

University of Southampton Research Repository ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON
SCHOOL OF CHEMISTRY

Computational methods for density
functional theory calculations on
insulators and metals based on localised
orbitals

by
Álvaro Ruiz-Serrano

Thesis for the degree of Doctor of Philosophy

August 2013

*Sueña el rico en su riqueza,
que más cuidados le ofrece;
sueña el pobre que padece
su miseria y su pobreza;
sueña el que a medrar empieza,
sueña el que afana y pretende,
sueña el que agravia y ofende,
y en el mundo, en conclusión,
todos sueñan lo que son,
aunque ninguno lo entiende.*

*Yo sueño que estoy aquí
de estas prisiones cargado,
y soñé que en otro estado
más lisonjero me vi.
¿Qué es la vida? Un frenesí.
¿Qué es la vida? Una ilusión,
una sombra, una ficción,
y el mayor bien es pequeño;
que toda la vida es sueño,
y los sueños, sueños son.*

La vida es sueño, acto segundo.
Pedro Calderón de la Barca, 1635.

Declaration of authorship

I, Álvaro Ruiz-Serrano, declare that the thesis entitled “Computational methods for density functional theory calculations on insulators and metals based on localised orbitals” and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
 - Pulay forces from localized orbitals optimized in situ using a psinc basis set. Á. Ruiz-Serrano, N. D. M. Hine, and C.-K. Skylaris. *J. Chem. Phys.* **136**(23), 234101 (2012).
 - A variational method for density functional theory calculations on metallic systems with thousands of atoms. Á. Ruiz-Serrano and C.-K. Skylaris. *J. Chem. Phys.* **139**(5), 054107 (2013).
 - Variationally-localized search direction method for constrained optimization of non-orthogonal, localized orbitals in electronic structure calculations. Á. Ruiz-Serrano and C.-K. Skylaris. *J. Chem. Phys.* **139**(16), 164110 (2013).
 - DIIS and Hamiltonian diagonalisation for total-energy minimisation in the ONETEP program using the ScaLAPACK eigensolver. Á. Ruiz-Serrano. *M. Sc. in High Performance Computing*. University of Edinburgh (2010).

Signed:

Date:

UNIVERSITY OF SOUTHAMPTON
SCHOOL OF CHEMISTRY

Doctor of Philosophy

Computational methods for density functional theory calculations on insulators
and metals based on localised orbitals
by Álvaro Ruiz-Serrano

Abstract

Kohn-Sham density functional theory (DFT) calculations yield reliable accuracy in a wide variety of molecules and materials. The advent of linear-scaling DFT methods, based on locality of the electronic matter, has enabled calculations on systems with tens of thousands of atoms. Localisation constraints are imposed by expanding the Kohn-Sham states in terms of a set of atom-centred, spherically-localised functions. Chemical accuracy is then achieved via a self-consistent optimisation using a high-resolution basis set. This formalism reduces the size of, and brings predictable sparsity patterns to, the matrices expressed in this representation, such as the Hamiltonian matrix. In this work, we used the ONETEP program for DFT calculations, which is based on the abovementioned principles. The vision behind our research is to advance the method by developing new and robust algorithms to enable novel applications based on localised orbitals.

We investigated the consequences of strict spatial localisation constraints on the energy minimisation problem. This part of our research led to new theoretical realisations and approaches to calculate the total energy and atomic forces. We show that the self-consistent energy minimisation process is in fact a *constrained* optimisation problem where the localisation constraints must be taken into account for consistency. With this important point in mind, we were able to show that the atomic forces must be corrected with extra terms that account for the localisation constraints. We developed the code to determine the correction terms, known as Pulay forces, which provide consistency between the total energy and atomic forces. Weakly-bound systems, such as protein-ligand complexes, are now accurately described after the total forces were corrected.

We took a step forward in the formulation of the optimisation problem with localisation constraints. Using an analytic manipulation of the equations that govern the energy minimisation procedure, we were able to design a new algorithm for constrained optimisation of non-orthogonal, localised functions. This new method is exact and accounts for the localisation constraints in a fully variational manner. We named this new algorithm the *variationally-localised search direction* (VLSD) method. We implemented a prototypical version of the VLSD method

in the ONETEP program. We were able to compare the convergence speed towards the solution between the VLSD and the standard method, based on an *unconstrained* optimisation process that does not account for localisation constraints in a variational fashion. Calculations with the VLSD method typically show faster convergence, especially in systems with a certain degree of expected natural delocalisation.

We not only studied the effects of localisation, we also used the localised-functions formalism to develop two new types of algorithms for calculations on insulators (various density mixing methods) and metals (a direct minimisation method for finite-temperature DFT). These methods use the SCALAPACK dense algebra package to distribute memory and computational requirements over a large number of processors. In these approaches, the Hamiltonian matrix must be diagonalised for consistency. The formalism based on localised functions makes this step efficient, despite the cubic-scaling cost, and allowed us to perform calculations on systems with thousands of atoms.

Density mixing methods are well-established techniques for total energy minimisation in self-consistent field approaches. Their implementation is very useful for profiling and benchmarking other approaches that are fundamental for achieving linear-scaling cost. Validation calculations on a set of insulators and semiconductors show that accuracy in the description of the electronic structure is consistently achieved with the density mixing schemes.

The new algorithm that we have implemented in the ONETEP program for calculations on metallic systems is an important new functionality for the code. Typically, standard DFT methods for metals, based on plane-waves or other forms of delocalised orbitals, are limited to a few tens of atoms. Calculations on thousands of atoms are possible at the expense of a high computational cost, requiring tens of thousands of processors. With our new approach, calculations on large metallic systems with thousands of atoms can be executed employing for the task only a few hundreds of processors. This method can be used in studies of metallic complexes that are not viable with traditional approaches, such as the kind of nanoparticles used in innovative catalytic processes.

