to use in the simulation has to be chosen according to the physics at hand. One can estimate from ATI-related considerations how many such partial waves to include, depending on how many ATI peaks one needs to resolve. From the energy conservation perspective, the number of photons g (an integer) which contribute to the process can be estimated as: $g=(E_f+U_p-E_i)/\omega$, where E_f is the final energy of the electron, $U_p=I/\omega^2$ is the ponderomotive energy of the electron in the laser field, with I being the light-field intensity, and E_f is the final energy of the electron. For the first ATI-peak to be resolved, one has to use at least $l_{\max}\approx(\omega+U_p-E_i)/\omega$. In practice I have taken many more such partial-waves for full-convergence. Similarly, for the j-th peak to be resolved, at least $l_{\max}\approx(j\omega+U_p-E_i)/\omega$ has to be in place, in practice much more than this being actually needed: often even one order of magnitude more.

The summation over the projection of the angular momentum along the $\hat{z}$ axis, that is the sum over the magnetic quantum number m, does not necessarily have to go from $-l$ to l, but one can truncate this expansion sooner.

The short-time propagator

$$U\left(t,\Delta t\right)=\exp\left\{-i\Delta t\left[\mathcal{H}_0+V\left(t+\Delta t/2\right)\right]\right\} \tag{3.6}$$

is used (Δt is the timestep of the simulation and is fixed throughout the whole simulation, i.e. no adaptive-time-stepping is used) and is approximated using the second-order split-operator formula:

$$\begin{aligned}U\left(t,\Delta t\right)=&\exp\left[-i\frac{\Delta t}{2}\mathcal{H}_0\right]\exp\left[-i\Delta tV\left(t+\frac{\Delta t}{2}\right)\right]\\&\exp\left[-i\frac{\Delta t}{2}\mathcal{H}_0\right].\end{aligned} \tag{3.7}$$

The numerical error of this splitting is bound from above by $(\Delta t)^3$, this quantity being multiplied by some scalar prefactors. From these prefactors, it is important to note that this upper bound also contains the E-field strength squared. Thus it is the case that for higher field strengths, the timestep to be set for the simulation has to be smaller relative to the timestep of the same simulation run with a lower field strength. This is the explanation (for the splitting scheme used in my work) for the well-known fact that higher-field strength TDSE simulations are more difficult to fully-converge. The field-free propagator $\exp\left[-i(\Delta t/2)\mathcal{H}_0\right]$ is time-independent (calculated as explained above using a diagonalisation routine) and each partial wave is propagated individually under its action [99]:

$$\exp\left[-i(\Delta t/2)\mathcal{H}_0\right] = \sum_{l=0}^{l_{\max}} \exp\left[-i(\Delta t/2)\mathcal{H}_0^l\right], \tag{3.8}$$

with

$$\exp\left[-i(\Delta t/2)\mathcal{H}_0^l\right] = \sum_{n=0} \exp\left[-i(\Delta t/2)\,\epsilon_{nl}\right]\,|\psi_{nl}\rangle\,\langle\psi_{nl}|\,. \tag{3.9}$$

The two parts of the evolution under the field-free Hamiltonian $\mathcal{H}_0$ are best performed on the spectral-representation of the time-dependent wavefunction.

However, the part of the evolution under the laser-field only is performed on the grid representation of the time-dependent wavefunction: any operator which is a multiplication operator is diagonal in the GPS method and thus is a computationally cheap array-array element-wise multiplication. That is, the $\exp\left[-i\Delta t V(t+\frac{\Delta t}{2})\right]$ part is performed on the grid-representation. The grid representation is obtained by means of a spherical harmonics transform operation.

Of course, at each timestep, an inverse spherical harmonics transform is needed to go back to the previous representation (that is, to change from the grid representation used for applying the laser field to the spectral representation used for applying the radial propagators).

The propagator from Eq. (3.7) is applied iteratively on the time-dependent wavefunction up until the end of the laser-pulse.

At the beginning of each timestep, the norm of the time-dependent wavefunction is calculated from its spectral representation. This operation is just

a simple summation over the modulus-squared of each data-element of this representation. At the beginning of the time-dependent simulation, this has to be close to 1, to machine-precision. Over the course of the simulation, an absorber (mask-function) is applied at each timestep. Its form is given by:

$$mask(r) = \begin{cases} \cos^{0.25}\left(\frac{\pi}{2}\frac{(r-R_a)}{(R_b-R_a)}\right) & \text{, if } R_a < r < R_b \\ 1 & \text{, if } 0 < r < R_a. \end{cases} \tag{3.10}$$

This means that if the time-dependent wavefunction is dragged enough (i.e. beyond R_a) by the laser-field towards the boundary of the computational domain, its norm will decrease according to the results obtaining from the multiplication of Eq. (3.10) with the grid-representation of the time-dependent wavefunction.

3.3.3 Photoelectron-momentum-distribution at the end of the simulation

It is known from scattering theory [122] that it is needed to project the time-dependent wavefunction onto the target-atom's continuum scattering states with correct boundary-conditions in order to retrieve the three-dimensional photoelectron-momentum-distribution (PMD). The projection is done onto states which have plane waves and incoming spherical waves in asymptotic forms at large distances [123]. This is routinely performed only for relatively-simple systems: the scattering states are available in analytical form only for the atomic Coulomb potential. For other potentials for the one-centre system, numerical solutions for these scattering states are needed. For systems different than atoms, but again having one active electron (say, 2-centre systems, such as the molecular ion H_2^+), the situation is even more complicated and the efforts to obtain these states are usually more significant.

There exists other methods to extract the PMD and considerable effort has been devoted over the years to improve them, in order to avoid calculating continuum states. The window-operator method is the first to be mentioned. It extracts contributions from the time-dependent wavefunction at the end of the simulation which are included in a bin of a given width in the energy-space.

However, by avoiding the direct projection, the correct boundary conditions for the scattering states are not included (not defined at all) in this method [123].

Partitioning the space into two regions and analytically propagating the time-dependent wavefunction under the influence of the Volkov-Hamiltonian in the second region has been also considered in great detail [124]. While many times very accurate, the PMD's obtained with the help of this method differ depending on where the separation into two parts is done and this can leave room for ambiguity. Also, this separation has to be done at a rather large distance from the core because the Volkov-Hamiltonian has as its fundamental characteristic the fact that it neglects the core's attraction. Provided that the TDSE-propagation method scales across processors, then this problem can be easily circumvented by adding more computational resources to the simulation. Nevertheless, it is clear that it will eventually become hard to accurately sample the real-space very far away from the origin of the coordinate system in the GPS-formalism.

The t-SURF [96] approach makes use of the wavefunction as it passes through an absorbing boundary (surface). It thus changes the projection onto the scattering states to a time-integral of the flux of the wavefunction through the wall. The i-SURF method [97] is an evolved version of t-SURF: it computes PMD's in a more computationally-efficient manner than the original.

In the current version of the solver presented in this chapter, only the projection onto the scattering states of atomic hydrogen is presented. This approach is then applied to studies of this atom in high-frequency laser-fields. It can also be adapted to other hydrogen-like ions (i.e. any atom having Z-protons and having had $Z-1$ electrons already ionised, for example He^+).

The continuum states are considered in a partial-waves expansion:

$$\psi_{\vec{k}}(\vec{r}) = \sqrt{\frac{2}{\pi}}\frac{1}{k}\sum_{l=0}^{l_{\max}}\sum_{m=-l}^{l} i^l e^{-i\Delta_l}\frac{y_l(k,r)}{r}Y_l^m(\Omega)\,Y_l^{m*}(\Omega_{\vec{k}}), \tag{3.11}$$

where i is the imaginary unit, Δ_l is the scattering phase-shift (here only equal to the Coulomb-phase shift $\sigma_l = \arg\Gamma(l+1-\frac{i}{k})$ with Γ being the gamma-function [121] and defined in the previous chapter) and $y_l(k,r)$ satisfying the

TISE and being the regular Coulomb function [121], $y_l(k,r) = F_l(\eta_C, kr)$, with $\eta_C := -Z/k = -1/k$ the Sommerfeld parameter. The spherical harmonics are functions of $\Omega = (\theta, \varphi)$ and $\Omega_{\vec{k}} = (\theta_k, \varphi_k)$.

The scattering wavefunctions are obtained using the `Fortran` routines from [125] and the phase-shifts are retrieved using `GSL` [1]. I compiled and run the workflow from [125] at the input values of interest for the physics I investigate.

Eq. (3.11) presents scattering states normalised on the momentum-grid, that is:

$$\left\langle \psi_{\vec{k}}(\vec{r}) \middle| \psi_{\vec{k}'}(\vec{r}) \right\rangle = \delta(\vec{k} - \vec{k}'). \tag{3.12}$$

Upon a careful investigation of the source codes (and the text of [125] itself), it is visible that the routines actually produce Landau-normalised scattering states (and not momentum normalised scattering states). A postprocessing of them is thus required if one wants to compare PMD's obtained using them against other results from the literature. All these are conventions which vary upon authors, and in this case it is all about a multiplicative factor of 2 (which, when calculating mod-probabilities squared, becomes a multiplicative factor of 4), and has to be taken into account if quantitative comparisons are needed.

Then the (one-dimensional radial-) projection onto the required scattering states is performed for each radial slice from the spectral-representation of the time-dependent wavefunction at the end of the TDSE simulation:

$$b_{l,m}(k) = \sqrt{\frac{2}{\pi}} \frac{1}{k} (-i)^l e^{i\Delta_l} \left\langle y_l(k,r) \middle| f_l^m(r, t = T_f) \right\rangle, \tag{3.13}$$

where $T_f = T$ and represents the duration of the laser pulse. The triply-differential PMD is then obtained as a forward spherical-harmonics transform:

$$\frac{dP}{k^2\, dk\, d\Omega_{\vec{k}}} = \left| \sum_l \sum_m b_{l,m}(k) Y_l^m(\Omega_{\vec{k}}) \right|^2. \tag{3.14}$$

This can be integrated over the angles θ_k and φ_k and the result, which is given in the literature as the photoelectron-energy-spectrum (PES), is obtained as follows:

$$\frac{dP}{dE_K} = \sqrt{2E_K} \sum_l \sum_m \left| b_{l,m}\left(\sqrt{2E_K}\right) \right|^2, \tag{3.15}$$

[1] GSL - GNU Scientific Library, `https://www.gnu.org/software/gsl/`

where the kinetic energy of the electron E_K is related to its momentum k as $E_K = k^2/2$.

3.4 Implementation

3.4.1 The "serial" workflow: parallelism exploited only inside one NUMA-region

Preliminary-calculations

In this subsection I describe how one obtains the ingredients for the radial propagators as from the diagonalisation of the field-free Hamiltonians, for each partial wave l.

Once the number of radial grid-points N_r is chosen, the Gauss-Lobatto nodes and weights are calculated using the `Python` functions from [126].

For values of N_r below (say) 2560, it only takes a few seconds on a machine equipped with an Intel® Core(TM) i7-10700 CPU @2.90 GHz chip to obtain the results.

For a Gauss-Lobatto quadrature with a total of $N_r + 2$ points, the two end-nodes are -1 and +1. The other N_r nodes are the roots of the derivative of the Legendre polynomial (denoted here as P) of order $N_r + 1$ (thus a polynomial of order N_r, having N_r zeroes), i.e. the roots of $P'_{N_r+1}(x)$. The Gauss-Lobatto weights (used at the end of the time-dependent simulation, for the photoelectron-momentum-distribution retrieval workflow) are given by:

$$w_i = \begin{cases} \dfrac{2}{(N_r+1)(N_r+2)} & \text{, if } x_i = \pm 1 \\ \dfrac{2}{(N_r+1)(N_r+2)\,(P_{N_r+1}(x_i))^2} & \text{, if } x_i \neq \pm 1 \end{cases} \tag{3.16}$$

Then the population of the field-free Hamiltonians (one for each partial wave l) is started. The same workflow is applied to each integer $l \in [0, ..., l_{max}]$, and these workflows are independent of each other. Each such $\underline{\underline{H_0^l}}$ is shaped (N_r, N_r) and contains along its (main) diagonal the elements:

$$\left(\underline{\underline{H_0^l}}\right)_{(i,i)} = \frac{(N_r+1)(N_r+2)}{6} \frac{1}{(r'(x_i))^2\,(1-x_i^2)} + \frac{l(l+1)}{2(r(x_i))^2} - \frac{1}{r(x_i)}, \tag{3.17}$$

for $i \in [0, ..., N_r - 1]$, for a fixed $l \in [0, ..., l_{max}]$, and where $r'(x_i)$ represents the derivative of the algebraic-mapping from Eq. (3.4), at the Gauss-Lobatto collocation point (the nodes for this type of quadrature) x_i.

The non-diagonal elements of the field-free Hamiltonian matrix are populated as:

$$\left(\underline{\underline{H_0^l}}\right)_{(i,j)} = \frac{1}{r'(x_i)} \frac{1}{r'(x_j)} \frac{1}{(x_i - x_j)^2} \ , \tag{3.18}$$

for $i \neq j$ and $i, j \in [0, ..., N_r - 1]$ and for a fixed $l \in [0, ..., l_{max}]$.

The `scipy.linalg.lapack` package from `Python` is used to invoke DSYEV and the $(l_{max} + 1)$ real, symmetric, matrices $\underline{\underline{H_0^l}}$ with $l \in [0, ..., l_{max}]$ are diagonalised to obtain their eigenvalues and eigenvectors, which are stored to disk and will later be input to the time-dependent simulation.

The wall-clock-time spent on this part of the preliminary calculations can become significant if large angular momenta values are needed (above, say, 1000) and a fine discretisation of the radial coordinate is required (above, say, 2560 grid-points). However, multiple threads can be used (i.e. the DSYEV to benefit from multiple threads) if this is the case. Even further parallelisation can be implemented if needed: `MPI` is used and each rank processes a given subset of l-values, as the diagonalisation of each partial wave field-free Hamiltonian is independent of the other such operations (for the other l-values). Inside each rank, the diagonalisation routine will benefit from using multiple threads, so resources are used in the most-optimum way.

The radial propagators are built at the beginning of the time-dependent simulation from the eigenvalues and eigenvectors of the field-free Hamiltonians for each partial wave l. This is a time-independent operation, and thus this assembly of the radial propagators is done only once and after they are created, they are kept in memory and used at every timestep. The eigenvectors and eigenvalues are loaded from memory using the `xtensor`-suite [2]. This framework is appealing since it allows `Numpy`-like manipulation of multi-dimensional arrays

[2] `xtensor is a C++| library meant for numerical analysis with multi-dimensional array expressions. See: https://xtensor.readthedocs.io/en/latest/index.html`

in a C++ program, but also loading in a C++ program any .npy-array stored on disk from a Python program (which is particularly useful here as the diagonalisation of the field-free Hamiltonian $\mathcal{H}_0$ is performed in Python).

Each radial propagator is the result of the multiplication of three matrices, denoted here as $\underline{\underline{V}}$, $\underline{\underline{D}}$, $\underline{\underline{V}}^T$, populated with different values for different l-partial waves.

For any l, $\underline{\underline{V}}$ is composed from the eigenvectors of $\underline{\underline{H_0^l}}$ as obtained from DSYEV.

For any l, $\underline{\underline{D}}$ is a diagonal matrix, having at its f diagonal element the value $\exp\left(-I\,\epsilon_f^l\,\Delta t/2\right)$, where I is the imaginary unit, i.e. $I^2 = -1$, and ϵ_f^l is the f-th eigenvalue of the field-free Hamiltonian for the partial wave l. f can range between 0 and an upper value equal to or less than $N_r - 1$. While choosing $N_r - 1$ includes all the physics, to improve the numerical stability of the time-dependent simulation, it is possible to chose a lower value than $N_r - 1$, say Q, with $Q > 0$. Then one would also ignore the last $(N_r - Q)$ eigenvectors, thus the shapes of the matrices would be as follows: for $\underline{\underline{V}}$ the shape would be (N_r, Q) and for $\underline{\underline{D}}$ the shape would be (Q, Q).

Irrespective to chosen value of Q, the product of the three matrices denoted here as $\underline{\underline{A}}^l$ is always of shape (N_r, N_r) and equals to the radial propagator for the partial wave l.

The radial propagators are stored in a long one-dimensional, dynamically allocated C-array of length $N_r \cdot N_r \cdot (l_{max} + 1)$, and having elements of the `std::complex<double>` type.

Time-propagation

A discretised wavefunction is introduced in the computer-code and propagated. In both its spectral and grid representations, the time-dependent wavefunction is represented as an array expressed as a C++ vector. This choice of the C++ vector over the classical C-array does not influence in any negative way the interaction with the MPI-library or cblas-library described later, as the elements are guaranteed to be contiguous in memory. The C++ vector also allows for

out-of-bounds checks at runtime, at the cost of a very modest degradation in performance in my studies, which is also compensated by the appealing feature that they go out of scope automatically, thus eliminating the potential leaks appearing due to forgetting to `delete` memory allocated with `new`.

Thus a long `std::vector<std::complex<double>>` holds the time-dependent wavefunction in its spectral representation: $f_l^m(X,t)$ for all l, m, at any X and at the current time t.

Another long `std::vector<std::complex<double>>` holds the time-dependent wavefunction in its grid representation: $f(X,\theta,\varphi,t)$, for all θ and φ points, at any X and at the current time t.

The number of radial grid-points is denoted as N_r. The number of angular θ-points is denoted as $nlat$. The number of azimuthal φ-points (physics-convention) is denoted as $nphi$. The number of elements of the grid representation is equal to $(N_r \cdot nlat \cdot nphi)$.

The spherical harmonics transform involves integration against the angular and azimuthal coordinates.

For the θ-integral, Gauss-Legendre quadrature is used [127] and in the current implementation the total number of l-values (i.e. $l_{max}+1$) is set to be equal to the number of θ-points. That is $nlat = l_{max}+1$.

For the φ-integral, Fast-Fourier-Transform (FFT) is used [40].

An unsigned-integer $mmax$ is chosen at the beginning of the time-dependent simulation to represent the maximum value which m (the angular momentum projection quantum number) can take. To include all m-channels, $mmax = l_{max}$, where l_{max} is fixed in the input-deck and independent of $mmax$. To save on computational resources, $mmax$ can be chosen as $mmax < l_{max}$, but shall be always the case that $mmax \geq 0$. If $mmax < l_{max}$, then for any l, the expansion in spherical harmonics (explained below) is done for m-values going from $-(\min\{mmax, l\})$ up until $+(\min\{mmax, l\})$ only.

The number of φ-points $nphi$ is equal to $2mmax+1$, and it is thus the case that a one-to-one-mapping between m-values and φ-points is in place.

The change between the grid representation and the spectral representation is done as follows [127]:

$$f_l^m(X,t) = \int_{\theta=0}^{\theta=\pi} \int_{\varphi=0}^{\varphi=2\pi} f(X,\theta,\varphi,t)\ Y_l^m(\theta,\varphi) \sin\theta\, d\theta\, d\varphi\ . \tag{3.19}$$

The change between spectral representation and the grid representation is done as follows [127]:

$$f(X,\theta,\varphi,t) = \sum_{l=0}^{l=l_{max}} \sum_{m=-m_{max}}^{m=m_{max}} f_l^m(X,t)\ Y_l^m(\theta,\varphi)\ , \tag{3.20}$$

where m_{max} is equal to $\min\{mmax, l\}$, where the signification of the unsigned integer $mmax$ has been explained above.

The field-free propagation is performed in the spectral-representation of the time-dependent wavefunction. That is, each slice of length N_r denoted as $f_l^m(X,t)$ is acted upon by the corresponding radial propagator ($\underline{\underline{A}}^l$, the radial propagator matrix for the angular momentum l). The slices for all the physically-correct m-values ($|m| \leq l_{max}$) which exist in the simulation for a given l are all acted upon by the same radial propagator (there is no reference to the m-quantum-number in $\underline{\underline{A}}^l$ or in its creation). Thus the result of this matrix-vector multiplication is an update of the $f_l^m(X,t)$ slice of the spectral-representation of the time-dependent wavefunction. The multiplication is performed with the help of a `cblas` library, in particular with its `cblas_zgemv(...)` function.

On the ARCHER2 supercomputer, I used the collection of hardware-optimised `Cray-Libsci` routines as this is a Cray® machine. I ran simulations using the solver also on the `SCARF` HPC-centre. There, I have installed myself `AMD-BLIS` and used the collection of routines from this library, as the SCARF nodes I ran the solver on were equipped with `AMD-EPYC 7502` processors. They are used either in their serial form, or on their threaded form. Usually, for a reasonably large N_r, a scalable way to use the parallelism of the HPC-cluster was to call the routines with the maximum number of threads from 1 NUMA region. For small values of N_r, this does not posses any computational-advantage with respect to the time it takes to execute (actually for a very small N_r value and as many

as 16 threads, the parallel `cblas_zgemv` call can potentially be slower than its serial counterpart), and the parallelism is used in a different way. An in-depth discussion about the choices related to how the solver is designed to make the best possible use of the multiple levels of parallelism of a given computer-architecture is delayed up until Sec. 3.4.2.

For now, the aim is to remember that the first and the second parts of the propagation in the field-free Hamiltonian $\mathcal{H}_0$, see Eq. (3.7), is achieved by matrix-vector multiplications of complex, double-precision, type. The vectors are radial-slices from the spectral representation of the time-dependent wavefunction and the matrices are the radial propagators for each partial wave l.

To apply the laser field, i.e. to perform the propagation from the middle part of Eq. (3.7), the spectral to grid transformation is undertaken. In the current implementation of the solver, this is achieved using the highly-optimised and flexible `shtns` `C`-library of Schaeffer [127, 128]. This library has support for GPU-offloading which is shown to improve its performance even further, becoming faster than the CPU-computations, which are already highly efficient. I have not made use of these advanced capabilities of `shtns`, but for very-long wavelengths studies the GPU-offloading might prove itself highly useful. As from Eq. (3.20), it is clear that the same manipulations have to be performed for each X-grid-point in order to perform the spectral to grid transform. For reasonable target-values for l_{max}, the computational-performance of the `shtns` commands is not dramatically improved when using a large number of threads when making one such call. The number of threads that `shtns` uses can be easily set at the beginning of the `C++` program. I set it to be equal to one. Thus it is the N_r transforms (each having identical code-logic/operations) which are split among threads, and each thread executes the required `shtns` commands to perform the spectral to grid transform in a serial fashion, isolated from the other threads. There are no data race-conditions because each thread updates the wavefunction stored at a different X-coordinate value. After all the threads

finished their iterations, the grid representation is obtained: $f(X, \theta, \varphi, t)$, for any X-grid point, at the current time t.

At this point, the grid-representation is ready to have the laser-field operator $\exp\left[-i\,\Delta t\,V\left(t+\frac{\Delta t}{2}\right)\right]$ applied onto it. This is a simple update of each element of the grid-representation `std::vector<std::complex<double>>`. For a grid point (r, θ, φ), the value of the product $\vec{r}\cdot\vec{E}(t+\Delta t/2)$ for the current timestep t, is multiplied by the negative of the imaginary unit, $-i$ (with $i^2 = -1$) and Δt and is stored in a temporary. Then e (where e denotes here the base of natural logarithms) is raised to the power denoted by this temporary, to be then multiplied by the corresponding value of the grid representation of the time-dependent wavefunction (the one at the same (r, θ, φ)-grid point). This repeats for all the real-space grid points (a total of $N_r \cdot nlat \cdot (2m_{max}+1)$ such points).

To perform the second part of the propagation in the field-free Hamiltonian $\mathcal{H}_0$, the grid to spectral transform has to be performed. Again, the exact same manipulations have to be applied for each X-grid point in order to get from (θ, φ) to (l, m) for that X-grid point. As mentioned above for the spectral to grid transform, for reasonable target-values for l_{max}, the computational-performance of the `shtns` commands is not dramatically improved when using a large number of threads when calling them. This is also the case here, for the grid to spectral transform. Thus here as well, each call to `shtns` uses its serial implementation. The N_r transforms are split among threads and each thread works on its own to compute the amount of work given to it. After all the threads finished their iterations, the spectral representation is obtained: $f_l^m(X, t)$, for pairs (l, m), at any X-grid point at the current time t.

The same matrix-vector multiplications are performed to propagate for the second half of the timestep in the field-free Hamiltonian $\mathcal{H}_0$.

The propagation for the current timestep has finished, but the absorber has not been applied on the time-dependent wavefunction yet. As shown in Eq. (3.10), this is a function of the radial coordinate only, and is applied on the grid-representation of the wavefunction. Another set of spherical harmonics

transforms needs to be performed: the first one gets to the grid-representation such that the mask can be applied, then the inverse transform converts back to the spectral-representation in order to be able to start the next timestep.

It will become clear later why it is instructive to estimate the computational complexity of the solver now. That is, the computational-time-complexity of one timestep of the simulation is analysed from first principles. The assumption I make for this estimate is that $mmax = l_{max}$. This is a more than reasonable assumption: many times one includes all the m-channels in the simulation. The multiplication between a square-matrix of shape (N_r, N_r) and a vector of length N_r is $O(N_r^2)$ in computation-time.

Thus, the propagation in the field-free Hamiltonian $\mathcal{H}_0$ is of time-complexity of $O(N_r^2 \cdot \texttt{total_num_of_LM_pairs})$.
It was explained above that `total_num_of_LM_pairs` is $nlat^2$ for when $mmax = l_{max}$. This means that the field-free propagation is $O(N_r^2 \cdot nlat^2)$.

The spectral to grid transformation performs the same operation N_r times. From [127], the time-complexity of `shtns` is $O(nlat^3)$. Thus the total complexity of the spectral to grid is $O(N_r \cdot nlat^3)$. The same argument holds for the grid to spectral transformation. Experimentally, it was observed by the author of `shtns` that the `shtns` operations are $O(N_r \cdot nlat^{2.7})$ due to the various optimisations made inside the library, but this is subject to various conditions on $nlat$ and $mmax$ and cannot be considered a general statement which is true in all possible cases of interest. For an estimate only, it can be safely assumed that the spherical-harmonics transforms are a fourth-order polynomial in the discretisation parameters.

The application of the laser-field is $O(N_r \cdot nlat \cdot mmax)$ and this is a third order polynomial in the discretisation parameters, similarly for the the absorber's operation on the time-dependent wavefunction.

Of course, there are prefactors present in front of these products between the discretisation parameters, for an exact consideration of all the floating point opera-

tions which come into play in the solver. They might prove themselves significant when the values of the discretisation parameters fall within certain ranges.

In total, the time-complexity of the GPS-TDSE solver can be estimated as $O(N_r^2 \cdot nlat^2 + N_r \cdot nlat^3)$ + lower-order-terms (l.o.t.).

A pictorial representation of how the various parts of the workflow connect for generating the solution to the physics problem which the described framework aims to solve is in Fig. 3.1.

3.4.2 MPI - OpenMP paralellisation strategy

For the field-free propagation it is necessary to operate on radial slices of the spectral representation of the time-dependent wavefunction. In the computer-memory, these blocks having N_r elements each are stored contiguously, in ascending order of l-values, for all valid m-values of that l-value. Each block belongs to a given l-value and each family of blocks from an l-value is propagated individually with the help of the propagator for that l-value. This is because the propagation in the field-free Hamiltonian does not depend on the m-quantum number.

It is then natural to consider how to distribute the l-values and the associated m-values across the `MPI`-ranks, in a way which will be discussed shortly. Each rank must then operate only on the radial slices belonging to its received l-values and m-values, so its needs access only to the propagators for some l-values (thus unless not using one rank only, each rank requires access to fewer radial propagators than the total number of l-values). There is no need to hold all the field-free Hamiltonians $\mathcal{H}_0^l$'s eigenvectors and eigenvalues in each rank's memory, but only those which are needed by each rank.

A similar argument is valid for the spectral representation of the time-dependent wavefunction as well: each rank needs to hold only those radial slices of the wavefunction belonging to the l-values and m-values which the program distributes to it at the start of the simulation.

For the spectral representation of the time-dependent wavefunction, the buffer of memory (an array) holding it contains $N_r \cdot$ `total_#_of_LM_pairs`

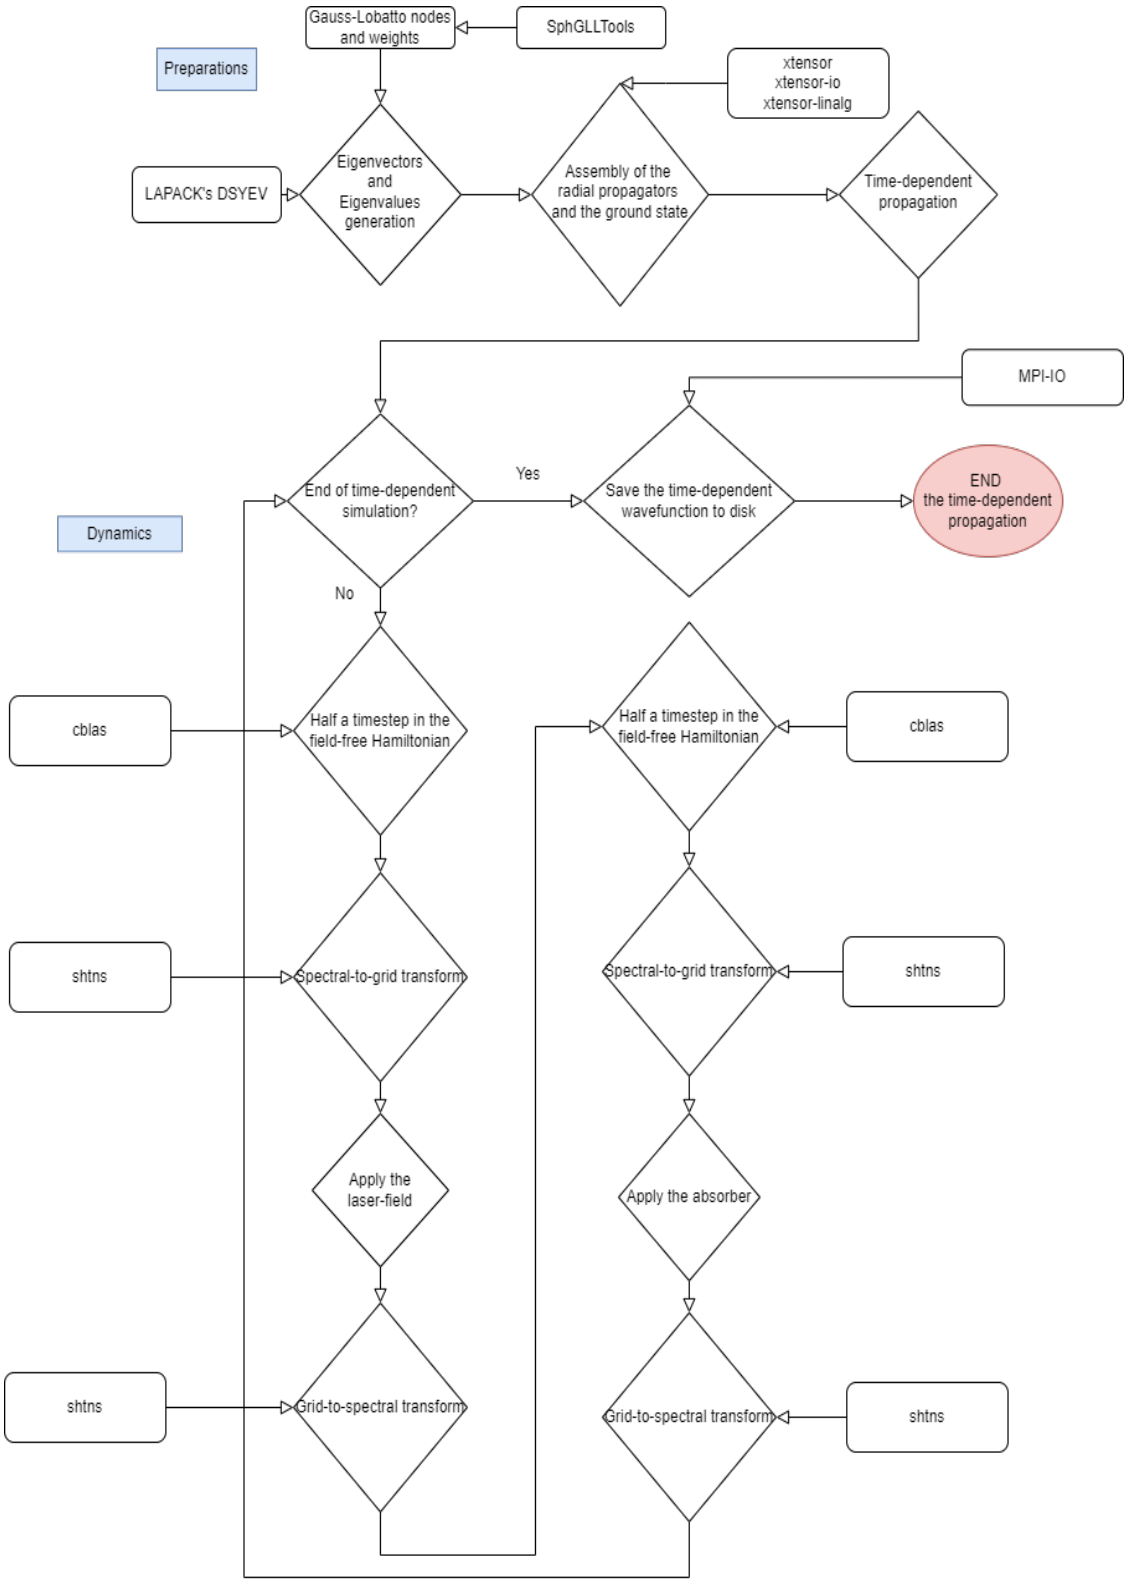


Figure 3.1: The workflow and pre-processing pipeline of the GPS-TDSE solver. The external libraries are pictured as rounded-edges boxes and the important locations inside the solver are pictured as diamonds. The separation into preliminary calculations, temporal dynamics and the end is done through coloured-shapes.

elements, and the splitting across ranks is based on the second operand from the product, i.e. `total_#_of_LM_pairs`. This `total_#_of_LM_pairs` is firstly calculated as in Alg. 1.

Algorithm 1 Calculation of the workload for the field-free propagation

```
total_#_of_LM_pairs ← 0
if mmax = l_max then
    for L ∈ [0, l_max] do
        total_#_of_LM_pairs ← total_#_of_LM_pairs + (2L + 1)
    end for
else
    for L ∈ [0, l_max] do
        upper ← min(mmax, L)
        total_#_of_LM_pairs ← total_#_of_LM_pairs + (2 · upper + 1)
    end for
end if
return total_#_of_LM_pairs
```

Then the ranks obtain their work across the (l, m) pairs as in Alg. (2).

Algorithm 2 Calculation of the workload for the field-free propagation

```
whole_LM ← total_#_of_LM_pairs / size
remainder_LM ← total_#_of_LM_pairs % size
if rank < remainder_LM then
    start_LM ← rank · (whole_LM + 1)
    end_LM ← start_LM + whole_LM
else
    start_LM ← rank · whole_LM + remainder_LM
    end_LM ← start_LM + (whole_LM - 1)
end if
return (start_LM, end_LM)
```

A pictorial representation of the radial slices held in memory by a communicator composed of three ranks (i.e. rank 0, rank 1 and rank 2) is shown in Fig. 3.2 (memory locations indices increment to the right and three consecutive dots represent memory locations not drawn on the diagram due to lack of space). As an example, it is assumed that this first rank holds the (l, m) pairs $(0, 0)$ and $(1, -1)$, the second rank holds the (l, m) pair $(1, 0)$ and the third rank holds the (l, m) pair $(1, 1)$, for when `total_#_of_LM_pairs` = 4 and *size* = 3 and $mmax = l_{max} = 1$.

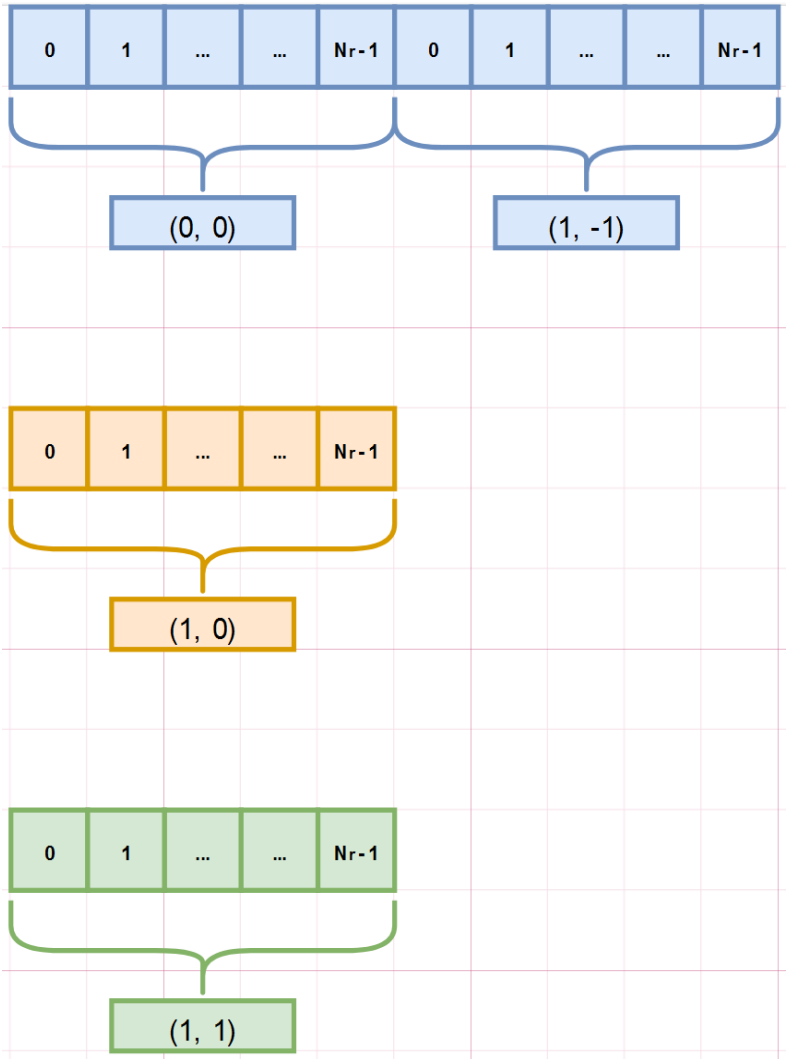


Figure 3.2: Schematic diagram showing three ranks holding the radial-slices of the time-dependent wavefunction. The zero-th rank (the top row) holds two (l, m) pairs, the first rank (the middle row) holds one (l, m) pair, and the second rank (the bottom row) holds one (l, m) pair. The memory locations increment to the right and the dots signify elements which are not drawn in the figure due to lack of space. In the figure, `total_#_of_LM_pairs` = 4.

Of course, for any combination of *size* and l_{max} and $mmax$, a roughly similar picture holds, however its details depend on the results obtained from Alg. (1) and Alg. (2).

In a 3D TDSE simulation, there is a non-trivial dependence on the φ coordinate. It was explained in Sec. 3.4.1 that the number of φ-points is $2\,mmax + 1$, thus for any l-value, there are $2M + 1$ blocks of N_r elements, where M is $min(mmax, l)$. The matrix-vector multiplications needed to propagate the time-dependent wave-

function for half a timestep in the field-free Hamiltonian $\mathcal{H}_0^l$ are performed as explained in Sec. 3.4.1. Each radial slice representing a given (l, m) pair has its position known and fixed inside the buffer of memory which holds it on the corresponding rank. Thus it is not hard to match any given (l, m) pair radial-slice with the corresponding radial propagator for that l-value.

In order to then apply the laser-field operator, the grid representation is needed. The spectral to grid transform works on (pseudo-2D) angular slices. That is, for any radial-grid point X, the same operation is applied.

This is an opportune moment to discuss the distribution of workload across the radial-coordinate r. The N_r slices are distributed uniformly across the ranks, just as the (l, m) pairs. In pseudocode, this is detailed in Alg. (3).

Algorithm 3 Calculation of the distribution of the radial-coordinate

```
whole_r ← N_r / size
remainder_r ← N_r % size
if rank < remainder_r then
    start_r ← rank · (whole_r + 1)
    end_r ← start_r + whole_r
else
    start_r ← rank · whole_r + remainder_r
    end_r ← start_r + (whole_r - 1)
end if
return (start_r, end_r)
```

Returning to the spectral to grid transform, it is clear that a modification of how `Psi_spec` is held in memory across ranks is needed before commencing the work on the transform. Each rank shall possess contiguous (l, m) slices for some i's, that is to say, fixing on a rank only, for some i's in a subset of $[0, ..., N_r - 1]$, that rank shall hold all the (l, m) pairs. Currently (after the half a timestep propagation in the field-free Hamiltonian $\mathcal{H}_0$) the ranks posses contiguous radial slices, that is to say, fixing on one rank only, it holds all the i's in $[0, ..., N_r - 1]$, for some (l, m) pairs.

A distributed-matrix transposition shall be performed to exchange the data as needed. This is an All-to-All-Variable (`MPI_Alltoallv(...)`) communication and it will be subject to an extensive discussion shortly.

Because the ranks do not hold the same number of radial slices and are not supposed to hold the same number of (l, m) slices, each rank sends to each other rank variable amounts of data.

A pictorial representation is shown in Fig. 3.3. There, the objective is to go from the **START** distribution to the **AIM** distribution. As an example, it is assumed that a total of 3 ranks are present in the communicator (that is, $size = 3$), $N_r = 8$ and the number of (l, m)-pairs each rank takes being $[2, 1, 1]$, where the first value from this example $[2, 1, 1]$ container refers to the number of (l, m)-pairs the zeroth rank takes, the second value refers to the number of (l, m)-pairs that the first rank takes and the third value refers to the number of (l, m)-pairs that the second rank takes. Indeed, it is shown that initially (START), rank 0 (i.e. $p(\text{rocess})_0$) holds two (the number of (l, m)-pairs it gets) radial slices, each of N_r elements. The buffers in computer memory are storing values vertically in the picture in the START distribution, that is, $0, 9, 18, ..., 54, 63, 1, 10, 19, ...$ etc, indicated by the blue arrow. Similarly for p_1 and p_2, with their associated number of (l, m)-pairs they hold (1 and 1, respectively). At the end of this transposition, the AIM distribution shall be obtained. As $size = 3$ and $N_r = 8$, the first two ranks shall hold three 2D-angular slices each (as from Alg. (3)), while the last one collects the remaining portion (as from Alg. (3)).

The re-distribution is achieved by means of an `MPI_Alltoallv(...)` call, followed by a re-shuffling of the received memory-buffers. Each rank has to fill with elements an initially empty (i.e. filled with whatever data was residing at those bytes from memory) memory-buffer for the communication. This buffer is understood by the `MPI`-function as being composed of a number of *size* blocks of elements. Each block will be sent to the rank corresponding to the position of that block in the buffer to be communicated. The length of the whole buffer depends on the rank on which the buffer is constructed. From Fig. 3.3 it is visible that this length is equal to ($N_r \cdot$ `#_of_lm_pairs_of_this_rank`), where `#_of_lm_pairs_of_this_rank` is equal to 2 for the zeroth rank, 1 for the first rank and 1 for the second rank.