Acknowledgements

The last few years of research activity have been very intense and productive, both professionally and personally. It is a pleasure for me to finally be able to see all our scientific contributions put together into a single document, to form this Ph. D. dissertation. I am grateful to all the people that contributed to make it a reality:

- To my thesis supervisor, Chris-Kriton Skylaris, for his direct participation in all our projects, for his much-appreciated scientific advise, and for his support.
- To the Engineering and Physical Sciences Research Council (EPSRC) for the financial support via a High End Computing Studentship through the UKCP consortium (Grant No EP/F038038/1).
- To the members of the Edinburgh Parallel Computing Centre (EPCC) for sharing their expertise with me during my studies at the University of Edinburgh. In particular, I would like to thank my M. Sc. dissertation supervisor, Bartosz Dobrzelecki, for his advice.
- To Nicholas Hine for his active involvement in the work regarding the calculation of the Pulay forces, as well as for being a reliable source of advise regarding high-standards computing programming.
- To the members of the Skylaris group, past and present, for creating a friendly environment to work and coexist. Ben Lowe, Chris Cave-Ayland, Chris Pittock, Chris Sampson, Jacek Dziejcz, Karl Wilkinson, Max Phipps, Nawzat Saadi, Nicholas Zonias, Quintin Hill, Peter Cherry, Stephen Fox and Valerio Vitale. It was great sharing the office with you.

- To those that I can proudly call my friends in Southampton, and whose company always made feel at home. Mike Carter, Yannis Haldoupis, Yuki Ikuno, Kathy Aston, Sam Golten, Carole Despiau, Stefanos Marangos, Ally Hu, Zohra Ouaray, Neus Campanyà Llovet and Javier Jaso Buzzoni. Cheers.
- To Antonia Minou, for her loving company and care, for making new ideas flourish in me, for being my partner in the exploration of the World, and for making me happy. *Ευχαριστώ πολύ.*

This Ph. D. dissertation is dedicated to my parents, whose efforts to provide the best education for their children ultimately made this thesis possible. I hope the reader will not take offence if I say a few words in Spanish.

Quiero dar las gracias a todos los que, desde la distancia, siempre me han deseado lo mejor en mi aventura hacia el doctorado.

- A mis amigos, a los que siempre tengo ganas de volver a ver cuando regreso a casa. César Castro Martínez, Silvia Vega Vega, Leo y Juliana Rodríguez Díaz, Diego Recio Velázquez, Sergio Fernández González, Raúl Fernández Cobos, Santi Fernández Carnicero, Chus Fernández Sainz e Iñaki López Ferreño.
- A mis abuelos, por su cariño sincero y por cuidar siempre de mí.
- A mi hermano Javi, por su fiel compañía desde que éramos pequeños.
- Y muy especialmente, a mis padres, Juan Antonio y María Isabel, quienes supieron brindarme la mejor educación y me prepararon para la vida adulta. Esta tesis no hubiera sido posible sin ellos, y por tanto, a ellos va dedicada.

Contents

Declaration of authorship	5
Abstract	7
Acknowledgements	9
List of Figures	15
List of Tables	27
Introduction	31
Thesis outline	32
List of related publications	35
1 Principles of quantum mechanics	37
1.1 The wavefunction space	37
1.2 The Schrödinger equation	41
1.3 Pure and mixed states	42
1.4 The variational principle of quantum mechanics	44
1.5 The Hellmann-Feynman theorem	47

2	Density functional theory	49
2.1	The molecular many-body problem	50
2.2	The Hohenberg-Kohn theorems	54
2.3	Finite-temperature DFT	59
2.4	Kohn-Sham mapping	62
2.5	The exchange-correlation functional	70
2.6	Basis set functions	73
2.7	Periodic systems and plane waves	75
2.8	Core electrons	79
2.9	Linear-scaling density functional theory	82
2.10	Localised orbitals	86
2.11	The ONETEP program	91
2.11.1	Calculation of the total energy	91
2.11.2	Inner loop optimisation	95
2.11.3	Outer loop optimisation	99
3	Pulay forces	103
3.1	Origin of the Pulay forces	104
3.2	Pulay forces in ONETEP	107
3.3	Convergence of the forces	111
3.4	Geometry optimisation calculations	113
3.4.1	Geometry optimization using NGWFs	114
3.4.2	Geometry optimization using PAOs	120
3.4.3	Geometry optimization of large systems	120
3.5	Summary	123

4	The variationally-localised search direction (VLSD) method	125
4.1	Localised functions	127
4.2	Analysis of the constraints	129
4.2.1	Orthonormality of the Kohn-Sham states	129
4.2.2	Localisation of the non-orthogonal functions	130
4.3	Steepest-descent optimisation of the non-orthogonal functions . .	133
4.4	Truncated search direction (TSD) method	138
4.5	The variationally-localised search direction (VLSD) method . . .	141
4.5.1	Mathematical principle of the VLSD method	141
4.5.2	The VLSD algorithm	146
4.6	Results and discussion	150
4.6.1	Description of the calculations	150
4.6.2	Analysis of the TSD method	152
4.6.3	Convergence properties	155
4.6.4	Validation tests	157
4.7	Summary	159
5	Density mixing and Hamiltonian diagonalisation methods	163
5.1	Direct inversion of the iterative subspace	165
5.2	Density kernel optimisation with DIIS methods	167
5.2.1	Linear mixing	169
5.2.2	ODA mixing	170
5.2.3	Pulay mixing	172
5.2.4	LiST mixing	173
5.3	NGWF optimisation	174
5.4	Numerical eigensolvers	175

5.5	Performance of the eigensolvers	178
5.6	Validation tests	183
5.7	Summary	188
6	Direct minimization method for metals	191
6.1	Implementation in the ONETEP program	193
6.2	Inner loop: energy eigenvalues optimisation	195
6.3	Outer loop: NGWF optimisation	197
6.4	Scaling with the system size	201
6.5	Results	204
6.6	Summary	208
7	Conclusions and final remarks	211
Appendix A	Spherical wave representation of the NGWFs	219
A.1	Theory	219
A.2	Generation of the SWs	221
A.3	Truncation rules for the expansion of the NGWFs in terms of SWs	223
A.4	Accuracy tests	227
Appendix B	A VLSD approach for QM/QM calculations with selective NGWF optimisation	231
B.1	Unsuccessful attempts	232
B.2	Complete VLSD approach for QM/QM calculations	235
Appendix C	VLSD method for spatial localisation of the density kernel	239
	Bibliography	245