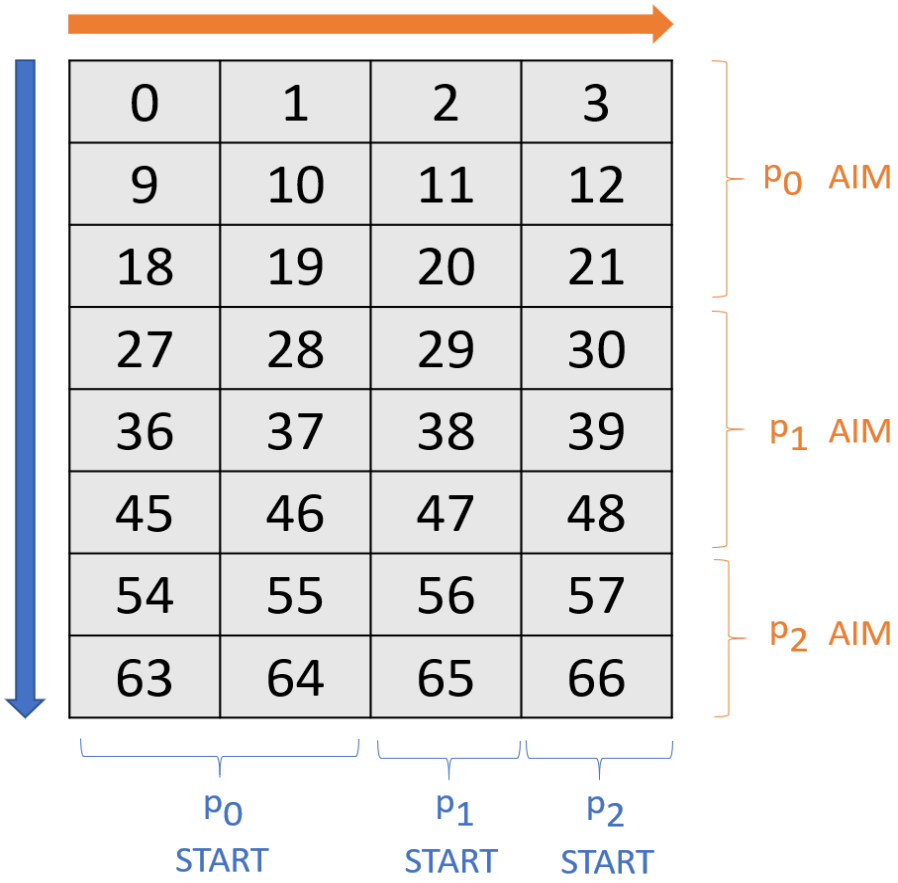


Figure 3.3: Pictorial representation of the starting distribution before the first `MPI_Alltoallv(...)` (denoted by blue arrows and braces) and the end distribution (denoted by orange arrows and braces). This AIM-distribution is to be obtained after the call and a necessary reshuffling have been performed. Three processes are shown ($size = 3$), $N_r = 8$ and `total_#_of_LM_pairs` = 4.

The buffers to be used in the `MPI_Alltoallv(...)` call are presented in Fig. 3.4. The top collection of three rows is the buffer as constructed on rank 0, the middle collection of three rows is the buffer as constructed on rank 1 and the last three rows represent the buffer as constructed on rank 2. Inside each collection, one row represents one block (as discussed above), and its destination is shown after the arrow. It is to be observed that the length of these buffers is not the same across the ranks.

After the communication, each process receives the needed elements, however not in the correct order. That is, one finds that they are not in contiguous 2D-angular slices, thus they need a re-shuffling (intra-rank only). In a pictorial representation, this is achieved as in Fig. 3.5.

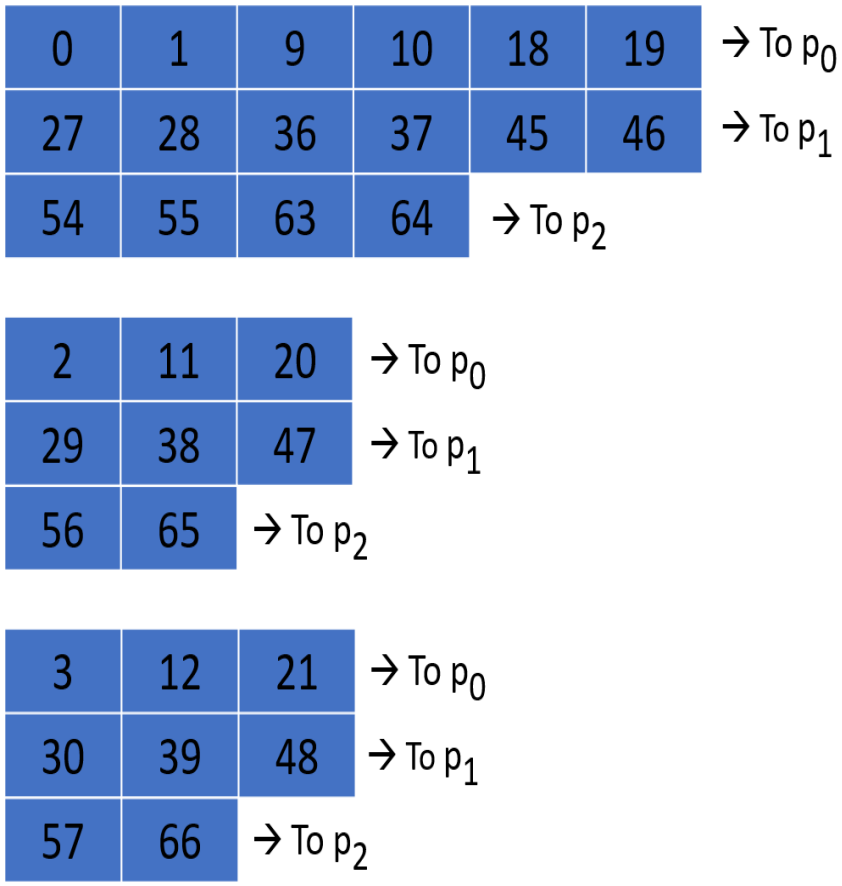


Figure 3.4: The memory buffers to be sent via the first `MPI_Alltoallv(...)` communication at each timestep of the solver. Top three rows: the buffers to be sent from process p_0, middle three rows: the buffers to be sent from process p_1, bottom three rows: the buffers to be sent from process p_2. For each process, the top line represents the data to be sent to p_0, the middle line represents the data to be sent to p_1 and the bottom line represents the data to be sent to p_2. It is observed that each process does not send the same amount of data to the other processes, including itself, thus the variable-communication patterns emerge.

Finally, the spectral to grid transform operation is applied on individual 2D-angular slices of the time-dependent wavefunction.

After this is completed, each process holds pieces of the grid-representation of the time-dependent wavefunction: `Psi_grid[i, :, :]`, for some `i`'s (in the example above, 3 `i`-values for rank 0, 3 `i`-values for rank 1 and 2 `i`-values for rank 2, i.e. just as the distribution of work across the radial-coordinate to ranks, for the spectral-representation).

The laser-field operator, the middle part of the three-parts propagation from Eq. (3.7), being a diagonal operator, is applied as a simple multiplication at

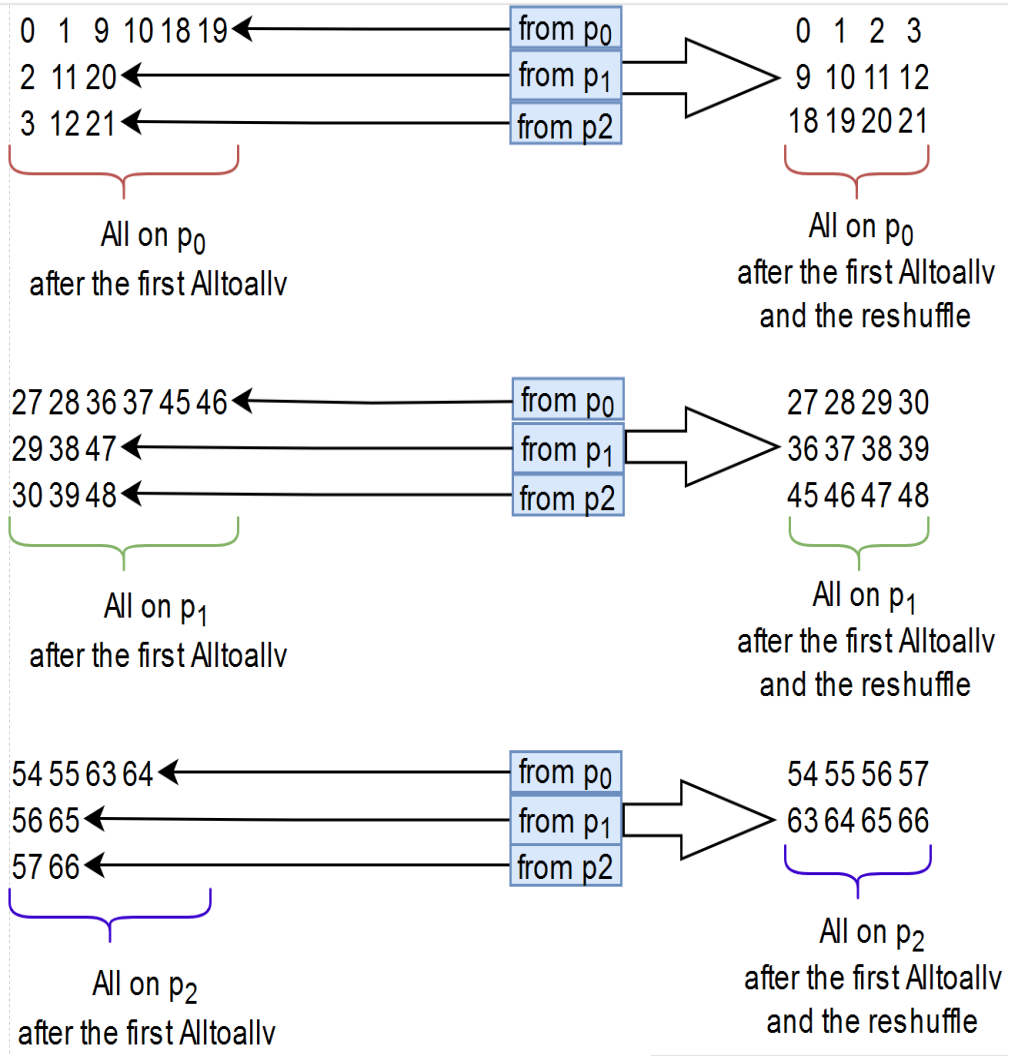


Figure 3.5: Data re-distribution after the first `MPI_Alltoallv(...)` call. Top section: the movement inside p_0, middle section: the movement inside p_1, bottom section: the movement inside p_2.

every grid-point from the real, 3D-space. As each rank knows what `i`-values it holds, there is no problem to identify the actual real coordinates of the datapoints represented in the `Psi_grid` buffer (distributed across the ranks).

Now the second part of the iterative algorithm starts, and the first objective is to transform from the grid-representation to the spectral-representation (such that the remaining half-a-timestep field-free propagation, i.e. in $\mathcal{H}_0$, is performed). This is done again with the help of the `shtns` C library, which, as expected, performs this inverse-transform with roughly the same API-commands as the forward-transform. However, this results in the ranks holding `Psi_spec[i, :, :]`

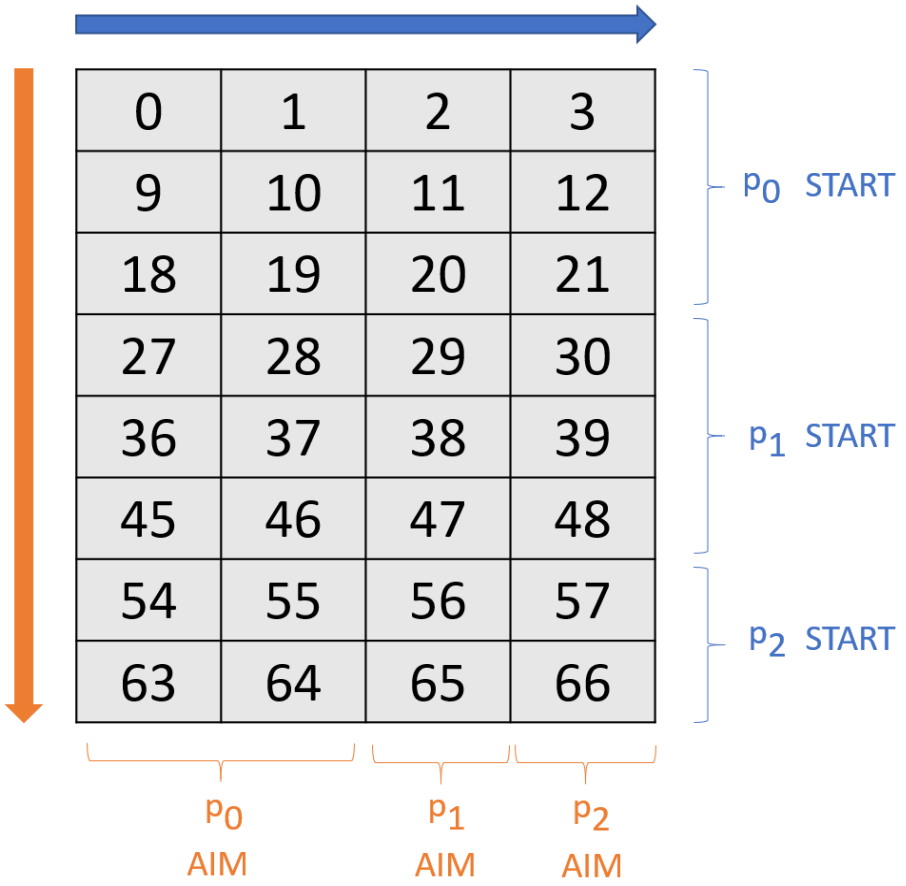


Figure 3.6: Pictorial representation of the starting distribution before the second `MPI_Alltoallv(...)` (denoted by blue arrows and braces) and the end distribution (denoted by orange arrows and braces). This AIM-distribution is to be obtained after the call and a necessary reshuffling have been performed. Three processes are shown ($size = 3$), $N_r = 8$ and `total_#_of_LM_pairs` = 4.

for some `i`'s (the same as before). Still, to apply the radial-propagators, one needs to have contiguous blocks of N_r elements, for some (l, m)-pairs. This again means that an all-to-all communication between the ranks needs to be performed. A pictorial representation is shown in Fig. 3.6.

Just as in Fig. 3.3, the initial distribution is marked with START, while the resulting distribution shall be the one marked with the AIM keyword. Now, in the START distribution, the memory buffers store elements as indicated by the blue arrow (horizontally to the right and at the end of one row going to the first element of the next row). At the end, the AIM distribution is signified by the orange arrow and generally, the orange colour. To obtain this pattern of data in the memory of the ranks, an `MPI_Alltoallv(...)` function call is performed,

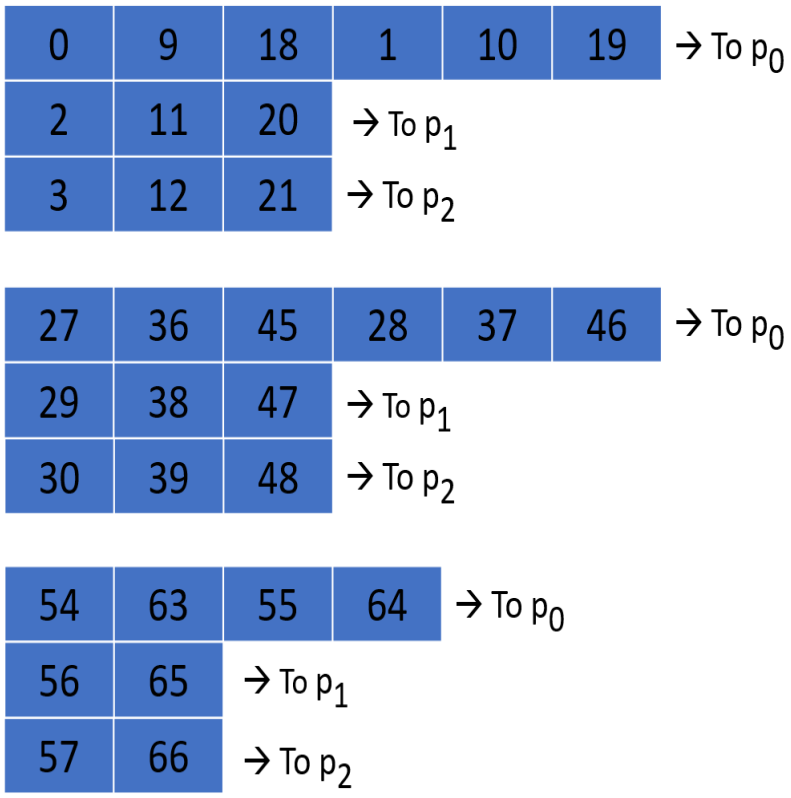


Figure 3.7: The memory buffers to be sent via the second `MPI_Alltoallv(...)` communication at each timestep of the solver. Top three rows: the buffers to be sent from process p_0, middle three rows: the buffers to be sent from process p_1, bottom three rows: the buffers to be sent from process p_2. For each process, the top line represents the data to be sent to p_0, the middle line represents the data to be sent to p_1 and the bottom line represents the data to be sent to p_2. It is observed that each process does not send the same amount of data to the other processes, including itself, thus the variable-communication patterns emerge.

followed by an intra-rank re-shuffling of the elements received by each rank. The buffers to be populated in each rank for the function call are illustrated in Fig. 3.7.

The intra-rank re-shuffling is pictorially represented in Fig. 3.8. Albeit with a different elements' transversal-pattern, the intuition is the same: this step makes the ranks to actually hold the data received from the all-to-all communication in the correct, contiguous order, needed for the propagation in the field-free Hamiltonian.

The actual steps for this propagation are the same as before, detailed above. Once this second half a timestep in the field-free Hamiltonian is completed, in between a pair of spectral-to-grid and grid-to-spectral transforms, the absorber

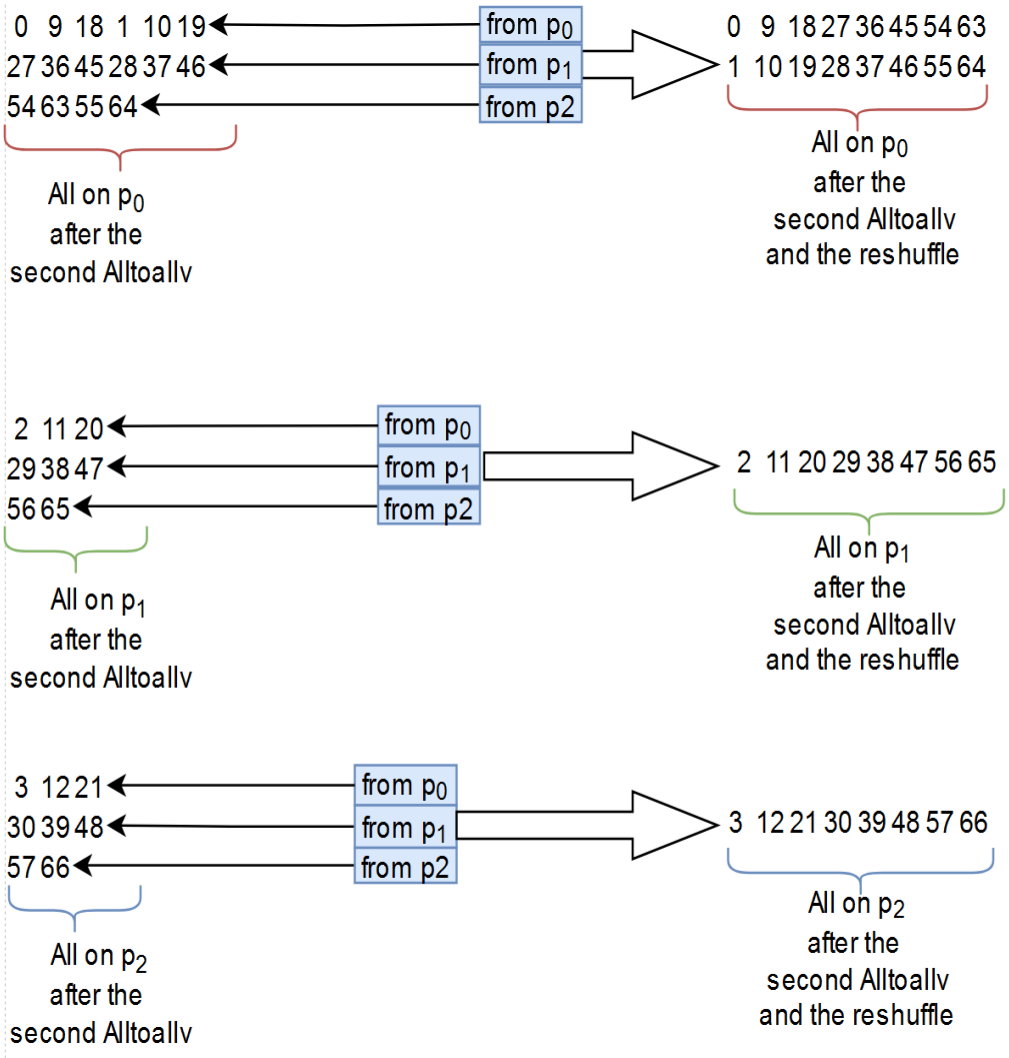


Figure 3.8: Data re-distribution after the second `MPI_Alltoallv(...)` call. Top section: the movement inside p_0, middle section: the movement inside p_1, bottom section: the movement inside p_2.

(mask) is applied on the time-dependent wavefunction. As the mask is a function of the radial coordinate only, and because the ranks know which 2D angular slices they hold (that is, which `i`'s from `Psi_grid[i, :, ;]`), it is an easy matter to identify the real coordinates in the 3D space of each datapoint contained in `Psi_grid`.

One full timestep of the propagation can thus be written as in the pseudocode listed in Alg. (4).

At the end of the simulation, the results are saved to disk using `MPI-IO`. This choice is preferred over classical `POSIX-IO`, as it does not serialise the solver and thus does not become the bottleneck when increasing the number of ranks to non-

Algorithm 4 One full timestep (Δt) of the propagation

1: Propagate for $\Delta t/2$ in the field-free Hamiltonian the l-values of this rank
2: Transform the spectral-representation to the grid-representation
3: Apply the laser field calculated at the grid points held by this rank onto the time-dependent wavefunction
4: Transform the grid-representation to the spectral-representation
5: Propagate for $\Delta t/2$ in the field-free Hamiltonian the l-values of this rank
6: Transform the spectral-representation to the grid-representation
7: Apply the absorber calculated at the grid points held by this rank onto the time-dependent wavefunction
8: Transform the grid-representation to the spectral-representation

negligibly small values. From the point of view of computer memory, the usage of `MPI-IO` does not eliminate the ability to work with very fine spatial-discretisations which would not fit into the RAM of only one `MPI`-process.

The data is saved to disk from its spectral representation, in the state in which the ranks hold slices of the type `Psi_spec[i, :, :]`, for some given `i`'s. This saving-to-disk workflow is executed in parallel by means of hardware-optimised `MPI-IO` function calls. In particular, the `MPI_File_write_at_all(...)` is used, and before that the `offset` of each rank is calculated in order to be fed as an input argument to this function. As it is not very CPU-time consuming, the saving of the time-dependent wavefunction to disk could also be done at each timestep, however for the PMD-retrieval method used in the current implementation, only the data at the end of the simulation is needed. Thus the storage (hard-disk)-requirements for running a simulation are drastically reduced in the current implementation of the workflow.

3.5 Benchmark calculations

3.5.1 Three-dimensional problem-setting

In order to test that the 3D TDSE solver based on the GPS discretisation produces correct results in full-dimensionality as well, two fully-3D problems well known in the literature are solved.

Elliptical polarisation

Apart from the case of a linearly polarised laser field, the azimuthal symmetry of $\mathcal{H}$ (the total Hamiltonian) is broken (because of the presence of the φ angle in $V(\vec{r},t) = \vec{r}\cdot\vec{E}(t) = x\cdot E_x + y\cdot E_y + z\cdot E_z = r\sin\theta\cos\varphi\cdot E_x + r\sin\theta\sin\varphi\cdot E_y + r\cos\theta\cdot E_z$) and the problem is 3D, requiring the solution of TDSE in full-dimensionality. For any ellipticity which is non-zero, the electron is ionised by the peak of the laser field along the major-axis and it is drifted towards the minor-axis. A rotation of the 2D PMD (the heatmap in the plane of the polarisation) shall be observed.

In Fig. 3.9 the $p_z - p_y$ distribution (for a polarisation in the $\hat{z}-\hat{y}$ plane, the major axis being along $\hat{z}$ and the laser propagating along $\hat{x}$) is shown for a 800 nm pulse of six optical-cycles ($\sin^2$ envelope) with ellipticity $\epsilon = 0.5$ and obtained by means of direct projection onto the scattering states at the end of the simulation. The interaction with the laser-field was treated in the length gauge in the DA ($V(\vec{r},t) = \vec{r}\cdot\vec{E}_{DA}(t)$), where $\vec{E}_{DA}(t) = -\partial\vec{A}(t)/\partial t$. The vector potential $\vec{A}(t)$ used in this elliptically-polarised-light simulation was:

$$\begin{aligned} A_x(t) &= 0, \\ A_y(t) &= A_{max}\sin^2\left(\frac{\pi t}{T}\right)\epsilon\cos(\omega t), \\ A_z(t) &= A_{max}\sin^2\left(\frac{\pi t}{T}\right)\sin(\omega t), \end{aligned} \tag{3.21}$$

with $A_{\max}$ above being set to 0.4421 a.u..

A "left-right" asymmetry is also visible, which shall be the case for any short laser pulse (below, say, 6-8 optical-cycles). This is a carrier-envelope phase (CEP) effect and disappears as the pulse duration increases, signalling the fact that a fine resolution for the PMD reveals dynamics-details which depend on this laser field parameter in the short-pulse regime. For this TDSE simulation, the numerical parameters were rather modest because the aim of this section is not to obtain fully-converged results, but to test that general features obtainable only in 3D problems are recovered: the number of radial grid points was 500, the

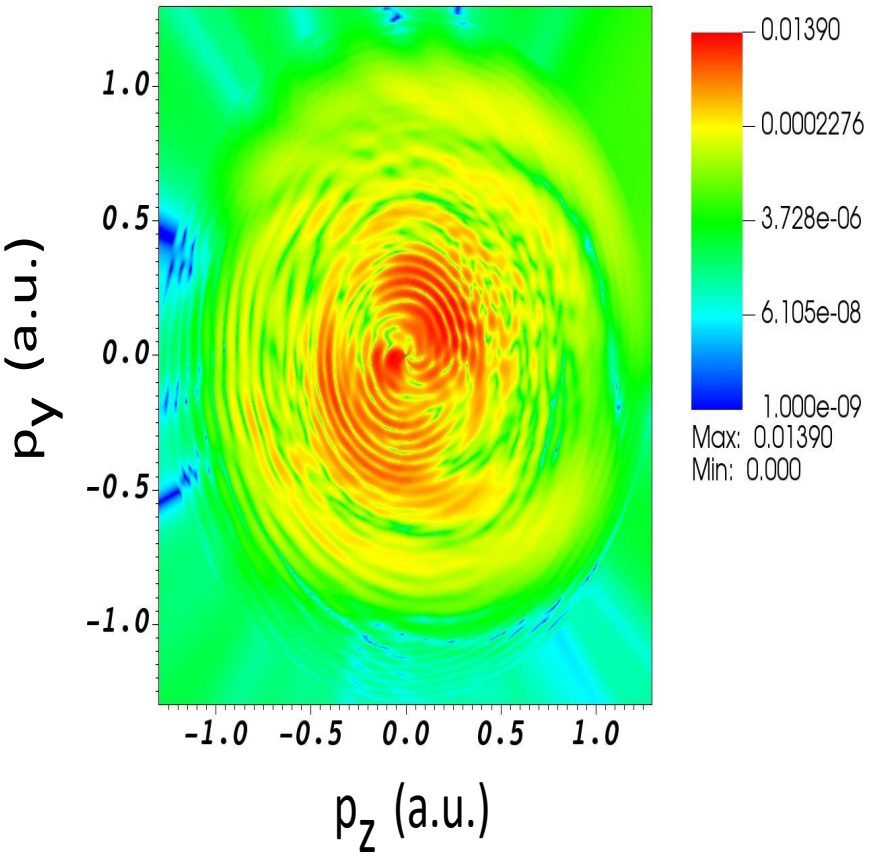


Figure 3.9: Two-dimensional heatmap showing the momenta in the laser polarisation plane p_z-p_y for an ellipticity of $\epsilon = 0.5$, at the wavelength $\lambda = 800$ nm, and peak intensity of 10^{14} W/cm^2, for a laser pulse of six optical-cycles in duration.

number of θ-points was 64 and all the m-channels were included. The temporal discretisation had 6000 timesteps per 1 optical cycle.

Attosecond angular-streaking: the "attoclock"

The elliptical polarisation of light can be applied in an experimental setting commonly known as the "attoclock": a means to achieve attosecond time-resolution, it helps resolving sub-cycle tunnel-ionisation dynamics, as it was very beautifully shown in [129] and other such studies. A circularly polarised laser field ($\epsilon = 1.0$) is used and the final momentum of the electron is nicely mapped back to its instant of ionisation, thus opening doors for studies in attosecond-science. In the experiment, the pulse was of two optical cycles in duration and here it was kept the same. The wavelength chosen was 400 nm. I mention that

the physics simulation setup is identical to the physics conditions from [98] and from [96], and these two closely match the experiments from the past. In my simulations, interaction with the laser-field was treated in the length gauge in the DA ($V(\vec{r},t)=\vec{r}\cdot\vec{E}_{DA}(t)$), where $\vec{E}_{DA}(t)=-\partial\vec{A}(t)/\partial t$. The vector potential $\vec{A}(t)$ used in this attoclock-simulation was:

$$\begin{aligned} A_x(t) &= 0, \\ A_y(t) &= A_{max}\sin^2\left(\frac{\pi t}{T}\right)\cos(\omega t), \\ A_z(t) &= A_{max}\sin^2\left(\frac{\pi t}{T}\right)\sin(\omega t), \end{aligned} \tag{3.22}$$

with $A_{\max}$ above being set to 0.4682 a.u.. For the numerical parameter of this simulation, the number of radial points was 1200, the number of θ-points was 90 and all the m-channels were included. The temporal discretisation was rather coarse, with 5000 timesteps per one optical-cycle.

One would expect that the PMD to be obtained in this attoclock-setting possesses some special features, as clearly detailed in [96].

It is known that the most probable instant of ionisation is when the E-field peaks in time. This happens after one optical-cycle for the pulse considered in this subsection. At that moment, $E_z(t)=-\partial A_z(t)/\partial t$ is negative, thus the position of the electron at the tunnel exit is in the $z>0$ region. When neglecting the attraction with the core, according to the three-step model (simple-man model) of ATI [9], the final drift of the electron in the continuum is given by the negative A-vector-potential at the tunnel exit time. For the pulse considered in this subsection, at the tunnel exit time, $-A_y$ is negative (A_y is positive) and $-A_z=0$. Thus, one would expect a rather large value of the PMD along $p_z=0$ and $p_y<0$. When the Coulomb potential is included in the analysis, this large value of the differential ionisation probability is shifted to negative p_z (maintaining the negative p_y characteristic) because the electron swings-by the nucleus due to its attraction (thus builds up negative-drift-momentum from the $z>0$ exit-point position). The above explanations can also be seen in [96], for example.

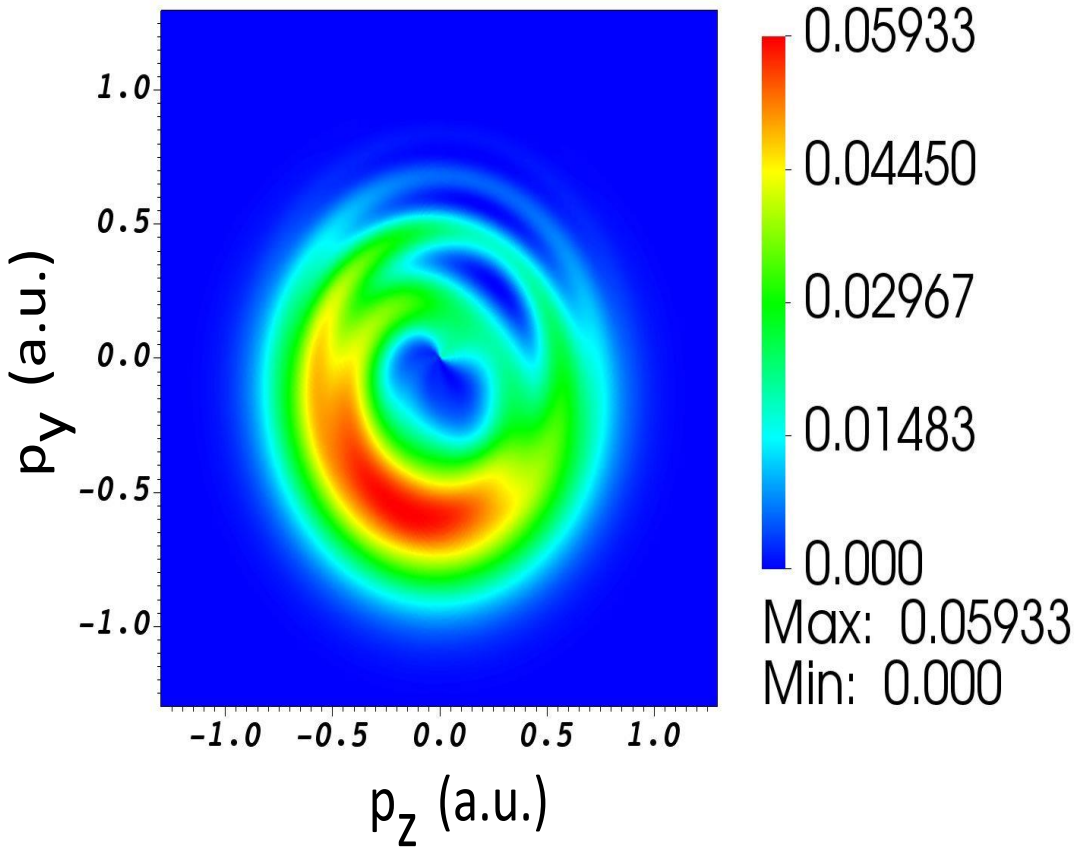


Figure 3.10: Two-dimensional heatmap showing the momenta in the laser polarisation plane p_z-p_y for circular polarisation (ellipticity $\epsilon = 1$), at the wavelength $\lambda = 400$ nm and at a peak intensity of 10^{14} W/cm^2 for a laser pulse of 2 optical-cycles in duration and a radial box-size of 240 a.u.. The well-known "attoclock" PMD-pattern is visible and the simulation results are in excellent quantitative agreement with panel a) of Fig. 4 from [98] and with panel a) of Fig. 5 from [96], all three of which match (both as the general shape and characteristics but also as the actual heatmap values) with each other.

The obtained PMD in the plane of polarisation is presented in Fig. 3.10. It is seen that the characteristic shape and features are obtained, thus confirming the validity of the statement that the solver can be used for 3D TDSE PMD-retrieval studies. This is because the simulation results quantitatively agree with the results presented in panel a) of Fig. 4 from [98] and in panel a) of Fig. 5 from [96]. Because those studies used the same physics conditions as my simulation, the agreement in the heatmap's data values is a strong evidence that I benchmarked my solver in 3D against well-known and verified results, presented in both distant, but also very recent literature.

3.6 Results and discussion: the parallel-scaling and performance of the solver

3.6.1 The improvement in the timings per one timestep as more computational resources are used

The solver has been tested in order to see how well the parallelisation strategy discussed in Sec. 3.4.2 scales in two different scenarios of importance in parallel-programming.

In the first scenario, known as the ***strong-scaling*** regime, the aim is to look at a fixed problem size (that is, a fixed spatial-3D-discretisation which implicitly dictates the number of elements in the spectral representation) and analyze the improvement in the wall-clock-time needed for executing one timestep of the propagation as the number of cores the solver can use is increased.

In the second scenario, the ***weak-scaling*** regime, the aim is different. Here the amount of work per core is fixed from the start to a given value, and the aim is to look at how the wall-clock-time for executing one timestep of the propagation behaves as the number of cores the solver uses is increased.

Starting with the first setup, for the ***strong-scaling*** tests I decided that the following discretisation is used: $N_r = 1255$, $nlat = 511 = l_{max} + 1$, $mmax = 10$. This discretisation will be different for the weak-scaling tests which I will present later. For the strong-scaling tests, I chose the discretisation to be one-off multiples of 2, such that the `MPI_Alltoallv(...)` calls do not decay to `MPI_Alltoall(...)` calls as it would happen when the number of used `MPI`-ranks divides the discretisation values. This is because the chosen numbers of `MPI`-ranks to be used in this testing-campaign are powers of 2, and is known that `MPI_Alltoall(...)` calls scale better than their `MPI_Alltoallv(...)` counterparts, as there are more ways to optimise the communication for the former than for the latter. Much effort has been dedicated to improve these two types of the all-to-all communication, e.g., [130–135]. This is because this all-to-all communication is at the core of

non-trivial scientific methods and applications, for example in FFTs as applied to large-problem-size quantum-chemistry simulations [136].

For the strong-scaling results, the number of ARCHER2 nodes used started at 8, each node being set to have 8 `MPI`-ranks, each rank having 16 cores set to be active (thus with a maximum of 16 `OpenMP` threads being possible to be spawn inside each rank) and ended up at 128 nodes. For each `MPI`-setup, the same parts of the solver have been timed with the C++'s `chrono`.

Inside the root `MPI`-rank, I timed the following parts of the workflow: **a)** the total wall-clock-time per one timestep of the simulation; **b)** the matrix-vector (M-v) products for propagating those (l, m)-pairs the rank has received as its work; **c)** the first `MPI_Alltoallv(...)` call; **d)** the spectral-to-grid transform; **e)** the grid-to-spectral-transform; **f)** the application of the absorber; **g)** the application of the laser-field operator; **h)** the second `MPI_Alltoallv(...)` call;

CPUs	1 tstep	M-v	Av 1	S-G	G-S	Mask	Laser	Av 2
1024	1091	373	79	37	38	0	22	1
2048	515	204	5	18	19	0	10	1
4096	285	96	22	9	9	0	5	1
8192	162	50	14	5	5	0	3	1
16384	105	28	12	3	3	0	2	3

Table 3.1: Strong scaling tests results: all timings are in milliseconds (ms). M-v signifies the matrix-vector operations, Av signifies an Alltoall-variable communication (with 1 and 2 being the first and the second communications), S-G signifies the spectral to grid transform and G-S signifies the grid to spectral transform. The discretisation was $N_r = 1255$, $nlat = 511$, $mmax = 10$. There were 16 cores per one `MPI` rank, thus 16 `OpenMP` threads per process. One NUMA region of ARCHER2 was populated by one `MPI` rank. The second column represents how quickly one timestep of the time-dependent simulation finishes, for various number of processors used (first column). The other columns break down the duration of one timestep into its time-consuming parts. There are other parts of the solver which have been not presented here as they are negligible in wall-clock-time duration. The matrix-vector multiplications (M-v column) are applied twice per timestep, similarly for the spectral to grid transform (fifth column) and the grid to spectral transform (sixth column). The absorber (mask column) and the laser-field are applied once per timestep.

The most-important metric is the wall-clock-time per one timestep of the simulation, as this dictates when the simulation finishes and thus the physics-data becomes available.

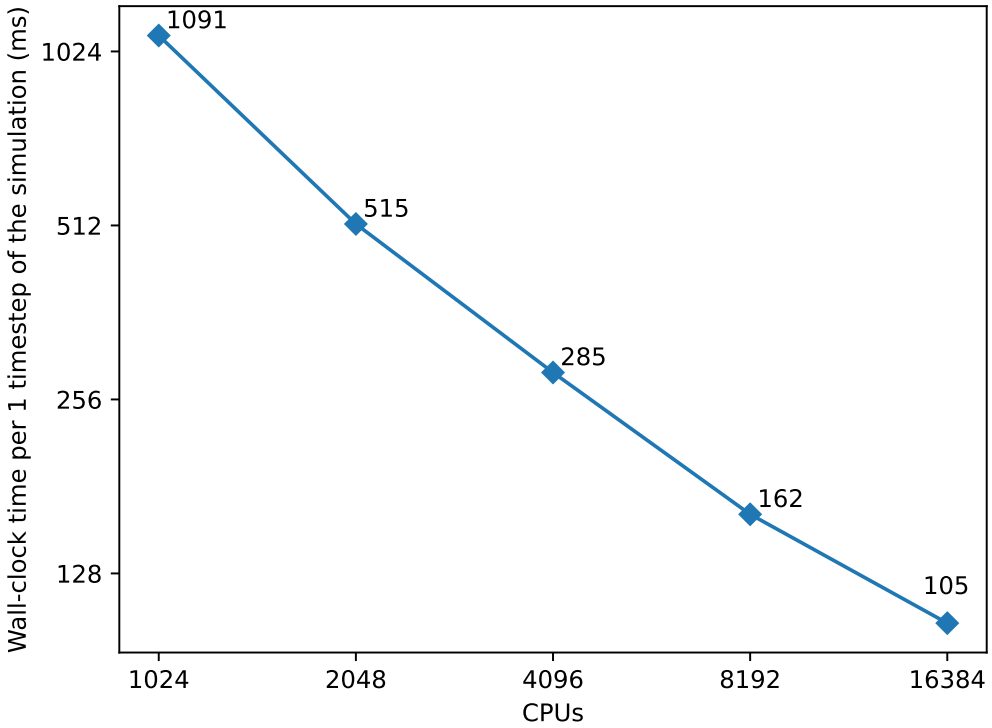


Figure 3.11: Strong-scaling (`OpenMP` enabled): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The spatial discretisation was fixed at $N_r = 1255$, $nlat = 511$, $mmax = 10$. While generally following a linear trend up to a large number of processors, the start of the development of a plateau can be noticed towards a very large number of processors, when the communication is too expensive for the amount of computation performed by each processor. One `MPI`-rank was composed of 16 cores, each core launching one `OpenMP` thread, and there has been one such rank per one NUMA region of ARCHER2. The datapoints present their associated wall-clock time value, for ease of reading the results.

The results for the strong-scaling tests are presented in Fig. 3.11.

The log-log (in base 2) plot shows across the vertical axis the wall-clock time per 1 timestep of the simulation, in milliseconds (ms). The horizontal axis represents the number of processors (cores) used. One node from ARCHER2 possesses 128 cores. No `AMD`-hyper-threading has been used in any run, thus a maximum of 1 `OpenMP` thread was possible to be launched per 1 (physical) core. One datapoint represents one run of the solver, having as computational resources the corresponding number of cores from the plot's horizontal-axis. It is visible that the decreasing trend is almost linear (close to the theoretical best-

scaling). As after a certain number of processors the communication becomes comparable to the computation, it is expected that the linear decreasing trend starts to plateau. For the tested problem size, there is a slight indication of this phenomenon happening at 64 nodes (8192 cores). Of course, the point at which degradation starts to happen depends on the problem size. For a larger problem size than the one from these tests (for example, more m-channels included in the simulation, as $mmax = 10$ is a rather modest value), the start of the degradation would be seen at a larger number of nodes than what is presented here. It is not a trivial task to obtain 256 (or more) full-compute-nodes on ARCHER2 for a normal `SLURM`-job (the time needed to be waited to be allocated such a big number of nodes is rather large), so I can conclude that for reasonable practical applications in laser-atom interaction studies, the solver can have more resources (CPUs) safely invested into it, without suffering from scalability losses.