List of Figures

2.1	Fermi-Dirac distribution for different values of the electronic temperature.	69
2.2	Pseudopotential approximation. The pseudopotential V_{pseudo} converges to $1/r$ after a given radial cut-off r_c . The pseudised wavefunction, ψ_{pseudo} , is smooth near the core and converges to the original one in the long-range. Image extracted from Ref. [40].	81
2.3	(a) Localised Kohn-Sham orbitals on a peptide. (b) Single NGWFs on the same peptide. Images taken from Ref. [4].	92
2.4	(a) Two-dimensional layout of the NGWFs in a uniform Cartesian grid, including their spherical localisation regions. (b) Three-dimensional representation of a psinc function. Images taken from Ref. [178].	92
2.5	Optimisation of a single NGWF on a carbon atom in ethene. Image taken from [178].	101
2.6	Self-consistent energy minimisation with two nested loops.	102

3.1	Convergence of the forces with respect to the kinetic energy cutoff and the radius, R , of the NGWF localisation spheres (assumed to be the same for all the NGWFs). Left: force acting on an oxygen atom in CO_2 along the covalent bond. Right column: force acting on a hydrogen atom in H_2O dimer along the hydrogen bond. The first row corresponds to the Hellmann-Feynman force, the second to the Pulay force and the third are the Pulay-corrected total forces (Hellmann-Feynman forces plus Pulay).	112
3.2	Eggbox effect on the energy (top) and force along the covalent bond (bottom) of a CO_2 molecule with SZ (red circles), DZP (blue squares) and TZDP (green diamonds) PAO basis sets and NGWFs (black triangles).	115
3.3	Convergence of the maximum absolute value of the force during geometry optimization of the adenine-thymine DNA base pair with NGWF radii of 3.70 \AA (top) and 7.94 \AA (bottom). The convergence threshold was 0.015 eV/\AA	116

3.4	Convergence of the binding energy in the adenine-thymine complex with increasing NGWF radii. The calculations that use the structures obtained with Pulay-corrected forces show systematic convergence (blue squares) with increasing NGWF radii. In contrast, the calculations that use structures obtained with uncorrected Hellmann-Feynman forces (red circles) lack of convergence and show large variations of the binding energy. The calculations represented with black crosses show convergence of the binding energy with increasing NGWF radii, based on fixed molecular geometries obtained with Pulay-corrected forces and NGWF radii of 3.70 and 7.94 Å, respectively.	118
3.5	Convergence of the maximum force during the BFGS geometry optimization of the tennis-ball dimer using only the Hellmann-Feynman forces (top left molecule) and Pulay-corrected forces (top right).	119
3.6	Convergence of the maximum value of the force during the geometry optimization of the DNA molecule using NGWFs and DZP and TZDP PAO basis sets.	123

4.1 Two-dimensional depiction of the real-space localisation regions in a calculation involving four non-orthogonal localised functions. The covariant functions $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$ and $\phi_4(\mathbf{r})$ are localised within different circles, which produces a distinctive sparsity pattern in $S_{\alpha\beta}$. In contravariant space, $\phi^1(\mathbf{r})$, $\phi^2(\mathbf{r})$ and $\phi^3(\mathbf{r})$ are localised within the same localisation region, made up from the union of the localisation regions of $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$, while $\phi^4(\mathbf{r})$ is localised in a circle, as a result of $\phi_4(\mathbf{r})$ not overlapping any other function. The volume covered by all the localisation regions is the same both in the covariant and contravariant representations. All the expansion coefficients associated to grid points outside of this volume are zero by definition of the localisation constraint. 132

4.2 Example of point-wise subset separation in a hypothetical calculation with a two-dimensional real space simulation cell, containing three covariant functions localised within circles. The grid point \mathbf{r}_a is inside LR($\phi_2(\mathbf{r})$), and outside LR($\phi_1(\mathbf{r})$) and LR($\phi_3(\mathbf{r})$). Therefore $\xi_a = \{2\}$ and $\chi_a = \{1, 3\}$. The grid point \mathbf{r}_b is inside LR($\phi_2(\mathbf{r})$) and LR($\phi_3(\mathbf{r})$), and outside LR($\phi_1(\mathbf{r})$). Therefore $\xi_b = \{2, 3\}$ and $\chi_b = \{1\}$. The grid point \mathbf{r}_c is inside all the localisation regions. Therefore $\xi_c = \{1, 2, 3\}$ and χ_c is an empty subset. The grid point \mathbf{r}_d is outside all the localisation regions. Therefore, ξ_d is an empty subset and $\chi_d = \{1, 2, 3\}$ 144

4.3 Schematic representation of the VLSD method. In this example, there are three mutually overlapping covariant functions, $\{\phi_\alpha\}$, localised within circles in a two-dimensional simulation cell. The algorithm begins in (A), calculating the derivative of the energy with respect to the variational parameters in Ω , imposing orthogonality of $\{\psi_i\}$ and applying occupancy and kinetic energy preconditioning Eq. (4.63). The result are the values of $\{d_p^\alpha \in \Upsilon'\}$, which correspond to the grid points confined within the circles. Then, a loop over all the grid points \mathbf{r}_p starts. Steps (B) and (C) represent the progress of the algorithm at a certain intermediate stage, where some, but not all, grid points have been scrutinised. In most situations, only the covariant search directions are needed to proceed with the steepest descent update. In such cases, the path through (B) and (C) calculates $\{d_{p1}\}$, $\{d_{p2}\}$ and $\{d_{p3}\}$ directly by applying Eq. (4.64) and Eq. (4.66) to the current grid point. Alternatively, the path through (D) and (E) calculates the remaining contravariant search direction coefficients using Eq. (4.65). The VLSD algorithm is capable of calculating the covariant and contravariant search direction coefficients simultaneously by running steps (B) and (D) as independent processes. At the conclusion of the look over \mathbf{r}_p , the resulting search directions are directly localised within the established localisation regions and duality between $\{d_{p\alpha}\}$ and $\{d_p^\alpha\}$ holds exactly. 149

4.4	Evolution of the parameter μ with the iteration counter l for different values of R_α . The solid lines are the Gaussian broadening curves used to fit the results.	153
4.5	Progression of μ_0 , l_0 and θ_0 , calculated after fitting a Gaussian broadening function through the results, for different values of R_α	155
4.6	Calculated value of the relative speed-up ζ for each test system and increasing R_α	157
5.1	Set of silicon nanorods benchmarks.	179
5.2	Set of benchmarks of fragments of an amyloid fibril protein.	180
5.3	Comparison of the scaling of the LNV and DIIS algorithms with the system size for the set of amyloid benchmarks. The systems have been simulated using 20 nodes and 4 cores per node on HEC-ToR and Iridis 3.	183
5.4	Calculations on $\text{Si}_{29}\text{H}_{36}$ silicon nanorod with DIIS. Pulay and LiST methods use a history of size $N_{mix} = 5$, and perform 2 leading linear mixing iterations with parameter $\lambda = 0.1$. (a) Convergence of the RMS of the commutator (2.122) during the execution of the inner loop. In this plot, the series corresponding to linear mixing with $\lambda = 0.1$ and ODA are overlapping. (b) Convergence of the RMS of the slope of the NGWF gradient along the search direction (2.127), with 5 iterations in the inner loop.	184

5.5	Calculations on $\text{Si}_{242}\text{H}_{140}$ silicon nanorod with DIIS. Pulay and LiST methods use a history of size $N_{mix} = 5$, and perform 5 leading linear mixing iterations with parameter $\lambda = 0.05$. (a) Convergence of the RMS of the commutator (2.122) during the execution of the inner loop. (b) Convergence of the RMS of the slope of the NGWF gradient along the search direction (2.127), with 10 iterations in the inner loop.	186
5.6	Flowchart of the DIIS implementation in ONETEP.	190
6.1	Time taken to complete one outer loop iteration with five inner loop iterations, in calculations on Au nanoparticles of increasing size. The plot shows the time taken by different parts of the algorithm. “Hamiltonian DD” and “Hamiltonian DI” must be interpreted as the density-dependent and density-independent terms of the Hamiltonian, respectively.	202
6.2	Density of states of Pt_{13} , obtained with ONETEP and CASTEP. Agreement is achieved for NGWF radii of 4.0 Å and above.	205
6.3	Lattice parameter stretching of bulk Cu. There are 4 atoms in the CASTEP simulation cell and 500 atoms in the ONETEP simulation cell, forming a $5 \times 5 \times 5$ supercell.	207
6.4	Convergence of the Helmholtz free energy functional, $A_{\mathcal{T}}$, with the number of outer loop (NGWF optimisation) iterations, for a set of Au cuboctahedral nanoparticles of increasing sizes. The structures of Au_{13} and Au_{2057} are also shown in the plot.	209
A.1	Constrained spherical Bessel function	225

A.2	NGWF for a atom of C in the ethene molecule with $l_{max} = 3, 4$ and 5, respectively and $n_{extra} = 0$. The black line and square marks represents the fully optimised NGWF before generating the SWs expansion.	229
A.3	NGWF for a atom of C in the ethene molecule with $l_{max} = 3$ and $n_{extra} = 2, 20$ and 50, respectively. The black line and square marks represents the fully optimised NGWF before generating the SWs expansion.	230
B.1	Multiscale QM/QM calculation example. The top left corner of the molecule in this space representation is the high-accuracy region, where both the density kernel and the NGWFs are optimised. The rest of the molecule is the lower-accuracy region, where the NGWFs are fixed and only the density kernel is optimised	232
B.2	QM/QM setup of benzene in solvation, with 200 water molecules in the box. The NGWFs in the benzene molecule are optimised, and the rest are fixed.	234

B.3 Example of point-wise subset separation in a hypothetical calculation with a two-dimensional real space simulation cell, containing three covariant functions localised within circles. The covariant function ϕ_3 (magenta border and white filling) is constrained to be fixed during the calculation. The grid point \mathbf{r}_a is inside $\text{LR}(\phi_2(\mathbf{r}))$, and outside $\text{LR}(\phi_1(\mathbf{r}))$ and $\text{LR}(\phi_3(\mathbf{r}))$. Therefore $\xi_a = \{2\}$ and $\chi_a = \{1, 3\}$. The grid point \mathbf{r}_b is inside $\text{LR}(\phi_2(\mathbf{r}))$ and $\text{LR}(\phi_3(\mathbf{r}))$, and outside $\text{LR}(\phi_1(\mathbf{r}))$, but ϕ_3 is fixed. Therefore $\xi_b = \{2\}$ and $\chi_b = \{1, 3\}$. The grid point \mathbf{r}_c is inside all the localisation regions, but ϕ_3 is fixed. Therefore $\xi_c = \{1, 2\}$ and $\chi_c = \{3\}$. The grid point \mathbf{r}_d is outside all the localisation regions. Therefore, ξ_d is an empty subset and $\chi_d = \{1, 2, 3\}$ 236

B.4 Schematic representation of the VLSD method for the QM/QM scheme. In this example, there are three overlapping covariant functions, $\{\phi_\alpha\}$, localised within circles in a two-dimensional simulation cell. The covariant NGWF ϕ_3 (magenta border and white filling) is constrained to be fixed during the calculation. The algorithm begins in (A), calculating the energy gradient with respect to the variational degrees of freedom. The result are the values of $\{d_p^\alpha\}$, which correspond to the grid points confined within the circles of the optimised NGWFs only. Then, a loop over all the grid points \mathbf{r}_p starts. Steps (B) and (C) represent the progress of the algorithm at a certain intermediate stage, where some, but not all, grid points have been scrutinised. The point-wise subset separation is repeated at every grid point, as in Eq. (B.6) and Eq. (B.7). The path through (B) and (C) is the most efficient one. The covariant search directions $\{d_{p1}\}$ and $\{d_{p2}\}$ are calculated, while $\{d_{p3}\}$ becomes zero by construction. Alternatively, the path through (D) and (E) calculates the remaining contravariant search direction coefficients using Eq. (4.65). The VLSD algorithm is capable of calculating the covariant and contravariant search direction coefficients simultaneously by running steps (B) and (D) as independent processes. At the conclusion of the look over \mathbf{r}_p , the resulting search directions are directly localised within the established localisation regions and duality between $\{d_{p\alpha}\}$ and $\{d_p^\alpha\}$ holds exactly, even if $\{d_{p3}\}$ are zero. 237