For the ***weak-scaling*** regime tests, I firstly made use of the approximated computational complexity of the solver. This is because in all such tests, one processor has to be devoted to a fixed, constant, workload. Thus in weak-scaling, the increase in the computational-resources has to come with an associated increase in the problem size which leaves each processor to have the same amount of work distributed to it.

As shown in Sec. (3.4.1), the estimated computational-complexity (time-complexity) of the solver is $O(N_r^2 \cdot nlat^2 + N_r \cdot nlat^3)$ for when $mmax = l_{max}$. This means that for any $p > 1$, where p refers to either the total number of processors for when no `OpenMP` is used, or to the number of `MPI`-ranks with 16 cores each when `OpenMP` is used, N_r and $nlat$ (and implicitly $mmax = l_{max}$) both have to grow as $p^{1/4}$ in order to maintain a constant workload per an element represented by p (i.e. constant work per either one core or per one `MPI`-process composed of 16 cores). This is because the computational complexity is a fourth-order polynomial in the discretisation parameters. Thus $N_r(p) = p^{1/4} N_r(1)$, and similarly, $nlat(p) = p^{1/4} nlat(1)$, together with $mmax(p) = p^{1/4} mmax(1)$, which

follows from the assumption that $mmax = l_{max}$ (i.e. all m-channels are included in the tests for the weak-scaling analysis), is needed such that the expression

$$\frac{N_r(p)^2 \cdot nlat(p)^2}{p} + \frac{N_r(p) \cdot nlat(p)^3}{p} \tag{3.23}$$

is constant with respect to p.

There have been two weak-scaling sets of tests performed, one in which each `MPI`-rank has 1 core only (thus the number of `OpenMP` threads was 1), and another one in which each `MPI`-rank has 16 cores (thus the number of `OpenMP` threads was 16 inside the rank). An initial discretisation was fixed, that is, $N_r(1)$ and $nlat(1)$, with $mmax(1)$ following from $nlat(1)$ as $mmax(1) = nlat(1) - 1$. Then, for each $p \neq 1$, $N_r(p)$ and $nlat(p)$ were calculated from $N_r(1)$ and $nlat(1)$, and $mmax(p)$ followed $nlat(p)$ as $mmax(p) = nlat(1) - 1$. For the first set of tests, the initial discretisation was $N_r(1) = 64$ and $nlat(1) = 32$. Then p was set to all the values from the set $\{16, 32, 64, 128, 256, 512\}$. The results for this first set of tests are presented in Fig. 3.12 and in Table 3.2.

CPUs	p	N_r	nlat	1 tstep	M-v	Av 1	S-G	G-S	Laser	Av 2
1	1	64	32	49	4	0	3	8	13	0
16	16	128	64	39	4	1	2	5	7	1
32	32	153	77	32	6	0	2	2	6	0
64	64	182	91	44	7	0	4	4	5	0
128	128	216	108	39	6	1	1	3	4	3
256	256	256	128	24	5	1	1	2	3	0
512	512	305	153	53	8	5	3	2	5	2

Table 3.2: Weak scaling tests results (no `OpenMP`): all timings are in milliseconds (ms). M-v signifies the matrix-vector operations, Av signifies an Alltoall-variable communication (with 1 and 2 being the first and the second communications), S-G signifies the spectral to grid transform and G-S signifies the grid to spectral transform. The initial discretisation was $N_r = 64$, $nlat = 32$, $mmax = 31$. The increasingly finer discretisations were calculated as $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks. There was one core per one `MPI` rank.

The log-log plot shows across the vertical axis the wall-clock-time for the execution of 1 timestep, in milliseconds (ms), while the horizontal-axis represents the number of cores used for that numerical experiment.

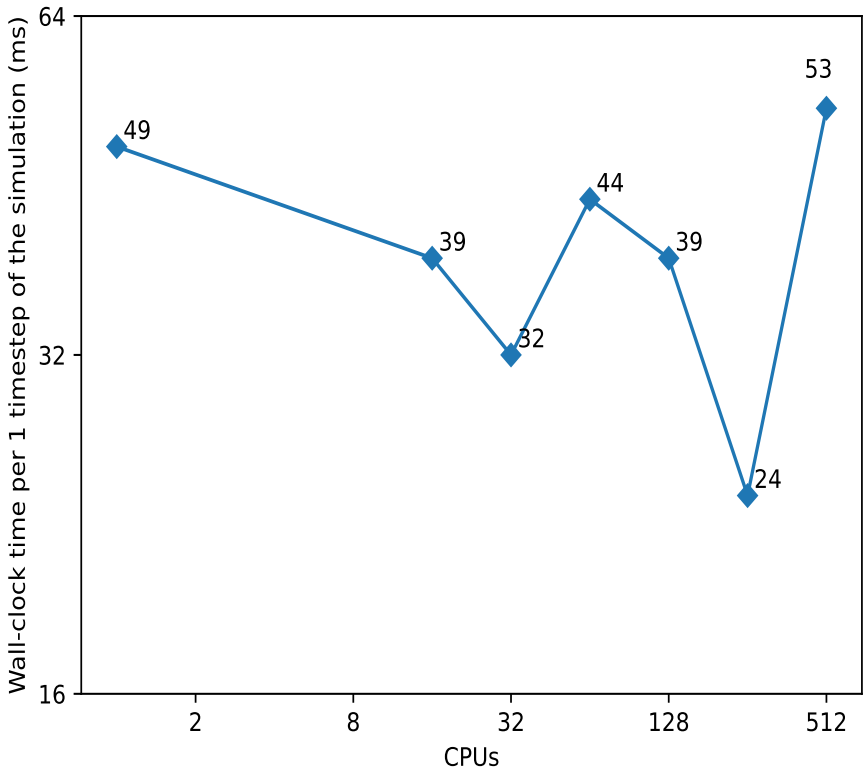


Figure 3.12: Weak-scaling (no `OpenMP`): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The initial discretisation was $N_r = 64$, $nlat = 32$, $mmax = 31$. The increasingly finer discretisation have been calculated as: $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks and one such `MPI`-rank was composed of one core. The datapoints present their associated wall-clock time value, for ease of reading the results.

It is known that weak-scaling is harder to be achieved in practice than the strong-scaling. In my case, the fact that only an estimation of the time-complexity has been made, together with the missing prefactors for the actual counts of the floating-point-operations happening in the iterative algorithm creates additional complications. For example, the application of the laser-field operator does not come into play in the estimated time-complexity, as it is a third-order polynomial in the discretisation parameters, but for rather low such parameters, there is not a large difference between a third-order polynomial in the discretisation parameters and a fourth-order one, thus inaccuracies (either small or large, depending on

the parameters) can easily appear. Depending on what type of calculations are performed for each point in the real 3D-space, obtaining the laser-field value at each grid point might start to show a non-trivial computational cost. This is just one aspect to consider, however the solver is rather involved, especially after the compiler optimises (many times aggressively) the workflow. Also, during one timestep of temporal propagation there are performed other operations which do not show up in the estimated computational complexity. All these make calculating the discretisations required for keeping the amount of work on each processor the same as for the initial discretisation not entirely straightforward: the calculated $N_r(p)$ and $nlat(p)$ can only be considered rough estimates. Still, the results I obtained, even in such a complex algorithmic-theory scenario, are showing that weak-scaling tends to be achieved rather well for the current implementation of the solver for its fully-`MPI` version, across a range of numbers of processors and across a rather large spectrum of real-space discretisation parameters. An oscillatory pattern around a linear trend is observed: for such coarse discretisations, the fact that the computational complexity of the solver is not accurate to exact numbers of floating-point-operations translates into timings which fluctuate around the mean trend, which is seem to follow a straight line.

For the weak-scaling tests for which `OpenMP`-threading was enabled, the initial discretisation was $N_r(1) = 256$ and $nlat(1) = 64$, again with $mmax(1) = nlat(1) - 1$. The p for this set of tests denotes one `MPI`-process and one such rank contained 16 cores. Numerical experiments have been performed for $p \in \{16, 32, 64, 128\}$, that is, up to 2048 cores. The results are presented in Fig. 3.13 and in Table 3.3.

Because the initial discretisation is finer than the one for the weak-scaling tests without `OpenMP`, the obtained timings are larger because the computational effort per processor is larger than for those tests. Still, the linear trend is again observed, this time a more pronounced such behaviour is seen: there is a rather constant plateau which extends up to the maximum number of cores used in these tests. Again, because of the various approximations made in estimating the

CPUs	p	N_r	nlat	1 tstep	M-v	Av 1	S-G	G-S	Laser	Av 2
16	1	256	64	3479	1669	3	5	11	14	3
256	16	512	128	3322	1606	26	2	5	6	23
512	32	609	153	3385	1617	51	5	5	9	2
1024	64	725	182	3288	1596	21	3	5	5	2
2048	128	862	216	3339	1579	90	5	8	4	2

Table 3.3: Weak scaling tests results (`OpenMP` enabled): all timings are in milliseconds (ms). M-v signifies the matrix-vector operations, Av signifies an Alltoall-variable communication (with 1 and 2 being the first and the second communications), S-G signifies the spectral to grid transform and G-S signifies the grid to spectral transform. The absorber (mask) is omitted from this table as its contribution to the timings was too small in all the experiments in the weak-scaling regime. The initial discretisation was $N_r = 256$, $nlat = 64$, $mmax = 63$. The increasingly finer discretisations were calculated as $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks. There were 16 cores per one `MPI` rank, thus 16 `OpenMP` threads per process.

computational (time-) complexity of the solver, the obtained datapoints (timings) are spread around the linear trend. The general conclusions drawn from these results is that a larger problem size can be solved in the same amount of wall-clock time as a smaller one, by allocating more processors to the finer simulation.

3.6.2 Serial-performance and insights

The roofline-model [137, 138] of a HPC-application is often preferred to provide insights about the general behaviour of its components when the application is run on one `MPI`-rank only (or on many ranks and on only one the cumulative result is analyzed at the end). Being a plot which presents across the vertical-axis the performance (usually measured in GFLOPs/sec) and across the horizontal-axis the arithmetic-intensity (defined here as the number of FLOPs per byte moved), it easily reveals how the solver could be optimised. That is, it shows what is most needed to make the solver reach the solution (end of simulation) in a smaller wall-clock time: either more memory-bandwidth (i.e. the program is memory-bound), or more-compute capabilities from the core (i.e. the program is compute-bound). Such conclusions can either be drawn as an average response across all the solver's parts, or as location-resolved, that is, the analyses are done per individual

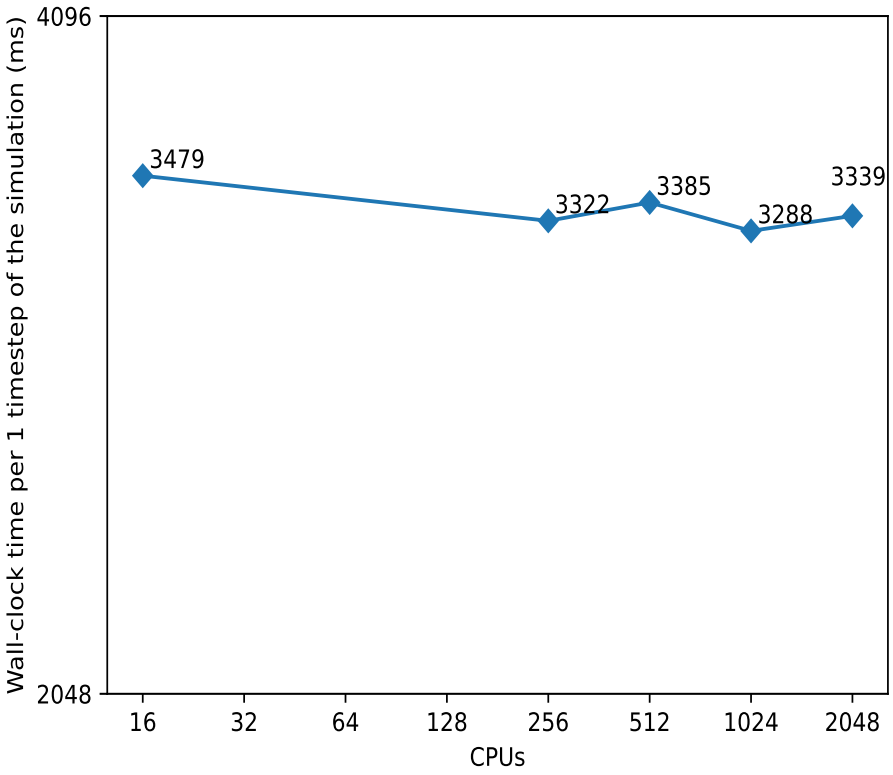


Figure 3.13: Weak-scaling (`OpenMP` enabled): the wall-clock time per one timestep of the simulation was measured as a function of the number of processors used. The initial discretisation was $N_r = 64$, $nlat = 32$, $mmax = 31$. The increasingly finer discretisations were calculated as $N_r(p) = p^{1/4} \cdot N_r(1)$, similarly for $nlat$ and $mmax(p) = nlat(p) - 1$, where p in this context stands for the number of ranks. There were 16 cores per one `MPI` rank, thus 16 `OpenMP` threads per process.

parts of the solver and then interpreted individually. I will call throughout this sub-section the functions of interest from within the solver and the loops which perform significant amount of compute-work as individual *kernels*.

To obtain the ingredients composing such a diagnostic, two independent measurements (pieces of information) are needed. The first one is the memory traffic (bandwidth, i.e. bytes/sec), and the second one is the FLOP rate (FLOPS/sec). Dividing the FLOP rate by the obtained bandwidth, one obtains the operational intensity (the value across the horizontal axis). The performance is just the FLOP rate (the value across the vertical axis). These allow to put one datapoint into an $x - y$ plot.

In order for this graph to be complete, there has to be one datapoint per region of interest (kernel), but also the graph needs the so-called "rooflines". These are theoretical guarantees above which the application's (solver's) performance cannot go above (i.e. it cannot be better than these, because of the hardware at hand). The first roofline is the maximum (or peak) performance, the second roofline is the maximum (or peak) data-throughput, and the point of their intersection is called the ridge-point. Anything which falls to the left of the ridge-point is denoted as memory-bound and would benefit from faster (memory) buses for data transportation. Sitting on the left of this point loosely means that the data movement takes a considerable amount of time and the arithmetic capabilities of the chip are not fully utilised (the computer waits for data to arrive, i.e. to fill the cache-misses). On the other side, i.e. to the right of the ridge-point, the location is denoted as being compute-bound, and anything residing there would benefit from faster computations. For these parts, the (main) bottleneck is the speed at which the FLOPS are executed, and the kernels located there have a rather large degree of data locality.

While the information for (all) the datapoint(s) for (the complete) plot can be collected manually (the rooflines can also be calculated theoretically using the vendor's reference manual), it can imply a lot of engineering effort and many times obtaining a correct such collection is not a trivial task to achieve. Initially I performed trials with the Intel® Advisor tool [3]. When the compilation and linking of the solver is done accordingly, and then a simulation is run "in-tandem" with this tool, the data necessary to create the roofline-model plot is collected automatically. The visualisation of the results is also relatively straightforward. It is to be mentioned that collecting the necessary metrics requires the HPC-application to be running more than how long it would take to complete the simulation without the executable being orchestrated with Advisor. Thus, these runs can quickly become expensive and the 3D real-space discretisations on

[3] Intel® Advisor is a design and analysis tool for developing performant code. See: `https://www.intel.com/content/www/us/en/developer/tools/oneapi/advisor.html`

which this tool can be used are rather modest. However, after running this tool, I observed that a few kernels presented performances which were above the automatically-generated rooflines, in visible contradiction to what I was expecting: the theoretical peak performance is the number (value) the vendor guarantees that cannot be exceeded by any application, yet I was obtaining that. I concluded that these measurement errors are due to the functions running for a very short time and the tool not having enough room for a proper judgement on their performance. This was later confirmed by a senior performance-engineer working on the Intel-Advisor product, who explained that this situation was known to the team and are working on finding ways to mitigate these issues, however I stopped using this tool way before I obtained this information.

This is because it was clear to me that a more manual collection might prove itself helpful. However, ARCHER2 nodes with their `Zen2` architecture implied that there has not been much room for the choices of the hardware counters to be used. In fact, `Zen2` has not evolved much since the `Bulldozer` (2011) in terms of what counters it offers. `PAPI` [139] as from the `cray-pat` module showed that there are as few as 17 events available on the compute nodes, from a total of 108 possible events.

Eventually, I compiled and used `likwid` [117] to obtain the roofline-model datapoints for my solver, on a `Debian`-based PC I had access to in my research-group and not on ARCHER2, as this tool requires non-trivial privileges to have all its capabilities accessible by the end-user. At the time of writing, the ARCHER2 helpdesk started to service my support-ticket and went on to allow `likwid` to be used by anyone on their super-machine. I have orchestrated the kernels of interest from the solver using the `Marker-API` from `likwid` and measured the traffic and (double-precision) FLOP-rate using `likwid-perfctr`. The peak performance and bandwidth have also been measured using `likwid`, by means of various `likwid-bench` commands and associated parameters.

For the peak performance measurement, a kernel is launched and inside it, the processor executes that many floating-point operations that anything else

which happens can be considered to have a negligible effect. Also, in order to minimise these unwanted-effects, the size of the array on which the operations (which can be part of, for example, a reduction) take place has to fully fit into the lowest-level cache of the target chip. In short, the aim of this kernel is to have a rather large number of in-register operations on a single data-element loaded from an array, before moving on to next array element.

For the peak-bandwidth measurement, `likwid` offers four alternatives to chose from, the simplest probably being a very large array being copied and the time taken for this to finish being measured. However, it is to be noted that the memory access pattern of the benchmark used (for example, the stream operation (`a[i] = b[i] + q * c[i]`), where `a`, `b` and `c` are standard vectors and `q` is a constant) is important for the obtained results, and this is due to complex effects from the prefetcher's interactions with the data.

The peak-performance measurements also depend on the number of double-precision elements of an array on which the operations are executed simultaneously. By means of an example, one can distinguish between scalar, SSE, and AVX, and all these dictate the operation-width of the instruction. There also exist some specific units inside the chip, the fused-multiply-add (FMA) parts of it, which, as their name suggests, glue together into a single (micro-)operation a multiplication which is accumulated into a target location. All these create multiple peak-performance rooflines and in order to go beyond (vertically above) any of this upper bounds, more of the next level of vectorisation is needed to be performed by the optimising compiler, potentially with the help of the application-developer by means of (`GCC`, for example) intrinsics and properly written code.

I have targeted the following parts of the solver: the spherical-harmonics transforms (both the forward transformation and the backwards transformation), the application of the laser-field and of the absorber, and the core of the propagation in the field-free Hamiltonian, i.e. the matrix-vector multiplications between the radial-propagators and the corresponding radial-slices from the wavefunction. Results are presented in Fig. 3.14.

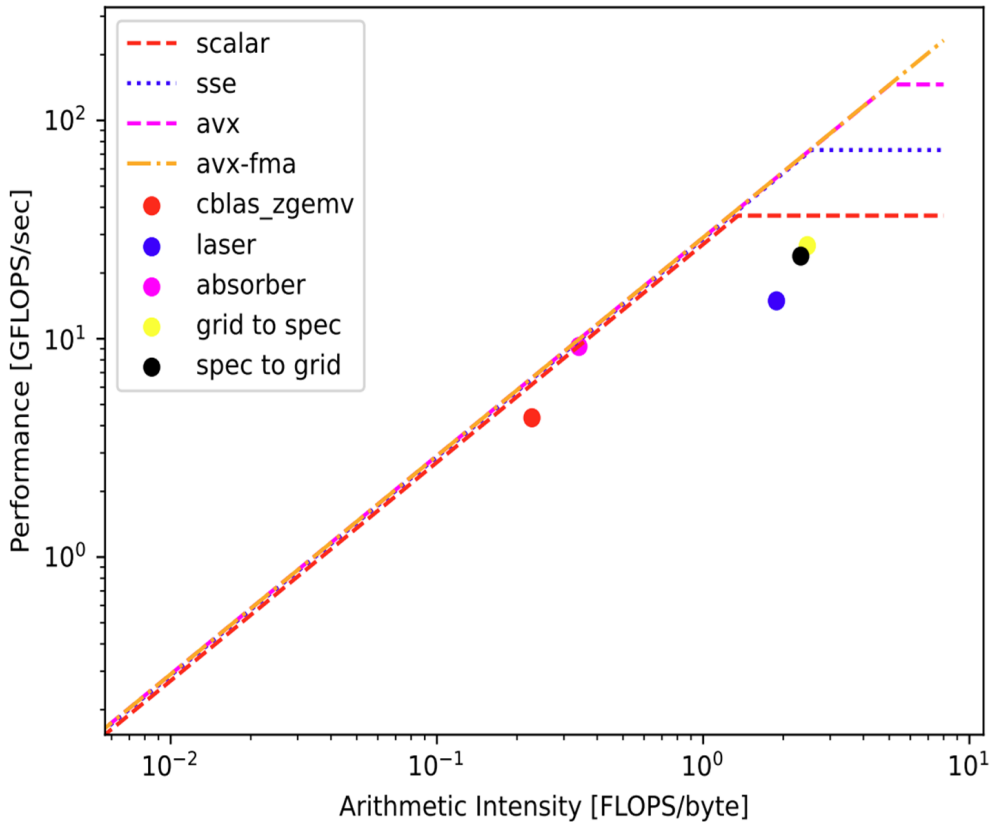


Figure 3.14: Results from applying the roofline-model to the TDSE solver by means of `likwid`. Across the vertical axis the number of floating point operations per second is pictured, while the degree of data locality, i.e. the number of floating point operations per byte moved, is shown on the horizontal axis. It is seen that some of the orchestrated (measured) parts of the solver are sitting to the left of the ridge-point, thus are memory-bound, while the others are compute-bound.

It is obtained that the matrix-vector multiplication, `cblas_zgemv`, is memory-bound, similar to the application of the absorber. This means that the way the time-dependent wavefunction is traversed when multiplying it with the mask-function needs to be made more cache-friendly for a higher performance, for example. For the `cblas_zgemv`, I used the vendor's optimised implementation of this operation and this is understood not to be able to be improved further. Because when applying the laser-field there are some non-trivial calculations to be performed to get the value of the E-field at that timestep, the datapoint for this part of the solver is moved to the right relative to the absorber. It is also visible that it is not very well optimised by the compiler. The whole workflow

for a spherical harmonics transform is obtained to be compute-bound and even although some parts of it are vectorised, its general behaviour is seen to mostly benefit from faster computations and a better usage of the single-instruction-multiple-data (SIMD) registers, for example.

3.6.3 Workload-distribution across the CPU-cores

The partitioning (decomposition) of data to processors is of crucial importance for a HPC-application. To reach the most optimum work-distribution approach, I performed various trials with different choices of the decomposition. After each re-engineering, measurements have been performed to validate assumptions and check results. Before the measurements, the application had to be instrumented accordingly. The data has been recorded by profiling techniques, as tracing campaigns can quickly go out of control due to the non-trivial amounts of storage needed to save their results. The analysis was post-mortem and not online, and mainly consisted of analysing wall-clock times, CPU-times, the amounts of bytes moved (sent and received), or synchronisation issues appearing due to work-imbalance. All these efforts have been possible due to the HPC-toolkit `ScoreP` [118]. For this section, the discretisation was $N_r = 251$, $nlat = 62$, $mmax = 61$ and the simulation ran for 3000 timesteps. The visualisation of the results is done with the help of the `Cube-GUI` [4]. This software is extremely appealing since it visually shows, via a colour-coded 2D-heatmap, how much work the processes (represented across the horizontal axis) and the threads spawned from each process (represented across the vertical axis) receive. It also automatically calculates some given scores (metrics) which are useful and quick to read indicators of how much the solver's performance benefits from the distribution of the work. A current limitation of this tool is that within one rank, `Score-P`, at least for the version used in this work (version 8.1), does not see any

[4] Cube, which is used as performance report explorer for `Scalasca| and Score-P|, is a generic tool for displaying a multi-dimensional performance space consisting of the dimensions (i) performance metric, (ii) call path, and (iii) system resource: https://www.scalasca.org/scalasca/software/cube-4.x/`

un-instrumented library routines. Thus, even if the matrix-vector multiplications for the field-free propagation part are performed using more than one thread when the matrices are not small, some scores reported for that non-trivial and important part of the solver are close to zero (i.e. very-poor scores) because the `OpenMP` threads are not all seen, but only the main thread which calls the `cblas_zgemv(...)` function. This is why a combination of HPC analysis-tools is needed to inspect such a solver (or any other HPC application).

From the methodology presented in Sec. 3.4.2, it is to be expected that the reported `MPI`-scores are rather high across all parts of the solver. This is because each rank receives a portion of the radial coordinate in the grid-representation and a portion of the total amount of (l, m)-pairs to be propagated in the spectral-representation, and these portions are seen, from Alg. 2 and Alg. 3, to be as close to even as possible.

Selected results as seen with the help of `Cube-GUI` are presented in Fig. 3.15, 3.16, 3.17, 3.18, 3.19 and 3.20.

The first two figures show the behaviour of the field-free propagation, i.e. the multiplication of the radial propagators with the corresponding slices from the time-dependent wavefunction. It is seen that the work distribution is close to perfection from the point of view of the `MPI`-ranks: the colour-coded topology shows the same colour for all the 8 ranks used in this simulation. Similarly, the automatically-calculated scores are close to the theoretical maximum, which is not perfectly achieved because the discretisation parameters for this simulation were especially chosen by myself such that they are not evenly divided by the number of ranks used here. This was done to showcase the general applicability of the solver for any given possible spatial discretisation.

Then, the next two figures are devoted to presenting the behaviour of the spectral-to-grid transform. The inverse transform, i.e. the grid to spectral one, follows the same trends and is not presented anymore as it does not add any new insights. The spectral to grid transform is efficiently distributed to all the cores. This is visible from Fig. 3.18. Indeed, each rank receives a given portion of the

number of radial grid-points and performs the same operation on each of them. Thus the routine for this is further parallelisable also inside one `MPI`-rank, and the `OpenMP`-threads are efficiently occupied. The colour-coded heatmap confirms this, and the reported performance-scores are very high. Again, these are not actually at the theoretical maximum due to, in this case, a small load-imbalance across the radial-coordinate, which translates into synchronisation issues at the end of the `#pragma omp parallel for` regions which encompass the logic for the spherical harmonics transform.

Finally, the last two figures show the distribution and the scores of the (second) `MPI_Alltoallv(...)` communication, that is, the one which helps establishing the correct data partitioning among the ranks for starting a new timestep of the simulation. While the load balance is good, the communication efficiency is only ranked as fair. Defined as a metric which identifies portions of code in which large amount of times are spent in communication rather than in performing useful computations, this communication efficiency is further decomposed into two sub-metrics: the serialisation efficiency (processes waiting) and the transfer efficiency (processes transferring amounts of data considered large relative to the network bandwidth). Depending on the simulation, the values of both of these can become non-trivial, lowering down the communication efficiency score.

3.7 Conclusion and outlook

This chapter started with a general review of the various computational techniques which have been invented in the past and used to numerically solve TDSE. The split-operator in the energy-representation is at the heart of the GPS method.

This was the focus of the next section, where I detailed the theoretical treatment of TDSE in the GPS-formalism. Before the start of any time-dependent simulation, the eigenvalues and the eigenvectors of the field-free Hamiltonian are needed to be obtained. These are a very important part of the recipe to create the radial propagators which dictate the temporal dynamics of the electron. I then finished the considerations by filling the gaps remaining in the discussion

about the time-evolution: I presented the spectral to grid and grid to spectral transforms, together with the application of the laser field. Having established the theory and numerical discretisation details, I looked at how the observable of interest in my case, the PMD, is obtained from the time-dependent wavefunction.

As the chapter continued, the focus shifted to the design and implementation of the 3D TDSE numerical solver in the GPS-formalism in a fully-distributed fashion. In order to build the understanding needed to obtain such a result, preliminary considerations were necessary (how is the parallelism exploited intra-node only), and I went on this route. With this in mind, I made the leap to the inter-node communication patterns required for the scaling across cores placed not necessarily inside the same node.

The whole workflow has been benchmarked against PMD's well-known in the literature.

I then analysed its performance using various tools and metrics which help judging the quality of a HPC-application. I obtained linear scaling in the strong-scaling regime and a-linear-trend scaling in the weak-scaling regime, then looked carefully at the workload distribution across the cores. All these show that more computational resources can be invested in a TDSE simulation and the (wall-clock-) time-to-solution will be shorter. Importantly, this respects, at least for the parameters I tested, a linear behaviour, making the fully-distributed 3D-TDSE-GPS approach a viable tool for studying ultra-intense laser-atom interactions.

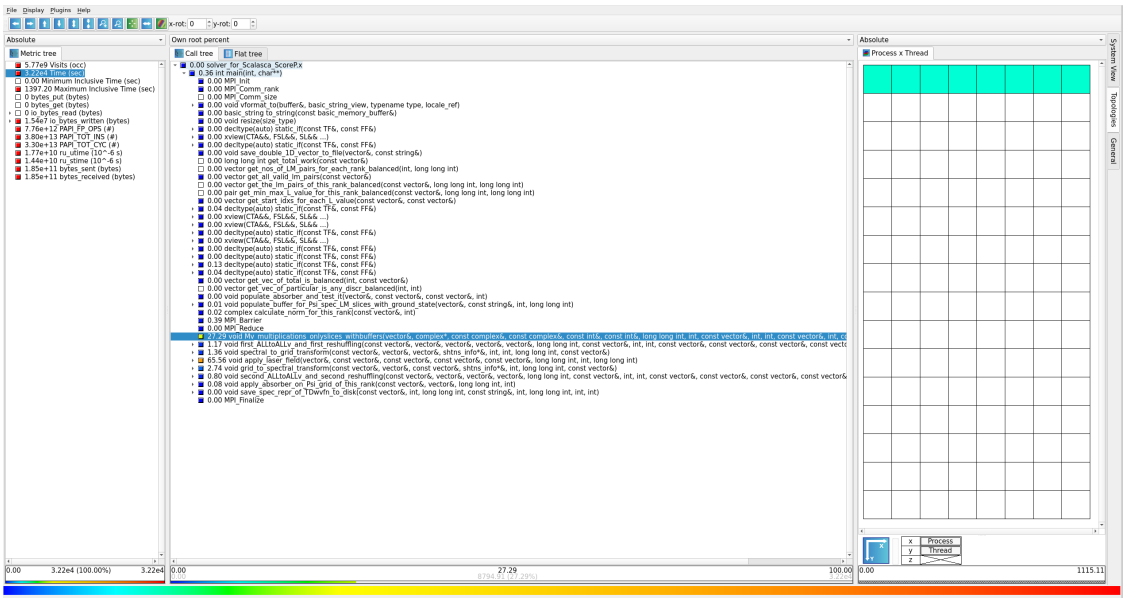


Figure 3.15: Workload distribution of the field-free propagation as shown by `Score-P`. The colourmap on the right shows across the horizontal axis the processes and along the vertical axis the threads spawn by each process. The colourbar for the data of the colourmap is presented on the bottom of the figure. As any external function which uses threads in its implementation is not currently seen as multi-threaded in `Score-P`, there are no threads visible, even if the `cblas_zgemv(...)` calls ran on multiple `OpenMP` threads.

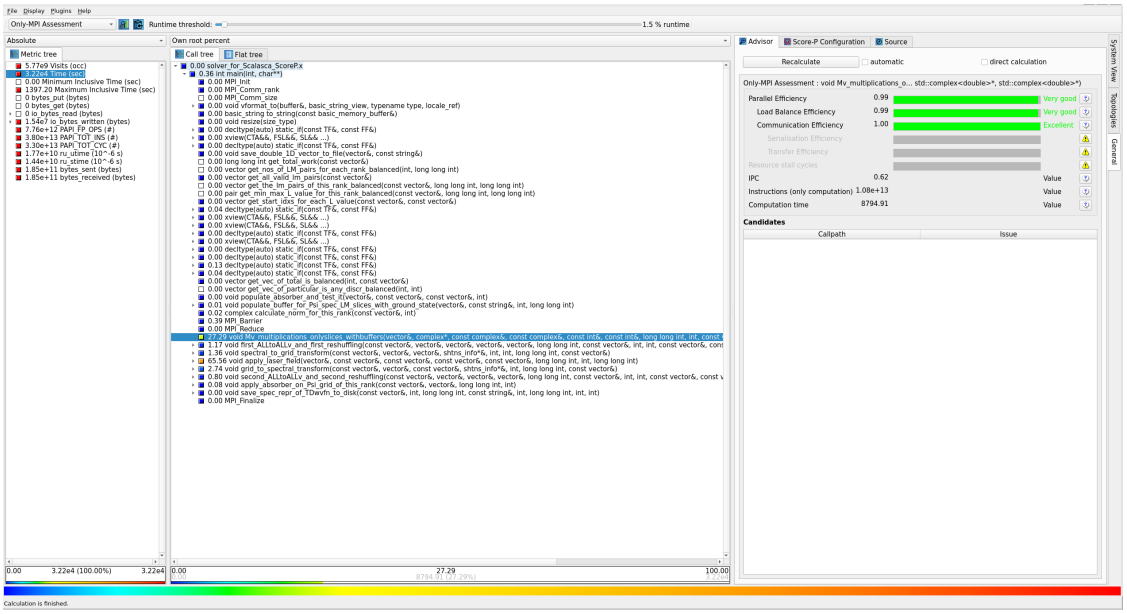


Figure 3.16: The parallelisation scores given by `Score-P` for the field-free propagation. As this is a fully and equally distributed part of the solver, the obtained numbers are close to perfection. The maximum score, 1.0, is not achieved as the spatial-discretisation for these examples was especially chosen as not to be evenly divisible by the number of processes.

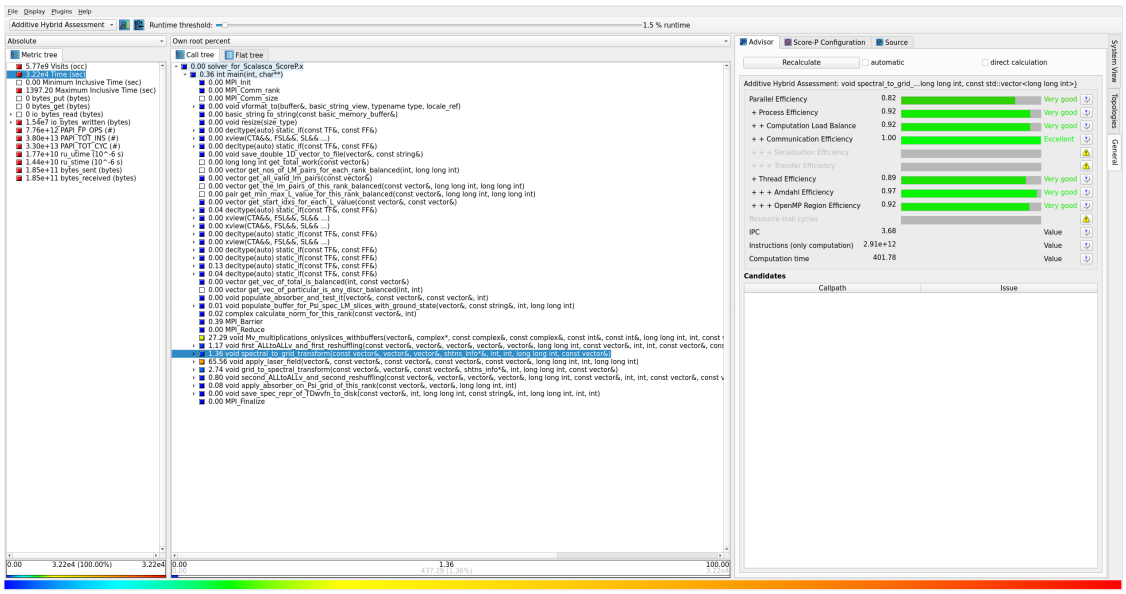


Figure 3.17: Spectral to grid transformation parallelisation scores as obtained using `Score-P`. As for each radial grid-point the same operations are performed, these grid-points can all be transformed in parallel. Thus the resulting scores for this operation are rather high, even for a grid-discretisation which is not evenly divisible by the number of the processors.

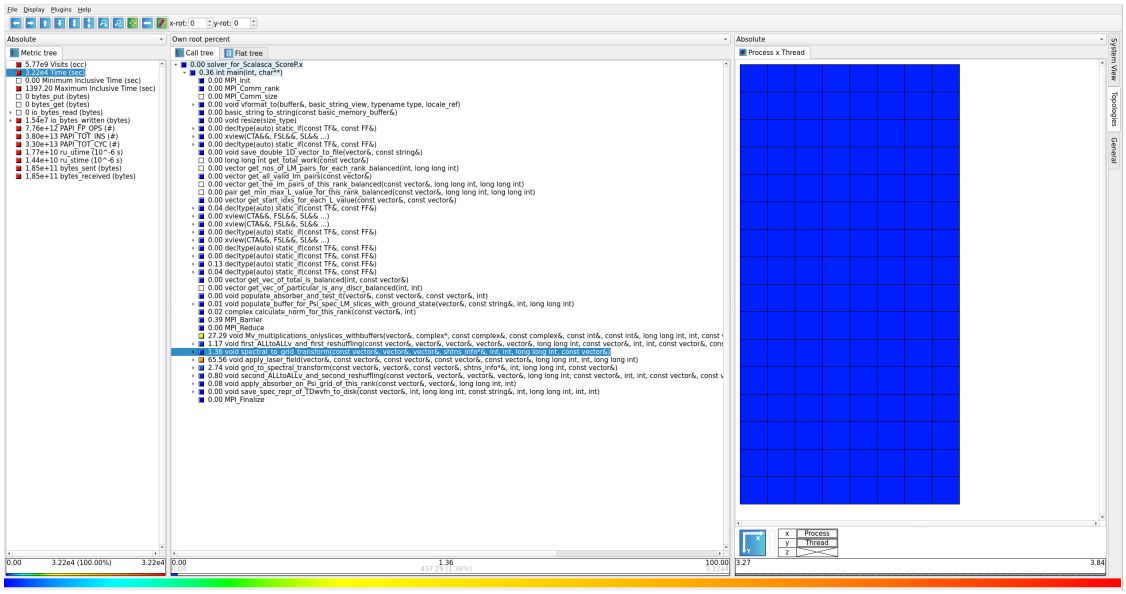


Figure 3.18: Spectral to grid transformation workload distribution as obtained using `Score-P`. Each `MPI`-process occupies one column from the heatmap on the right. The `OpenMP`-threads spawned by each process are visible across the vertical axis from the heatmap. The heatmap shows very similar colours across all its locations, indicating a rather well-balanced workload-distribution for the spectral to grid transformation. The colourbar for the data of the heatmap is presented on the bottom of the figure.

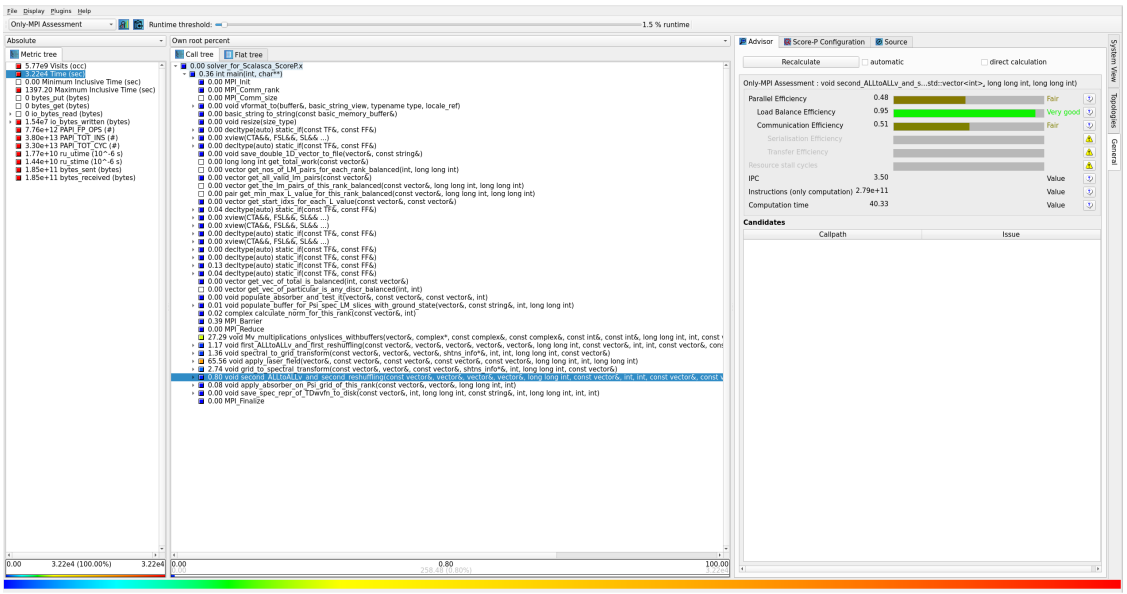


Figure 3.19: The second all-to-all (variable) communication parallelisation-scores obtained using `Score-P`. While the load-balance seems to approach a good one, the communication efficiency is only a fair one. This is because the costs associated with such a massive exchange of data across the processors are not trivial.

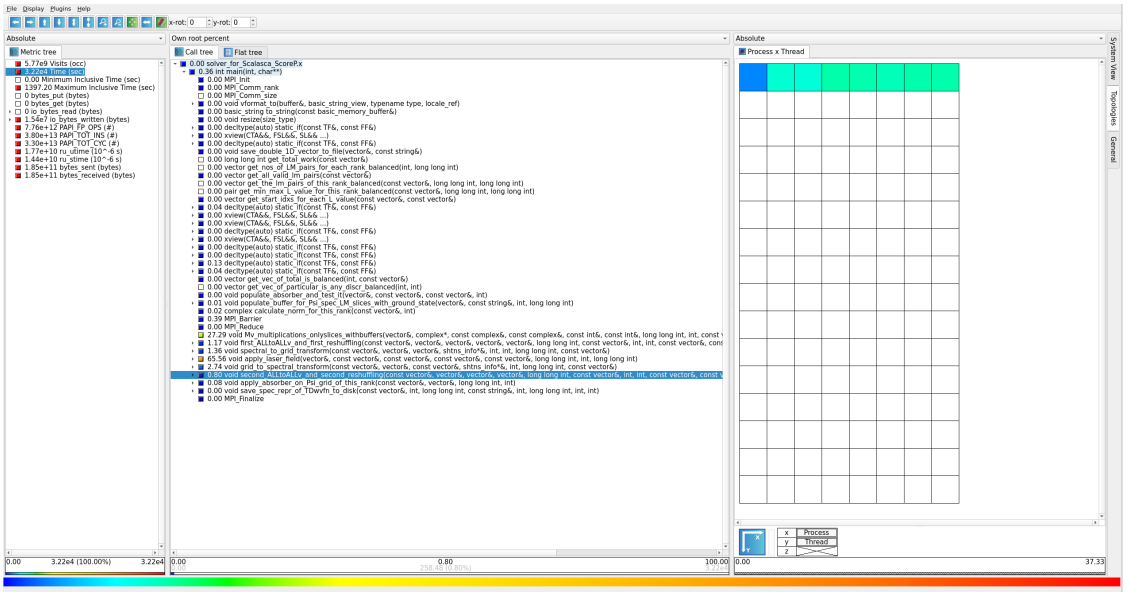


Figure 3.20: The workload distribution for the second all-to-all (variable) communication obtained using `Score-P`. As this communication is (potentially highly) non-uniform as each processor sends and receives (potentially highly) different amounts of data, the colours shown on the heatmap on the right are not the same across the processors, showcasing that this operation is rather costly and not perfectly (i.e. linearly) scalable across cores. The colourbar for the data of the heatmap is presented on the bottom of the figure.

4

Beyond-dipole-approximation electron quantum-dynamics investigations as performed using a length-gauge Hamiltonian correct up to first order in $1/c$

"The purpose of computing is to gain insight."

— Richard Hamming in "Numerical Methods for Scientists and Engineers"

"All my computing experience tells me, if something is wrong with the results, do not blame the computer, compiler, math library etc. Take a closer look at your own code."

— Dmitry

Contents

4.1 Introduction 112
4.2 The length gauge Hamiltonian, including non-dipole corrections up to first order in $1/c$ 115
4.3 Results and discussion 120
4.3.1 Ionisation probabilities at the end of the pulse 122
4.3.2 Angle-integrated dipole versus non-dipole photoelectron-momentum-distributions 126

4.3.3 Angle and momentum-resolved two-dimensional $p_z - p_y$ heatmaps 129
4.4 Conclusion and outlook 136

4.1 Introduction

Light is one of the most important means to study matter and understand the complex intricacies which come into play when considering the time-resolved dynamics of the very fast particles called electrons. There has been tremendous effort invested into making light sources more performant, adaptable and easy to use in experiments aimed at observing new phenomena or confirming theoretical predictions. In particular, a class of lasers which allow making real steps towards, for example, making a molecular movie [140], are the X-ray free-electron lasers [141].

The EuXFEL [142] has started user operation in 2021 and delivers astonishing beam-pulses of light: high-brilliance, low divergence, very tunable wavelengths and shapes. At the FLASH/FLASH-II facility many successful experiments have been carried out [143]. FERMI operates in the extreme ultra violet (EUV)-soft X-ray regime and produced scientific results over more than 10 years now and its upgrade plans to version 2.0 have been published [144]. LCLS [145] and the future LCLS-II enable scientific and technological progress: there are end-stations for atomic-molecular optics (AMO) and ultrafast X-ray science. With its hard X-rays of the uttermost quality, LCLS allows physicists to produce value in, e.g. material science or drug discovery, highlighting the complexity of the real world and the large effort needed to understand it from a fundamental point of view. The SwissFEL at the Paul Scherrer Institute in Villigen, Switzerland is also to be mentioned among all these experimental facilities.

High harmonic generation (HHG) [4, 7, 80, 81] from noble-gases can also produce pulses of light suitable for ultra-advanced electron imaging techniques.

A means to convert infra-red (IR) light to a soft X-ray laser or an attosecond-pulse, it has, and continues to attract broad interest [146, 147], receiving the 2023 Physics Nobel Prize award.

HHG originating from the plasma created at the interaction of a high-power optical-laser with a solid target could also be used as the source of temporally short, EUV radiation which could be then put to interact with the quantum-system under consideration [148, 149]. Recently, a series of two large-scale experiments, one at the ORION laser-facility and one at the Central-Laser-Facility (CLF), have been performed by the leader of my research group, Professor Peter Norreys and his collaborators from the Queen's University Belfast (Professors Dromey and Yueng) and AWE plc. The pulses of light seen in the CLF experiment could be focused to a spot and be the high-frequency driver of a laser-atom interaction.

HHG from a bulk crystal is also a potential feasible route for attosecond light sources [150].

Numerical simulations are needed to compare the theory with the results from the interaction chamber, irrespective to the driver with which the experimental data has been generated.

Laser-atom interaction is studied at its most fundamental level by means of numerically solving the time-dependent Schrödinger equation (TDSE) for the non-relativistic regime or the time-dependent Dirac-equation (TDDE) if the peak value of the electrical field (E-field) incident onto the target is too large for the non-relativistic regime, or if the target itself requires a relativistic description (very highly-charged ions).

The previous chapter tried to mention different approaches to solve TDSE. All these are unique on their own and have attracted major interest and improvements over the years. Still, many of those mentioned solvers employ laser fields in the so-called dipole-approximation (DA) (Ch. 1), one of the most widely used simplifications of the real Physics in the Atomic-Molecular-Optical-Physics (AMO) community. It is indeed a very accurate approximation for many cases of interest and at its core, it neglects the spatial dependence of the

laser field and only consider its time dependence. This in turn means that no magnetic B-field is present in the interaction, and that the photon momentum is completely neglected (set to zero).

However, beyond dipole approximation (BDA) corrections need to be added when the laser pulse is very intense, long in duration, or its photons are very energetic. As the wavelength of the radiation decreases, the size of the system under consideration stops being way smaller than it, thus violating even the textbook criterion of the applicability of the DA. The signatures of the breakdown of the dipole approximation in the high-frequency regime are the change in the ionisation probability and the appearance in the photoelectron-momentum-distribution (PMD) of the lobe against the propagation direction of the laser field predicted in [151]. Since then, advanced numerical simulations and theoretical considerations have been targeted at improving the understanding of the changes in the dynamics due to going BDA [152–155].

Even in the long-wavelength regime, one needs to go BDA if an accurate treatment is sought-after: indeed, as the laser-induced velocity of the electron scales as the field strength over the frequency of the laser-field, it can become significant for low (e.g. mid-infrared) frequencies [156]. Nondipole effects in the long wavelength regime of laser-atom interaction have been experimentally measured in the pioneering experiments by Smeenk et al. and Ludwig et al. [157, 158], and more recently [159–161]. They have been considered also in the laser-molecule interaction [162–164].

Using a DA Hamiltonian when solving TDSE offers important advantages, enabling, for example, propagation by means of operators solely dependent on the spatial coordinates or the momentum coordinates: there are no products of non-commuting terms in the DA-Hamiltonians. Thus a (relatively) straightforward propagation algorithm such as the Fourier split-operator can be used for solving the TDSE with DA-Hamiltonians [84]. Also, when using the DA-Hamiltonian, a linearly-polarised pulse (along the $\hat{z}$-axis) is fully and correctly described in a two-dimensional (2D) spherical-coordinates problem setup rather

than in the full three-dimensional (3D) space. Solving the TDSE in 2D reduces the memory requirements and computation time drastically compared to the full 3D case. For the BDA-Hamiltonian, the problem becomes full-3D even when considering a linearly-polarised pulse along the $\hat{z}$-axis.

Given the wide interest in treatments of quantum systems exposed to extreme pulses of light, this chapter shows that correct BDA results can be obtained using a Hamiltonian of the same structure as a DA-Hamiltonian, but having included in it corrections up to first order in $1/c$, where c is the speed of light. Unless otherwise stated, atomic units are used throughout this chapter ($e = m_e = \hbar = 1/(4\pi\epsilon_0) = 1$). The theoretical treatment of this highly-elegant Hamiltonian has its roots in [165]. The scope of this chapter is to show that *ab initio* predictions obtained by numerically solving 3D TDSE with this Hamiltonian in the *length gauge* by means of the GPS method are accurate and computationally efficient. It is also touched upon the fact that quantum-mechanical observables should be the same in the length and in the velocity-gauges, if both the simulations are converged to some degree. This paves the way to theoretical investigations of non-dipole effects in TDSE using solvers which natively employ DA Hamiltonians in the length-gauge, making it an appealing route to running full-dimensionality simulations aimed at regions of the parameters space for which the DA is questionable.

4.2 The length gauge Hamiltonian, including non-dipole corrections up to first order in $1/c$

The TDSE

$$i\frac{\partial \Psi(\vec{r},t)}{\partial t} = \mathcal{H}\,\Psi(\vec{r},t) \tag{4.1}$$

is considered in spherical coordinates, where $\vec{r} = (r,\theta,\varphi)$ and the interaction Hamiltonian is $\mathcal{H} = \mathcal{H}_0 + V(\vec{r},t)$. The time-dependent part of $\mathcal{H}$ is the coupling with the laser E-field expressed here in the length gauge $V(\vec{r},t) = \vec{r}\cdot\vec{E}(t)$. The time-independent part $\mathcal{H}_0$ represents the field-free (unperturbed) Hamiltonian and is composed from the attraction with the core ($V_C(r) = -1/r$ for atomic

hydrogen), the centrifugal part and the radial, second order derivative (together in the $-\nabla^2/2$).

For a laser pulse described purely in the DA (Ch. 1), the form of $\vec{E}(t)$ can be either linearly-polarised (along $\hat{z}$-axis and the problem becomes two-dimensional (2D): there will be no φ-angle dependence because the system possesses rotational symmetry around the $\hat{z}$-axis) or elliptically-polarised (and the problem is known to become fully 3D). The E-field for this DA-Hamiltonian will be denoted as $\vec{E}_{\mathrm{DA}}(t)$ and it has one non-zero component for a linearly-polarised laser-pulse and two non-zero components for an elliptically-polarised laser-pulse when expressed in 3D Cartesian coordinates.

In this section theoretical aspects related to an effective E-field which is a function only of the temporal parameter (just as a DA-E-field), but includes non-dipole corrections up to first order in $1/c$ and which will be denoted as $\vec{E}'_{\mathrm{BDA}}(t)$, are presented. This effective E-field is suitable for numerically solving the TDSE (Eq. (4.1)) beyond the dipole-approximation, by substituting the usual interaction with the laser field $V(\vec{r},t) = \vec{r}\cdot\vec{E}_{\mathrm{DA}}(t)$ with $V(\vec{r},t) = \vec{r}\cdot\vec{E}'_{\mathrm{BDA}}(t)$ and applying exactly the same iterative algorithm (time-stepping) as for the DA-E-field case (previous chapter of this thesis).

The end-result which is obtained from the calculations of this subsection is that a DA-E-field, linearly-polarised along $\hat{z}$, "acquires" a component along its propagation direction $\hat{x}$ and thus turns into a pseudo-E-field (which resembles an elliptically-polarised one) when it has the non-dipole corrections up to first order in $1/c$ added to it.

This makes sense: the electron, through the magnetic-force from the laser, is moved along the propagation direction of the laser: this movement is encompassed in the non-zero term which is obtained for the $\hat{x}$-component of the pseudo-E-field. This force along the propagation direction of the laser is the signature for non-dipole effects up to first order in $1/c$ and it resembles the radiation pressure of the light.

In this section, the linear polarisation along $\hat{z}$ and the propagation direction along $\hat{x}$ are assumed for the laser pulse.

Any spatially and temporally dependent vector-potential $\vec{A}(\vec{r},t)$, can be expressed through a dependence on $\eta := (\omega t - \omega x/c)$ as:

$$\vec{A}(\vec{r},t) = A_0(\eta) \begin{bmatrix} 0 \\ 0 \\ \sin(\eta) \end{bmatrix}, \tag{4.2}$$

where $A_0(\eta)$ is the envelope function. After a Taylor expansion up to first order in $1/c$ of the full vector potential, one recovers the usual dipole term, related to the DA-E-field as $\vec{E}_{\mathrm{DA}}(t) := -\dfrac{\partial \vec{A}^{\,0}(t)}{\partial t}$:

$$\vec{A}^{\,0}(t) = A_0(t) \begin{bmatrix} 0 \\ 0 \\ \sin(\omega t) \end{bmatrix}, \tag{4.3}$$

and two terms which depends on the distance along the propagation direction x:

$$\vec{A}_{\mathrm{c}}(\eta) = A_0(t) \begin{bmatrix} 0 \\ 0 \\ \dfrac{-x\omega}{c}\cos(\omega t) \end{bmatrix}, \tag{4.4}$$

and

$$\vec{A}_{\mathrm{e}}(\eta) = -\partial_\eta A_0(\eta)|_{\eta=\omega t} \begin{bmatrix} 0 \\ 0 \\ \dfrac{x\omega}{c}\sin(\omega t) \end{bmatrix}, \tag{4.5}$$

where the subscripts c and e mean the carrier variation and the envelope variation, respectively.

For a $\sin^2$ envelope:

$$A_0(\eta) = A_{\max} \sin^2\left(\frac{\pi}{T}\frac{\eta}{\omega}\right), \tag{4.6}$$

where T is the duration of the laser pulse ($T = N \cdot 2\pi/\omega$, with N the number of optical-cycles in the laser pulse). Similarly,

$$A_0(t) = A_{\max} \sin^2\left(\frac{\pi t}{T}\right). \tag{4.7}$$

It is then easy to calculate $-\partial_\eta A_0(\eta)\,|_{\eta=\omega t}$ as:

$$-\partial_\eta A_0(\eta)\,|_{\eta=\omega t} = -A_{\max}\,\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right). \tag{4.8}$$

Eq. (4.5) becomes:

$$\vec{A}_{\mathrm{e}}(\eta) = -A_{\max}\,\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right)\begin{bmatrix}0\\0\\\frac{x\omega}{c}\sin(\omega t)\end{bmatrix}. \tag{4.9}$$

Now $\nabla\times\vec{A}_{\mathrm{c}}$ and $\nabla\times\vec{A}_{\mathrm{e}}$ are calculated:

$$\begin{aligned}\nabla\times\vec{A}_{\mathrm{c}} &= A_0(t)\begin{vmatrix}\hat{\imath}&\hat{\jmath}&\hat{k}\\\partial_x&\partial_y&\partial_z\\0&0&\frac{-x\omega}{c}\cos(\omega t)\end{vmatrix} =\\ &= A_0(t)\begin{bmatrix}0\\\frac{\omega}{c}\cos(\omega t)\\0\end{bmatrix},\end{aligned} \tag{4.10}$$

$$\begin{aligned}\nabla\times\vec{A}_{\mathrm{e}} &= -\begin{vmatrix}\hat{\imath}&\hat{\jmath}&\hat{k}\\\partial_x&\partial_y&\partial_z\\0&0&\frac{x\omega}{c}\sin(\omega t)\end{vmatrix}\cdot A_{\max}\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right) =\\ &= -A_{\max}\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right)\begin{bmatrix}0\\-\frac{\omega}{c}\sin(\omega t)\\0\end{bmatrix}.\end{aligned} \tag{4.11}$$

In total, $\left(\nabla\times\vec{A}_{\mathrm{e}}+\nabla\times\vec{A}_{\mathrm{c}}\right) \coloneqq \vec{B}(t)$ is the spatially-independent magnetic B-field:

$$\vec{B}(t) = \begin{bmatrix}0\\A_0(t)\frac{\omega}{c}\cos(\omega t) + A_{\max}\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right)\frac{\omega}{c}\sin(\omega t)\\0\end{bmatrix}. \tag{4.12}$$

In [165], the authors argue that accounting for the Lorentz force only along the laser-induced dipole trajectory is an accurate approximation for the effect of the full (spatially and temporally-dependent) B-field on the electron's motion. This approximation is technically valid only for long laser pulses (longer than, e.g. 5-6 optical-cycles) which are not circularly-polarised. The quality of this approximation becomes better as the pulse is longer and closer to being linearly polarised. This way of modelling the electron dynamics gives rise to a mixed-gauge Hamiltonian (having one term dependent on the momentum operator and

one term dependent on the coordinate-space operator) which is then cast onto the length-gauge (thus obtaining the effective E-field dependent only on the time parameter and including the non-dipole corrections up to first order in $1/c$).

I present here how this effective E-field is obtained, including all the (non-trivial) details omitted in [165]. Apart from the calculations which have been derived above, it is needed to take the cross product between $-\vec{A}^{\,0}(t)$ and $\vec{B}(t)$:

$$\begin{aligned} &-\vec{A}^{\,0}(t)\times\vec{B}(t)= \\ &\begin{vmatrix} \hat{\imath} & \hat{\jmath} & \hat{k} \\ 0 & 0 & -A_0(t)\sin(\omega t) \\ 0 & \frac{\omega}{c}\left(h_1+h_2\right) & 0 \end{vmatrix}, \end{aligned} \tag{4.13}$$

where $h_1 \coloneqq A_0(t)\cos(\omega t)$ and $h_2 \coloneqq A_{\max}\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right)\sin(\omega t)$, so the result of Eq. (4.13) is equal to:

$$\begin{aligned} &\vec{A}^{\,M}(t) \coloneqq -\vec{A}^{\,0}(t)\times\vec{B}(t)= \\ &= \hat{\imath}\,A_0(t)\sin(\omega t)\cdot \\ &\cdot\left(A_0(t)\frac{\omega}{c}\cos(\omega t)+A_{\max}\frac{\pi}{T\omega}\sin\left(\frac{2\pi t}{T}\right)\frac{\omega}{c}\sin(\omega t)\right) \end{aligned}, \tag{4.14}$$

where it has been defined that this cross product between $-\vec{A}^{\,0}(t)$ and $\vec{B}(t)$ to be denoted by $\vec{A}^{\,M}(t)$. The effective E-field $\vec{E}^{\,'}_{\mathrm{BDA}}(t)$ is shown in [165] to be equal to

$$\vec{E}^{\,'}_{\mathrm{BDA}}(t) = -\left(\frac{\partial\vec{A}^{\,M}(t)}{\partial t}\right)-\left(\frac{\partial\vec{A}^{\,0}(t)}{\partial t}\right). \tag{4.15}$$

The prime sign is not related to a temporal derivative and is just a notation to signify that it is an effective E-field. It acquires a non-zero component along the propagation direction of the laser pulse ($\hat{x}$):

$$\vec{E}^{\,'}_{\mathrm{BDA}}(t) = \begin{bmatrix} -\frac{\omega}{c}A_{\max}^2\left(E_x^{(a)}+E_x^{(b)}\right) \\ 0 \\ -\frac{\partial A_0(t)}{\partial t}\sin(\omega t)-A_0(t)\cos(\omega t)\,\omega \end{bmatrix}, \tag{4.16}$$

where $E_x^{(a)} \coloneqq \sin^4\left(\frac{\pi t}{T}\right)\frac{\sin(2\omega t)}{2}$, $E_x^{(b)} \coloneqq \frac{\pi}{T\omega}\sin^2\left(\frac{\pi t}{T}\right)\sin\left(\frac{2\pi t}{T}\right)\sin^2(\omega t)$, $A_{\max}$ from Eq. (4.6) and $A_0(t)$ from Eq. (4.7).

The E-field from Eq. (4.16) is the ingredient of the length gauge, BDA-Hamiltonian, which includes corrections up to first order in $1/c$: $\mathcal{H}_0+\vec{r}\cdot$

$\vec{E}^{\,'}_{\mathrm{BDA}}(t)$. It can be readily propagated with the numerical-technique detailed in the previous chapter because it has the structure of an E-field dependent only on the temporal parameter, but includes non-dipole corrections. It does not possess rotational symmetry around the $\hat{z}$-axis: the φ-angle is present because of the $x \cdot E^{'}_{\mathrm{BDA},\,x}(t)$ which is non-zero because $E^{'}_{\mathrm{BDA},\,x}(t)$ is non-zero.

The usual DA E-field, linearly polarised along $\hat{z}$ and propagating along $\hat{x}$ is obtained by setting the $\hat{x}$-component of the E-field from Eq. (4.16) to zero.

Thus, this DA E-field can be written as:

$$\vec{E}_{\mathrm{DA}}(t) = \begin{bmatrix} 0 \\ 0 \\ -\frac{\partial A_0(t)}{\partial t}\sin(\omega t) - A_0(t)\cos(\omega t)\,\omega \end{bmatrix}. \tag{4.17}$$

This expression can also be obtained by calculating $-\partial \vec{A}^{\,0}(t)/\partial t$ with $\vec{A}^{\,0}(t)$ from Eq. (4.3). This linearly-polarised E-field can be solved in a 2D problem-setting: the $x \cdot E_{\mathrm{DA},\,x}(t)$ term is 0 in this DA-case because $E_{\mathrm{DA},\,x}(t) = 0$ and thus there is rotational-symmetry around the $\hat{z}$-axis for the DA-case.

Again, the interpretation of Eq. (4.16) is that while only an E-field is used in the TDSE-solver and there is no explicit B-field, the E-field is an effective one and it includes non-dipole effects up to first order in $1/c$, visible from the radiation pressure like force along the laser propagation direction.

These theoretical calculations which have led to the effective E-field which includes non-dipole corrections up to first order in $1/c$ allow for the usage of general propagation algorithms, designed for dipole-approximated fields, to study non-dipole effect in the laser-atom interaction. The effects obtained using highly-specialised solvers (see e.g. [90]) can be recovered by means of a length gauge 3D-TDSE solver which employs this "effective" E-field dependent only on the time parameter.

4.3 Results and discussion

The laser frequency of $\omega = 3.5$ a.u. (central pulse wavelength equals to approximately 13 nm) is considered in this section, i.e. the laser-photons have their

energy equal to 95 eV. This is at the border of entering the soft X-ray regime, technically still being in the EUV regime.

The general idea of the results presented in this section is valid also for (high-) frequencies greater or lower in magnitude than this, depending on the field strength and temporal duration.

Here, linear polarisation along $\hat{z}$ is assumed, a $\sin^2$ envelope is chosen, the temporal durations are between five and fifteen optical-cycles depending on the simulation and the peak E-field strength E_0 depends on the simulation as well, being between a rather low value (2 a.u.) and 60 a.u. (corresponding to 1.3×10^{20} W/cm^2 as obtained from the conversion formula I [a.u.] = E_0^2 [a.u.]).

As in this chapter the TDSE is solved, care has to be taken not to enter the fully-relativistic regime of laser-atom interaction (where TDDE shall be solved instead). Considering that a rough-estimate for the boundary for entering such a domain is that the classical velocity of the electron in the laser-field is below 10% of the speed of light $c \approx 137.0$, the highest-field calculations just start to touch beyond this boundary: the maximum classical-speed is $E_0/\omega = 60/3.5 = 17.1$ a.u.. Given that the most important (first-order) relativistic corrections are related to the B-field (the second-order corrections are associated with relativistic kinematic effects such as, e.g., the change in mass of the electron because of its very fast movement in the laser field), such an inclusion of the uniform magnetic field, as it is done in the present chapter, is justified, for the laser field parameters used here. In fact, numerical calculations of the responses given by both TDSE and time-dependent-Dirac-equation (TDDE) from [166], showcase this. There, they used the exact same physics parameters. While they do conclude that relativistic effects appear due to the mass increase, they consider them starting from $E_0 = 58 - 60$ a.u. only. This is where my numerical simulations finish, and for such field-values, the differences due to the mass increase are rather small. This discussion can be seen at their page 7, on the right-hand-side column (also presented in their Eq. (41)). The quantitative evidence is presented in Fig. 5 from the same study. Overall, the effect of the relativistic effects is to decrease

the ionisation probability compared to the non-relativistic case, as discussed at the start of Sec. IV B from the mentioned paper, and this effect increases with an increase in the peak field-strength.

4.3.1 Ionisation probabilities at the end of the pulse

The ionisation probability P_{i} is defined in this chapter to be the total contribution obtained from the projections onto the unbound eigenstates (the eigenvectors of Eq. (3.2) for which $\epsilon_{nl} > 0$ a.u.) of the time-dependent wavefunction at the end of the pulse. Accordingly, the probability to remain in a bound state is $P_{\mathrm{b}} = 1 - P_{\mathrm{i}}$. This can be thought of as a more unambiguous consideration of the potentially vague concept of ionisation probability than, e.g., the norm of the wavefunction at the end of the simulation (the temporal dependence of which is used in the present chapter to monitor the stability of the simulations and how much of the wavefunction is absorbed at the boundary of the computational domain). It is to be mentioned that in all the simulations from this chapter for which PMD's are calculated, the final norm of the time-dependent wavefunction is very close to unity (the initial value). The lowest value of the norm at the end of the simulation in the present chapter is obtained, as expected, for the highest field strength and the longest pulse-duration considered, that is 60 a.u. and fifteen optical-cycles.

The classical excursion length of the electron in the laser field is E_0/ω^2 and for the highest field strength of 60 a.u., it equals to 4.9 a.u. for the frequency considered in this section (3.5 a.u.), thus a box size of 80 a.u. correctly encompasses the dynamics during the simulation.

For this box size of $R_{\mathrm{b}} = 80.0$ a.u., the absorber started at $R_{\mathrm{a}} = 60.0$ a.u. and was a $\cos^{0.25}$ mask function [120]. For the most demanding laser field strengths and longest pulses, the GPS discretisation needed 400 radial points and 500 angular momenta values. Between 4100 and 7000 timesteps per 1 optical cycle have been used, depending on the simulation. Such a large number of timesteps per one optical cycle is motivated by the fact that the numerical error in the split-operator approach of solving the TDSE grows as the field strength

raised to second power. Thus a more intense laser field requires a finer temporal discretisation, and this effect is nonlinear.

For a fixed duration of 10 optical-cycles, I ran simulations to observe if any difference in the unbound populations appears between the usual, DA-Hamiltonian and the BDA-Hamiltonian.

The DA-Hamiltonian the TDSE is solved for in the results section (i.e. Sec. 4.3) is:

$$\mathcal{H}_{\mathrm{DA}} = \mathcal{H}_0 + \vec{r} \cdot \vec{E}_{\mathrm{DA}}(t), \tag{4.18}$$

with $\vec{E}_{\mathrm{DA}}(t)$ from Eq. (4.17) and $\mathcal{H}_0$ being the field-free (unperturbed) Hamiltonian, composed from the attraction with the core $(-1/r)$ and the $\nabla^2/2$ operator.

The BDA-Hamiltonian the TDSE is solved for in this results section (i.e. Sec. 4.3) is:

$$\mathcal{H}_{\mathrm{BDA}} = \mathcal{H}_0 + \vec{r} \cdot \vec{E}^{\,\prime}_{\mathrm{BDA}}(t), \tag{4.19}$$

with $\vec{E}^{\,\prime}_{\mathrm{BDA}}(t)$ from Eq. (4.16) and again $\mathcal{H}_0$ the field-free Hamiltonian, composed from the attraction with the core $(-1/r)$ and the $\nabla^2/2$ operator.

In all the simulations from this section, the relationship between A_{max} from Eq. (4.16) or Eq. (4.17) and E_0 is simply $E_0 = \omega A_{\mathrm{max}}$ because of the linear polarisation of the pulse considered in this section.

If these differences between the DA-BDA simulations exist, the discrepancy shall increase as the peak field strength increases.

Results for the duration of 10 optical-cycles at the laser-frequency of $\omega = 3.5$ a.u. and at a range of peak E-field values (one value per one TDSE simulation) are presented in the top panel of Fig. 4.1.

One would expect that with a higher peak-field strength, the electron has a higher chance to populate an unbound state.

This indeed happens for relatively modest field strengths, up to approximately 12 a.u. Then, a local maximum in the ionisation probability versus peak E-field strength dependence is seen. After that, the phenomenon of atomic stabilisation develops [167].

The atomic stabilisation is counter-intuitive and does not respect the expectation that a more intense laser pulse will ionise the atom more. Instead, in both the DA and BDA cases, even if the peak field strength keeps increasing, the atom does not keep the rate of increase in the ionisation probability seen for the abrupt slope from the modest field strengths region. The probability of ionisation actually drops a little after reaching the local maximum and somehow stabilises (remains roughly the same) for quite a broad range of peak E-field values before it starts to significantly increase again.

Visible differences between DA and BDA start to appear for a peak field strength of around $38 - 40$ a.u. for the pulse duration of 10 optical-cycles. These differences are more significant as the peak value of the field increases, as expected by looking at Eq. (4.16) where the x-component of the BDA, effective, E-field, depends on the A_{max}^2 term. The BDA Hamiltonian is shown to have a detrimental effect on the atomic stabilisation mentioned above: it hinders this phenomenon.

In the middle and bottom panels of Fig. 4.1, results for similar simulations are presented for pulse durations of 15 optical cycles and 5 optical cycles. For 15 optical cycles, the trends observed are the same as for the 10 optical cycles simulations: after the initial rapid rise in ionisation probability, atomic stabilisation in the super-strong, high-frequency field appears. Visible differences between DA and BDA Hamiltonians start to show up from values of the peak E-field around $36 - 38$ a.u.. Again, the BDA simulations clearly inhibit the degree of atomic stabilisation the atom shows.

For the 5 optical-cycles simulations, the atomic stabilisation does not fully-develop anymore. There is a change in the slope of the increase of ionisation probability versus the peak field-strength and a very narrow region in which this former quantity is roughly constant, but the overall trend shows a visible, clear increase throughout all the simulated peak-field values. It is thus the case that the atomic stabilisation depends on the pulse duration. To be complete and exact,

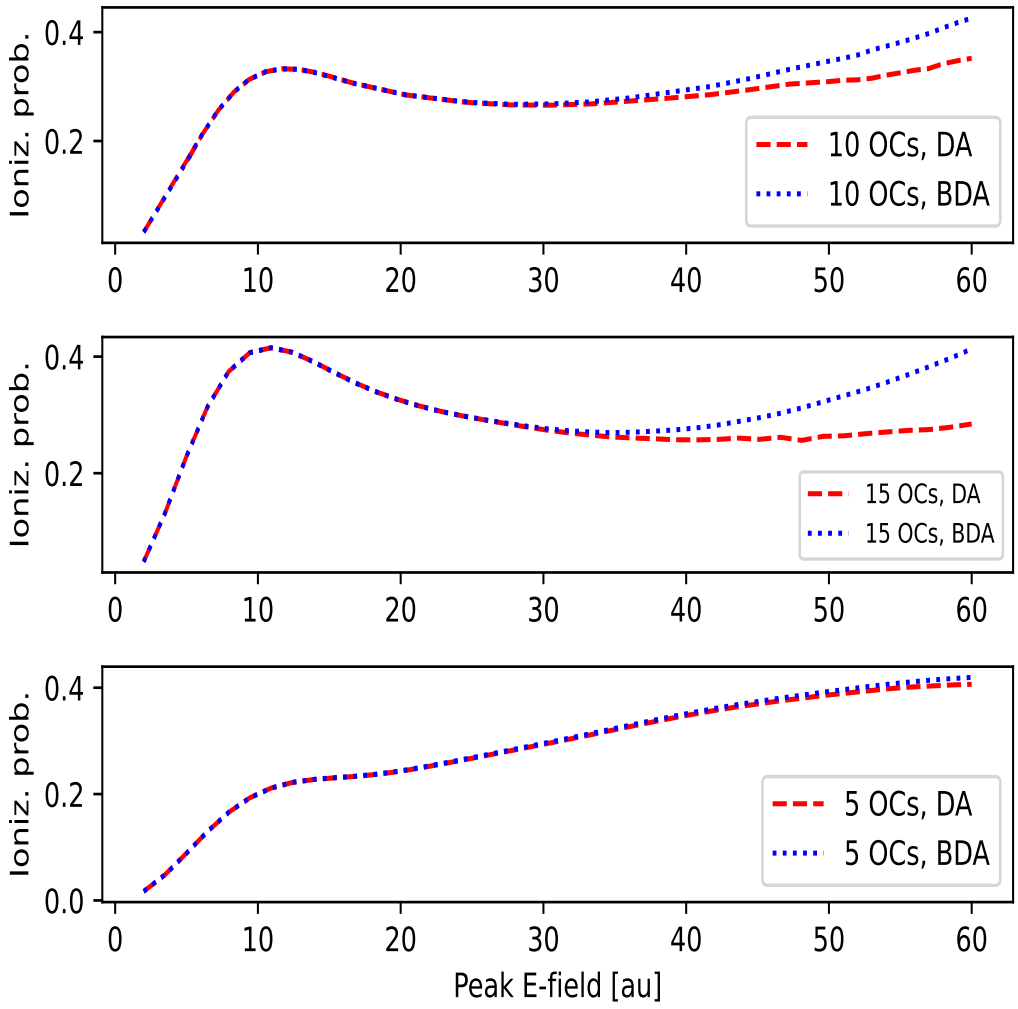


Figure 4.1: Total ionisation probabilities at the end of dipole-approximation (DA) and beyond dipole-approximation (BDA) TDSE-simulations with a laser-pulse of frequency $\omega = 3.5$ a.u. and various peak E-field strengths, for a $\sin^2$ temporal profile of 10 optical-cycles (upper panel), 15 optical-cycles (middle panel) and 5 optical-cycles (lower-panel).

it depends on the temporal shape of the pulse, in particular on the wings of the envelope (i.e. how quickly the envelope drops to zero) [168].

4.3.2 Angle-integrated dipole versus non-dipole photoelectron-momentum-distributions

The photoelectron energy spectrum (PES) is calculated as from Eq. (3.15), making use of the direct projection onto the scattering states (with proper asymptotic conditions) of the target atom at the end of the laser pulse. PES (or the angle-integrated PMD) can be measured experimentally and can reveal important information about the quantum-dynamics.

For this subsection, the results are presented in Fig. 4.2, where there are showcased, down one column from top to bottom, temporal durations of 10 and 15 optical cycles, for two columns, representing peak field values of 30 and 45 a.u., in order from left to right.

In all spectra, the first ATI peak is seen at 3.0 a.u. of energy and this value is equal to $\omega - I_\mathrm{p} = 3.5$ a.u. -0.5 a.u. for atomic hydrogen.

This is explained by the fact that the ATI peaks follow the energy-conservation relationship

$$E_{\mathrm{K},n} = E_{\mathrm{K},0} + n\omega, \tag{4.20}$$

where n represents the number of photon absorbed by the atom ($n = 1$ for the first ATI peak) and $E_{\mathrm{K},0}$ is the (unperturbed) ground state-energy. That is, they are spaced one photon energy apart due to the $n\omega$ term.

The fact that the peaks are not perfectly sharp (i.e. a Dirac-delta function centered at $E_{\mathrm{K},n}$) is because of the finite laser pulse duration (i.e. the pulses used in this chapter are not monochromatic, i.e. they are temporally-short).

One can also observe oscillations in the spectra inside the ATI peaks. They are due to the electron re-colliding with the nucleus: interferences appear due to these multiple re-scattering events [90]. These oscillations are faster (their period is smaller, i.e. more sub-peaks appear inside the ATI peak) for a longer pulse length.

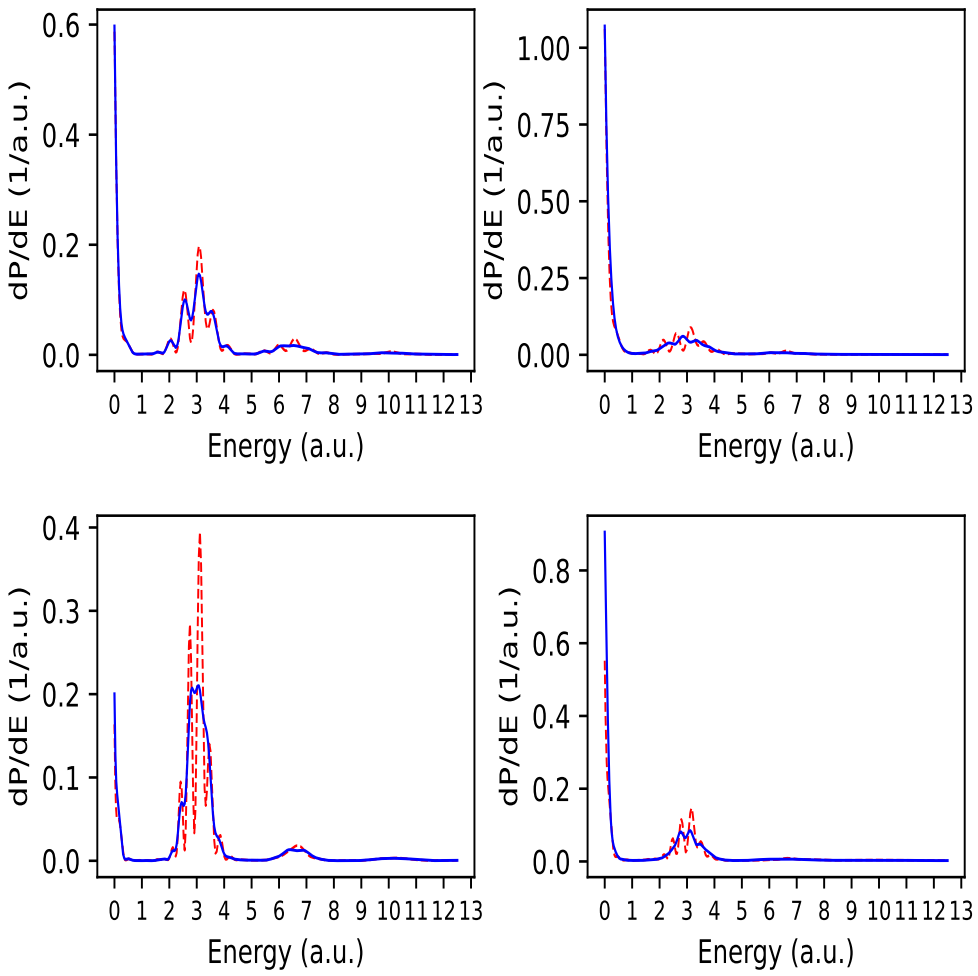


Figure 4.2: Photoelectron energy spectrum (PES) plots as obtained at the end of TDSE-simulations. Top panels: the pulse was of 10 optical-cycles in duration. Bottom panels: the pulse was of 15 optical-cycles in duration. Left column: the pulse had a peak field strength of 30.0 a.u. Right column: the pulse had a peak field strength of 45.0 a.u. In all panels, the dashed red curve represents the dipole-approximated (DA) Hamiltonian's results and the solid, blue curve represents the beyond-dipole-approximation (BDA) Hamiltonian's results. The first ATI peak is seen at 3.0 a.u. of energy and the oscillations inside it (due to recolissions with the core) are diminished when going BDA.

The LES seen here (the part of the spectrum located around very low kinetic-energies) is not the same as the one observed with mid-IR laser fields [169]. In the case of super-strong, high-frequency fields (as it is the case here), it can be explained by a two-photon (Raman-type) process: one photon is absorbed from the laser field and another one is emitted into the laser-field and their energy difference is transferred to the electron [170].

One expects that the oscillations seen in the ATI peaks become less intense when going BDA. This is because the BDA dynamics do not favour the electron to recollide with the core so many times anymore [90]. The explanation is that when going BDA, the force along the propagation direction pushes the electron to the continuum and re-encountering the nucleus is less likely because of this radiation pressure-like force (seen in the x-component of the effective E-field from Eq. (4.16)). The oscillations diminish as expected for all the field parameters considered. The most drastic change in their amplitude is obtained for the longest pulse-duration considered in this work: 15 optical-cycles (bottom panels of Fig. 4.2).

It is instructive to compare and contrast the angle-integrated spectra with those obtained from solutions of the TDDE numerically solved in full dimensionality. At least from the point of view of CPU-time, the effort which shall be devoted to solve this equation, rather than the TDSE, is significantly larger: the matrices involved in the calculations are 4 times larger than those of the TDSE workflow. As the operations from BLAS level 2 (matrix-vector) and BLAS level 3 (matrix-matrix) all scale non-linearly with respect to the input size ($\mathrm{O}(N^2)$ and $\mathrm{O}(N^3)$, respectively, where (N, N) is the shape of the input matrix), thus the computational cost increases significantly, even for a modest number of grid points. Still, there have been successful attempts to obtain such results for the X-ray regime. In [171], the first ATI peak is obtained for the singly-ionised helium at $E_{K,0} = 12.0$ a.u., and after re-scaling (by a division of Z^2, with $Z = 2$ is the nucleus charge for atomic He), this is equivalent to atomic hydrogen at $E_{K,0} = 3.0$ a.u., similar to the results presented in this chapter. It is important

thus to realise that with a significantly reduced cost and a much easier computer-implementation effort, one can obtain electron spectra which agree with the most general formalism one can use when gauging about a single electron in an ultra-strong laser field, at least for the parameters targeted in this chapter.

4.3.3 Angle and momentum-resolved two-dimensional $p_z - p_y$ heatmaps

Pseudocolour 2D momentum distribution plots are presented in this section. They are cross sections at $p_y = 0$, that is, the momenta along the laser polarisation direction ($\hat{z}$) and along the laser propagation direction ($\hat{x}$) are shown.

Over the history, depending on the laser pulse parameters, notable PMD characteristics which signify specific laser-atom interaction processes have been discovered: the ring structure (circles spaced one photon energy apart) [1], the low energy structure (LES) [169], the very low energy structure (VLES) [172], an interference pattern due to electrons emitted at different times [173], or the fork structure [174], to name some of them. These are all mentioned only and not discussed further here, in order to showcase that important information can be extracted from the patterns observed in the (measured or simulated) PMDs.

When in the DA, the PMD's do not depend on the φ-angle, and thus are rotationally-symmetric around the $\hat{z}$-axis because there is no other non-zero component in the DA-E-field than the $E_z(t)$. This means that the top-bottom half-planes are identical if they emerge from a DA-TDSE simulation.

The left-right half-planes are not necessarily identical: they become perfectly equal to each other only in the limit of an infinitely long pulse. For any real pulse, this is not the case, and the difference between these half-planes is a carrier-envelope phase (CEP) effect. Still, depending on the laser characteristics, the left-right half-planes can become roughly (visually) the same for as low as 10-15 optical-cycles pulses.

The upper panels from Fig. 4.3 and from Fig. 4.4 present DA-results for peak-field values of $E_0 = 30$ a.u. for the former figure and $E_0 = 45$ a.u. for the later

figure, respectively. Both figures showcase heatmaps on two columns, for 10 optical-cycles pulse duration on the left column and for 15 optical-cycles pulse duration on the right column.

In all these plots, the electron is ionised mainly along the laser polarisation direction (that is, there is a strong maximum at $p_x = 0$ a.u.).

The ATI-circles are also visible: along these curved lines the electron is shown to be ionised with quantised kinetic energies (see Eq. (4.20)).

The oscillations seen inside one ATI-peak from the previous sub-section are also seen in these angle and momentum-resolved heatmaps: they are the closely-spaced circles which encapsulate the main, higher in magnitude, circle.

There is a strong minimum, visible in all the plots, along the laser propagation direction $\hat{x}$, for these DA-heatmaps.

When going BDA, the rotational symmetry around the laser polarisation axis $\hat{z}$ is broken (the problem is full 3D because of the appearance of the uniform magnetic field, or equivalently, because of the $E'_{\mathrm{BDA},x}(t)$ component from the BDA, effective, E-field from Eq. (4.16)). It shall be thus the case that the presented heatmaps are not symmetric w.r.t. the horizontal line passing through $p_x = 0$ a.u. for the BDA-simulations.

In the high-frequency regime specifically, the appearance of a pronounced lobe against the propagation direction of the laser pulse (in this chapter the propagation direction is the positive $\hat{x}$-direction) manifests itself when going BDA in numerical TDSE simulations. This lobe is not seen in DA TDSE simulations. The lobe appears due to the fact that the combined E and B fields, when in the BDA case, are creating a radiation pressure-like force on the electron, pushing the wavepacket along the propagation direction of the pulse.