C.1 Example of pair-wise subset separation in a hypothetical calculation with four NGWFs that give rise to a 4×4 density kernel matrix. The localisation constraints with a finite radial cut-off R_c impose some off-diagonal elements being zero. The elements over the coloured background belong to the subset Λ , while the elements over the white background belong to the subset Ξ 241

List of Tables

3.1	Geometry optimization of adenine-thymine using ONETEP for different localisation radii R_α with Hellmann-Feynman forces (HF) and Pulay-corrected forces (PC). Results show the hydrogen bond lengths, the maximum absolute value of the force, $ F _{\max}$, and the number of BFGS steps required.	117
3.2	Length of the two hydrogen bonds in the adenine-thymine complex obtained with different PAO multiple- ζ basis sets of different localisation radii of size R , given in Å. The total forces converged to 0.015 eV/Å.	120
3.3	Structural parameters of the optimised 1045-atom DNA fragment as optimised by ONETEP. N. . .H(4), N. . .H(8) and N. . .H(12) correspond to the hydrogen bond of the fourth, eighth and twelfth pairs, respectively, that involve a nitrogen atom.	122
4.1	Number of iterations required to achieve self-consistent convergence in calculations with the TSD and VLSD methods. The values refer to the ratio l_{TSD}/l_{VLSD}	156
4.2	Calculated equilibrium bond length, L_0 , and total electric dipole moment at the equilibrium geometry, p_0 , of NaCl and H ₂ O.	159

5.1	Summary of the main features of the DIIS schemes implemented in ONETEP. The table shows the range of the $\{\lambda_j\}$ coefficients and the number of matrices required by each method.	169
5.2	Set of benchmarks corresponding to silicon nanorods of different size. The NGWF radii $\{R_\alpha\}$ are 7.0 Bohr for silicon atoms and 6.0 Bohr for hydrogen. The density kernel is truncated beyond a cut-off radius R_c of 25 Bohr. The table shows the number of atoms, N_{at} , number of NGWFs, N_ϕ , and plane-wave kinetic energy cut-off, E_c , of each calculation.	179
5.3	Set of benchmarks corresponding to a the amyloid protein of increasing size. The NGWF radii $\{R_\alpha\}$ are 7.0 Bohr for all species. The density kernel is truncated beyond a cut-off radius R_c of 20 Bohr. The table shows the number of atoms, N_{at} , number of NGWFs, N_ϕ , and plane-wave kinetic energy cut-off, E_c , of each calculation.	180
5.4	Successful calculations using LAPACK and ScaLAPACK eigensolvers on HECToR and Iridis 3. In all cases, 4 cores per node have been used. The asterisk * indicates that a smaller number of nodes has not been tested and it is possible for the calculation to fit on less nodes than that. For saving computational time, the protein systems have not been simulated using LAPACK on Iridis 3. . . .	182
5.5	Dipole moment of $\text{Si}_{29}\text{H}_{36}$	185
5.6	Dipole moment of $\text{Si}_{242}\text{H}_{140}$	187

6.1	BFGS geometry optimisation of Pt ₁₃ with ONETEP and CASTEP. The table shows the optimised value of the distance to the nearest-neighbour Pt atom.	205
6.2	Bulk modulus, B , and equilibrium lattice parameter, L_0 , of bulk Cu, calculated with CASTEP and ONETEP. The value of χ^2 corresponding to the fitting of the results to the third-order Birch-Murnaghan equation is also shown.	207
A.1	Iterations needed to converge the NGWFs of the ethene molecule to the tolerance value.	228

Introduction

The role of atomistic simulations of matter is to model the fundamental interactions that take place at the nano-scale and to produce reliable theoretical techniques capable of producing accurate predictions. It is by itself a vast and rapidly-evolving field with implications for many branches of nano-technology, from materials science [1] to bio-medicine [2].

Most of the relevant properties of atomic systems can be obtained from the detailed characterisation of their electronic structure, which, inevitably, requires the level of theory of quantum mechanics. It is well known that exact solutions of the quantum equations can only be obtained for very simple systems. Hence, realistic systems must be studied using well-established approximations. The complexity of solving the approximate problem is nevertheless enormous. Only using robust numerical algorithms and state-of-the-art computing techniques can simulations on many-atom systems be achieved.

Over the last decades, a wealth of scholarly research on electronic structure theory has been produced. The theory now known as density functional theory (DFT) [3] is an exact reformulation of the quantum-mechanical approach to condensed-matter physics in which the electronic density takes the central role of being the defining element of a given atomic system. DFT methods have been shown to be remarkably accurate and have enabled the research community to develop ever-

more sophisticated techniques for calculations on a wide range of materials. The concept of locality or “nearsightedness” of the electronic matter has led to the design of linear-scaling DFT methods that expand the size of the systems that can be simulated to tens of thousands of atoms.

Our research focused on the development of new approaches for DFT simulations on a wide range of systems, from organic molecules to semi-conductors and metals. We used the formalism of strictly localised orbitals, where the Kohn-Sham molecular orbitals are expanded as a linear combination of atom-centred functions localised within spheres. The ground-state is found using a self-consistent, variational approach to optimise the set of degrees of freedom that describe the system. We used the ONETEP program [4] as the software platform in which our developments were implemented.

In the course of our investigations, we were able to produce advances in our understanding of the theory behind the methods that use strict localisation constraints to reduce the computational cost. We developed new algorithms that are capable of treating localisation in an exact and fully-consistent manner, as an integral part of the optimisation process. Furthermore, we used the technique of localised functions to devise new methods for insulators and metals capable of calculations on thousands of atoms. The new software capabilities for atomistic simulations that resulted from our research are available as part of the ONETEP program and are accessible to other researchers.

Thesis outline

The present Ph. D. dissertation is organised in seven chapters. The first two chapters aim to give a comprehensive background of our work. Chapter 1 contains

a summary of the quantum theory and principles that form the theoretical background of our research. Chapter 2 contains a comprehensive literature review on atomistic calculations with DFT methods, including pivotal concepts such as the Born-Oppenheimer approximation, the Hohenberg-Kohn theorems, Kohn-Sham density functional theory and the fundamentals of computational approaches to quantum chemistry. We also include a description of the ONETEP method for linear-scaling density functional theory calculations. The results of our research are presented and discussed in Chapters 3 to 6.

In Chapter 3, we show that a contribution to the total atomic forces, known as Pulay forces, are present and must be calculated if strict localisation constraints are requested. This correction to the total force is essential for achieving consistency between the total energy and the total atomic forces. We show that, once the Pulay corrections have been taken into account, geometry optimisation calculations on weakly-bound systems produce correct relaxed structures.

Then, in Chapter 4 we describe the principles of a newly-developed algorithm, named the *variationally-localised search direction* method for constrained energy minimisation. This is a novel theoretical realisation of the optimisation problem with localisation constraints in a non-orthogonal representation, which shows that strict spatial localisation can be enforced exactly and in a fully variational manner. The VLSD method avoids the addition of an uncontrollable error factor arising from truncation of the search direction vectors, common in unconstrained optimisation schemes.

Chapter 5 describes the implementation of density mixing methods based on Hamiltonian diagonalisation. This development allowed us to assess the applicability of parallel eigensolvers to diagonalise the Hamiltonian operator as the basic

self-consistent step. Density mixing techniques are well-established methods for total energy minimisation which can be used to benchmark other techniques (such as direct minimisation schemes). Our results show that calculations in systems with thousands of atoms are possible, despite of the cubic-scaling properties of the diagonalisation operation. Accuracy in the description of the electronic structure is demonstrated with a set of validation calculations on silicon nanorods.

In Chapter 6, we outline the principles of a new direct free-energy minimisation technique that allows calculations on large metallic systems with thousands of atoms. The results show that the new algorithm substantially widens the range of metallic materials that can be simulated at the DFT level of theory at the same time that it reduces the computational cost. Our method has the potential to be used in studies of large-scale compounds (such as metallic nanoparticles) of direct relevance to industry.

There are many ramifications of our work that can be continued as future research projects. In Chapter 7 we summarise what we have achieved with our investigations and we outline the most relevant extensions of our work for future developments. Additionally, Appendices A, B and C show extra results that could potentially be used in future developments, and which we obtained during the course of the main research projects. Appendix A shows some results regarding the use of spherical waves as a method to store the result of a calculation in an effective manner. Appendix B shows an extension to the VLSD method for developing a multiple-accuracy QM/QM scheme based on DFT. Finally, Appendix C shows how the VLSD method can be used for the very challenging problem of density matrix optimisation with localisation constraints, which is central in linear-scaling approaches.

List of related publications

- The theory and results presented in Chapter 3 have been published as:

Pulay forces from localized orbitals optimized in situ using a psinc basis set.

Á. Ruiz-Serrano, N. D. M. Hine, and C.-K. Skylaris. *J. Chem. Phys.* **136**(23), 234101 (2012).

- The theory and results presented in Chapter 4 have been published as:

Variationally-localized search direction method for constrained optimization of non-orthogonal, localized orbitals in electronic structure calculations.

Á. Ruiz-Serrano and C.-K. Skylaris. *J. Chem. Phys.* **139**(16), 164110 (2013).

- Some, but not all, of the theory and results presented in Chapter 5 have been published as:

DIIS and Hamiltonian diagonalisation for total-energy minimisation in the ONETEP program using the ScaLAPACK eigensolver.

Á. Ruiz-Serrano. *M. Sc. in High Performance Computing*. University of Edinburgh (2010).

- The theory and results presented in Chapter 6 have been published as:

A variational method for density functional theory calculations on metallic systems with thousands of atoms.

Á. Ruiz-Serrano and C.-K. Skylaris. *J. Chem. Phys.* **139**(5), 054107 (2013).

Chapter 1

Principles of quantum mechanics

Quantum mechanics is a very complex branch of modern physics that has extraordinary repercussions in all the natural sciences as well as in engineering, mathematics and philosophy. In our research, we applied the equations of quantum mechanics to atomistic modelling and solid state physics. In this chapter we wish to provide a brief account of some of the principles of the quantum theory which are of special relevance to our project. For an in-depth study of the quantum theory, we refer the reader to Refs. [5–7] of well-known textbooks on the matter.

1.1 The wavefunction space

The predominant view of quantum mechanics is mostly due to the formalism developed by Schrödinger [8], termed *wave mechanics*, which makes the wavefunction, $\Psi(t)$, the defining element of a given quantum system. An alternative formulation of the quantum theory is due to Heisenberg, Born and Jordan, [9–11], termed *matrix mechanics*, which nominates the observable properties of a given quantum system as the elements that effectively define that system. Observables are the measurable quantities, either with experimental or theoretical techniques, that can

be obtained by probing the system. The interpretation of the quantum theory given by *matrix mechanics* is based on the action of operators over the elements of a Hilbert space, known as the *wavefunction space*. It is due to Dirac [12] the demonstration that these two visions of quantum mechanics are in fact equivalent. The rules of the so-called “new” quantum mechanics are enumerated in the Dirac-von Neumann axioms [12, 13].