Classical dynamics Monte-Carlo simulations revealed that the fraction of the wavepacket that is not ionised in a DA-simulation and is liberated when in the BDA-case is responsible for the lobe [151]. Thus, the lobe formation mechanism can be explained from classical considerations, keeping in mind that for such ultra-strong laser fields the electron essentially follows the trajectory of a

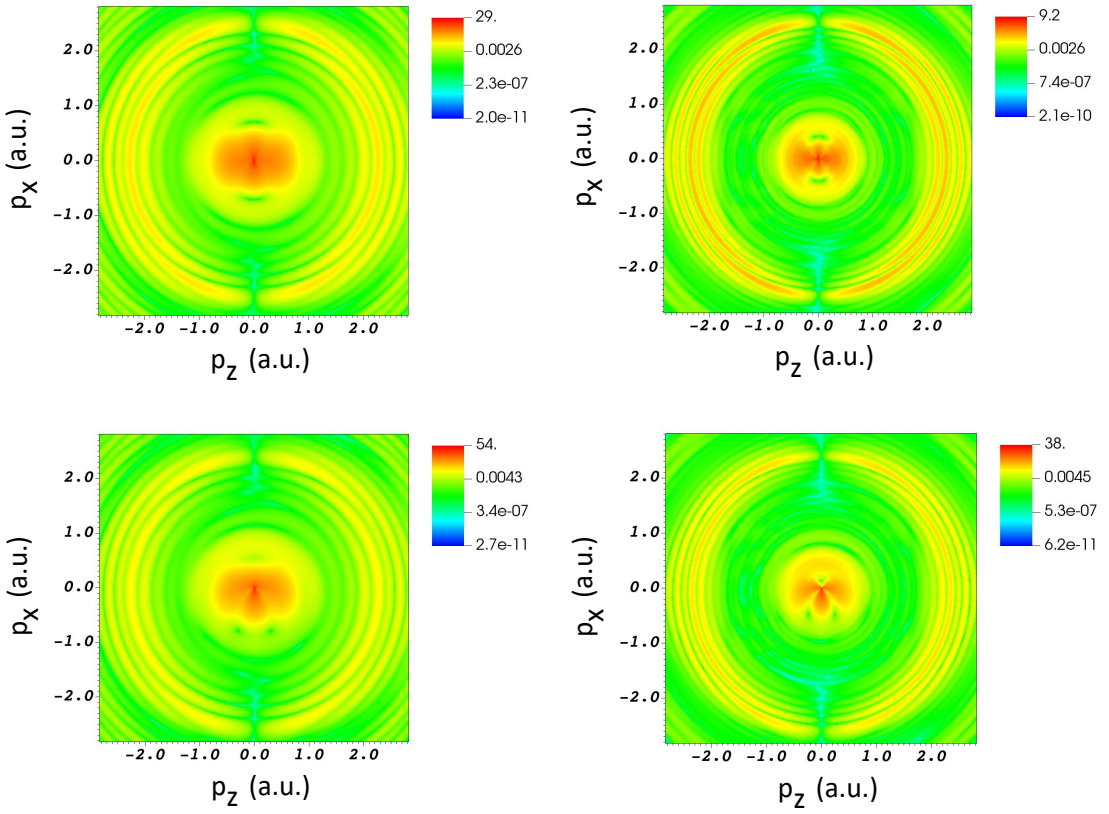


Figure 4.3: Two-dimensional heatmaps showing the momenta along the laser polarisation direction ($\hat{z}$) and along the laser propagation direction ($\hat{x}$). The peak laser field was 30 a.u. Top row: DA-simulations, bottom row: BDA-simulations. Left panels: 10 optical-cycles laser field duration, right panels: 15 optical-cycles laser field duration. The emergence of the characteristic lobe against the propagation direction is clearly seen in the BDA-simulations.

free particle in the laser field. It also builds up momentum from the scattering events with the atomic core. In the BDA simulations, because of the radiation pressure-like force along the laser propagation direction, the electron's trajectory is displaced along this propagation direction. The momentum which has built up (from the re-collisions with the core) has thus the effect to transfer an impulse along the negative propagation direction of the laser pulse in the BDA simulations. In the DA simulations, there is no radiation pressure-like force, and the electron oscillates along the laser polarisation direction without being shifted in any direction along the propagation-of-the-laser-field-axis. Thus, in DA, the up-down half-regions of the $p_z - p_x$ distributions shall be identical, whereas when going BDA, an up-down asymmetry shall be present if non-dipole corrections are strong

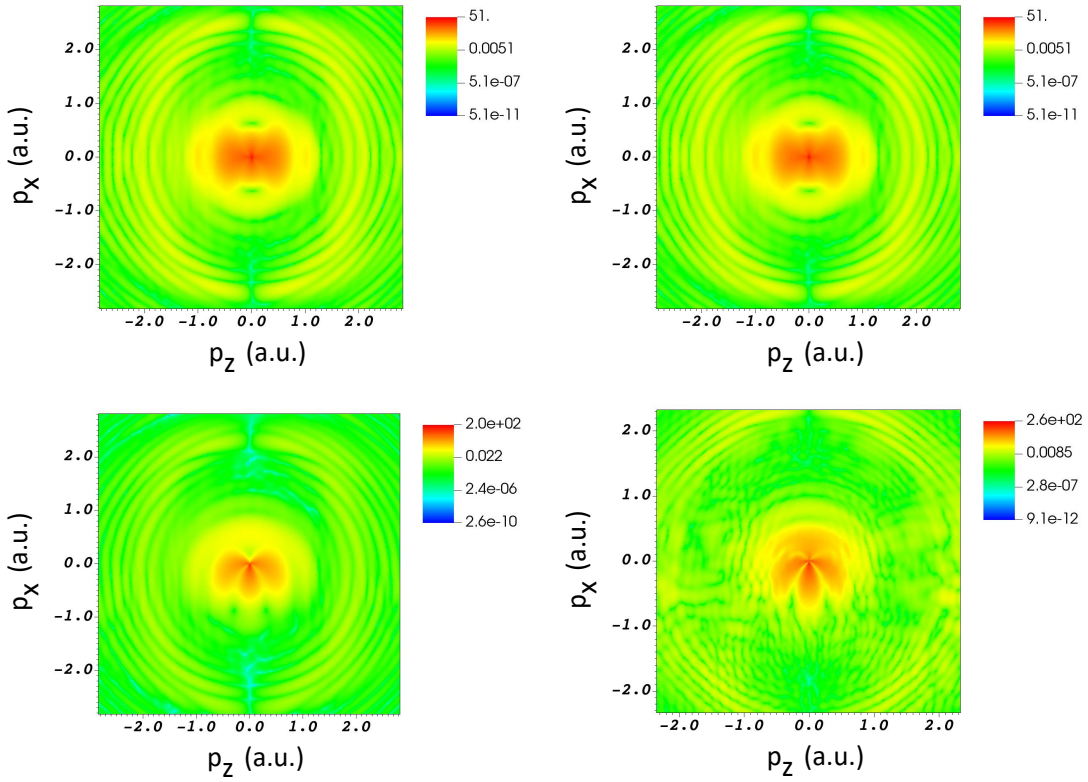


Figure 4.4: Two-dimensional heatmaps showing the momenta along the laser polarisation direction ($\hat{z}$) and along the laser propagation direction ($\hat{x}$). The peak laser-field was 45 a.u. Top row: DA-simulations, bottom row: BDA-simulations. Left panels: 10 optical cycles laser field duration, right panels: 15 optical cycles laser field duration. The emergence of the characteristic lobe against the propagation direction is clearly seen in the BDA-simulations.

enough (that is, if the field-strength is large enough).

The same two figures, that is Fig. 4.3 and Fig. 4.4, present BDA-results on the bottom panels. Again, the left column in each of these figures represents simulations performed for 10 optical-cycles laser pulse duration and the right column represents simulations performed for 15 optical-cycles laser pulse duration. For the BDA-simulations the peak field strengths follow the same arrangements in the panels as for the DA-simulations: the former figure showcases results for $E_0 = 30$ a.u. and the later figure showcases results for $E_0 = 45$ a.u..

All the BDA heatmaps exhibit the top-down asymmetry mentioned above. It is seen that the LES possesses indeed the lobe originating against the laser propagation direction, but also that the side-lobes are bent downwards. Be-

cause of the need to include a number of m-channels different than one in the BDA-simulations, CPU-time limit the full convergence capabilities of the solver for the highest peak field strengths and durations (45 a.u. and 15 optical cycles), this being reflected in the nonphysical artifacts visible both near the LES and near the first ATI-circle. These disappear with a more refined simulation (higher angular momentum maximum value, higher number of m-channels and a smaller timestep).

This shows that higher field strength and longer pulse duration simulations are computationally more expensive to be run, however this is not a problem for the current considerations as the trends seen are the same in all the previous simulations: with the help of the effective E-field dependent only on time and composing the length gauge BDA Hamiltonian, one can study in great detail the BDA TDSE dynamics in full dimensionality.

It is important to understand that the effective E-field can be cast into an effective A-vector-potential which can be used with 3D TDSE numerical solvers which employ the velocity gauge laser-atom interaction form. This can be done by means of integrating with respect to the time parameter Eq. (4.15). One then obtains an A-vector-potential dependent only on the time parameter, but including non-dipole corrections up to first order in $1/c$. It is known that simulations run with the velocity gauge formalism converge faster with respect to the number of partial waves included in the spherical harmonics expansion of the time-dependent wavefunction [175]. The comparison between the two formalisms' results can be a test of the correctness and accuracy of the obtained heatmaps, PES and total ionisation probabilities.

It is to be mentioned that the 2D heatmaps I obtained made use of Landau-normalised scattering states as obtained from running the routines from [125]. If a comparison is intended to be done against other literature results which used momentum-normalised such states, a slightly different post-processing of these continuum states would be needed.

More importantly, the triply-differential momentum-distributions as from fully-converged TDDE simulations from [171] also show the characteristic lobe against the laser's propagation direction. This is re-assuring: the resolved simulation-diagnostic (the PMD) matches a significantly more general formalism's results, however the current results are being obtained in a way more efficient manner.

It would be instructive to also compare and contrast my results with the results from [90]: the middle panel of my Fig. 4.1 is compared against Fig. 1 from the mentioned study. It is clear that both sets of simulation results agree. The peak (beyond which atomic stabilization develops) raises up to around 0.41 (slightly higher (around 3% increase) in [90] due to better convergence properties of their velocity-gauge solver) in both studies. Discrepancies in the total ionization probabilities at the end of the simulation between the Dipole and NonDipole simulations start to appear at 36-38 a.u. of field-strength in [90] and at around 37-38 a.u. of field strength in my figure. Finally, it is visible that these discrepancies increase as the field-strength increases. To be noted is that there might be some not-fully converged simulations towards extreme field-strengths in my case (the slightly oscillatory pattern in the Dipole-Approximated curve in my figure). This can be resolved with finer simulations, although at higher associated computational costs.

Now the bottom right-hand-side panel of my Fig. 4.2 is compared against the bottom right-hand side panel of Fig. 2 from [90]. It was visible from the 2D heatmaps (in particular, the bottom right-hand side panel of Fig. 4.4) that my simulation was not that converged for $E_0 = 45.0$ a.u. and $T = 15$ optical-cycles. I thus run additional simulations at the physics parameters from the bottom right-hand-side panel of my Fig. 4.2.

In Fig. 4.5 it is presented the PES of a rather well-converged such simulation.

The simulation's physics input parameters are identical (i.e. the laser field and box's size) to the conditions of the bottom right-hand-side panel of Fig. 2 from [90].

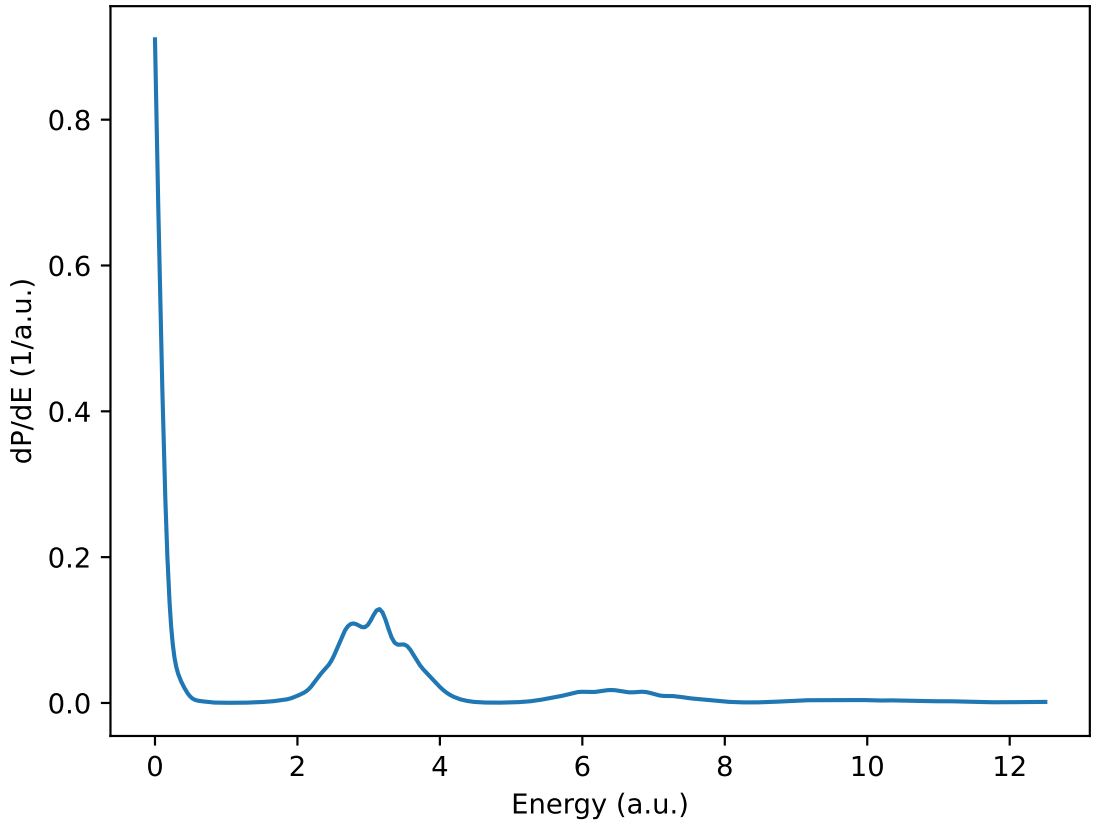


Figure 4.5: Photoelectron energy spectrum (PES) plot as obtained at the end of a TDSE-simulations. The laser field had its central frequency $\omega = 3.5$ a.u., it was 15 optical-cycles in duration with a sin-squared temporal envelope, and the laser peak-field strength was $E_0 = 45.0$ a.u., in a box extending up to $R_b = 300.0$ a.u..

The low-energy parts of the spectra tend to agree well. Due to numerical inaccuracies in the scattering-states generation (I used 300 k-points spanning a range extending from 0.005 a.u. to 5.0 a.u. of momentum), there are slight differences between the values of the spectra around $k = 0$ a.u. It is also unclear what is the actual minimum k-value [90] used in their bottom right-hand-side panel of Fig. 2. It is thus unclear what I shall compare against at $k = 0$. It is also known that it can become problematic to generate accurate scattering states at this type of extremely-small k-values. For example, when I wanted to generate such a state for $k = 0.002$ (so just slightly lower than my current minimum), the Fortran workflow resulted in "divide-by-0 IEEE" errors, similarly for even lower k-values. Thus I was limited by what value of the momentum/energy I could put/set as the first datapoint in the probability against energy plots. Overall,

the agreement is rather good. Moving on to the first ATI peak, it is visible that this is very well reproduced by my simulations when compared to [90]. Both its height and structure matches the [90]'s plot. The differences between the two do not go beyond a few tens of percent. After the $E = 3.0$ a.u. peak, in addition to what was presented in the plot from bottom right-hand-side panel of Fig. 2 in [90], I go to higher values of electron energy, showing the second ATI peak at $E = 6.5$ a.u. $= 2\omega - I_p = 7.0 - 0.5 = 6.5$ a.u. Although of smaller height (which is to be expected due to ATI/MPI considerations), the resolving power of the simulation manages to capture this second peak. To capture the next peaks, one would need to increase the number of angular-momentum l-channels even more than what I had and to try to obtain converged results for the third / fourth / ... peak. This clearly indicates that the BDA-Hamiltonian I used is accurate and can be used to capture Non-Dipole quantum dynamics, although at a non-trivial computational cost.

4.4 Conclusion and outlook

In this chapter, a BDA-Hamiltonian correct up to first order in $1/c$ has had all its terms calculated analytically and has been cast in the length-gauge, being dependent only on the time parameter. Then, numerical-simulations employing this non-dipole Hamiltonian in the 3D GPS TDSE-solver have been run for various laser pulses having their centre-frequency equal to 3.5 a.u.. Simulations with the same EUV laser pulses have also been run employing the DA Hamiltonian. Attention has been paid to the differences in the ionisation probabilities at the end of the pulse between these two sets of simulations. Moreover, the photo-electron energy-spectra have been calculated by means of direct projection of the time-dependent wavefunction at the end of the pulse onto the Coulomb scattering-states. It was found that the oscillations seen in the ATI-peaks due to the recollisions with the core are suppressed when going BDA and that the LES is increased (that is, more electrons with very small velocities are ejected) when going BDA, relative to the DA-simulations. Finally, the characteristic lobe

against the propagation direction is seen in the 2D BDA-heatmaps showcasing the momentum-distribution along the polarisation direction and along the propagation direction of the pulse. This appears only in the BDA simulations and is explained to originate from the fraction of the wavepacket which is not ionised in the DA simulations, but is detached from the core in the BDA simulations. All these are strong indications that the non-dipole, effective, E-field which is correct up to first order in $1/c$ is a perfect candidate to produce *ab initio* predictions in a computationally-efficient manner.

As mentioned in the end of Sec. 4.3.3, the effective E-field can be cast into an effective A-vector-potential. Using a velocity-gauge 3D-TDSE numerical solver, one can obtain results for the PMDs and PES observables as predicted by a velocity-gauge Hamiltonian. This will very clearly and directly touch upon the fact that the results shall be, at least in principle, gauge-invariant. Still, when performing numerical simulations, this perfect matching is rarely fully achieved (to machine-precision) because both the simulations have to be totally converged, this being very demanding in terms of CPU-time, especially if many simulations have to be performed in order to finely sample all the focal volume of the laser beam-pulse. It would be interesting and perhaps useful to analyse the results of such a quantitative comparison and this could be a direction for a future, self-contained, study.

The next chapter will use the TDSE solver in a different setting: it will go from a one-electron atom (a hydrogenic-atom) to a multi-electron atom, by means of a mean-field theory approach: the time-dependent-density-functional-theory (TDDFT).

5

Mean-field approaches to atomic ionisation

"Life is like riding a bicycle. To keep your balance, you must keep moving."

— Albert Einstein to his son Eduard

Contents

5.1 Introduction 138
5.2 Time-dependent density-functional-theory (TDDFT) and its application to strong-field atomic ionisation 141
5.2.1 Ground-state density-functional-theory 141
5.2.2 Time-dependence 154
5.3 Hartree-Fock formalism 173
5.3.1 Results, discussion and further work 176
5.4 Conclusions and outlook 178

5.1 Introduction

Having solved the time-dependent Schrödinger equation (TDSE) in the single-active-electron (SAE) approximation, it is natural to think on how to obtain the ionisation probabilities of all the electrons from the atom irradiated by the laser-pulse. When sequential ionisation prevails, consecutive such events are calculated

using the PPT model for the ionisation rate: this has been used in chapter 2 of this thesis. However, while it is fast to compute them in this way, one might want to look into ways to treat more than one-electron dynamics. As discussed in the previous chapters, it is unfeasible to exactly solve for the dynamics of all the electrons of the atom (apart from the helium one). Alternatives to circumvent the necessity to solve at each timestep for a scalar-field depending on a tuple of, e.g., 18 three-dimensional electron's coordinates (for atomic argon, which possess 18 electrons), have been in place. In this chapter the focus is on **mean-field theories**.

Probably one of the most used one such mean-field approach nowadays is the **Kohn-Sham-density-functional-theory** (KS-DFT), for which the aim here is to use its time-dependent extension (**KS-TDDFT**), as it is always the case in my interaction scenario that a laser field illuminates an ultra-rarefied noble-gas. It is computationally feasible to run such KS-TDDFT simulations and these incorporate, although in an approximate way, the response of all the electrons to the time-dependent perturbation. The limitations of this method will be discussed at length, but in short many of its disadvantages arise from the approximation which has to be made to the exchange-correlation (XC) energy functional (or potential), as it is unknown in practice, but represents a core, fundamental, part of the DFT-formalism.

Another way is to use the **Hartree-Fock** (HF) approach, again including the time-dependence due to the external laser-field (**TDHF**) [176]. A further extension of this is the multi-configuration time-dependent Hartree-Fock (MCTDHF), a more accurate treatment than its predecessor [177]. The limitations of the (TD)HF are diverse, however the most stringent one is that it does not include the correlation of the electrons. Post-HF methods have been derived and, while they come at a way greater computational cost (sometimes that large that it simply becomes unfeasible to run simulations based on these approaches) relative to their predecessors, they can partially (and many times successfully) resolve the issues with the lack of the correlation between the electrons.

Including more and more computational resources helps in two different ways based on the approach used. In (Kohn-Sham-TD)DFT, the aim is to obtain a better representation (closer to reality) of the exchange-correlation functional, not to directly improve the gap between the simulation wavefunction and the real many-electron wavefunction. This is because the simulation wavefunction is just an auxiliary construct which helps predicting the quantities of interest without resorting to the many-electron wavefunction. In TDHF, more compute time and memory are helpful in obtaining a more accurate representation of the real many-body (electrons in this chapter) wavefunction by means of the actual simulation wavefunction. Here, the simulation wavefunction tries to directly mimic the many-electron wavefunction. While the numerical ways to solve the equations of these two methods can be considered roughly similar (in a very crude approximation), there is a distinction (albeit subtle) on how more CPU-core clock-cycles and memory help predicting the experimental observables.

It is important to note that these mean-field methods can be applied to systems larger than a single atom, i.e. molecules, possibly many of them organised in a structure, in order to obtain the optical and dynamical responses of a real material or biological sample of interest.

In this chapter the focus is on TDDFT, leaving TDHF as a possible future direction arising from the work from this thesis. The Kohn-Sham-TDDFT equations I solved are firstly shown and explained: in short, they are coupled one-electron TDSE's which are propagated together in real-time and real-space. Results of two-dimensional (2D) TDDFT numerical simulations based on the generalized-pseudospectral-grids (GPS) discretisation of TDSE are presented. The effort to solve these equations increases with the number of Kohn-Sham orbitals used. Thus large box sizes cannot be trivially used, so I mainly dealt with very-high intensity pulses (up to 4.0 - 5.0 $\times$ 10^{19} W/cm^2) having their centre-frequency in the region of extreme ultraviolet (EUV) to soft-X-ray regime. An additional complication which appears when running TDDFT simulations relative to the one-electron TDSE simulations is that any numerical inaccuracy in the total

electron density on the grid is carried in time and this makes the propagation unstable as these numerical errors in the approximation grow over subsequent timesteps. Thus a finer spatial and temporal discretisation is needed to have the simulations stable during the time-propagation, relative to the one-electron case. It might even be the case that the timestep has to be one order of magnitude smaller than for the one-electron case. In order to simulate the focal-volume effect of a laser-gas jet interaction in the laboratory, many such TDDFT simulations have to be run to finely sample the laser focus intensity region, constraining the pulse durations which can be used in the numerical modelling.

5.2 Time-dependent density-functional-theory (TDDFT) and its application to strong-field atomic ionisation

5.2.1 Ground-state density-functional-theory

It is instructive to start without having any time-dependence in the problem considered, i.e. there is no dynamics involved as the laser-field is not included for the moment. The focus is thus on the ground-state properties alone, and these will be the initial input to the time-dependent considerations which will follow. That is, the current interest is to derive properties of the unperturbed system, without solving exactly the N-electron time-independent Schrödinger equation (TISE). The Born-Oppenheimer approximation (i.e. the nuclei are fixed in place) is made from the start in this chapter, and the TISE reads:

$$\hat{H}\,\Psi(\vec{x}_1, \vec{x}_2, ..., \vec{x}_N) = E\,\Psi(\vec{x}_1, \vec{x}_2, ..., \vec{x}_N), \tag{5.1}$$

where $\vec{x}_i = (\vec{r}_i, \sigma_i)$ is a tuple of four real values: the three spatial coordinates of the electron number i, and its spin value. The wavefunction Ψ has to be correctly anti-symmetrised in order to respect Pauli's exclusion principle. Above, the Hamiltonian $\hat{H}$ is composed from the kinetic energy of the electrons, the attraction between the nucleus and all the electrons (called the external potential,

which in the time-dependent case will also include the interaction between the laser and the electrons), and the electron-electron interaction,

$$\hat{H} = -\sum_{i=1}^{N} \frac{\nabla_i^2}{2} + \sum_{i=1}^{N} V_{ext}(\vec{r}_i) + \frac{1}{2}\sum_{i=1}^{N}\sum_{\substack{j=1 \\ j\neq i}}^{N} \frac{1}{|\vec{r}_i - \vec{r}_j|}, \tag{5.2}$$

with $V_{ext}(\vec{r}) = -Z/|\vec{r}|$ where Z is the nucleus charge and the $1/2$ multiplicative prefactor from the last term avoids double-counting the pairs (for example, $(i = 3; j = 5)$ with $(i = 5; j = 3)$).

As pictured above and from the previous chapters, Eq. (5.1) is an eigenvalue problem. The lowest energy eigenvalue, E_0, is the ground-state energy. The eigenvector of this eigenvalue is called the ground-state and is denoted by Ψ_0. It is normalised to 1.0 and related to the ground-state total electronic-density $n_0(\vec{r})$, a function of only three spatial coordinates (r, θ, φ), as follows:

$$n_0(\vec{r}) = N \sum_{\sigma} \int d\vec{x}_2 \int d\vec{x}_3 \ldots \int d\vec{x}_N \, |\Psi_0(\vec{x}, \vec{x}_2, \vec{x}_3, \ldots, \vec{x}_N)|^2. \tag{5.3}$$

The physical interpretation of the total electron-density is that it represents the probability to find any electron at position $\vec{r}$, and this probability is normalised to the total number of electrons composing the many-electron wavefunction $\Psi_0(\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N)$.

It turns out that for the ground state, knowing the electron-density amounts to having information about the system's energy. This is because of the **first theorem of Hohenberg and Kohn [178]**, which states that the ground-state electron-density uniquely determines the static external potential (up to an additive constant). This in turn means that the Hamiltonian is determined and one can, in principle, exactly extract any observable property of the many-electron system. Thus the one-body density can be used as the element which generates the system's characteristics.

This theorem can be alternatively stated in the following way: for the lowest energy state, the energy of the system of interacting electrons is a functional of the electron density (a functional is a mapping between a function and a number, i.e. a function of functions). Thus the knowledge of a function of three

coordinates, $n_0(\vec{r})$, suffices for obtaining the observable properties of a many-electron system. There exists a caveat here, which is the biggest difficulty of DFT: it is known that the functional exists and, for any given external potential, is unique (ignoring the possible problems which might appear with densities which do not respect the v-representability problem [179]), but no-one knows its exact form and approximations to it have to be made in practice. The quality of these approximations directly dictate the quality of the obtained results for the many-electron system. There is an ever increasing effort devoted to developing more and more accurate such approximations, potentially with the aid of statistical techniques in the form of machine-learning and deep-learning incorporated in quantum-chemistry workflows such as in [180–185]. In [180], the authors nicely and intelligently used a neural network to represent the unknown of interest, the exchange-correlation functional. Then experimental data has been fed into the constructed framework and the physics of the quantum interactions inside the materials (in particular, diatomic molecules) has been learned. This helped predict properties of unseen samples (not seen during the training of the network), with a strikingly magnificent accuracy. Interestingly, the network has been trained on as few as eight examples, but managed to infer the properties of as much as 104 new molecules. This shows that this might definitely be a path forward to more accurate representations of the core of DFT. This paradigm of keeping the parts of the model which are known to be accurate as they are and replacing the parts which are unknown (or are not accurate) with neural networks gains more and more attraction [186].

Hohenberg and Kohn formulated a second theorem [178], which gives a variational principle and can be stated as follows: for any correctly normalised trial electron-density $n_{trial}(\vec{r})$ which is not the ground-state density, the energy resulting from it is higher than the ground-state energy.

Their two theorems together express the fact that the ground-state density and energy are the minimum of some (again, unique, that is, being the same for any many-electron interacting system) energy functional, under the constraint

that the density is correctly normalised (i.e. gives the total number of electrons when integrated across the real three-dimensional space). Again, this functional is unknown in practice - had it been known, one could have just minimised it by any means which achieves the purpose and found the ground-state energy and density for any system under consideration.

Looking at the three elements of the Hamiltonian from Eq. (5.2), there is only one of them for which it is known the exact form of the functional of the density: the external potential. The full electron-electron interaction and the full kinetic energy of the interacting electrons do not have a closed-form expression known as of today and these are the reasons why DFT, even if designed as an exact theory, has to be solved approximately in practice.

The Kohn-Sham auxiliary-system

The method to solve this and the way in which (TD)DFT is used in practice is to move the attention towards **an auxiliary system of non-interacting electrons**, which possesses the exact same electron density as the original, interacting electrons system. This system is called **the Kohn-Sham system**, and it is mathematically described by a single Slater-determinant in the electron orbitals. Of course, the electrons from the two systems behave differently, however for any real system, its Kohn-Sham variant is exact and well-defined, and it is an aid (i.e. a tool) to obtain the quantities of interest of the real system. For this new system the density can be obtained as simple sum over the orbitals' moduli squared magnitudes, so an addition of all the single-electron densities from the Kohn-Sham system. The KS-orbitals are denoted here as $\{\phi_1(\vec{r}), \phi_2(\vec{r}), ..., \phi_N(\vec{r})\}$. Also, the kinetic energy functional can be expressed as the sum of the kinetic energy functionals of each electron (orbital) and it turns out that this represents a great part of the total, real, kinetic energy of the interacting electrons. It is because of the non-interacting nature of this auxiliary system that one can express its kinetic energy as a sum over individual kinetic energies, each obtained from the non-interacting wavefunction (i.e. the corresponding KS-orbitals from it). Finally,

the electron-electron interaction functional is written as the classical Coulomb interaction between the electrons (the Hartree potential). It is now understood that a remainder has to be in place - the error in the energy functional made by: **a)** treating the kinetic energy as from a system of non-interacting electrons, rather than from the real system plus **b)** treating the electron-electron interaction only classically. The remainder is called **the exchange-correlation energy functional** $E_{xc}[n(\vec{r})]$, where the square parentheses [ ] denote a functional. It is important to underline that it contains a part of the exact kinetic energy, even if its name does not suggest this. In HF-theory, the exchange and correlation integrals (for the energy contributions) are not related at all to the kinetic-energy.

At this point, equipped with these four functionals which together produce the total energy functional, one can apply the variational principle resulting from the two theorems of Hohenberg and Kohn (which also apply to this Kohn-Sham system: it can be always viewed as an interacting system with the electron-electron interaction strength set to exactly zero, thus in the realm of the applicability of the theorems) and obtain **the Kohn-Sham equations** (KS-equations) for the orbitals $\{\phi_1(\vec{r}), \phi_2(\vec{r}), ..., \phi_N(\vec{r})\}$:

$$\left(-\frac{\nabla^2}{2} + V_{ext}(\vec{r}) + V_H(\vec{r}) + V_{xc}(\vec{r})\right)\phi_i(\vec{r}) = \varepsilon_i\,\phi_i(\vec{r}),\ \ i = 1, 2, ..., N. \tag{5.4}$$

Above, the (XC) potential $V_{xc}(\vec{r})$ (a multiplicative operator on the grid representation of the orbitals) is the functional derivative of the exchange-correlation energy functional:

$$V_{xc}(\vec{r}) = \frac{\delta E_{xc}[n(\vec{r})]}{\delta n(\vec{r})}. \tag{5.5}$$

It is a good moment to discuss which avenue has been taken to approximate the exchange-correlation energy functional in the work presented in this chapter.

Exchange-correlation energy functionals

The **local density approximation (LDA)** is a very simple functional approximation which can be used in practice. Maybe surprisingly, it tends to offer an accuracy on the ground state energies which is many times not particularly unsatisfactory, thus the focus is on this approach. In brief, LDA means that the density-functional of the exchange-correlation energy is approximated by the following integral across the real volume:

$$E_{xc}\left[n(\vec{r})\right] \approx E_{xc}^{LDA}\left[n(\vec{r})\right] = \int d\vec{r}\; n(\vec{r})\, \varepsilon_{xc}^{HEG}(n(\vec{r})), \tag{5.6}$$

where $\varepsilon_{xc}^{HEG}(n(\vec{r}))$ is the exchange-correlation energy density (energy per electron) of the uniform (homogeneous) electron gas (HEG) with density $n(\vec{r})$. The uniform electron gas is a system in which all the electrons feel the same potential and can be characterised by a single real value, the electron density n^{HEG}, which is the same throughout all the volume of this gas (i.e. $n^{HEG} \neq n^{HEG}(\vec{r})$). This system can be almost exactly solved. As the name suggests, in LDA the energy density $\varepsilon_{xc}(n(\vec{r}))$ is a function of only the local value of $n(\vec{r})$. Within this LDA-formalism, it is customary to split ε_{xc} into two parts: the exchange $\varepsilon_x(n(\vec{r}))$ plus the correlation $\varepsilon_c(n(\vec{r}))$:

$$\varepsilon_{xc}(n(\vec{r})) := \varepsilon_x(n(\vec{r})) + \varepsilon_c(n(\vec{r})). \tag{5.7}$$

The exchange part can derived [187] and reads:

$$\varepsilon_x(\vec{r}) = -\frac{3}{4}\left(\frac{3}{\pi}\, n(\vec{r})\right)^{1/3}, \tag{5.8}$$

thus the associated exchange potential is:

$$V_x(\vec{r}) = -\left(\frac{3}{\pi} n(\vec{r})\right)^{1/3}. \tag{5.9}$$

The correlation part $\varepsilon_c(\vec{r})$ cannot be derived fully analytically and is approximated (interpolated, or generally speaking, fitted) based on: **a)** very accurate (for example, using quantum-Monte-Carlo approaches) numerical techniques which target parts of the space of the density-values of interest and on **b)** some

limiting cases in which the results are known analytically. The simulations which have been run for this chapter have used the **Perdew and Wang (1992) correlation PW92**: $\varepsilon_c^{PW92}(n(\vec{r}))$:

$$\varepsilon_c^{PW92}(n(\vec{r})) = -2a\,n(\vec{r})\,(1+\alpha_1 r_s(\vec{r}))\cdot \\ \cdot \ln\left(1+\frac{1}{2a\left(\beta_1\sqrt{r_s(\vec{r})}+\beta_2 r_s(\vec{r})+\beta_3(r_s(\vec{r}))^{3/2}+\beta_4(r_s(\vec{r}))^2\right)}\right),$$

where the constants are $a = 0.0310907, \alpha_1 = 0.21370, \beta_1 = 7.5957, \beta_2 = 3.5876, \beta_3 = 1.6382, \beta_4 = 0.49294$, and $r_s(\vec{r}) = \left(\frac{4\pi n(\vec{r})}{3}\right)^{-1/3}$ is proportional to the inverse cubic root of the electron density $n(\vec{r})$.

It might be instructive to list another very well known correlation energy-density functional: the **Vosko, Wilk and Nusair (VWN)** one, which is as below:

$$\varepsilon_c^{VWN}(n(\vec{r})) = \\ \frac{A}{2}\left(\ln\left(\frac{z}{Z(z)}\right)+\left(\frac{2b}{Q}\right)\tan^{-1}\left(\frac{Q}{2z-b}\right)\right) - \tag{5.10}$$

$$\frac{A}{2}\frac{bz_0}{Z(z_0)}\left(\ln\left(\frac{(z-z_0)^2}{Z(z)}\right)+\frac{2(b+2z_0)}{Q}\tan^{-1}\left(\frac{Q}{2z+b}\right)\right), \tag{5.11}$$

with the definitions: $z = z(\vec{r}) := \sqrt{r_s(\vec{r})},\ Z(z) := z^2+bz+c,\ Q := \sqrt{4c-b^2}$, and the constants $A = 0.0621814, z_0 = -0.409286, b = 13.0720, c = 42.7198$.

The usage of LDA acknowledges the fact that the binding energies of the core electrons are obtained to be larger than their real-values (thus these electrons are obtained to be under-bound to the atom) [188]. This is because the fact that the LDA-exchange-correlation potential decays exponentially, so it does not posses the correct asymptotic behaviour (that of the Coulomb potential, $-1/r$). For negative ions, the electrons can even become unbound because of this. The fundamental reason why this happens is that the electrons shall not self-interact. However, for the LDA exchange-correlation energy functional, a self-interaction (SI) energy contribution appears. This nonphysical SI behaviour can be corrected in a number of ways. Perdew and Zunger theoretically added a **self-interaction-correction (SIC)** [189] to the LDA-functional. Still, in practice, the usage of this new form leads to the KS-potential being different for different KS-orbitals,

thus it leads to them becoming non-orthogonal. **Optimised-effective-potential (OEP)** [190] techniques were generally derived to help increasing the accuracy of E_{xc} functional. Pure-OEP theory leads to an integral equation that is costly to be solved numerically, so the **Krieger-Li-Iafrate (KLI)** approximation to this equation [191, 192] is preferred (this way, a linear system of equations has to be solved and this is achievable in practice). Tong and Chu showed that the OEP-KLI formalism can help alleviating the SIC-issue of LDA and the numerical implementation of this is based on the GPS-discretisation [193]. In this chapter, such advanced considerations have not been made use of. Instead, a simple restoration of the asymptotic behaviour has been put in place, as from [194]. It is achieved by subtracting from the XC-potential the following quantity:

$$V_{correction}\left(r\right) = \frac{1}{r}\left(1 - e^{-Ar^2}\right), \tag{5.12}$$

where the constant A has to be chosen such that not to disturb too much the original LDA potential next to the core and to still have a considerable effect at large radii r, the location of space where it shall act in the most powerful way. A similar technique is briefly mentioned and applied in, e.g. [110].

Also, some major difficulties appear to be visible (and most pronounced) when going to the molecular targets, as it turns out the errors resulting from calculations done using LDA increase with the size of the quantum-system under consideration [188] (and a molecule is indeed much more extended than an isolated, single atom). Still, such discrepancies between numerical simulations and true experimental results in the case of molecules will not be further discussed in this chapter, but it is to be mentioned that creation or rupture of chemical bonds is not accurately treated within the LDA-realm [188], or that molecules are generally over-bound in LDA [188].

LDA is not perfectly accurate in either of its parts (exchange or correlation), however one is computed to be larger than in reality, while the other is computed to be smaller than in reality, and their sum totals a relatively accurate total exchange-correlation energy. LDA has its merits by providing accurate results for

many quantum-systems by employing a simple functional form which trivially respects many exact conditions [195]. As quite a few of these are rather technical in their nature, they will not be discussed further, however they can be a solid contender for the top spot in the ranking of the reasons why LDA is definitely not wrong in practical calculations for which the system is trivially seen not to be uniform. When constructing (ab-initio or by fitting expressions to datasets of interest) non-local or semi-local functionals, there might be cases that some exact conditions are violated or hard to be exactly satisfied.

In the search for functionals which are more accurate than LDA, one can move up across the rungs of the Jacob's ladder [196], arriving at **generalised-gradient-approximations (GGA)** for the XC-energy-functional. This means that now the functionals can depend not only on the local electron density, but also on its gradients (which are trivially non-zero for non-uniform media), and even on second-order gradients (Laplacians). This way a (potentially highly) non-flat electron-density can have its characteristics taken into account when calculating $E_{xc}[n(\vec{r})]$. The first efforts to make the transition from a simple local functional of the electron-density to this new, potentially improved functional dependence resulted in unsatisfactory results for the experimental observables calculations (the accuracy was worse than with pure LDA, even if the computational cost was higher). This was mainly because they did not respect quite a few exact conditions. Further efforts managed to reduce the errors relative to those obtained when using LDA, many times significantly. Mathematically, GGA's can be written as (and the notation also includes the so-called **meta-GGA**'s):

$$E_{xc}^{GGA}\left[n(\vec{r}), |\nabla n(\vec{r})|, \nabla^2 n(\vec{r})\right] = \int d\vec{r}\, f\left[n(\vec{r}), |\nabla n(\vec{r})|, \nabla^2 n(\vec{r})\right], \tag{5.13}$$

where the form of $f\left[n(\vec{r}), |\nabla n(\vec{r})|, \nabla^2 n(\vec{r})\right]$ depends on which (meta-)GGA-functional is actually used.

Here will be mentioned three of the most well-known GGA's, namely the BLYP, the PBE, and the PW91 ones. The BLYP is composed of the Becke (1988)

[197] exchange functional plus the Lee, Yang and Parr correlation functional [198]:

$$E_{xc}^{BLYP}\left[n(\vec{r}),|\nabla n(\vec{r})|,\nabla^2 n(\vec{r})\right] = E_x^{B88}\left[n(\vec{r}),|\nabla n(\vec{r})|\right] + E_c^{LYP}\left[n(\vec{r}),|\nabla n(\vec{r})|,\nabla^2 n(\vec{r})\right],$$

with the forms of the two parts reading as below:

$$E_x^{B88}\left[n(\vec{r}),|\nabla n(\vec{r})|\right] = \int d\vec{r}\,\varepsilon_x^{B88}\left[n(\vec{r}),|\nabla n(\vec{r})|\right], \tag{5.14}$$

$$E_c^{LYP}\left[n(\vec{r}),|\nabla n(\vec{r})|,\nabla^2 n(\vec{r})\right] = \int d\vec{r}\,\varepsilon_c^{LYP}\left[n(\vec{r}),|\nabla n(\vec{r})|,\nabla^2 n(\vec{r})\right], \tag{5.15}$$

with the generating quantities:

$$\varepsilon_x^{B88}\left[n(\vec{r}),|\nabla n(\vec{r})|\right] = \varepsilon_x^{LDA}[n(\vec{r})]\left(1-\frac{\beta}{2^{1/3}A_x}\frac{X^2}{1+6\beta X\sinh^{-1}(X)}\right), \tag{5.16}$$

and with the definition $X := 2^{1/3}\dfrac{|\nabla n(\vec{r})|}{(n(\vec{r}))^{4/3}}$ and constants $A_x := 0.75\left(\dfrac{3}{\pi}\right)^{1/3}$, $\beta = 0.0042$, and

$$\varepsilon_c^{LYP} = \frac{-a}{1+M\,(n(\vec{r}))^{-1/3}} \left(n(\vec{r}) + b\,(n(\vec{r}))^{-2/3}\left(C_F\,n(\vec{r})^{5/3} - 2t_w + \frac{1}{9}\left(t_w + 0.5\nabla^2 n(\vec{r})\right)\right)e^{-c(n(\vec{r}))^{-1/3}}\right),$$

with the definitions $t_w := \frac{1}{8}\left(\frac{|\nabla n(\vec{r})|^2}{n(\vec{r})} - \nabla^2 n(\vec{r})\right)$ and $C_F = 0.3\left(3\pi^2\right)^{2/3}$, $a = 0.04918$, $b = 0.132$, $c = 0.2533$, $M = 0.349$.

The PBE is composed from the exchange and correlation of Perdew, Burke and Ernzerhof and has been derived in 1996 [199]. The PW91 is composed from the exchange and correlation of Perdew and Wang, and as the name suggests, have been derived in 1991 [200, 201].

Finally, there have been derived the so-called **hybrid-functionals**, which try to **merge the DFT theory with the HF theory**, in a way in which the good parts of the two formalisms complement each other in a functional which is envisioned to be more accurate than its predecessors. Adding HF means removing the self-interaction which is experienced by the electrons in pure DFT. This self-interaction problem can be easily seen for the KS-auxiliary system, where the

one-body potential depends on all the orbitals, and when it acts on a given orbital, it also includes its contribution into it, this meaning a self-interaction is in place in pure DFT. In HF theory, it is known that there is no such problem because the two integrals (exchange integral and classical Coulomb interaction integral) add up for each orbital and cancel each other's contributions from the given orbital.

This merging of DFT and HF is done for the exchange functional, and it is almost never the case that the whole (100%) DFT-exchange is replaced by the exact Fock exchange. Depending on the hybrid used, there are different percentages which are swapped, for example, the PBE0 hybrid [202] uses 25% of exact exchange, while the Becke one splits the exchange in half, meaning a 50% exact exchange is used [203]. In the KS-auxiliary system, the recipe to calculate the exact Fock exchange is based on the KS-orbitals (i.e. one replaces the HF-orbitals with the KS-orbitals in the exchange sum-integral formula), and this is how this calculation of the exchange is implemented in practice. The disadvantage of this merging-method is that the Fock operator is non-local in space, thus the computational effort required to obtain the exact exchange and use a percentage of it is rather high, scaling non-linearly with the number of electrons in the system, potentially making it unfeasible for rather large systems.

The self-consistent (SC) iterations

The sum of the Hartree potential and the XC potential is sometimes collectively called **the Kohn-Sham potential** $V_{KS}(\vec{r}) := V_H(\vec{r}) + V_{xc}(\vec{r})$. The **Hartree potential** is the integral across all the real volume of the total electron-density divided by the corresponding distance:

$$V_H(\vec{r}) = \int \frac{n(\vec{r})}{|\vec{r} - \vec{r}\,'|} d\vec{r}\,'. \tag{5.17}$$

In practice, this is obtained by solving its associated Poisson equation (otherwise, one of the difficulties which might appear when numerically calculating it is the divergence at $\vec{r} = \vec{r}\,'$, for example), or by performing the integrals against the Laplace's expansion of the inverse distance instead of the problematic $1/|\vec{r} - \vec{r}\,'|$.

It is clear by now that the single-particle potential encompassed by the parentheses in Eq. (5.5) depends on the total electron-density $n(\vec{r})$, which in the KS-formalism is a sum over all the orbitals' densities:

$$n(\vec{r}) := \sum_{i=1}^{N} |\phi_i(\vec{r})|^2. \tag{5.18}$$

However, this is the quantity which is sought-after by solving the KS-equations. A **self-consistent-iterations (SC) procedure** shall be in place. That is, an initial guess for the set of orbitals shall be made at the beginning of the first iteration. I have found that I achieve self-consistency (convergence in the SC-iterations) with the pure-hydrogen attraction as the initial guess for the orbitals' form (that is, using a GPS-like discretisation of the TISE equation for atomic hydrogen and diagonalisation of the Hamiltonian matrix, as in the previous chapters). Then the equations are solved for the new set of orbitals originating from the guess made previously. Once obtained, this new set of orbitals generates a new density, which is in turn fed as input again to the KS equations. The procedure repeats up until the obtained density does not change significantly over a range of consecutive iterations of the workflow (i.e., convergence has been obtained when iteratively solving the KS equations). The fact that these equations produce densities that converge depends on the initial guesses for the orbitals, on the exchange-correlation potential, but also on the mixing scheme used at each iteration. The mixing scheme refers to how the obtained density is fed as input into the KS equations. While it might be the case that a naïve output-becomes-the-new-input routine works in practice, **linear-mixing** is many times used. There exist other mixing-schemes which are more advanced and offer better convergence properties, e.g. [204, 205], but will not be discussed and have not been made use of in this chapter. Linear-mixing simply means that at each iteration, a fraction of the density of the previous iteration is kept as a constituent part of the new input to the equations, in a linear-combination with the output from the current iteration. The coefficients of the linear-combination are real and are constrained to add up to 1.0.

Results and discussion

As the preliminary step to running a TDDFT simulation, the converged ground state (here in the KS-formalism) is needed. I have selected the argon atom as the target, a closed-shell atom composed of 18 electrons in its neutral state. In a spin-restricted formalism, this means that only 9 (spatial) orbitals have to worked with, the others evolving the same way during the TD-simulation. As will be shown in Sec. 5.2.2, due to the laser being linearly polarised, even less orbitals had to be actually propagated, due to the symmetry of the problem. The initial guesses for these orbitals have been the $Z = 1$ pure-hydrogen-attraction ones. That is, the TISE has been solved for the hydrogen atom and the $n = 1, l = 0$, $n = 2, l = 0$, $n = 2, l = 1$, and so on, up to $n = 3, l = 1$ eigenstates have been saved, where the n and l have the usual meaning in the hydrogenic-atoms theory (i.e. n is not an electron-density in this case). These represent the initial input to the self-consistent iterations involved in the ground-state KS-DFT workflow. The mixing scheme was set to have a fraction equal to 0.382. The simulation's radial extent was fixed to a rather small value $R_b = 100$ a.u., as the frequency of the laser pulse from the TD-simulation was rather high, thus the electrons do not move to far distances as the classical oscillator radius is equal to E_0/ω_{laser}^2. The smoothing parameter for restoring the correct asymptotic behaviour of the LDA XC-potential was set to $\eta = 5.0 \times 10^{-3}$. The discretisation which was chosen was $N_r = 150$. This offers a balance between accuracy and the speed of computation in the TD-case, especially when multiple such TDDFT simulations have to be run.

Other ground-state calculations have been run, each of them aimed at analysing the behaviour of the converged-eigenenergies as one of the discretisation parameters changes. It is concluded that any discretisation from $N_r = 150$ up to $N_r = 500$, for box sizes ranging between $R_b = 50$ a.u. and $R_b = 100$ a.u., gives eigenenergies which converge towards the same numerical values at the end of the SCF-iterations. For $R_b = 100$ a.u., the SC-iterations converged in 15 iterations.

Table 5.1 lists the obtained eigenenergies for $N_r = 150$ and $R_b = 100$ a.u..

	LDA KS-DFT	Experimental data
1s	-113.81	-117.83
2s	-10.80	-12.00
2p	-8.45	-9.13
3s	-0.89	-1.08
3p	-0.39	-0.58

Table 5.1:
Eigenenergies in atomic units (a.u.) after the convergence of the self-consistent-field (SCF) iterations. The ground-state KS-DFT calculation was performed with the LDA XC-energy functional and the real-space discretisation had a number of radial points $N_r = 150$. The box size was set to $R_b = 100$ a.u.. The experimental values are from [206].

The orbitals obtained at the end of the SC-iterations procedure are the initial input to the time-dependent calculations which dictate their dynamics under the influence of the laser-field.

5.2.2 Time-dependence

The extension of DFT to TDDFT has as its basis **the Runge-Gross theorem** [207], which can be stated as follows: for a many-body system evolving from a given fixed initial state, the one-body time-dependent electron-density $n(\vec{r}, t)$ uniquely defines the external, time-dependent potential $V_{ext}(\vec{r}, t)$. This external potential, as in the time-independent case, completely specifies the (now time-dependent) Hamiltonian operator $\hat{H}(t)$ acting on the many-electron system, which in turn shows that all the observables can be derived from the electron-density alone, by-passing the need to compute the full many-electron wavefunction at each timestep. That is, there is a one-to-one correspondence between the electron-density and the external potential, at each timestep. The proof of this theorem is more involved than the one from the time-independent case. Again, the practical way to solve the TDDFT dynamics is to make use of the KS-system, now with orbitals which depend on both time and real-space. Describing these temporal-dynamics, the KS-equations have to be upgraded to **the time-dependent KS-**

equations (TDKS-equations), all of which are similar in form to the TDSE:

$$i\frac{\partial}{\partial t}\phi_i(\vec{r},t) = \left(-\frac{\nabla^2}{2} + V_{ext}(\vec{r},t) + V_H(\vec{r},t) + V_{xc}(\vec{r},t)\right)\phi_i(\vec{r},t)\,, \quad i = 1,2,...,N, \tag{5.19}$$

where all the time-dependencies have been explicitly showcased, and where now $V_{ext}(\vec{r},t) = -Z/|\vec{r}| + \vec{r}\cdot\vec{E}(t)$ for a dipole-approximated laser-field $\vec{E}(t)$. These equations are numerically-solved mainly the same as a TDSE for a one-electron (SAE) target using the GPS discretisation from the previous chapters. Still, there are additional technical complications which appear for the TDKS-equations: the need to calculate the Hartree potential and the XC-potential at each timestep. The XC-potential is not a very complicated matter in the approximation I have chosen to use, i.e. the Adiabatic-Local-Density-Approximation (ALDA). The Hartree potential, however, involves supplementary computational effort at each timestep, as this potential is obtained by solving its associated Poisson equation, discretised again using the GPS technique. It is to be noted that the propagation used in this chapter is different than both the Casida approach and the Sternheimer approach for deriving optical properties of quantum-systems [208].

The Hartree potential

Having completely tackled the XC-potential, the last remaining part is the Hartree potential defined in Eq. (5.17), a functional of the total electron-density $n(\vec{r},t)$. In this chapter, in all the time-dependent simulations which have been run, the laser-field was linearly-polarised. This means that all the dynamics are completely determined in a 2D spherical geometry. Also, the magnetic quantum number m of each KS-orbital $\phi_i(\vec{r},t)$ is conserved during the time-propagation and the electron-density of each orbital does not depend on the azimuthal angle φ. The expansion in the spherical harmonics $Y_l^m(\theta,\phi)$ from previous chapters becomes here a simpler expansion in the Legendre polynomials only, and this is applied to the electron density of each KS-orbital and to the Hartree potential of each KS-orbital as well. Having calculated the Hartree potential of each KS-orbital,

their sum (updated at each timestep) is the required Hartree-potential which will be a multiplicative operator on the real-space grid for each orbital [209]. Thus:

$$V_H[n](r,\theta,t) = \sum_{i=1}^{N} V_H[n_i](r,\theta,t) \ , \tag{5.20}$$

where $V_H[n_i](r,\theta,t)$ is expanded onto the Legendre polynomials at each timestep, as below.

$$V_H[n_i](r,\theta,t) = \sum_{l=0}^{l_{max}} V_l^H[n_i](r,t)\sqrt{\frac{2l+1}{4\pi}}P_l(\cos\theta). \tag{5.21}$$

Introducing $U_l(r,t) \coloneqq rV_l^H[n_i](r,t)$, this function respects the following ordinary differential equation:

$$\left(-\frac{d^2}{dr^2} + \frac{l\,(l+1)}{r^2}\right) U_l(r,t) = 4\pi\, r\, n_i^l(r,t), \tag{5.22}$$

which is easily recognised as a Poisson-like equation with the term $n_i^l(r,t)$ given by the expansion of the electron density of the orbital number i in Legendre polynomials:

$$n_i(r,\theta,t) = \sum_{l=0}^{l_{max}} n_i^l(r,t)\sqrt{\frac{2l+1}{4\pi}}P_l(\cos\theta). \tag{5.23}$$

The boundary-conditions for Eq. (5.22) are [209]:

$$U_l(r=0,t) = 0, \tag{5.24}$$

$$U_l(r=R_{max},t) = \frac{4\pi}{(2l+1)\,R_{max}^l}\int_{r=0}^{R_{max}} n_i(r,t)\,r^{l+2}\,dr. \tag{5.25}$$

It is important to note that while in the GPS method the wavefunctions (and electron-density, accordingly) are imposed to be 0 at the grid boundaries, the Hartree potential shall not be 0 at the grid boundary. This is taken care of by also including the last Lagrange basis-function in the Hartree potential expansion on the GPS-grid, unlike for the wavefunctions.

The Poisson-like equation from Eq. (5.22) is solved firstly constructing the (operator-)matrix of the second-order derivative with respect to the radial coordinate and the contribution from the centrifugal term, then inverting it using `LAPACKE_dgetrf` and `LAPACKE_dgetri`, to be then further used in matrix-vector multiplication operations.

The time-dependent exchange-correlation potential

The exchange-correlation potential $V_{xc}(\vec{r},t)$ from Eq. (5.19) is a functional of the time-dependent electron-density $n(\vec{r},t)$, and all the values of the electron-density at the previous times, $n(\vec{r},t')$, with $0 \leq t' < t$ [210]. This rather subtle dependence on the history of the electron-density up to the current time t is almost always ignored altogether in all practical calculations by using the so-called **adiabatic-approximation**. This means that $V_{xc}(\vec{r},t)$ has the same expression as that from $t = 0$, but with the time-dependent electron density $n(\vec{r},t)$ replacing $n(\vec{r})$. This can translate to any functional-type from those presented earlier, however the time-dependent simulations presented in this chapter all used the **ALDA**. Unfortunately, as time passes during the simulation, the memory of the previous timesteps does not disappear, but continuously accumulates and consequently becomes harder and harder to completely ignore it. There is no simple mechanism to help cleaning up this dependence. The accuracy of the numerical predictions based on this approximation can only be judged when compared against real, experimentally-measured, results, for the interaction-scenario of interest. However, it has been already established that it gives extremely accurate results for low-frequency laser-fields [211], and can be successfully quantitatively compared and contrasted to lowest-order-perturbation-theory (LOPT) predictions for very high-frequency laser-fields [212].

Atomic ionisation modelling: the ionisation probabilities

Having explained the numerical aspects of the time-propagation, the focus is now on how to get information from the KS-orbitals after they have been numerically-propagated in time under the influence of the laser-field. By information, in this chapter, I refer to predictions on how many electrons have left the atom after the passage of the laser pulse (i.e. the charge state of the atom). There are still open questions onto this direction, i.e. on how to retrieve ionisation probabilities from the time-dependent KS-wavefunction (orbitals), and the results which will be presented in this chapter followed the **statistical-approach firstly**

proposed by Ullrich [213]. There, an approximate formula for the probability of ejection of k electrons from the atom has been derived based on the single-particle orbitals from the KS-auxiliary system. Two approximations are highlighted in the derivation: **a)** considering the interaction between the laser and the atom to take place in a finite simulation domain (which grid methods do, with the GPS, being a discrete-variable-representation-type method employing both a grid and a basis, falling into this category), the norm inside this region is taken to represent the probability to be bound and **b)** the replacement of the true many-body wavefunction with the KS-system wavefunction. The approximate ionisation probabilities formulae obtained and used to extract information from the TDDFT simulation are thus based on the usual simplifications encountered in practice. It is instructive to briefly go through the derivation, with all the details to be found in [213].

Let a part of the simulation domain volume be denoted by ν, with $\bar{\nu}$ being what would be beyond ν. For spherical-coordinates, the radial coordinate r defines ν and $\bar{\nu}$. By means of example, any point in the real space with its radial coordinate $r \leq R$ is in ν and any other point, i.e. with $r > R$, is in $\bar{\nu}$. Then it is the case that the integral across $(\nu + \bar{\nu})$ of the real, many-electron wavefunction shall always equal 1.0, because ν and $\bar{\nu}$ together extend from $r = 0$ to infinity. Mathematically, this means that for any time t:

$$1.0 = \int d\vec{r}_1 ... \int d\vec{r}_N \, |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2. \tag{5.26}$$

Then each of these spatial integrals is partitioned into two components, one across ν and the other across $\bar{\nu}$:

$$1.0 = \left(\int_\nu d\vec{r}_1 + \int_{\bar{\nu}} d\vec{r}_1 \right) ... \left(\int_\nu d\vec{r}_N + \int_{\bar{\nu}} d\vec{r}_N \right) |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2. \tag{5.27}$$

Remembering that the formula of the binomial coefficient is:

$$\binom{b}{c} = \frac{b!}{c! \, (b-c)!}, \tag{5.28}$$

Eq. (5.27) becomes:

$$
\begin{aligned}
1.0 = & \int_{\nu} d\vec{r}_1 ... \int_{\nu} d\vec{r}_N \, |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2 + \\
& \binom{N}{1} \int_{\bar{\nu}} d\vec{r}_1 \int_{\nu} d\vec{r}_2 ... \int_{\nu} d\vec{r}_N \, |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2 + \\
& \binom{N}{2} \int_{\bar{\nu}} d\vec{r}_1 \int_{\bar{\nu}} d\vec{r}_2 \int_{\nu} d\vec{r}_3 ... \int_{\nu} d\vec{r}_N \, |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2 + \\
& ... \\
& ... \\
& ... \\
& \binom{N}{N-1} \int_{\bar{\nu}} d\vec{r}_1 ... \int_{\bar{\nu}} d\vec{r}_N \, |\Psi(\vec{r}_1, ..., \vec{r}_N, t)|^2. \qquad (5.29)
\end{aligned}
$$

Denoting the probability that the atom is in charge state Q, where $Q = 0$ signifies a neutral atom, $Q = 1+$ a singly-ionised atom, and so on, up to $Q = N+$, where N is the atomic number of the target, which represents a fully-ionised atom, to be $P^Q(t)$, it is clear that the sum $\sum_{Q=0}^{N} P^Q(t)$ also equals to 1.0. This is easily understood from the fact that the atom has to be in a given charge state, at any given time, and because there are $(N + 1)$ such charge states, their summed probabilities equal to 1.0 (the certainty in probability theory). One then makes the identification (based on how many integrals are performed across ν and how many across $\bar{\nu}$) that the first term on the right of the equal sign from Eq. (5.29) is $P^0(t)$, the second term is $P^{1+}(t)$, the third term is $P^{2+}(t)$ and so on, up until the last term is $P^N(t)$. It is visible that the last term has all its integrals performed across $\bar{\nu}$. The **first approximation** is now made: the separation between ν and $\bar{\nu}$ is equated to the separation between bound and unbound (continuum) states. This assumes that ν includes all the bound states relevant for the interaction-scenario under consideration, and all the electron density which is not in ν is considered to be ionised. In practice, this is done with the help of an absorber (mask function) placed at the end of the simulation domain, and the integral across ν decreases in time. Of course, the spatial domain has to be large enough to include all the temporal-dynamics relevant to the laser and target characteristics which are input to the simulation. At all timesteps, the difference from 1.0 is what would be

contained in $\bar{\nu}$, i.e. the unbound-parts. With this explanation at hand, the **second approximation** is performed: the many-electron wavefunction $\Psi(\vec{r}_1, \vec{r}_2, ..., \vec{r}_N, t)$ is replaced with the KS-system wavefunction (a single Slater-determinant in the KS-orbitals $\phi_i(\vec{r}, t)$ where $i = 1, ..., N$). For the time-evolution of each of these orbitals an absorber is placed at the end of the simulation domain (all absorbers being identical to each other) and the integrals across ν can be readily calculated. In a numerical implementation, one almost always wants an automatic way to handle any target (atom). This is because when the atomic number of the element starts being large, it becomes impossible to input all the contributions in Eq. (5.29) by hand. Thus, denoting $M_1(t) := \int_\nu d\vec{r}\, |\phi_1(\vec{r}, t)|^2$, $M_2(t) := \int_\nu d\vec{r}\, |\phi_2(\vec{r}, t)|^2$, $\bar{M}_1(t) := \int_{\bar{\nu}} d\vec{r}\, |\phi_1(\vec{r}, t)|^2$, $\bar{M}_2(t) := \int_{\bar{\nu}} d\vec{r}\, |\phi_2(\vec{r}, t)|^2$, and so on, Eq. (5.29) can be reformulated as below, noting that $M_j(t) + \bar{M}_j(t) = 1.0$ for any t:

$$\sum_{Q=0}^{N} P^Q(t) = \prod_{j=1}^{N} \left(M_j(t) + \bar{M}_j(t) \right). \tag{5.30}$$

This form is helpful because the probability for no-ionisation $P^0(t)$ is obtained as the sum of all the terms which contain no $\bar{M}$-factors, but only M-factors (i.e. no bars involved at all). Similarly for the probability for single-ionisation, the aim is to collect all the terms which contain one $\bar{M}$-factor, and $(N-1)$ M-factors. The pattern is the same for all charge states, up to the fully-ionised atom $P^N(t)$, which is the sum of those terms containing only $\bar{M}$-factors, i.e. no M-factors. I have made use of `Mathematica`© and a high-level language (`Python`) to programatically obtain the expressions for the ionisation probabilities of all charge states for argon ($N = 18$). In short, the binomial coefficients are calculated and exported to a file on the disk using `Mathematica`©, and the postprocessing of the norms of all the KS-orbitals from the simulation and their assembly in sums as in Eq. (5.30) is done in the scripting-language. I developed and applied this workflow to calculate the probabilities at the end of the simulation only. However, if the probabilities throughout the simulation (i.e. at each timestep of a given simulation) are needed, then a compiled language has to be used for this task. Potentially, this has to include a distributed-computing (message-passing)

library such as `MPI`, as the same exact postprocessing operations are applied for each timestep, fully independent of each other, thus being trivially parallelisable.

A very thorough study on X-rays illuminating atoms has been done in the work from [212]. There, the simulation framework OCTOPUS in its Cartesian-3D mode was used to run (real time propagation) TDDFT calculations and look at ion yields through the statistical approach of Ullrich. Overall, in [212] it is found that TDDFT offers quantitatively correct (valid) predictions when compared against lowest-order-perturbation-theory (LOPT) results, and that the agreement depends on the type of XC-kernel used.

In their simulations, they used a Dipole-Approximated laser pulse with a gaussian temporal envelope, for which they present the full-width-at-half-maximum-intensity (FWHM) duration.

It is instructive to compare and contrast my results (obtained for roughly the same set of physics input parameters) against their results, to benchmark my TDDFT solver and workflow.

I set my central laser frequency the same as theirs ($\omega = 3.85$ a.u.), but kept my original sin-squared temporal profile.

I targeted the same atom (argon) and performed a parameter scan across the peak laser-field value, for a fixed temporal length which the temporal envelope spawns, just as they did. It was nice to discover closer and closer agreement with their results as I increased the number of optical-cycles from my sin-squared envelope. That is, I actually run three sets of TDDFT simulations, each set having a fixed value of the laser pulse's length (10 fs, 30 fs, 40 fs).

My results for the 40 fs duration set of simulations are presented in Fig. 5.1. The box size was $R_b = 50.0$ a.u., the number of radial grid points was $N_r = 150$, $l_{max} = 31$, and the set of simulations consisted of 40 different ones. There were 1000 timesteps per one optical-cycle, for each such simulation. The set of simulations spawned an intensity range between 10^{13} - 10^{16} W/cm^2.

It is known that if the FWHM of a gaussian is a given number of optical-cycles, a sin-squared envelope spawning the same number of optical-cycles does not

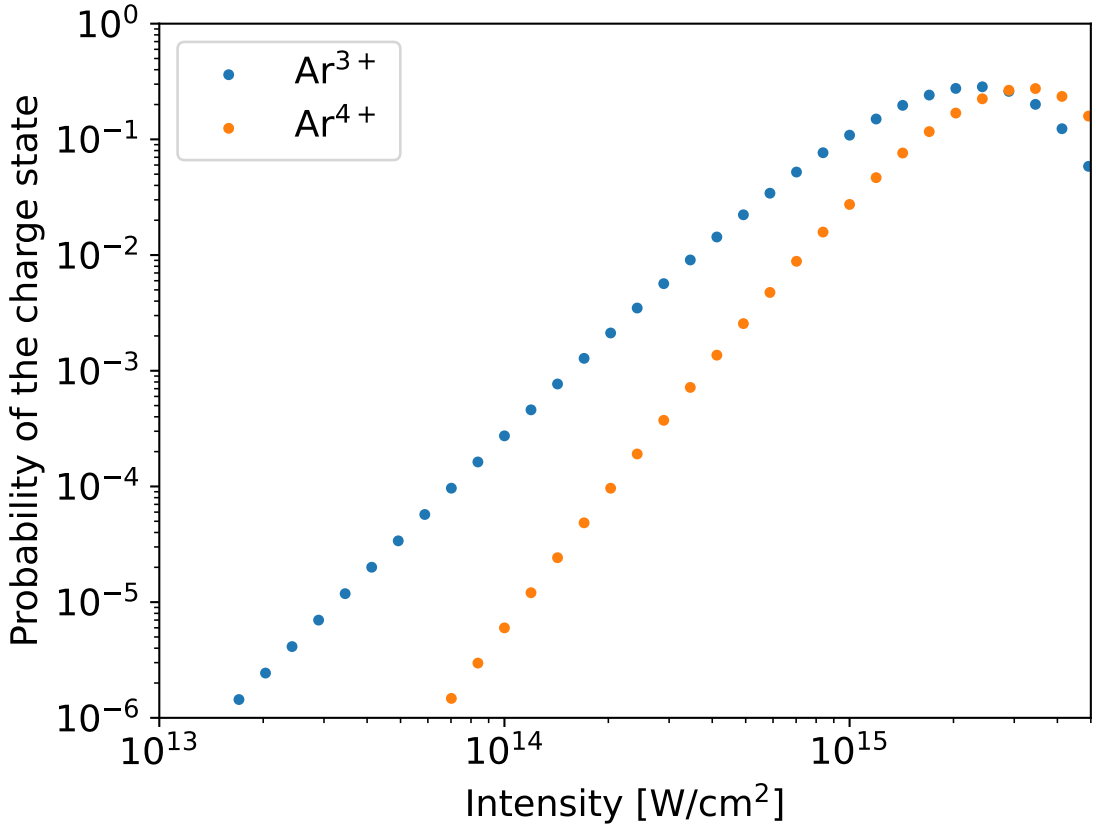


Figure 5.1: Triply ionised and four-times ionised argon atom probabilities at the end of TDDFT simulations, as a function of the peak intensity of the laser field for the simulation, at a fixed central frequency $\omega = 3.85$ a.u., for a sin-squared temporal envelope of 40 fs duration. The box size was $R_b = 50$ a.u. and there were 1000 timesteps per one optical-cycle for all simulations.

produce the same dynamics. It is only when the sin-squared spawns way more optical cycles than the FWHM-value the dynamics become roughly similar. This is because a gaussian does not actually drop to 0.0 intensity beyond the FWHM window, but the sin-squared actually drops to exactly 0.0 after the number of optical-cycles it spawns finishes.

Without loss of generality, I chose to focus on Ar^{3+} and Ar^{4+} charge states and compare and contrast against the results from panel b) of Fig. 4 of [212]. The curves representing LOPT results are drawn with solid lines: a green line for Ar^{3+} and a blue line for Ar^{4+}. The same colour scheme was used for their TDDFT results in the same figure, but with dashed lines.

The horizontal and vertical axes spawn the same range (from 10^{13} to 5.0 $\times$

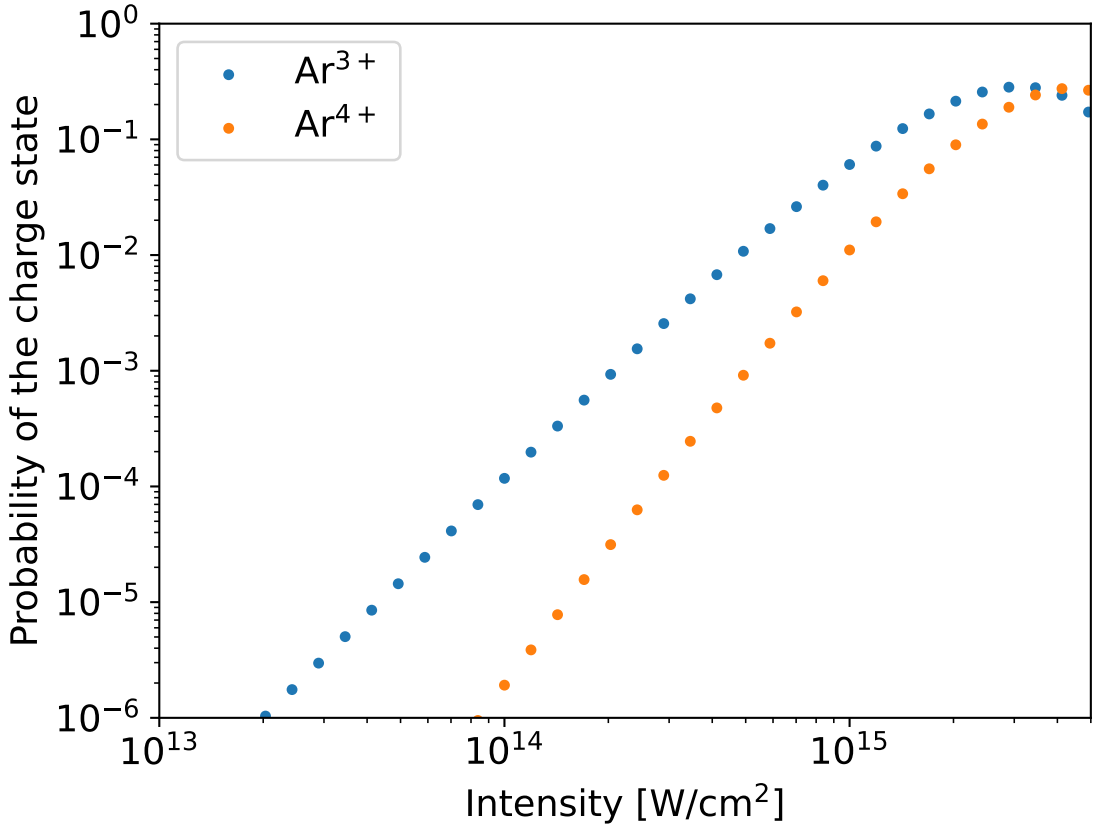


Figure 5.2: Triply ionised and four-times ionised argon atom probabilities at the end of TDDFT simulations, as a function of the peak intensity of the laser field for the simulation, at a fixed central frequency $\omega = 3.85$ a.u., for a sin-squared temporal envelope of 30 fs duration. The box size was $R_b = 50$ a.u. and there were 1000 timesteps per one optical-cycle for all simulations.

10^{15} W/cm^2 across the x-axis and from 10^{-6} to 1.0 across the y-axis) in both their panel and my plots, and both axes are logarithmic in both plots.

It is visible that the agreement between my simulations' results and LOPT results is very good, for both charge states. The agreement gets better for the 40 fs duration sin-squared laser pulse than it is for 30 fs duration sin-squared laser pulse, for reasons explained above.

The results for the 30 fs duration are presented in Fig. 5.2.

Different than what is presented in the paper, I used the LDA XC-functional. From their study, it is visible that even with the way more performant LB94 XC-potential, there exist some differences with respect to LOPT. It is thus to be expected that my results do not perfectly match LOPT, and this is clearly seen.

Still, the agreement is visible. Another sources of inaccuracy in my simulations would be that 10 fs of a Gaussian (what they used) is surely not perfectly matched by either 30 fs or 40 fs of a sin-squared envelope, but also that 150 radial grid points are not necessarily enough for full-convergence for my TDDFT simulations. Still, because these are parameter scans across intensity, I am limited by the computational resources and time. Perhaps around 500 radial grid-points for each simulation would improve the results even more, together with a decrease of the simulation timestep by a factor of 4 or even 10.

Overall, this shows a very good agreement with results from literature, benchmarking my TDDFT solver based on the GPS discretisation.

Results and discussion

Having had run the ground-state SCF calculations, the obtained orbitals from these numerical simulations were the starting point for the time-dependent simulations. Thus the grid discretisation (real-space-domain size and number of grid points) and the XC-potential parameters (for ALDA, the parameter for restoring the Coulomb-asymptotic behaviour) were kept the same. The laser pulse temporal profile was chosen in all simulations to have a $\sin^2$-envelope:

$$envelope(t) = \sin^2\left(\frac{\pi t}{T}\right) \tag{5.31}$$

The central frequency was fixed to $\omega_{laser} = 3.5$ a.u., and the temporal durations varied depending on the simulation, from $T = 10$ optical-cycles (OC's) to $T = 100$ OC's, where one OC is equal to $\frac{2\pi}{\omega_{laser}}$ a.u. of time. For the current ω_{laser}, one OC is equal to 1.80 a.u., or equivalently, 43.08 as. The absorber was placed to start at $R_a = 90$ a.u. and was chosen to be a mask-type function depending on the radial coordinate r:

$$mask(r) = \begin{cases} \cos^{0.25}\left(\frac{\pi}{2}\frac{(r-R_a)}{(R_b-R_a)}\right) & \text{, if } R_a < r < R_b \\ 1 & \text{, if } 0 < r < R_a. \end{cases} \tag{5.32}$$

In all the simulations, the maximum angular momentum appearing in the expansion of the orbitals and potentials was set to be $l_{max} = 62$.

Because the electric field of the laser pulse was linearly polarised (along the $\hat{z}$-axis) in all the simulations, the initial symmetry of the problem is conserved throughout all the time-dependent simulation. It is also the case that the magnetic quantum-number m is a good quantum-number and thus is conserved throughout all the temporal evolution. A check for correctness of the implementation of the KS-TDDFT equations is to observe if the dynamics of the single-particle orbitals which are equivalent from the point of view of the laser-field are actually similar. This can be done, for example, observing the $3p_{-1}$ and the $3p_{+1}$ single-particle orbitals of the argon target. Their magnetic quantum numbers are -1 and $+1$, respectively, and for a laser oscillating purely vertically (linear polarisation) across the $\hat{z}$-axis, these are equivalent and accordingly, the simulation results show this as well.

The spatial integrals across the volume ν of the orbitals are calculated throughout the whole simulation domain, at each timestep, as this is an efficient operation when using the spectral representation of the single-particle wavefunctions and does not carry a great computational effort. These values all started at 1.0 and decreased throughout the simulation as the absorber was acting on each single-particle orbital.

Various laser pulse temporal durations have been analysed. For each temporal duration, one set of 128 TDDFT simulations has been run. Each simulation from a given set had a different peak E-field strength, but apart from this, they all have been identical, in particular all had the same laser-field implemented. The aim was to observe the distribution of the populations of the charge states of the atom at the end of the laser pulse, as the peak intensity was varied. The range of luminous-intensities which has been (logarithmically) sampled was from 10^{16} W/cm^2 to close to 10^{20} W/cm^2. This workflow has been repeated for the following durations: 10 OC's, 20 OC's, 30 OC's, 50 OC's and 100 OC's. Equivalently, the durations were: 430.85 as, 861.69 as, 1292.54 as, 2154.24 as, 4308.47 as. Of course, longer durations meant that the simulations had to be let to run for longer wall-clock-times.

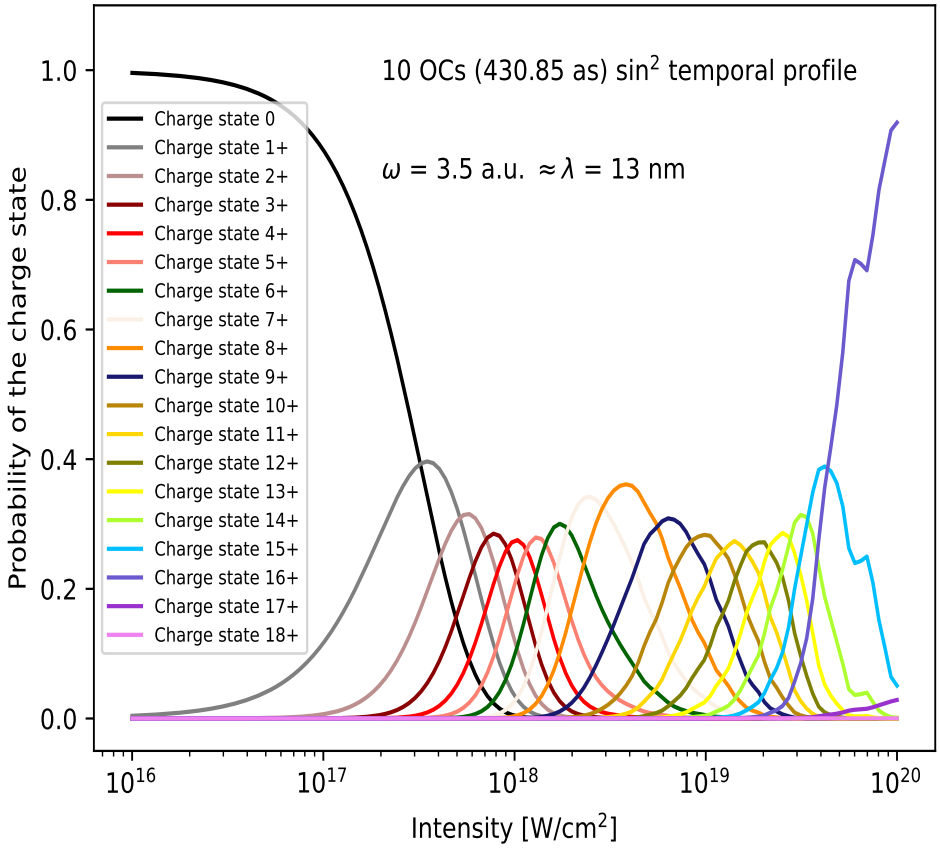


Figure 5.3: Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 10 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. In particular, the simulation domain extended up to 100 a.u. in the radial direction and the $\cos^{0.25}$ absorber started at 90 a.u. of radial extent from the origin. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The numerical discretisation for the governing equations has been performed with the aid of the generalized-pseudospectral-grids (GPS) formalism. For all the simulations, the number of radial-points was $N_r = 150$ and the number of angular-points was $N_\theta = 63$.

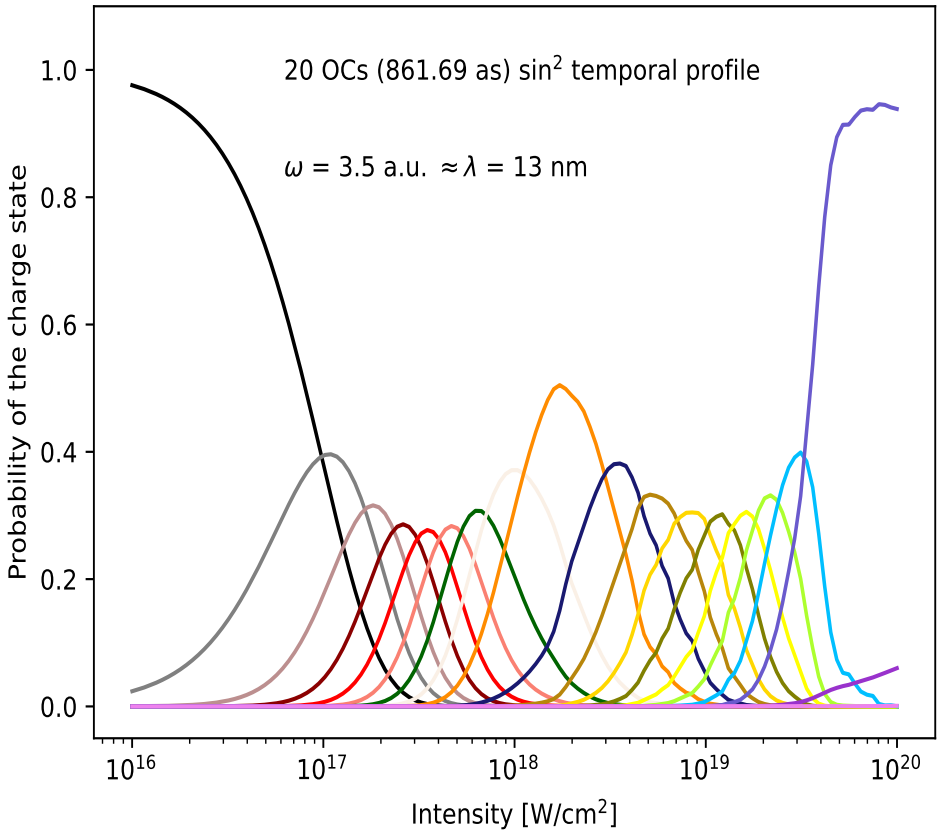


Figure 5.4: Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 20 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here.

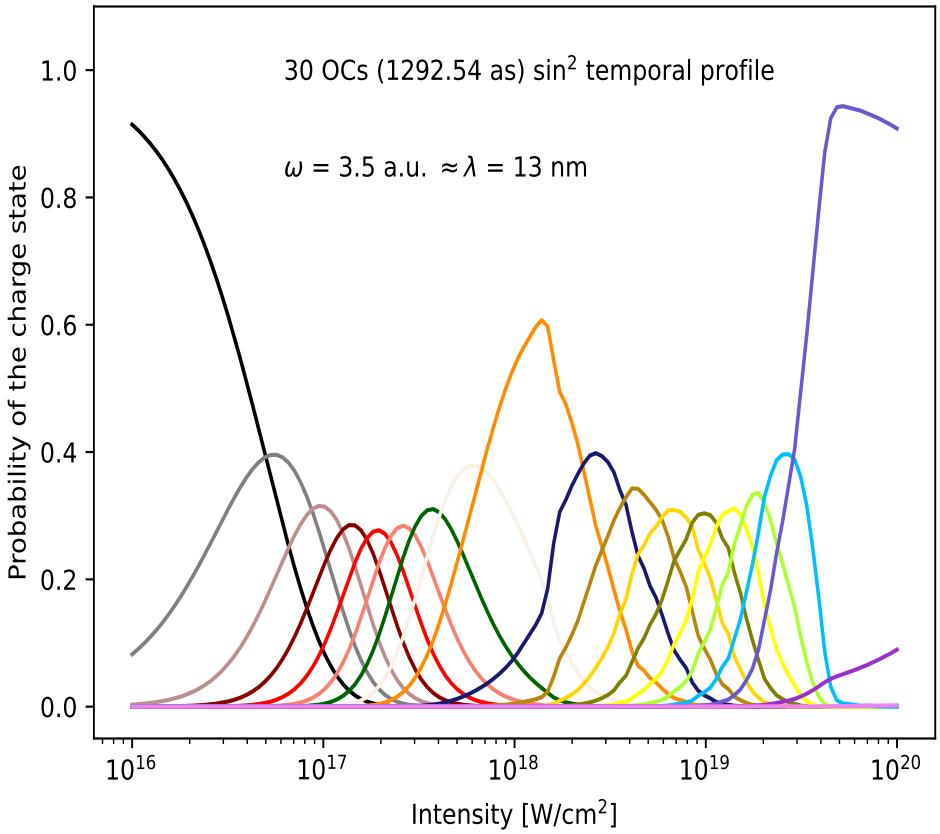


Figure 5.5: Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a sin^2-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 30 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here.

The first two sets of simulations (10 OC's (430.85 as) and 20 OC's (861.69 as)) are shown in Figures 5.3 and 5.4, each figure having 128 ticks across its horizontal axis (the intensity-axis). Focusing on the 10 OC's, it is clear that the higher charge states get populated as the peak laser-field of the simulation increases.

The source of this is the neutral atom state, i.e. the charge state 0, for which its population always gets lower and lower as the laser-field peak intensity increases.

This is to be expected, more and more electrons are to be ejected as the electrical-force from the laser E-field acting on the electron is more powerful. For very high field-strengths, the atom is almost surely ionised. The atom cannot be fully ionised for any of these laser-field peak-intensities. It is expected that a longer pulse duration, for a given fixed peak field-strength, is able to ionise the atom more. This is confirmed in the second set of simulations, all which had a temporal envelope of 20 OC's. While the same behaviour of consecutive charge states being populated as the intensity is increased is clearly visible, the curves are all shifted to the left relative to the 10 OC's case. The speed of the decrease in the neutral-atom state is more pronounced, and the range of intensities from the lower end of the simulated domain (close to 10^{16} W/cm^2) which in the 10 OC's case did not ionise the atom at all almost disappear in the 20 OC's case. There is also a slight increase in the charge state 17+ population (one electron remaining in the atom) relative to the 10 OC's case, however it is clear that even for 20 OC's the atom cannot be fully ionised for this frequency and these simulated peak-field strengths. It seemed a good avenue to explore if longer pulse lengths, that is, 30 OC's (1292.54 as), 50 OC's (2154.24 as) and 100 OC's (4308.47 as) manage to populate the final charge states.