Notation in quantum mechanics is important. The notation due to Dirac [14], called the “bra-ket” notation, is predominant in quantum theory for its consistency and simplicity. Every wavefunction $\Psi(t)$ suitable of describing a quantum system is a unique point of a complex Hilbert space \mathcal{H} , and must be single-valued and square-integrable. The wavefunction space \mathcal{H} has infinite dimensions. However, an infinite subset $\mathcal{B} \subset \mathcal{H}$ can always be found such as any other element $\Psi(t) \in \mathcal{H}$ can be written as a linear combination of elements in \mathcal{B} . The subset \mathcal{B} is said to be a basis set of \mathcal{H} . Seen as points in a Hilbert space, the wavefunctions must be, by definition, independent of their representation. For example, in literature, $\Psi(\mathbf{r}, t)$ often refers to the representation of $\Psi(t)$ in the real tridimensional space, while $\Psi(\mathbf{k}, t)$ often refers to the representation of $\Psi(t)$ in the reciprocal tridimensional space (or \mathbf{k} -space). It is obvious that $\Psi(\mathbf{r}, t) \neq \Psi(\mathbf{k}, t)$, but, at the same time, the difference between these two expressions is not in the actual wavefunction $\Psi(t)$, but in its representation in the \mathbf{r} -space and \mathbf{k} -space, respectively. The bra-ket notation, addresses this important point by assigning a bra $\langle \Psi(t) |$ and a ket $|\Psi(t)\rangle$ to every wavefunction $\Psi(t) \in \mathcal{H}$ in a one-to-one correspondence via Hermitian conjugacy, so that $|\Psi(t)\rangle = (\langle \Psi(t) |)^\dagger$. In this notation, the inner product of two dynamical states $\Psi_1(t)$ and $\Psi_2(t)$ is symbolised by $\langle \Psi_1(t) | \Psi_2(t) \rangle$. The inner product has the following properties:

- Linearity:

$$\langle \lambda_1 \Psi_1(t) + \lambda_2 \Psi_2(t) | \Psi_3(t) \rangle = \lambda_1 \langle \Psi_1(t) | \Psi_3(t) \rangle + \lambda_2 \langle \Psi_2(t) | \Psi_3(t) \rangle.$$

- Conjugate symmetry: $\langle \Psi_1(t) | \Psi_2(t) \rangle = \langle \Psi_2(t) | \Psi_1(t) \rangle^\dagger$.
- Positive definiteness: $\langle \Psi_1(t) | \Psi_1(t) \rangle \geq 0$, with $\langle \Psi_1(t) | \Psi_1(t) \rangle = 0$ if and only if $\Psi_1(t) = 0$.

Representations of dynamical states are correctly managed with the bra-ket notation via inner products. To follow with the examples put before, the real-space representation corresponds to $\Psi(\mathbf{r}, t) = \langle \mathbf{r} | \Psi(t) \rangle$, while the reciprocal-space representation corresponds to $\Psi(\mathbf{k}, t) = \langle \mathbf{k} | \Psi(t) \rangle$. The conjugates of these expressions are $\Psi^\dagger(\mathbf{r}, t) = \langle \Psi(t) | \mathbf{r} \rangle$ and $\Psi^\dagger(\mathbf{k}, t) = \langle \Psi(t) | \mathbf{k} \rangle$.

For every observable \mathcal{A} , there is a linear Hermitean operator $\hat{A}(t)$ and an associated eigenvalue equation:

$$\hat{A}(t) |\chi_n(t)\rangle = a_n(t) |\chi_n(t)\rangle, \quad (1.1)$$

such as $\{a_n(t)\}$ take real values and $\{\chi_n(t)\}$ form a basis set of \mathcal{H} . Eigenvectors associated to different eigenvalues are mutually orthogonal by construction. An orthonormal basis set such as $\langle \chi_n(t) | \chi_m(t) \rangle = \delta_{nm}$ can always be constructed, even if one or more eigenstates are degenerate. Assuming orthonormality of the basis set, any wavefunction $\Psi(t)$ can be written as a linear combination of $\{\chi_n(t)\}$ as:

$$|\Psi(t)\rangle = \sum_n |\chi_n(t)\rangle c_n(t). \quad (1.2)$$

where $c_n(t)$ are the expansion coefficients:

$$c_n(t) = \langle \chi_n(t) | \Psi(t) \rangle. \quad (1.3)$$

The basis set formed by $\{\chi_n(t)\}$ gives rise to the closure relation:

$$\sum_n |\chi_n(t)\rangle \langle \chi_n(t)| = \hat{I}, \quad (1.4)$$

where \hat{I} is the identity operator. The postulates of quantum mechanics dictate that when a measurement of the observable \mathcal{A} is taken, the result is one, and only one, of the eigenvalues of $\hat{A}(t)$, and the system collapses to the state given by the corresponding eigenvector. The probability of registering the eigenvalue $a_n(t)$ after a measurement of \mathcal{A} is¹:

$$\mathcal{P}_n(t) = \frac{|c_n(t)|^2}{\langle \Psi(t) | \Psi(t) \rangle}. \quad (1.5)$$

This probabilistic interpretation of the laws of nature is intrinsic to the quantum theory. The expectation value, or mean value, of $\hat{A}(t)$, can be calculated as:

$$\langle \hat{A}(t) \rangle = \frac{\langle \Psi(t) | \hat{A}(t) | \Psi(t) \rangle}{\langle \Psi(t) | \Psi(t) \rangle} = \sum_n \mathcal{P}_n(t) a_n(t). \quad (1.6)$$

Equations (1.5) and (1.6) can be simplified by requiring the system's wavefunction to be normalised, i.e., $\langle \Psi(t) | \Psi(t) \rangle = 1$. This choice of normalisation does not alter the value of the observables, as it will be demonstrated in Section 1.4.

Two Hermitean operators $\hat{A}(t)$ and $\hat{B}(t)$ are simultaneously observable if their commutator:

¹If the eigenspectra of \hat{A} is degenerate, then $P_n(t) = \frac{\sum_j |c_{jn}|^2}{\langle \Psi(t) | \Psi(t) \rangle}$, where j labels the degenerate eigenvectors corresponding to $a_n(t)$. The system collapses to a state confined to the subspace of \mathcal{H} spanned by $\{\chi_j(t)\}$.

$$\left[\hat{A}(t), \hat{B}(t) \right] = \hat{A}(t) \hat{B}(t) - \hat{B}(t) \hat{A}(t), \quad (1.7)$$

is zero, meaning that $\hat{A}(t)$ and $\hat{B}(t)$ share the same set of eigenvectors. A typical example of operators that do not commute are the position and momentum operators, which are observable separately, but not simultaneously. This is the basis of the uncertainty principle as originally formulated by Heisenberg [15].