The 128 simulations which had a laser-field of 30 OC's duration are presented in Fig. 5.5. All the simulations from this set have had the atom ionised in a considerable fashion at the end of the propagation. That is, even for 10^{16} W/cm^2, the neutral atom population is clearly different than 1.0, in contrast to the 20 and 10 OC's simulations at the same peak intensity. Towards the other side of the intensity range, the appearance of the charge state 17+ is a bit more pronounced, however this not dramatically changed from the results from the other sets.

It might be interesting to note the behaviour of the charge state 8+. Its population at the end of the pulse keeps increasing as the duration increases. For 10 OC's case, its maximum value is slightly above 0.3. For the 20 OC's case, its maximum value has increased. There is an even more increase for the 30 OC's case. Looking at the ionisation energies for argon [214], there is a big jump at the

one changing the atom from 8+ to 9+. That is, the 5+ to 6+ energy is listed to be 91.29 eV, the 6+ to 7+ one is 124.41 eV, the 7+ to 8+ one is 143.46 eV, while the 8+ to 9+ one is 422.60 eV, a considerably larger change relative to the differences between the consecutive previous ones. This might be a potential reason to explain the population increase of that 8+ charge state as the pulse duration increases. Still, the peak is also shifted to the left as the duration increases, that is, it appears at a lower intensity value, so this complicates and inhibits a clear explanation for this observed phenomenon.

Finally, the last two sets of 128 simulations, one of which had the laser-field of 50 OC's duration and the other which had the laser-field of 100 OC's duration, are shown in Fig. 5.6 and in Fig. 5.7. These sets confirm the behaviour of the charge state 8+ inferred from the previous sets, i.e. its behaviour as the length of the pulse duration increases. Interestingly, the behaviour of charge state 16+ is roughly similar to that of 8+, and this starts to be revealed in these 50 OC's simulations. Looking at the ionisation energies of argon, it is striking that to get from 14+ to 15+ one needs 855.5 eV, from 15+ to 16+, 918 eV, but from 16+ to 17+ one needs 4120 eV. This is quite the same pattern as for the sequence 6+ to 7+, 7+ to 8+ and then the jump at 8+ to 9+.

It is also to be noted that some numerical noise can be visible more and more as the durations of the laser pulses increases. This has to do with the numerical discretisations used in the simulations. All the sets kept the same number of grid points and timestep duration between them. Still, a 50 OC's simulation is significantly harder to converge than a 10 OC's one, not to mention a 100 OC's one. This is because in the GPS method, there exists a numerical error for the split-operator in the energy representation time-propagation scheme, and this is in place at each timestep of the propagation. To reduce it, one can decrease the timestep duration, or decrease the field strength (which here, for a parameter scan across the intensity, is not possible to be done). Unfortunately, as all the orbitals have to be propagated simultaneously in this type of simulations, a quite stringent

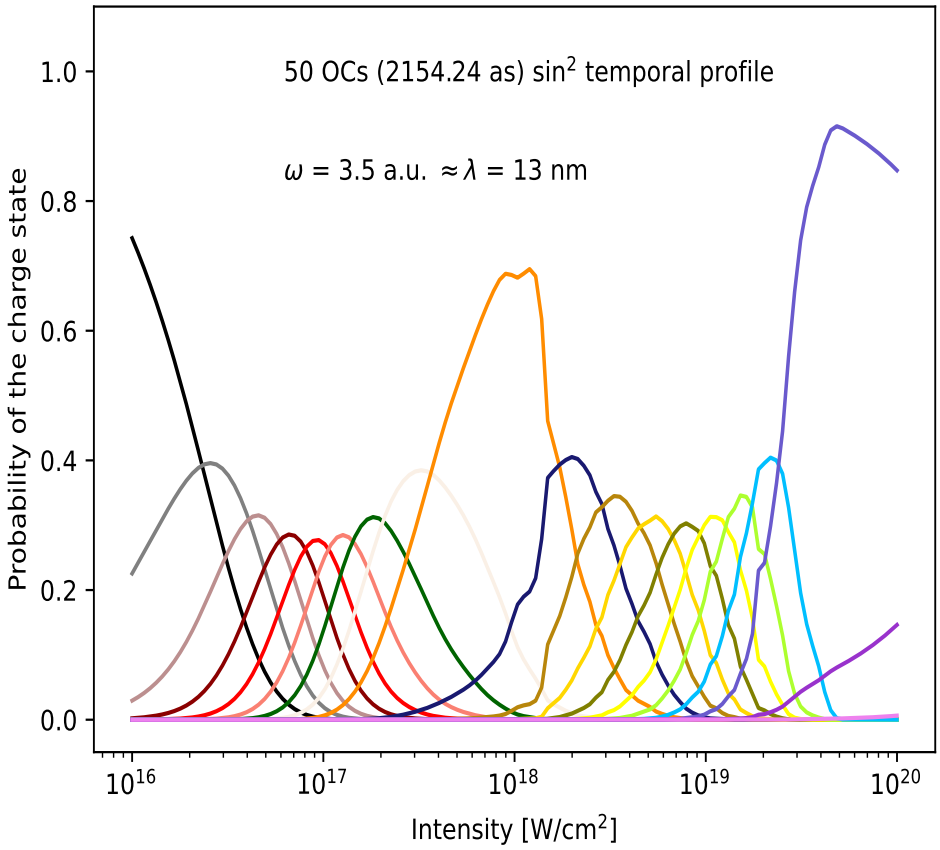


Figure 5.6: Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 50 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here.

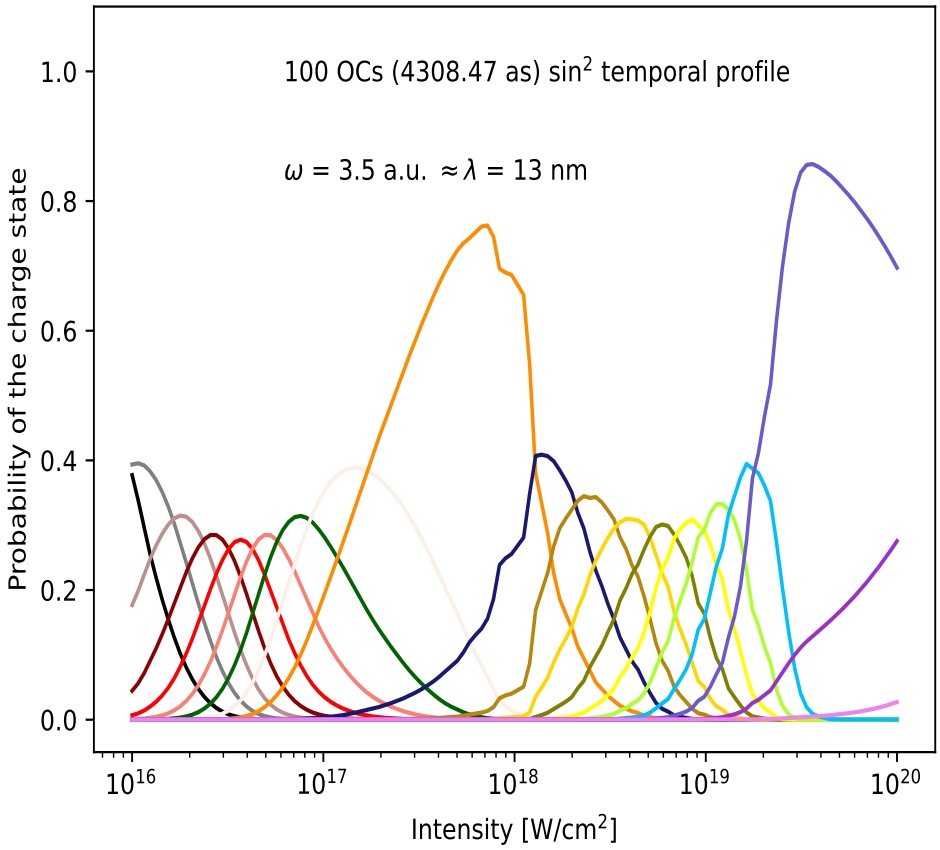


Figure 5.7: Kohn-Sham Time-Dependent-Density-Functional-Theory (KS-TDDFT) results for the ionisation probabilities at the end of a given simulation for an argon atom (having its atomic number $Z = 18$) illuminated by a $\sin^2$-temporal-envelope laser pulse of central frequency 3.5 a.u. and of 100 optical-cycles in duration. Different simulations had different laser peak intensities and the horizontal axis has been sampled with 128 such simulations. In all other aspects, the simulations have been set-up in an identical fashion. Each simulation produced one datapoint on each of the 19 curves presented. The curves signify the charge states of the atom: the charge state denoted by 0 means a neutral atom, the charge state denoted by 1+ means a singly ionised atom, and so on, up to 18+ which represents a fully ionised atom. The mapping between the colour of the curve and the actual charge state is the same as in the label of Fig. 5.3 and has been omitted here.

upper bound on the maximum possible timestep which can be used such that the simulations are actually feasible (run in a reasonable wall-clock time) is in place.

5.3 Hartree-Fock formalism

It would be interesting to try to gain insight into the atomic ionisation process through the TDHF formalism as well. In this approach, one Slater-determinant is again used to represent the correctly anti-symmetrised many-electron wave-function. The determinant is composed of orbitals, similarly to how the KS-wavefunction is constructed, however in TDHF the orbitals are intended to represent real electrons: there is no auxiliary-system to resort to. Still, again the problem is reduced to solving coupled one-electron TDSE-like equations, but the physics arguments motivating these equations possess differences with respect to TDKS. In short, the focus is shifted **from** non-interacting electrons under the influence of a mean-field potential composed from **a)** the classical Coulomb potential and **b)** a local exchange-correlation potential **to** again, non-interacting electrons, but under the influence of a mean-field potential composed from **a)** the classical Coulomb potential and **b')** a non-local exchange potential.

It is clear that a major disadvantage in pure-HF is that there is no correlation included. In TDDFT, this is accounted for, albeit in an approximate way. **Post-Hartree methods** can help solving this issue of the lack of correlation in HF, however they will not be discussed here at great length. The first method which could be mentioned is the **multi-configuration TDHF** (MCTDHF) formalism. This approximates the multi-electron wavefunction to be a linear combination of Slater-determinants and the orbitals of all these determinants evolve in time, i.e. a generalisation of the single Slater-determinant which evolves in time approach. In MCTDHF also the coefficients multiplying each determinant are time-dependent and change throughout the simulation. Thus there are two sets of coupled differential equations: one set for the determinants (i.e. the orbitals), and one set for the coefficients. These have to be solved (propagated) together by a time-integrator routine, many times the (adaptive) Runge-Kutta(-Felhberg)

method being preferred for its ease of implementation [215], however other, more advanced time-stepping routines can be used and some have been implemented in production-level-softwares such as `GAUSSIAN` [216]. While improving upon the accuracy of the results, considering many determinants in the expansion of the many-electron wavefunction leads to an increase in the computational effort, which for this approach can be demonstrated to be non-linear. Another post-Hartree method which can be used to determine many-electron wavefunctions is the **Configuration-Interaction** (CI), which also quickly expands to unfeasible computational costs in its full-CI form, thus truncations are used in practice, such as the CI singles-doubles (CISD), where only one or two ground state orbitals from the reference-determinant are replaced with other orbitals in the terms which compose the sum over the determinants. Finally, the **Coupled-Cluster** (CC) [217] methods are many times the choice for accurately obtaining the many-electron wavefunction of the quantum-system of interest.

Returning to TDHF, it is to be mentioned that including the full-exchange is non-trivial on its own because of its spatial non-locality, even if there is no correlation involved at all. In practice, the non-locality translates into the need that a non-trivial number of real-volume-integrals to be performed in a computationally tractable way. Efficient methods have been developed for this, many of them exploiting the mathematical properties of the basis functions used to represent the atomic orbitals and making use of very optimised computer-numerical-routines. Using only basis functions, that is, resorting to a spectral representation only to describe the atom, possesses both advantages and disadvantages. The radial integrals can be calculated efficiently, and there has been specialised effort invested to this over time. Ionisation, in particular, is many times poorly defined (or cannot be accurately treated) in this context, for a non-trivial variety of basis functions. This is why a real-space-grid is many times needed and used for real-time propagation. The combination between these two approaches leads to the

discrete-variable-representation (DVR)-type discretisations: the finite-element-method (FEM)-DVR combination has been observed to perform rather well, at reasonable computational costs [218].

To derive the (ground-state) HF-equations, the HF-ansatz is made for the many-electron wavefunction Ψ:

$$\Psi_{HF}(\vec{r}_1, \vec{r}_2, ..., \vec{r}_N) = \frac{1}{\sqrt{N!}} \begin{vmatrix} \psi_1(\vec{r}_1) & \psi_2(\vec{r}_1) & ... & \psi_N(\vec{r}_1) \\ \psi_1(\vec{r}_2) & \psi_2(\vec{r}_2) & ... & \psi_N(\vec{r}_2) \\ ... & ... & ... & ... \\ ... & ... & ... & ... \\ \psi_1(\vec{r}_N) & \psi_2(\vec{r}_N) & ... & \psi_N(\vec{r}_N) \end{vmatrix}, \tag{5.33}$$

where the ψ_i's with $i = 1, 2, ..., N$ represent the single-particle orbitals and the vertical bars signify a determinant. This is plugged in Eq. (5.1) and with the Hamiltonian from Eq. (5.2) one obtains that the total energy $E_{total} = \langle \Psi_{HF} | \hat{H} | \Psi_{HF} \rangle$ is equal to:

$$E_{total} = \int \sum_i^N \psi_i^*(\vec{r}) \left(-\frac{1}{2}\nabla_i^2 + V_{ext} \right) \psi_i(\vec{r})\, d\vec{r} + \tag{5.34}$$

$$\frac{1}{2} \sum_{i,j}^N \iint \frac{\psi_i^*(\vec{r}_1)\psi_i(\vec{r}_1)\psi_j^*(\vec{r}_2)\psi_j(\vec{r}_2)}{|\vec{r}_i - \vec{r}_j|} d\vec{r}_1 d\vec{r}_2 - \tag{5.35}$$

$$\frac{1}{2} \sum_{i,j}^N \iint \frac{\psi_i^*(\vec{r}_1)\psi_j(\vec{r}_1)\psi_i(\vec{r}_2)\psi_j^*(\vec{r}_2)}{|\vec{r}_i - \vec{r}_j|} d\vec{r}_1 d\vec{r}_2. \tag{5.36}$$

The terms from the second row are the Coulomb integrals, totalling an energy contribution which can be explained classically from the point of view of electrostatics, and the terms from the third row are the exchange integrals, totalling an energy contribution which cannot be explained classically, but only quantum-mechanically: these exchange-operators arise purely from the Pauli exclusion principle.

Applying the variational principle to this expression of the HF-total energy leads to the HF-equations, similar in form to the KS-DFT equations, but with a different mean-field potential (a non-local one in this case):

$$\left(-\frac{\nabla^2}{2} + V_{ext}(\vec{r}) + V_H(\vec{r}) \right) \psi_i(\vec{r}) + \int V_x(\vec{r}, \vec{r}\,') \psi_i(\vec{r}\,') d\vec{r}\,' = \varepsilon_i \psi_i(\vec{r}), \;\; i = 1, 2, ..., N. \tag{5.37}$$

Again, the mean-field potential depends on the orbitals, so an iterative procedure to solving Eq. (5.37) is performed, just as for the KS-DFT Eq. (5.4). The initial guess for the orbitals can, in principle, be of various forms, however a simple choice is the hydrogen-like atoms' analytical one-electron wavefunctions.

As mentioned, many times the equations are usually solved in a basis (i.e. the spectral-approach). That is, the orbitals are expanded in a linear combination of fixed basis-functions, and the HF-equations become the Roothaan–Hall-equations [219]. Then the aim is to find those set of coefficients which compose the expansion generating the minimum total-energy, in a self-consistent way. This is obtained by diagonalizing the so-called Fock-matrix at each SCF-iteration. The size of this matrix is directly related the accuracy of the end-results of the orbitals' energies, as the inclusion of more basis-functions in the expansion sums lead to an increased dimension of this matrix and thus a greater computational effort.

Usually the basis functions which are used are the Gaussian ones or variants of this type, choices including a wide range of complexities built upon the pure-Gaussians. This is because integrals involving this type of functions can be solved with the help of well-known formulae (an example being the product of Gaussians) and an exact, analytical result can be obtained and used.

There also exist pieces of work for which Slater basis functions have been employed: they are similar to analytical hydrogen-like-atoms TISE-solutions, but have no nodes (the functions do not touch zero at any point), however they do keep the asymptotical exponential decay from the origin.

5.3.1 Results, discussion and further work

Ground-state

I have used the open-source simulation framework PySCF [220] to obtain the HF-approximation to the ground-state wavefunction of an argon atom. The initial guess for the orbitals has been the one-electron analytical forms which solve the TISE for a hydrogen-like atom. Although more advanced initial-guesses exist in PySCF, and are known to produce more accurate end-results, this simple choice of

the form of the orbitals at the beginning of the SCF-iterations produced reasonable values for the orbital energies, in agreement with the published literature [221]. The convergence was fast, as it is expected for a single closed-shell atom having a relatively low atomic number Z. Still, the speed with which the iterations finished was also motivated by the special acceleration techniques used in PySCF [222]. This has been briefly touched upon in Sec. 5.2.1. The basis functions used were from the Dunning's correlation consistent basis sets, in particular, the cc-pVDZ ones. The orbitals energies obtained are presented in Table 5.2, together with the same experimental results from Table 5.1.

It is instructive to compare and contrast the HF results with the ground-state KS-DFT results. It is clear that there is an improvement upon the LDA-KS-DFT eigenenergies, even if the HF-formalism does not include correlation. The HF orbital energies are, depending on the orbital considered, closer, or very much closer to the experimentally determined values. This is because the quality of the KS-DFT calculations crucially depend on the XC-energy functional parametrisation used, so even if correlation is taken into account, the fact that its inclusion is an approximation to it only, and that for this atom, the correlation energy is rather small, explains the differences observed in the presented energies. Actually, with a more performant XC-potential, the LB94, the KS-DFT eigenenergies which are numerically obtained for the argon atom have a better chance to be considered on par with the HF results (Table 1 from [104]).

The estimation of the correlation energy has been obtained again using PySCF by running a post-Hartree-method ground-state simulation: this yielded $E_{correlation} = -0.14$ a.u., to be compared with the total energy yielded by the same calculation, $E_{total} = -526.80$ a.u. but also to the individual orbital energies from Table 5.2.

Time-dependence: potential future work

While costing a non-trivial amount of CPU-time, a possible direction to pursue for including the time-dependence is using the OCTOPUS [223] simulation-

	Py-SCF HF	Literature HF	Experimental data
1s	-118.61	-118.61	-117.83
2s	-12.32	-12.32	-12.00
2p	-9.57	-9.57	-9.13
3s	-1.27	-1.28	-1.08
3p	-0.59	-0.59	-0.58

Table 5.2:
Hartree-Fock orbital energies for argon obtained with PySCF using the cc-pVDZ basis set and one-electron hydrogen-like analytical solutions as initial guesses. The literature-HF reference results are from [221]. The experimental values are from [206]. All are presented in a.u..

framework. Unfortunately, it does not exploit the spherical symmetry of the atoms and the simulations are performed in full Cartesian coordinates (x, y, z) rather than in a two-dimensional setting, i.e. the (r, θ) coordinates, and this adds up to the cost of calculating the exchange-integrals and managing to have a fine-enough spatial discretisation to resolve (at least approximately) the inner shells. Still, because OCTOPUS is fully distributed across cores by using both `MPI` and `OpenMP`, using more computational resources the time for the simulations to complete can be made shorter.

5.4 Conclusions and outlook

In this chapter, the TDSE solver presented before in this thesis has been used to simulate the laser-field induced ionisation dynamics of all the electrons from an argon atom by means of a mean-field-theory, the time-dependent density-functional-theory. Ionisation probabilities have been presented for all the charge states of argon at the end of TDDFT simulations, for various laser pulses of high-frequency. It was found that the charge states are populated sequentially and that for the simulated field-parameters, even pulses of 100 optical-cycles cannot fully ionise the argon atom.

It would be interesting to see, as a further direction, what is the peak field-strength for which the atom is fully ionised in a simulation involving laser fields

similar to the ones from this chapter.

Perhaps even more beneficial for the community would be to change the target and compare and contrast the results obtained with the current ones, taking into account the increased (or decreased) number of electrons of the new atomic element illuminated by the laser-field.

On another note, the elephant in the room is the exchange-correlation functional type: surely the results obtained with a different such choice would be amendable to a detailed and thorough investigation, possibly in the view of the present (A)LDA-type results.

The basis functions could be changed from being fully numerical and simulations could be run faster given that the new elements are of the Gaussian type, where closed-form formulae exist for the product of such functions integrated across the real space. Still, the difficulty arising if one pursues this approach is the non-trivial effort to describe ionisation (so large-radial-distance quantum-dynamics) using such orbitals.

6
Conclusions and outlook

"Learn from yesterday, live for today, hope for tomorrow! The important thing is not to stop questioning!"

— Albert Einstein

Contents

6.1 Conclusions 180
6.2 Possible further work 182

6.1 Conclusions

This thesis started by presenting how atomic laser-field induced ionisation can potentially help inferring the peak intensity of an ultra-intense beam-pulse of light. The interaction scenario was quantitatively-modelled through state-of-the-art laser-field ionisation-rates and put into an experimental setup, considering how counting the detected ionised atoms of various charge states could reveal the beam's electrical-field peak-value. The numerical simulations have been carried out using SMILEI, a fully-relativistic particle-in-cell plasma-simulation code. I have added a new functionality (ionisation rate) to the codebase. This accurately

models the ionisation regimes of atoms across a very wide range of peak electric-field values: it is a piece-wise defined, general, field-ionisation rate, composed of three different rates. The workflow allows selecting the most optimum one from the three, at each timestep, based on the electric-field value at the location of the atom being ionised. The choice is always made by analysing how large is that value compared to some reference quantities for the atom and its charge state. This naturally allowed the extension into the barrier-suppression-ionisation (BSI) regime (rather large field strengths), but now also very accurately models the problematic transition between pure tunnelling and BSI, which usually warrants special attention, but many times is not treated and considered at all in PIC codes. With this state-of-the-art modelling of field-induced ionisation at hand, I ran numerical simulations of ultra-rarefied xenon gases illuminated by ultra-intense laser beam-pulses and focused, specifically, on the ionisation yields of each charge state.

Then the atomic laser-field induced ionisation has been analysed at its most fundamental level by numerical solutions to the time-dependent-Schrödinger-equation. Initially theoretical and numerical considerations about the equation have been presented and these led to a fully-distributed solver based on the generalized-pseudospectral-grids (GPS) numerical discretisation. This first time full-parallelisation has had its scalability and performance measured using dedicated tools and instruments and after I collected this data, I analysed and carefully thought of the obtained results. This solver was used to search for various potential differences between a hydrogenic-atom's electron quantum-dynamics as dictated by a dipole-approximated Hamiltonian and by a beyond-dipole-approximation Hamiltonian. These differences have been found, are non-trivial, and the two results have been compared and contrasted in full.

Finally, the multi-electron nature of a noble-gas irradiated by a laser pulse has been taken into account in the Schrödinger equation by means of a mean-field approach: the time-dependent-density-functional-theory formalism. As it

consisted of numerically solving coupled one-electron TDSE, one for each Kohn-Sham orbital which was propagated, this extension to multi-electron targets of the one-electron solver mentioned previously only needed to include a single new additional non-trivial part: the Hartree potential which the electrons feel. Results of simulations carried out in this setting, for an argon atom, have been presented and analysed in the context of atomic-ionisation due to pulses of light of very high-frequency. One such simulation takes a wall-clock-time which increases worse than linear with respect to the number of orbitals, this restricting the targets one can use when wanting to obtain solutions in reasonable amounts of compute-time. Similarly for the peak field-strengths and the laser pulse's temporal duration, any increase in these requires an associated upgrade for the grid resolution. Using a statistical approach to derive the charge states' populations, the TDDFT results now nicely predict the ionisation dynamics and experimental observables of interest: the ion yields.

6.2 Possible further work

I discuss below some possible further directions arising from the ideas previously presented in the thesis.

The peak-intensity atomic-diagnostic could be put on a Bayesian framework. That is, an inverse-problem setting shall be created and the aim to identify the shape of the focal spot becomes a statistical-inference effort: what are the most probable beam's characteristics which have created the experimental observables as they are seen by the detectors? It is not trivial how one defines the family of functions to be searched in, i.e. the functional form of the beam's electric-field, because in a real experiment, a beam can be very fragmented and non-uniform: it can have very intense peaks in rather random locations from across the focal volume. Also, if the distribution of beam's conditions results as (is) bi-modal (or multi-modal) the problem complicates and special attention has to be paid to this. The results can be narrowed down in terms of their associated errors by

constraining the experimental conditions based on prior-knowledge about the laser pulse's generation and its movement along the beamlines.

A search which could be worth investigating is when does the relativistic self-focusing actually become important? It could be useful if an error bound could be placed on the obtained results based on the amount of beam's modification (degradation) due to this plasma-effect. This would be a non-trivial endeavour, but then higher gas-densities can be used in practice, practically making the experiment less demanding in terms of the requested technical parameters. It will then be a trade-off between the accuracy of the obtained intensity-distribution and the ease of practical implementation and the researchers could decide on when to stop increasing (or decreasing) the gas density.

One could potentially insert "exact" laser-field ionisation rates in the SMILEI routines which calculate the probabilities of ionisation at each timestep. This means that fully-converged TDSE simulations have to be carrier out in advance and the field-ionisation rates to be extracted accurately from them. These numbers will then be fed as input-parameters to the SMILEI's solvers.

The last point above is very much dependent on being able to run many TDSE simulations at rather fine spatial and temporal discretisations. One possible way to go after increased performance for the TDSE solver is to make it efficiently use GPU's. As advanced statistical techniques become more and more prevalent nowadays, it could be interesting to make the solver to be able to propagate (operational) gradients backwards in order to make it "differentiable". This could be combined with the GPU usage, as most modern machine-learning/deep-learning (ML/DL) frameworks harness, behind the top-level wrappers on matrix-matrix multiplications, the compute-power of GPU's.

Also, surrogate modelling could be a rewarding path forward. It could be a means to approximate the numerical solutions obtained using the solver but in a much faster way. This could be achieved by training a ML/DL model on the real datasets obtained from the solver, then querying the model to make predictions at unseen input examples, rather than running a full TDSE simulation. As the

core characteristic of the GPS discretisation is that it employs non-uniform grids, convolutional-neural-networks (CNN's) which traditionally operate on equidistant pixels from an image are not seen to be a plug-and-play solution for a neural-network which learns the quantum electron dynamics. Continuous-convolutions exist [224] and can potentially provide the sought-after generalisation, although in a non-trivial manner. Still, the most flexible grid is, in general, a graph, so graph-neural-networks (GNN's) with their message-passing paradigm might be the way forward to capture the irregularities of the GPS (two-dimensional or three-dimensional) grid. This gains traction nowadays, for example in [225], the authors nicely and accurately modelled very complicated plasma-dynamics by means of a well-trained GNN regressor. With clever algorithmic designs and implementations of the information flowing through the graph, the rather high scaling of the network with respect to the number of graph nodes (vertices) can be made smaller and thus its training to become feasible and there has been recent effort into this direction [226].

It will be important to find a way to filter through the obtained predictions and see if they are correct: a (robust) means to quantify the uncertainty (UQ) in these predictions shall be in place.

The excitation and ionisation dynamics of atoms illuminated by a chirped pulse of light (for which the instantaneous frequency is deemed to be dependent on time) could be also analysed by means of numerical TDSE or time-dependent-density-functional-theory (TDDFT) simulations.

A different expansion direction could be to enable the solver to accept molecular targets (for example, two-centre systems, i.e. diatomic-molecules). This could be done using the prolate spheroidal coordinates, which can be discretised with the GPS method, in a manner similar to how this was done for the spherical coordinates and presented in this thesis. A different workflow shall be constructed as the scattering states for such quantum-systems can easily become cumbersome to be numerically-calculated in an accurate fashion. The response

of such molecules when irradiated with X-rays could be tested by means of the non-dipole Hamiltonian presented previously.

Similarly, the response of an excited state of a hydrogenic-atom could be of interest to be simulated numerically.

One could go on the long-wavelength (3400 nm, say) avenue and search for changes in the photoelectron-momentum-distributions due to the inclusion of non-dipole corrections to the usual dipole-approximated Hamiltonian. This long-wavelength regime would mean that the electron has an oscillation amplitude which is rather large, compared to the X-ray regime. Thus the simulation domain has to be sampled accurately also at large radial distances. A modification of the non-linear mapping of the radial coordinate could be useful to achieve this with a reasonable number of radial grid-points. In short, the density of the grid points across the radial axis shall be made more uniform. A step-like structure of this mapping could prove itself the most-optimum: finely-sampled next to the core, less finely-sampled at intermediate distances from the core, and constant, but rather loosely-sampled from there up until towards the infinity. Of course, a finite-differences grid (a family of Dirac delta-functions) has a constant grid-spacing across the radial coordinate and thus can easily sample the distances towards infinity in an efficient manner. It would be worth first trying to answer if such a grid can accurately capture the fast oscillations of the wavefunctions next to the core, at least in a one-electron setting.

A whole field of research is devoted to finding very accurate approximations to the exchange-correlation functional. Although far from trivial, it could be a path forward to test the numerical predictions obtained using a ML/DL-fitted such functional [180] placed inside the TDDFT solver. A hybrid-type exchange-correlation functional could produce predictions which might offer interesting insights into the ionisation process.

A different mean-field approach, the time-dependent-Hartree-Fock (TDHF) formalism, could be implemented and the numerical predictions obtained using

it be compared and contrasted with the TDDFT ones for X-rays illuminating a noble-gas atom.

References

[1] P. Agostini et al. “Free-Free Transitions Following Six-Photon Ionization of Xenon Atoms”. In: *Phys. Rev. Lett.* 42 (17 Apr. 1979), pp. 1127–1130. URL: https://link.aps.org/doi/10.1103/PhysRevLett.42.1127.

[2] S. Augst et al. “Laser Ionization of Noble Gases by Coulomb-barrier Suppression”. In: *JOSA B* 8.4 (Apr. 1991), pp. 858–867.

[3] M Ferray et al. “Multiple-harmonic conversion of 1064 nm radiation in rare gases”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 21.3 (Feb. 1988), p. L31. URL: https://dx.doi.org/10.1088/0953-4075/21/3/001.

[4] X. F. Li et al. “Multiple-harmonic generation in rare gases at high laser intensity”. In: *Phys. Rev. A* 39 (11 June 1989), pp. 5751–5761. URL: https://link.aps.org/doi/10.1103/PhysRevA.39.5751.

[5] G G Paulus et al. “Rescattering effects in above-threshold ionization: a classical model”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 27.21 (Nov. 1994), p. L703. URL: https://dx.doi.org/10.1088/0953-4075/27/21/003.

[6] Juan J. Carrera, X. M. Tong, and Shih-I Chu. “Creation and control of a single coherent attosecond xuv pulse by few-cycle intense laser pulses”. In: *Phys. Rev. A* 74 (2 Aug. 2006), p. 023404. URL: https://link.aps.org/doi/10.1103/PhysRevA.74.023404.

[7] Ferenc Krausz and Misha Ivanov. “Attosecond physics”. In: *Rev. Mod. Phys.* 81 (1 Feb. 2009), pp. 163–234. URL: https://link.aps.org/doi/10.1103/RevModPhys.81.163.

[8] C Blaga et al. “Imaging ultrafast molecular dynamics with laser-induced electron diffraction”. In: *Nature* 483 (2012), pp. 194–197.

[9] P. B. Corkum. “Plasma perspective on strong field multiphoton ionization”. In: *Phys. Rev. Lett.* 71 (13 Sept. 1993), pp. 1994–1997. URL: https://link.aps.org/doi/10.1103/PhysRevLett.71.1994.

[10] D. N. Fittinghoff et al. “Observation of nonsequential double ionization of helium with optical tunneling”. In: *Phys. Rev. Lett.* 69 (18 Nov. 1992), pp. 2642–2645. URL: https://link.aps.org/doi/10.1103/PhysRevLett.69.2642.

[11] B. Walker et al. “Precision Measurement of Strong Field Double Ionization of Helium”. In: *Phys. Rev. Lett.* 73 (9 Aug. 1994), pp. 1227–1230. URL: https://link.aps.org/doi/10.1103/PhysRevLett.73.1227.

[12] Andreas Becker and Farhad H M Faisal. “Mechanism of laser-induced double ionization of helium”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 29.6 (Mar. 1996), p. L197. URL: https://dx.doi.org/10.1088/0953-4075/29/6/005.

[13] Enam A. Chowdhury and Barry C. Walker. "Multielectron Ionization Processes in Ultrastrong Laser Fields". In: *Journal of the Optical Society of America B: Optical Physics* 20.1 (Jan. 2003), pp. 109–112.

[14] Enam A. Chowdhury, C. P. J. Barty, and Barry C. Walker. ""Nonrelativistic" Ionization of the L-shell States in Argon by a "Relativistic" 10^{19}W/Cm^2 Laser Field". In: *Physical Review A* 63.4 (Mar. 2001), p. 042712.

[15] A. Sagisaka et al. "Production of highly charged ions of argon by optical field ionization in a relativistic laser field". In: *JAERI-Conf 2001-011*. 2001.

[16] H Bachau et al. "Applications of B-splines in atomic and molecular physics". In: *Reports on Progress in Physics* 64.12 (Nov. 2001), p. 1815. URL: https://dx.doi.org/10.1088/0034-4885/64/12/205.

[17] Kane Yee. "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media". In: *IEEE Transactions on Antennas and Propagation* 14.3 (1966), pp. 302–307.

[18] J. Derouillat et al. "Smilei : A Collaborative, Open-Source, Multi-Purpose Particle-in-Cell Code for Plasma Simulation". In: *Computer Physics Communications* 222 (Jan. 2018), pp. 351–373.

[19] R. Nuter et al. "Field Ionization Model Implemented in Particle In Cell Code and Applied to Laser-Accelerated Carbon Ions". In: *Physics of Plasmas* 18.3 (Mar. 2011), p. 033107.

[20] I C E Turcu et al. "High Field Physics and QED Experiments at ELI-NP". In: (), p. 87.

[21] K Homma et al. "COMBINED LASER GAMMA EXPERIMENTS AT ELI-NP". In: (), p. 42.

[22] J.-P. Chambaret et al. "Extreme Light Infrastructure: Laser Architecture and Major Challenges". In: *Proc.SPIE*. Vol. 7721. May 2010.

[23] F. Lureau et al. "High-Energy Hybrid Femtosecond Laser System Demonstrating 2 × 10 PW Capability". In: *High Power Laser Science and Engineering* 8 (2020), e43.

[24] S. Weber et al. "P3: An Installation for High-Energy Density Plasma Physics and Ultra-High Intensity Laser–Matter Interaction at ELI-Beamlines". In: *Matter and Radiation at Extremes* 2.4 (July 2017), pp. 149–176.

[25] J. Nees et al. "ZEUS: A National Science Foundation Mid-Scale Facility for Laser-Driven Science in the QED Regime". In: Conference on Lasers and Electro-Optics. Optical Society of America, 2020, JW2B.9.

[26] Jin Woo Yoon et al. "Realization of laser intensity over 1023W/cm2". In: *Optica* 8.5 (May 2021), pp. 630–635. URL: https://opg.optica.org/optica/abstract.cfm?URI=optica-8-5-630.

[27] F. Mirani et al. "Integrated quantitative PIXE analysis and EDX spectroscopy using a laser-driven particle source". In: *Science Advances* 7.3 (2021), eabc8660. eprint: https://www.science.org/doi/pdf/10.1126/sciadv.abc8660. URL: https://www.science.org/doi/abs/10.1126/sciadv.abc8660.

[28] B. Le Garrec and the Apollon team. "Design Update and Recent Results of the Apollon 10 PW Facility". In: Conference on Lasers and Electro-Optics. Optical Society of America, 2017, SF1K.3.

[29] A. R. Bell and John G. Kirk. "Possibility of Prolific Pair Production with High-Power Lasers". In: *Physical Review Letters* 101.20 (Nov. 2008), p. 200403.

[30] C. P. Ridgers et al. "Dense Electron-Positron Plasmas and Ultraintense \gamma-Rays from Laser-Irradiated Solids". In: *Physical Review Letters* 108.16 (Apr. 2012), p. 165006.

[31] D.F. Gordon, B. Hafizi, and M.H. Helle. "Solution of relativistic quantum optics problems using clusters of graphical processing units". In: *Journal of Computational Physics* 267 (2014), pp. 50–62. URL: https://www.sciencedirect.com/science/article/pii/S002199911400148X.

[32] I. I. Artemenko and I. Yu. Kostyukov. "Ionization-Induced Laser-Driven QED Cascade in Noble Gases". In: *Physical Review A* 96.3 (Sept. 2017), p. 032106.

[33] G. Pariente et al. "Space–Time Characterization of Ultra-Intense Femtosecond Laser Beams". In: *Nature Photonics* 10.8 (Aug. 2016), pp. 547–553.

[34] S. Gales et al. "The Extreme Light Infrastructure—Nuclear Physics (ELI-NP) Facility: New Horizons in Physics with 10 PW Ultra-Intense Lasers and 20 MeV Brilliant Gamma Beams". In: (Jan. 2018).

[35] Rick Trebino et al. "Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating". In: *Review of Scientific Instruments* 68.9 (Sept. 1997), pp. 3277–3295. eprint: https://pubs.aip.org/aip/rsi/article-pdf/68/9/3277/8811505/3277_1_online.pdf. URL: https://doi.org/10.1063/1.1148286.

[36] C. Iaconis and I.A. Walmsley. "Self-referencing spectral interferometry for measuring ultrashort optical pulses". In: *IEEE Journal of Quantum Electronics* 35.4 (1999), pp. 501–509.

[37] Miguel Miranda et al. "All-optical measurement of the complete waveform of octave-spanning ultrashort light pulses". In: *Opt. Lett.* 44.2 (Jan. 2019), pp. 191–194. URL: https://opg.optica.org/ol/abstract.cfm?URI=ol-44-2-191.

[38] V. Yanovsky et al. "Ultra-High Intensity- 300-TW Laser at 0.1 Hz Repetition Rate." In: *Optics Express* 16.3 (Feb. 2008), pp. 2109–2114.

[39] Zhaoyang Li et al. "Degradation of Femtosecond Petawatt Laser Beams: Spatio-temporal/Spectral Coupling Induced by Wavefront Errors of Compression Gratings". In: *Applied Physics Express* 10.10 (Sept. 2017), p. 102702.

[40] M. Frigo and S.G. Johnson. "The Design and Implementation of FFTW3". In: *Proceedings of the IEEE* 93.2 (2005), pp. 216–231.

[41] E. S. Sarachik and G. T. Schappert. "Classical Theory of the Scattering of Intense Laser Radiation by Free Electrons". In: *Physical Review D: Particles and Fields* 1.10 (May 1970), pp. 2738–2753.

[42] Szu-yuan Chen, Anatoly Maksimchuk, and Donald Umstadter. "Experimental Observation of Relativistic Nonlinear Thomson Scattering". In: *Nature* 396.6712 (Dec. 1998), pp. 653–655.

[43] C. Z. He et al. “Towards an in Situ, Full-Power Gauge of the Focal-Volume Intensity of Petawatt-Class Lasers”. In: *Optics Express* 27.21 (Oct. 2019), pp. 30020–30030.

[44] Omri Har-Shemesh and Antonino Di Piazza. “Peak Intensity Measurement of Relativistic Lasers via Nonlinear Thomson Scattering”. In: *Optics Letters* 37.8 (Apr. 2012), pp. 1352–1354.

[45] Wenchao Yan et al. “High-Order Multiphoton Thomson Scattering”. In: *Nature Photonics* 11.8 (Aug. 2017), pp. 514–520.

[46] O. E. Vais, S. G. Bochkarev, and V. Yu. Bychenkov. “Nonlinear Thomson Scattering of a Relativistically Strong Tightly Focused Ultrashort Laser Pulse”. In: *Plasma Physics Reports* 42.9 (Sept. 2016), pp. 818–833.

[47] C. Z. He et al. “Towards an in situ, full-power gauge of the focal-volume intensity of petawatt-class lasers”. In: *Opt. Express* 27.21 (Oct. 2019), pp. 30020–30030. URL: `https://opg.optica.org/oe/abstract.cfm?URI=oe-27-21-30020`.

[48] Felix Mackenroth, Amol R Holkundkar, and Hans-Peter Schlenvoigt. “Ultra-Intense Laser Pulse Characterization Using Ponderomotive Electron Scattering”. In: *New Journal of Physics* 21.12 (Dec. 2019), p. 123028.

[49] O. E. Vais et al. “Angular Distribution of Electrons Directly Accelerated by an Intense Tightly Focused Laser Pulse”. In: *Quantum Electronics* 47.1 (Feb. 2017), p. 38.

[50] O. E. Vais et al. “Characterizing Extreme Laser Intensities by Ponderomotive Acceleration of Protons from Rarified Gas”. In: *New Journal of Physics* 22.2 (Feb. 2020), p. 023003.

[51] O. E. Vais and V. Yu Bychenkov. “Complementary Diagnostics of High-Intensity Femtosecond Laser Pulses via Vacuum Acceleration of Protons and Electrons”. In: *Plasma Physics and Controlled Fusion* 63.1 (Dec. 2020), p. 014002.

[52] K. Krushelnick et al. “Diagnosis of Peak Laser Intensity from High-Energy Ion Measurements during Intense Laser Interactions with Underdense Plasmas”. In: *Laser and Particle Beams* 18.4 (Oct. 2000), pp. 595–600.

[53] Anthony Link et al. “Development of an in Situ Peak Intensity Measurement Method for Ultraintense Single Shot Laser-Plasma Experiments at the Sandia Z Petawatt Facility”. In: *Review of Scientific Instruments* 77.10 (2006), 10E723. eprint: `https://doi.org/10.1063/1.2336469`.

[54] M. F. Ciappina et al. “Progress toward Atomic Diagnostics of Ultrahigh Laser Intensities”. In: *Physical Review A* 99.4 (Apr. 2019), p. 043405.

[55] M. F. Ciappina and S. V. Popruzhenko. “Diagnostics of Ultra-Intense Laser Pulses Using Tunneling Ionization”. In: *Laser Physics Letters* 17.2 (Jan. 2020), p. 025301.

[56] M. F. Ciappina, E. E. Peganov, and S. V. Popruzhenko. “Focal-Shape Effects on the Efficiency of the Tunnel-Ionization Probe for Extreme Laser Intensities”. In: *Matter and Radiation at Extremes* 5.4 (July 2020), p. 044401.

[57] L. V. Keldysh. “Ionization in the Field of a Strong Electromagnetic Wave”. In: *J. Exp. Theor. Phys.* 20.5 (1965), pp. 1307–1314.

[58] Vladimir S. Popov. “Tunnel and Multiphoton Ionization of Atoms and Ions in a Strong Laser Field (Keldysh Theory)”. In: *Physics-Uspekhi* 47.9 (Sept. 2004), p. 855.

[59] A M Perelomov, V S Popov, and M V Terent'Ev. "IONIZATION OF ATOMS IN AN ALTERNATING ELECTRIC FIELD". In: *JETP* 23.5 (1966).

[60] I. Yu Kostyukov and A. A. Golovanov. "Field Ionization in Short and Extremely Intense Laser Pulses". In: *Physical Review A* 98.4 (Oct. 2018), p. 043407. arXiv: 1808.06890.

[61] X. M. Tong and C. D. Lin. "Empirical Formula for Static Field Ionization Rates of Atoms and Molecules by Lasers in the Barrier-Suppression Regime". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 38.15 (July 2005), pp. 2593–2600.

[62] D. Bauer and P. Mulser. "Exact Field Ionization Rates in the Barrier-Suppression Regime from Numerical Time-Dependent Schrödinger-equation Calculations". In: *Physical Review A* 59.1 (Jan. 1999), pp. 569–577.

[63] A A Golovanov and I Yu Kostyukov. "Formula for the Ionisation Rate of an Atom or Ion in a Strong Electromagnetic Field for Numerical Simulation". In: *Quantum Electronics* 50.4 (Apr. 2020), pp. 350–353.

[64] W Han. *Elementary Numerical Analysis Course*. https://homepage.math.uiowa.edu/~whan.

[65] Michael Tipping. "Sparse Bayesian Learning and Relevance Vector Machine". In: *Journal of Machine Learning Research* 1 (Jan. 2001), pp. 211–244.

[66] Tianqi Chen and Carlos Guestrin. "XGBoost: A Scalable Tree Boosting System". In: Aug. 2016, pp. 785–794.

[67] Takuya Akiba et al. "Optuna: A next-Generation Hyperparameter Optimization Framework". In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 2623–2631.

[68] A. P. Dempster, N. M. Laird, and D. B. Rubin. "Maximum Likelihood from Incomplete Data via the EM Algorithm". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 39.1 (1977), pp. 1–38.

[69] Christopher M Bishop. *Pattern Recognition and Machine Learning*. New York : Springer, [2006] ©2006.

[70] F. Llorente et al. "Marginal Likelihood Computation for Model Selection and Hypothesis Testing: An Extensive Review". In: *SIAM Review* 65.1 (2023), pp. 3–58. eprint: https://doi.org/10.1137/20M1310849. URL: https://doi.org/10.1137/20M1310849.

[71] C. M. Pooley and G. Marion. "Bayesian model evidence as a practical alternative to deviance information criterion". In: *Royal Society Open Science* 5.3 (2018), p. 171519. eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rsos.171519. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsos.171519.

[72] H. Akaike. "A new look at the statistical model identification". In: *IEEE Transactions on Automatic Control* 19.6 (1974), pp. 716–723.

[73] Gideon Schwarz. "Estimating the Dimension of a Model". In: *The Annals of Statistics* 6.2 (1978), pp. 461–464. URL: https://doi.org/10.1214/aos/1176344136.

[74] E Fong and C C Holmes. "On the Marginal Likelihood and Cross-Validation". In: *Biometrika* 107.2 (Jan. 2020), pp. 489–496. eprint: https://academic.oup.com/biomet/article-pdf/107/2/489/33217997/asz077.pdf.

[75] Jean-Pierre Berenger. "A perfectly matched layer for the absorption of electromagnetic waves". In: *Journal of Computational Physics* 114.2 (1994), pp. 185–200. URL: https://www.sciencedirect.com/science/article/pii/S0021999184711594.

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

[77] B. Fryxell et al. "FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes". In: 131.1 (Nov. 2000), pp. 273–334.

[78] Ulf Zastrau et al. "The High Energy Density Scientific Instrument at the European XFEL". In: *Journal of Synchrotron Radiation* 28.5 (Sept. 2021), pp. 1393–1416. URL: https://doi.org/10.1107/S1600577521007335.

[79] H R Reiss. "The tunnelling model of laser-induced ionization and its failure at low frequencies". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 47.20 (Oct. 2014), p. 204006. URL: https://dx.doi.org/10.1088/0953-4075/47/20/204006.

[80] A. McPherson et al. "Studies of multiphoton production of vacuum-ultraviolet radiation in the rare gases". In: *J. Opt. Soc. Am. B* 4.4 (Apr. 1987), pp. 595–601. URL: https://opg.optica.org/josab/abstract.cfm?URI=josab-4-4-595.

[81] Pierre Agostini and Louis F DiMauro. "The physics of attosecond light pulses". In: *Reports on Progress in Physics* 67.6 (May 2004), p. 813. URL: https://dx.doi.org/10.1088/0034-4885/67/6/R01.

[82] F H M Faisal. "Collision of electrons with laser photons in a background potential". In: *Journal of Physics B: Atomic and Molecular Physics* 6.11 (Nov. 1973), p. L312. URL: https://dx.doi.org/10.1088/0022-3700/6/11/003.

[83] Howard R. Reiss. "Effect of an intense electromagnetic field on a weakly bound system". In: *Phys. Rev. A* 22 (5 Nov. 1980), pp. 1786–1813. URL: https://link.aps.org/doi/10.1103/PhysRevA.22.1786.

[84] M.D Feit, J.A Fleck, and A Steiger. "Solution of the Schrödinger equation by a spectral method". In: *Journal of Computational Physics* 47.3 (1982), pp. 412–433. URL: https://www.sciencedirect.com/science/article/pii/0021999182900912.

[85] Mark R. Hermann and J. A. Fleck. "Split-operator spectral method for solving the time-dependent Schrödinger equation in spherical coordinates". In: *Phys. Rev. A* 38 (12 Dec. 1988), pp. 6000–6012. URL: https://link.aps.org/doi/10.1103/PhysRevA.38.6000.

[86] André D. Bandrauk and Hai Shen. "Exponential split operator methods for solving coupled time-dependent Schrödinger equations". In: *The Journal of Chemical Physics* 99.2 (July 1993), pp. 1185–1193. eprint: https://pubs.aip.org/aip/jcp/article-pdf/99/2/1185/11042089/1185_1_online.pdf. URL: https://doi.org/10.1063/1.465362.

[87] J. T. Lin and T. F. Jiang. "Photodissociation of H_2^+ in intense chirped laser fields". In: *Phys. Rev. A* 63 (1 Dec. 2000), p. 013408. URL: https://link.aps.org/doi/10.1103/PhysRevA.63.013408.

[88] Hengtai Yu, Tao Zuo, and André D Bandrauk. "Intense field ionization of molecules with ultra-short laser pulses - enhanced ionization and barrier-suppression effects". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 31.7 (Apr. 1998), p. 1533. URL: https://dx.doi.org/10.1088/0953-4075/31/7/018.

[89] J. Crank and P. Nicolson. "A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 43.1 (1947), pp. 50–67.

[90] Morten Førre and Aleksander Skjerlie Simonsen. "Nondipole ionization dynamics in atoms induced by intense XUV laser fields". In: *Phys. Rev. A* 90 (5 Nov. 2014), p. 053411. URL: https://link.aps.org/doi/10.1103/PhysRevA.90.053411.

[91] M. Nurhuda and F. H. M. Faisal. "Numerical solution of time-dependent Schrödinger equation for multiphoton processes: A matrix iterative method". In: *Phys. Rev. A* 60 (4 Oct. 1999), pp. 3125–3133. URL: https://link.aps.org/doi/10.1103/PhysRevA.60.3125.

[92] Wei-Chao Jiang and Xiao-Qing Tian. "Efficient Split-Lanczos propagator for strong-field ionization of atoms". In: *Opt. Express* 25.22 (Oct. 2017), pp. 26832–26843. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-25-22-26832.

[93] H.G. Muller. "An efficient propagation scheme for the time-dependent Schrödinger equation in the velocity gauge". In: *Laser Physics* 9.1 (1999), pp. 138–148. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-0000289329&partnerID=40&md5=6d8fa294c41b9c0544549b71d273dc01.

[94] Serguei Patchkovskii and H.G. Muller. "Simple, accurate, and efficient implementation of 1-electron atomic time-dependent Schrödinger equation in spherical coordinates". In: *Computer Physics Communications* 199 (2016), pp. 153–169. URL: https://www.sciencedirect.com/science/article/pii/S001046551500394X.

[95] Dieter Bauer and Peter Koval. "QPROP: A Schrödinger-solver for intense laser–atom interaction". In: *Computer Physics Communications* 174.5 (2006), pp. 396–421. URL: https://www.sciencedirect.com/science/article/pii/S0010465505005825.

[96] Volker Mosert and Dieter Bauer. "Photoelectron spectra with QPROP and t-SURFF". In: *Computer Physics Communications* 207 (2016), pp. 452–463. URL: https://www.sciencedirect.com/science/article/pii/S0010465516301801.

[97] Vasily Tulsky and Dieter Bauer. "QPROP with faster calculation of photoelectron spectra". In: *Computer Physics Communications* 251 (2020), p. 107098. URL: https://www.sciencedirect.com/science/article/pii/S0010465519304035.

[98] Zhao-Han Zhang et al. "QPC-TDSE: A parallel TDSE solver for atoms and small molecules in strong lasers". In: *Computer Physics Communications* 290 (2023), p. 108787. URL: https://www.sciencedirect.com/science/article/pii/S0010465523001327.

[99] Xiao-Min Tong and Shih-I Chu. "Theoretical study of multiple high-order harmonic generation by intense ultrashort pulsed laser fields: A new generalized pseudospectral time-dependent method". In: *Chemical Physics* 217.2 (1997). Dynamics of Driven Quantum Systems, pp. 119–130. URL: https://www.sciencedirect.com/science/article/pii/S0301010497000633.

[100] Xiao-Min Tong. "A three-dimensional time-dependent Schrödinger equation solver: an application to hydrogen atoms in an elliptical laser field". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 50.14 (June 2017), p. 144004. URL: https://dx.doi.org/10.1088/1361-6455/aa77a0.

[101] Jingyan Wang, Shih-I Chu, and Cecil Laughlin. "Multiphoton detachment of H^-. II. Intensity-dependent photodetachment rates and threshold behavior—complex-scaling generalized pseudospectral method". In: *Phys. Rev. A* 50 (4 Oct. 1994), pp. 3208–3215. URL: https://link.aps.org/doi/10.1103/PhysRevA.50.3208.

[102] Guanhua Yao and Shih-I Chu. "Generalized pseudospectral methods with mappings for bound and resonance state problems". In: *Chemical Physics Letters* 204.3 (1993), pp. 381–388. URL: https://www.sciencedirect.com/science/article/pii/000926149390025V.

[103] Dmitry A. Telnov and Shih-I Chu. "Multiphoton detachment of H^- near the one-photon threshold: Exterior complex-scaling–generalized pseudospectral method for complex quasienergy resonances". In: *Phys. Rev. A* 59 (4 Apr. 1999), pp. 2864–2874. URL: https://link.aps.org/doi/10.1103/PhysRevA.59.2864.

[104] Dmitry A. Telnov et al. "Exterior complex scaling method in time-dependent density-functional theory: Multiphoton ionization and high-order-harmonic generation of Ar atoms". In: *Phys. Rev. A* 87 (5 May 2013), p. 053406. URL: https://link.aps.org/doi/10.1103/PhysRevA.87.053406.

[105] Dmitry A. Telnov and Shih-I Chu. "Above-threshold-ionization spectra from the core region of a time-dependent wave packet: An ab initio time-dependent approach". In: *Phys. Rev. A* 79 (4 Apr. 2009), p. 043421. URL: https://link.aps.org/doi/10.1103/PhysRevA.79.043421.

[106] D. A. Tumakov et al. "Photoelectron spectra after multiphoton ionization of Li atoms in the one-photon Rabi-flopping regime". In: *Phys. Rev. A* 100 (2 Aug. 2019), p. 023407. URL: https://link.aps.org/doi/10.1103/PhysRevA.100.023407.

[107] Dmitry A. Telnov and Shih-I Chu. "Minima in low-energy above-threshold-ionization spectra induced by electronic structure". In: *Phys. Rev. A* 100 (4 Oct. 2019), p. 043423. URL: https://link.aps.org/doi/10.1103/PhysRevA.100.043423.

[108] Mitsuko Murakami and Shih-I Chu. "Photoelectron momentum distributions of the hydrogen atom driven by multicycle elliptically polarized laser pulses". In: *Phys. Rev. A* 93 (2 Feb. 2016), p. 023425. URL: https://link.aps.org/doi/10.1103/PhysRevA.93.023425.

[109] Mitsuko Murakami, Oleg Korobkin, and Marko Horbatsch. "High-harmonic generation from hydrogen atoms driven by two-color mutually orthogonal laser fields". In: *Phys. Rev. A* 88 (6 Dec. 2013), p. 063419. URL: https://link.aps.org/doi/10.1103/PhysRevA.88.063419.

[110] Mitsuko Murakami et al. "Photoelectron momentum distribution of ground- and excited-state lithium atoms induced by extreme-ultraviolet photon absorption". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 51.17 (Aug. 2018), p. 175002. URL: https://dx.doi.org/10.1088/1361-6455/aad1ad.

[111] Dmitry A. Telnov and Shih-I Chu. "Ab initio study of high-order harmonic generation of H_2^+ in intense laser fields: Time-dependent non-Hermitian Floquet approach". In: *Phys. Rev. A* 71 (1 Jan. 2005), p. 013408. URL: https://link.aps.org/doi/10.1103/PhysRevA.71.013408.

[112] Dmitry A. Telnov and Shih-I Chu. "Ab initio study of the orientation effects in multiphoton ionization and high-order harmonic generation from the ground and excited electronic states of H_2^+". In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 043412. URL: https://link.aps.org/doi/10.1103/PhysRevA.76.043412.

[113] Dmitry A. Telnov and Shih-I Chu. "Effects of multiple electronic shells on strong-field multiphoton ionization and high-order harmonic generation of diatomic molecules with arbitrary orientation: An all-electron time-dependent density-functional approach". In: *Phys. Rev. A* 80 (4 Oct. 2009), p. 043412. URL: https://link.aps.org/doi/10.1103/PhysRevA.80.043412.

[114] K Nasiri Avanaki, Dmitry A Telnov, and Shih-I Chu. "Exploration of the origin of anomalous dependence for near-threshold harmonics in H_2^+ on the ellipticity of driving laser fields". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 49.11 (May 2016), p. 114002. URL: https://dx.doi.org/10.1088/0953-4075/49/11/114002.

[115] Dmitry A. Telnov, John Heslar, and Shih-I Chu. "Effect of nuclear vibration on high-order-harmonic generation of aligned H_2^+ molecules". In: *Phys. Rev. A* 90 (6 Dec. 2014), p. 063412. URL: https://link.aps.org/doi/10.1103/PhysRevA.90.063412.

[116] Lin Zhu et al. "Revisiting the generalized pseudospectral method: Radial expectation values, fine structure, and hyperfine splitting of confined atom". In: *International Journal of Quantum Chemistry* 120.14 (2020), e26245.

[117] Thomas Gruber et al. *LIKWID*. Version v5.3.0. Nov. 2023. URL: https://doi.org/10.5281/zenodo.10105559.

[118] Andreas Knüpfer et al. "Score-P: A Joint Performance Measurement Run-Time Infrastructure for Periscope,Scalasca, TAU, and Vampir". In: *Tools for High Performance Computing 2011*. Ed. by Holger Brunst et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 79–91.

[119] E. Angerson et al. "LAPACK: A portable linear algebra library for high-performance computers". In: *Supercomputing '90:Proceedings of the 1990 ACM/IEEE Conference on Supercomputing*. 1990, pp. 2–11.

[120] Dmitry A Krapivin and Dmitry A Telnov. "Anomalous dependence of ionization probability and electron angular distributions on orientation of molecular axis in photoionization of H_2^+: effect of two-center interference". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 54.20 (Nov. 2021), p. 205601. URL: https://dx.doi.org/10.1088/1361-6455/ac34de.

[121] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. 1st ed. New York: Dover, 1972.

[122] R.G. Newton. *Scattering Theory of Waves and Particles*. International series in pure and applied physics. McGraw-Hill, 1966. URL: https://books.google.co.uk/books?id=QXzAk-GrN-kC.

[123] B. Feti ć, W. Becker, and D. B. Milo ševi ć. "Extracting photoelectron spectra from the time-dependent wave function: Comparison of the projection onto continuum states and window-operator methods". In: *Phys. Rev. A* 102 (2 Aug. 2020), p. 023101. URL: https://link.aps.org/doi/10.1103/PhysRevA.102.023101.

[124] X. M. Tong, K. Hino, and N. Toshima. "Phase-dependent atomic ionization in few-cycle intense laser fields". In: *Phys. Rev. A* 74 (3 Sept. 2006), p. 031405. URL: https://link.aps.org/doi/10.1103/PhysRevA.74.031405.

[125] Liang-You Peng and Qihuang Gong. "An accurate Fortran code for computing hydrogenic continuum wave functions at a wide range of parameters". In: *Computer Physics Communications* 181.12 (2010), pp. 2098–2101. URL: https://www.sciencedirect.com/science/article/pii/S0010465510003413.

[126] Caio Ciardelli et al. "SphGLLTools: A toolbox for visualization of large seismic model files based on 3D spectral-element meshes". In: *Computers Geosciences* 159 (2022), p. 105007. URL: https://www.sciencedirect.com/science/article/pii/S0098300421002909.

[127] Nathanael Schaeffer. "Efficient spherical harmonic transforms aimed at pseudospectral numerical simulations". In: *Geochemistry, Geophysics, Geosystems* 14.3 (2013), pp. 751–758.

[128] Keiichi Ishioka. "A New Recurrence Formula for Efficient Computation of Spherical Harmonic Transform". In: *Journal of the Meteorological Society of Japan* 96.2 (2018), pp. 241–249.

[129] P. Eckle et al. "Attosecond angular streaking". In: *Nature Physics* 4 (2008), pp. 565–570. URL: https://www.nature.com/articles/nphys982.

[130] Evelyn Namugwanya et al. *Collective-Optimized FFTs*. 2023. arXiv: 2306.16589 [cs.MS].

[131] Rahul Kumar, Amith Mamidala, and D. K. Panda. "Scaling alltoall collective on multi-core systems". In: *2008 IEEE International Symposium on Parallel and Distributed Processing*. 2008, pp. 1–8.

[132] Chen-Chun Chen et al. "Highly Efficient Alltoall and Alltoallv Communication Algorithms for GPU Systems". In: *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 2022, pp. 24–33.

[133] Michael Hofmann and Gudula Rünger. "An In-Place Algorithm for Irregular All-to-All Communication with Limited Memory". In: *Recent Advances in the Message Passing Interface*. Ed. by Rainer Keller et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 113–121.

[134] Cong Xu et al. "SLOAVx: Scalable LOgarithmic AlltoallV Algorithm for Hierarchical Multicore Systems". In: *2013 13th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing*. 2013, pp. 369–376.

[135] Qinghua Zhou et al. "Accelerating MPI All-to-All Communication with Online Compression on Modern GPU Clusters". In: *High Performance Computing*. Ed. by Ana-Lucia Varbanescu et al. Cham: Springer International Publishing, 2022, pp. 3–25.

[136] Weile Jia, Lin-Wang Wang, and Lin Lin. "Parallel transport time-dependent density functional theory calculations with hybrid functional on summit". In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. SC '19. Denver, Colorado: Association for Computing Machinery, 2019. URL: https://doi.org/10.1145/3295500.3356144.

[137] Samuel Williams, Andrew Waterman, and David Patterson. "Roofline: an insightful visual performance model for multicore architectures". In: *Commun. ACM* 52.4 (Apr. 2009), pp. 65–76. URL: https://doi.org/10.1145/1498765.1498785.

[138] Georg Ofenbeck et al. "Applying the roofline model". In: *2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 2014, pp. 76–85.

[139] Daniel Barry et al. "Memory Traffic and Complete Application Profiling with PAPI Multi-Component Measurements". In: *2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 2023, pp. 393–402.

[140] K Pande et al. "Femtosecond structural dynamics drives the trans/cis isomerization in photoactive yellow protein". In: *Science* 352.6286 (2016), pp. 725–729. eprint: https://www.science.org/doi/pdf/10.1126/science.aad5081. URL: https://www.science.org/doi/abs/10.1126/science.aad5081.

[141] B. McNeil and N. Thompson. "X-ray free-electron lasers". In: *Nature Photon* 4 (2010), pp. 814–821. URL: https://www.nature.com/articles/nphoton.2010.239#citeas.

[142] T. Tschentscher and R. Feidenhans'l. "Starting User Operation at the European XFEL". In: *Synchrotron Radiation News* 30.6 (2017), pp. 21–28. eprint: https://doi.org/10.1080/08940886.2017.1386995. URL: https://doi.org/10.1080/08940886.2017.1386995.

[143] Bart Faatz et al. "The FLASH Facility: Advanced Options for FLASH2 and Future Perspectives". In: *Applied Sciences* 7.11 (2017). URL: https://www.mdpi.com/2076-3417/7/11/1114.

[144] E. Allaria et al. *FERMI 2.0, Conceptual Design Report*.

[145] R. W. Schoenlein et al. "The Linac Coherent Light S ource: Recent Developments and Future Plans". In: *Applied Sciences* 7.8 (2017). URL: https://www.mdpi.com/2076-3417/7/8/850.

[146] Anne L'Huillier and Ph. Balcou. "High-order harmonic generation in rare gases with a 1-ps 1053-nm laser". In: *Phys. Rev. Lett.* 70 (6 Feb. 1993), pp. 774–777. URL: https://link.aps.org/doi/10.1103/PhysRevLett.70.774.

[147] M. Hervé et al. "On-the-fly investigation of XUV excited large molecular ions using a high harmonic generation light source". In: *Sci Rep* 12 (13191 2022). URL: https://www.nature.com/articles/s41598-022-17416-4#citeas.

[148] P. A. Norreys et al. "Efficient Extreme UV Harmonics Generated from Picosecond Laser Pulse Interactions with Solid Targets". In: *Phys. Rev. Lett.* 76 (11 Mar. 1996), pp. 1832–1835. URL: https://link.aps.org/doi/10.1103/PhysRevLett.76.1832.

[149] B. Dromey et al. "High harmonic generation in the relativistic limit". In: *Nat. Phys.* 2 (2006), pp. 456–459. URL: https://www.nature.com/articles/nphys338.

[150] S. Ghimire and D. A. Reis. "High-harmonic generation from solids". In: *Nat. Phys.* 15 (2019), pp. 10–16. URL: https://www.nature.com/articles/s41567-018-0315-5.

[151] M. Førre et al. "Nondipole Ionization Dynamics of Atoms in Superintense High-Frequency Attosecond Pulses". In: *Phys. Rev. Lett.* 97 (4 July 2006), p. 043601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.97.043601.

[152] Aleksander Skjerlie Simonsen and Morten Førre. "Magnetic-field-induced enhancement of atomic stabilization in intense high-frequency laser fields". In: *Phys. Rev. A* 92 (1 July 2015), p. 013405. URL: https://link.aps.org/doi/10.1103/PhysRevA.92.013405.

[153] Morten Førre and Aleksander Skjerlie Simonsen. "Generalized velocity-gauge form of the light-matter interaction Hamiltonian beyond the dipole approximation". In: *Phys. Rev. A* 93 (1 Jan. 2016), p. 013423. URL: https://link.aps.org/doi/10.1103/PhysRevA.93.013423.

[154] Morten Førre and Sølve Selstø. "Schrödinger formulation of the nondipole light-matter interaction consistent with relativity". In: *Phys. Rev. A* 101 (6 June 2020), p. 063416. URL: https://link.aps.org/doi/10.1103/PhysRevA.101.063416.

[155] Zhongyuan Zhou and Shih-I Chu. "Multiphoton above-threshold ionization in superintense free-electron x-ray laser fields: Beyond the dipole approximation". In: *Phys. Rev. A* 87 (2 Feb. 2013), p. 023407. URL: https://link.aps.org/doi/10.1103/PhysRevA.87.023407.

[156] H. R. Reiss. "Dipole-approximation magnetic fields in strong laser beams". In: *Phys. Rev. A* 63 (1 Dec. 2000), p. 013409. URL: https://link.aps.org/doi/10.1103/PhysRevA.63.013409.

[157] C. T. L. Smeenk et al. "Partitioning of the Linear Photon Momentum in Multiphoton Ionization". In: *Phys. Rev. Lett.* 106 (19 May 2011), p. 193002. URL: https://link.aps.org/doi/10.1103/PhysRevLett.106.193002.

[158] A. Ludwig et al. "Breakdown of the Dipole Approximation in Strong-Field Ionization". In: *Phys. Rev. Lett.* 113 (24 Dec. 2014), p. 243001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.113.243001.

[159] J. Maurer et al. "Probing the ionization wave packet and recollision dynamics with an elliptically polarized strong laser field in the nondipole regime". In: *Phys. Rev. A* 97 (1 Jan. 2018), p. 013404. URL: https://link.aps.org/doi/10.1103/PhysRevA.97.013404.

[160] Nida Haram et al. "Relativistic Nondipole Effects in Strong-Field Atomic Ionization at Moderate Intensities". In: *Phys. Rev. Lett.* 123 (9 Aug. 2019), p. 093201. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.093201.

[161] A. Hartung et al. "Magnetic fields alter strong-field ionization". In: *Nat. Phys.* 15 (2019), pp. 1222–1226. URL: https://www.nature.com/articles/s41567-019-0653-y.

[162] Szczepan Chelkowski and André D. Bandrauk. "Photon-momentum transfer in molecular photoionization". In: *Phys. Rev. A* 97 (5 May 2018), p. 053401. URL: https://link.aps.org/doi/10.1103/PhysRevA.97.053401.

[163] Simon Brennecke and Manfred Lein. "High-order above-threshold ionization beyond the electric dipole approximation: Dependence on the atomic and molecular structure". In: *Phys. Rev. A* 98 (6 Dec. 2018), p. 063414. URL: https://link.aps.org/doi/10.1103/PhysRevA.98.063414.

[164] R. Mishra, D. Kalita, and A. Gupta. "Breakdown of dipole approximation and its effect on high harmonic generation". In: *Eur. Phys. J. D* 66 (169 2012). URL: https://link.springer.com/article/10.1140/epjd/e2012-20750-0#citeas.

[165] Simon Vendelbo Bylling Jensen, Mads Middelhede Lund, and Lars Bojer Madsen. "Nondipole strong-field-approximation Hamiltonian". In: *Phys. Rev. A* 101 (4 Apr. 2020), p. 043408. URL: https://link.aps.org/doi/10.1103/PhysRevA.101.043408.

[166] Tor Kjellsson, Sølve Selstø, and Eva Lindroth. "Relativistic ionization dynamics for a hydrogen atom exposed to superintense XUV laser pulses". In: *Phys. Rev. A* 95 (4 Apr. 2017), p. 043403. URL: https://link.aps.org/doi/10.1103/PhysRevA.95.043403.

[167] Mihai Gavrila. "Atomic stabilization in superintense laser fields". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 35.18 (Sept. 2002), R147. URL: https://dx.doi.org/10.1088/0953-4075/35/18/201.

[168] M. Dondera, H. G. Muller, and M. Gavrila. "Observability of the dynamic stabilization of ground-state hydrogen with superintense femtosecond laser pulses". In: *Phys. Rev. A* 65 (3 Feb. 2002), p. 031405. URL: https://link.aps.org/doi/10.1103/PhysRevA.65.031405.

[169] C. Blaga et al. "Strong-field photoionization revisited". In: *Nature Phys* 5 (2009), pp. 335–338. URL: https://www.nature.com/articles/nphys1228.

[170] A M Popov, O V Tikhonova, and E A Volkova. "Strong-field atomic stabilization: numerical simulation and analytical modelling". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 36.10 (Apr. 2003), R125. URL: https://dx.doi.org/10.1088/0953-4075/36/10/201.

[171] Dmitry A. Telnov and Shih-I Chu. "Relativistic ionization probabilities and photoelectron distributions of hydrogenlike ions in superstrong electromagnetic fields". In: *Phys. Rev. A* 104 (2 Aug. 2021), p. 023111. URL: https://link.aps.org/doi/10.1103/PhysRevA.104.023111.

[172] C. Y. Wu et al. "Characteristic Spectrum of Very Low-Energy Photoelectron from Above-Threshold Ionization in the Tunneling Regime". In: *Phys. Rev. Lett.* 109 (4 July 2012), p. 043001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.109.043001.

[173] Diego G. Arbó, Emil Persson, and Joachim Burgdörfer. "Time double-slit interferences in strong-field tunneling ionization". In: *Phys. Rev. A* 74 (6 Dec. 2006), p. 063407. URL: https://link.aps.org/doi/10.1103/PhysRevA.74.063407.

[174] M. Möller et al. "Off-axis low-energy structures in above-threshold ionization". In: *Phys. Rev. A* 90 (2 Aug. 2014), p. 023412. URL: https://link.aps.org/doi/10.1103/PhysRevA.90.023412.

[175] E Cormier and P Lambropoulos. "Optimal gauge and gauge invariance in non-perturbative time-dependent calculation of above-threshold ionization". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 29.9 (May 1996), p. 1667. URL: https://dx.doi.org/10.1088/0953-4075/29/9/013.

[176] Kenneth C. Kulander. "Time-dependent Hartree-Fock theory of multiphoton ionization: Helium". In: *Phys. Rev. A* 36 (6 Sept. 1987), pp. 2726–2738. URL: https://link.aps.org/doi/10.1103/PhysRevA.36.2726.

[177] Jürgen Zanghellini et al. "An MCTDHF approach to multi–electron dynamics in laser fields". In: *Laser Physics* 13 (Aug. 2003), pp. 1064–1068.

[178] P. Hohenberg and W. Kohn. "Inhomogeneous Electron Gas". In: *Phys. Rev.* 136 (3B Nov. 1964), B864–B871. URL: https://link.aps.org/doi/10.1103/PhysRev.136.B864.

[179] W. Kohn. "v-Representability and Density Functional Theory". In: *Phys. Rev. Lett.* 51 (17 Oct. 1983), pp. 1596–1598. URL: https://link.aps.org/doi/10.1103/PhysRevLett.51.1596.

[180] M. F. Kasim and S. M. Vinko. "Learning the Exchange-Correlation Functional from Nature with Fully Differentiable Density Functional Theory". In: *Phys. Rev. Lett.* 127 (12 Sept. 2021), p. 126403. URL: https://link.aps.org/doi/10.1103/PhysRevLett.127.126403.

[181] Muhammad F. Kasim, Susi Lehtola, and Sam M. Vinko. "DQC: A Python program package for differentiable quantum chemistry". In: *The Journal of Chemical Physics* 156.8 (Feb. 2022), p. 084801. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0076202/18277993/084801_1_5.0076202.pdf. URL: https://doi.org/10.1063/5.0076202.

[182] Chuin Wei Tan, Chris J. Pickard, and William C. Witt. "Automatic differentiation for orbital-free density functional theory". In: *The Journal of Chemical Physics* 158.12 (Mar. 2023), p. 124801. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0138429/18281199/124801_1_5.0138429.pdf. URL: https://doi.org/10.1063/5.0138429.

[183] Kyle Bystrom and Boris Kozinsky. "CIDER: An Expressive, Nonlocal Feature Set for Machine Learning Density Functionals with Exact Constraints". In: *Journal of Chemical Theory and Computation* 18.4 (2022). PMID: 35235322, pp. 2180–2192. eprint: https://doi.org/10.1021/acs.jctc.1c00904. URL: https://doi.org/10.1021/acs.jctc.1c00904.

[184] Etienne Cuierrier et al. "The fourth-order expansion of the exchange hole and neural networks to construct exchange–correlation functionals". In: *The Journal of Chemical Physics* 157.17 (Nov. 2022), p. 171103. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0122761/16552263/171103_1_online.pdf. URL: https://doi.org/10.1063/5.0122761.

[185] R. Nagai, R. Akashi, and O. Sugino. "Completing density functional theory by machine learning hidden messages from molecules". In: *npj Comput Mater* 6 (43 2020). URL: https://www.nature.com/articles/s41524-020-0310-0.

[186] Karim K. Alaa El-Din et al. "STEP: extraction of underlying physics with robust machine learning". In: *Royal Society Open Science* 11.6 (2024), p. 231374. eprint: https://royalsocietypublishing.org/doi/pdf/10.1098/rsos.231374. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rsos.231374.

[187] P. A. M. Dirac. "Note on Exchange Phenomena in the Thomas Atom". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 26.3 (1930), pp. 376–385.

[188] NM Harrison. "An introduction to density functional theory". In: *Nato Science Series Sub Series III Computer and Systems Sciences* 187 (2003), pp. 45–70.

[189] J. P. Perdew and Alex Zunger. "Self-interaction correction to density-functional approximations for many-electron systems". In: *Phys. Rev. B* 23 (10 May 1981), pp. 5048–5079. URL: https://link.aps.org/doi/10.1103/PhysRevB.23.5048.

[190] James D. Talman and William F. Shadwick. "Optimized effective atomic central potential". In: *Phys. Rev. A* 14 (1 July 1976), pp. 36–40. URL: https://link.aps.org/doi/10.1103/PhysRevA.14.36.

[191] J. B. Krieger, Yan Li, and G. J. Iafrate. "Systematic approximations to the optimized effective potential: Application to orbital-density-functional theory". In: *Phys. Rev. A* 46 (9 Nov. 1992), pp. 5453–5458. URL: https://link.aps.org/doi/10.1103/PhysRevA.46.5453.

[192] J. B. Krieger, Yan Li, and G. J. Iafrate. "Construction and application of an accurate local spin-polarized Kohn-Sham potential with integer discontinuity: Exchange-only theory". In: *Phys. Rev. A* 45 (1 Jan. 1992), pp. 101–126. URL: https://link.aps.org/doi/10.1103/PhysRevA.45.101.

[193] Xiao-Min Tong and Shih-I Chu. "Density-functional theory with optimized effective potential and self-interaction correction for ground states and autoionizing resonances". In: *Phys. Rev. A* 55 (5 May 1997), pp. 3406–3416. URL: https://link.aps.org/doi/10.1103/PhysRevA.55.3406.

[194] V. A. Yerokhin, A. N. Artemyev, and V. M. Shabaev. "QED treatment of electron correlation in Li-like ions". In: *Phys. Rev. A* 75 (6 June 2007), p. 062501. URL: https://link.aps.org/doi/10.1103/PhysRevA.75.062501.

[195] Lucas O. Wagner, Zeng-hui Yang, and Kieron Burke. "Exact Conditions and Their Relevance in TDDFT". In: *Fundamentals of Time-Dependent Density Functional Theory*. Ed. by Miguel A.L. Marques et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 101–123. URL: https://doi.org/10.1007/978-3-642-23518-4_5.

[196] John P. Perdew and Karla Schmidt. "Jacob's ladder of density functional approximations for the exchange-correlation energy". In: *AIP Conference Proceedings* 577.1 (July 2001), pp. 1–20. eprint: https://pubs.aip.org/aip/acp/article-pdf/577/1/1/12108089/1_1_online.pdf. URL: https://doi.org/10.1063/1.1390175.

[197] A. D. Becke. "Density-functional exchange-energy approximation with correct asymptotic behavior". In: *Phys. Rev. A* 38 (6 Sept. 1988), pp. 3098–3100. URL: https://link.aps.org/doi/10.1103/PhysRevA.38.3098.

[198] Chengteh Lee, Weitao Yang, and Robert G. Parr. "Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density". In: *Phys. Rev. B* 37 (2 Jan. 1988), pp. 785–789. URL: https://link.aps.org/doi/10.1103/PhysRevB.37.785.

[199] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. "Generalized Gradient Approximation Made Simple". In: *Phys. Rev. Lett.* 77 (18 Oct. 1996), pp. 3865–3868. URL: https://link.aps.org/doi/10.1103/PhysRevLett.77.3865.

[200] John P. Perdew and Yue Wang. "Accurate and simple analytic representation of the electron-gas correlation energy". In: *Phys. Rev. B* 45 (23 June 1992), pp. 13244–13249. URL: https://link.aps.org/doi/10.1103/PhysRevB.45.13244.

[201] John P. Perdew et al. "Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation". In: *Phys. Rev. B* 46 (11 Sept. 1992), pp. 6671–6687. URL: https://link.aps.org/doi/10.1103/PhysRevB.46.6671.

[202] Carlo Adamo and Vincenzo Barone. "Toward reliable density functional methods without adjustable parameters: The PBE0 model". In: *The Journal of Chemical Physics* 110.13 (Apr. 1999), pp. 6158–6170. eprint: https://pubs.aip.org/aip/jcp/article-pdf/110/13/6158/19068890/6158_1_online.pdf. URL: https://doi.org/10.1063/1.478522.

[203] Axel D. Becke. "A new mixing of Hartree–Fock and local density-functional theories". In: *The Journal of Chemical Physics* 98.2 (Jan. 1993), pp. 1372–1377. eprint: https://pubs.aip.org/aip/jcp/article-pdf/98/2/1372/19125465/1372_1_online.pdf. URL: https://doi.org/10.1063/1.464304.

[204] D.R Bowler and M.J Gillan. "An efficient and robust technique for achieving self consistency in electronic structure calculations". In: *Chemical Physics Letters* 325.4 (2000), pp. 473–476. URL: https://www.sciencedirect.com/science/article/pii/S0009261400007508.

[205] G. P. Kerker. "Efficient iteration scheme for self-consistent pseudopotential calculations". In: *Phys. Rev. B* 23 (6 Mar. 1981), pp. 3082–3084. URL: https://link.aps.org/doi/10.1103/PhysRevB.23.3082.

[206] Kenneth D. Sevier. "Atomic electron binding energies". In: *Atomic Data and Nuclear Data Tables* 24.4 (1979), pp. 323–371. URL: https://www.sciencedirect.com/science/article/pii/0092640X79900123.

[207] Erich Runge and E. K. U. Gross. "Density-Functional Theory for Time-Dependent Systems". In: *Phys. Rev. Lett.* 52 (12 Mar. 1984), pp. 997–1000. URL: https://link.aps.org/doi/10.1103/PhysRevLett.52.997.

[208] David A. Strubbe et al. "Response Functions in TDDFT: Concepts and Implementation". In: *Fundamentals of Time-Dependent Density Functional Theory*. Ed. by Miguel A.L. Marques et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 139–166. URL: https://doi.org/10.1007/978-3-642-23518-4_7.

[209] Mitsuko Murakami, G. P. Zhang, and Shih-I Chu. "Multielectron effects in the photoelectron momentum distribution of noble-gas atoms driven by visible-to-infrared-frequency laser pulses: A time-dependent density-functional-theory approach". In: *Phys. Rev. A* 95 (5 May 2017), p. 053419. URL: https://link.aps.org/doi/10.1103/PhysRevA.95.053419.

[210] Eberhard K. U. Gross and Neepa T. Maitra. "Introduction to TDDFT". In: *Fundamentals of Time-Dependent Density Functional Theory*. Ed. by Miguel A.L. Marques et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 53–99. URL: https://doi.org/10.1007/978-3-642-23518-4_4.

[211] M. Thiele, E. K. U. Gross, and S. Kümmel. "Adiabatic Approximation in Nonperturbative Time-Dependent Density-Functional Theory". In: *Phys. Rev. Lett.* 100 (15 Apr. 2008), p. 153004. URL: https://link.aps.org/doi/10.1103/PhysRevLett.100.153004.

[212] A. Crawford-Uranga et al. "Time-dependent density-functional theory of strong-field ionization of atoms by soft x rays". In: *Phys. Rev. A* 90 (3 Sept. 2014), p. 033412. URL: https://link.aps.org/doi/10.1103/PhysRevA.90.033412.

[213] C.A. Ullrich. "Time-dependent Kohn–Sham approach to multiple ionization". In: *Journal of Molecular Structure: THEOCHEM* 501-502 (2000), pp. 315–325. URL: https://www.sciencedirect.com/science/article/pii/S016612809900442X.

[214] A. Kramida et al. NIST Atomic Spectra Database (ver. 5.11), [Online]. Available: https://physics.nist.gov/asd [2024, August 7]. National Institute of Standards and Technology, Gaithersburg, MD. 2023.

[215] Skattum Sigve Bøe. "TIME EVOLUTION IN QUANTUM DOTS Using the Multiconfiguration Time-Dependent Hartree-Fock Method". In: (2013). URL: https://core.ac.uk/download/pdf/30814239.pdf.

[216] Xiaosong Li et al. "A time-dependent Hartree–Fock approach for studying the electronic optical response of molecules in intense fields". In: *Phys. Chem. Chem. Phys.* 7 (2 2005), pp. 233–239. URL: http://dx.doi.org/10.1039/B415849K.

[217] Jiří Čížek. "On the use of the cluster expansion and the technique of diagrams in calculations of correlation effects in atoms and molecules". In: *Advances in chemical physics* 14 (2007), pp. 35–89.

[218] David Hochstuhl and Michael Bonitz. "Two-photon ionization of helium studied with the multiconfigurational time-dependent Hartree–Fock method". In: *The Journal of Chemical Physics* 134.8 (Feb. 2011), p. 084106. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.3553176/14905108/084106_1_online.pdf. URL: https://doi.org/10.1063/1.3553176.

[219] C. C. J. Roothaan. "New Developments in Molecular Orbital Theory". In: *Rev. Mod. Phys.* 23 (2 Apr. 1951), pp. 69–89. URL: https://link.aps.org/doi/10.1103/RevModPhys.23.69.

[220] Qiming Sun et al. "Recent developments in the PySCF program package". In: *The Journal of Chemical Physics* 153.2 (July 2020), p. 024109. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0006074/16722275/024109_1_online.pdf. URL: https://doi.org/10.1063/5.0006074.

[221] C.F. Bunge, J.A. Barrientos, and A.V. Bunge. "Roothaan-Hartree-Fock Ground-State Atomic Wave Functions: Slater-Type Orbital Expansions and Expectation Values for Z = 2-54". In: *Atomic Data and Nuclear Data Tables* 53.1 (1993), pp. 113–162. URL: https://www.sciencedirect.com/science/article/pii/S0092640X8371003X.

[222] Péter Pulay. "Convergence acceleration of iterative sequences. the case of scf iteration". In: *Chemical Physics Letters* 73.2 (1980), pp. 393–398. URL: https://www.sciencedirect.com/science/article/pii/0009261480803964.

[223] Nicolas Tancogne-Dejean et al. "Octopus, a computational framework for exploring light-driven phenomena and quantum dynamics in extended and finite systems". In: *The Journal of Chemical Physics* 152.12 (Mar. 2020), p. 124119. eprint: https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/1.5142502/16712083/124119_1_online.pdf. URL: https://doi.org/10.1063/1.5142502.

[224] David M Knigge et al. "Modelling Long Range Dependencies in ND: From Task-Specific to a General Purpose CNN". In: *The Eleventh International Conference on Learning Representations*. 2023. URL: https://openreview.net/forum?id=ZW5aK4yCRqU.

[225] S. Zhang et al. "A Graph-based machine learning approach to accelerate plasma simulations". In: *Bulletin of the American Physical Society* (2024). URL: https://meetings.aps.org/Meeting/DPP24/Session/UO09.4.

[226] Hamed Shirzad et al. “Exphormer: Sparse transformers for graphs”. In: *International Conference on Machine Learning*. PMLR. 2023, pp. 31613–31632.