1.2 The Schrödinger equation

The evolution in time of the system's wavefunction is determined by the non-relativistic Schrödinger equation [8]:

$$i\hbar \frac{\partial}{\partial t} |\Psi(t)\rangle = \hat{H}(t) |\Psi(t)\rangle, \quad (1.8)$$

Equation (1.8) can be developed by establishing analogies with the Euler-Lagrange and Hamilton variational approaches to classical dynamics [16]. It is a first-order differential equation (a condition derived from the concept of continuous time of classical mechanics), and has the property of linearity, which implies that any linear combination of solutions is also a valid solution. In classical mechanics, evaluating the Hamiltonian at a given set of canonical coordinates returns the value of the energy of the system at time t . In quantum mechanics, the energy is given by the expectation value of the Hamiltonian operator acting on the system's wavefunction:

$$E(t) = \langle \Psi(t) | \hat{H}(t) | \Psi(t) \rangle, \quad (1.9)$$

where normalisation of Ψ is assumed.

If the Hamiltonian is independent of t , the energy eigenvalues and their associated eigenvectors will also be independent of t . The energy is said to be a constant of motion if it is independent of t , and in such cases, the system is said to be conservative. Conservative systems are very important for the study of quantum mechanical systems, for they link to the familiar principle of conservation of the total energy. Under this assumption, the solutions of (1.8) are the stationary states:

$$|\Psi(t)\rangle = \sum_n |\Psi_n\rangle e^{iE_n t/\hbar}, \quad (1.10)$$

where the time-independent wavefunctions $\{\Psi_n\}$ are the solutions of the eigenvalue equation of the Hamiltonian operator:

$$\hat{H} |\Psi_n\rangle = E_n |\Psi_n\rangle. \quad (1.11)$$

Equation (1.11) is referred to as the time-independent Schrödinger equation.

1.3 Pure and mixed states

A quantum system is said to be in a pure state if it can be completely described by a wavefunction. On the other hand, a quantum system is said to be in a mixed state if it cannot be fully described by a wavefunction. Such cases correspond to systems where a set of different pure states $\{\Psi_u(t)\}$ coexist with certain probabilities $\{p_u\}$, with $0 \leq p_u \leq 1$ and $\sum_u p_u = 1$. Mixed states must be described via the density operator $\hat{\rho}$, defined as [5]:

$$\hat{\rho}(t) = \sum_u p_u |\Psi_u(t)\rangle \langle \Psi_u(t)|. \quad (1.12)$$

The probabilities and expectation values of observables are fully characterised by the product of the density operator and other Hermitian operators, thus correctly describing the mixed state. For example, the probability of registering the eigenvalue $a_n(t)$ and the eigenvector $\chi_n(t)$ after a measurement of $\mathcal{A}(t)$, is:

$$\mathcal{P}_n(t) = \hat{\rho}(t) \hat{P}_n(t), \quad (1.13)$$

where $\hat{P}_n(t)$ is the projection operator onto the space spanned by $\chi_n(t)$ ²:

$$\hat{P}_n(t) = |\chi_n(t)\rangle \langle \chi_n(t)|. \quad (1.14)$$

The expectation value of $\hat{A}(t)$ is:

$$\langle \hat{A}(t) \rangle = \text{tr} \left[\hat{\rho}(t) \hat{A}(t) \right] = \sum_n \mathcal{P}_n(t) a_n(t). \quad (1.15)$$

The density operator has the following properties:

- Idempotency: $\hat{\rho}(t) = \hat{\rho}^2(t)$ if, and only if, the system is in a pure state.
- Hermitean: $\hat{\rho}(t) = \hat{\rho}^\dagger(t)$.
- Normalisation: $\text{tr} [\hat{\rho}(t)] = \sum_u p_u = 1$.

The time-dependent Schrödinger equation (1.8) can be reformulated in terms of $\hat{\rho}(t)$ as:

$$i\hbar \frac{d}{dt} \hat{\rho}(t) = \left[\hat{H}(t), \hat{\rho}(t) \right]. \quad (1.16)$$

In conservative systems, $\hat{\rho}$ is independent of the time. Under this condition, the above equation yields

²Assuming non-degeneracy of $a_n(t)$.

$$\left[\hat{H}, \hat{\rho} \right] = 0, \quad (1.17)$$

and so the Hamiltonian and the density operator commute. A common set of wavefunctions that are eigenstates of both $\hat{\rho}$ and \hat{H} can be found. From now onwards, our study will be restricted to conservative systems only.

1.4 The variational principle of quantum mechanics

The variational principle of quantum mechanics states that solving the eigenvalue equation of the Hamiltonian $\hat{H} |\Psi\rangle = E |\Psi\rangle$ is identical to solving the variational problem $\delta E [\Psi] = 0$, where Ψ is the system's wavefunction. To begin with the proof, consider an infinitesimal variation of Ψ and its effect on the total energy of the system, calculated as the expectation value of \hat{H} :

$$\begin{aligned} \delta E [\Psi] &= E [\Psi + \delta\Psi] - E [\Psi] & (1.18) \\ &= \frac{\langle \Psi + \delta\Psi | \hat{H} | \Psi + \delta\Psi \rangle}{\langle \Psi + \delta\Psi | \Psi + \delta\Psi \rangle} - \frac{\langle \Psi | \hat{H} | \Psi \rangle}{\langle \Psi | \Psi \rangle} \\ &\dots \text{rearranging terms} \dots \\ &= \frac{\langle \delta\Psi | \hat{H} | \Psi \rangle + \langle \Psi | \hat{H} | \delta\Psi \rangle}{\langle \Psi | \Psi \rangle} - \frac{\langle \Psi | \hat{H} | \Psi \rangle}{\langle \Psi | \Psi \rangle^2} (\langle \delta\Psi | \Psi \rangle + \langle \Psi | \delta\Psi \rangle) + \mathcal{O}(\delta\Psi^2) \\ &= \frac{1}{\langle \Psi | \Psi \rangle} \left[\langle \delta\Psi | \hat{H} - E | \Psi \rangle + \langle \Psi | \hat{H} - E | \delta\Psi \rangle \right] + \mathcal{O}(\delta\Psi^2), \end{aligned}$$

where normalisation of Ψ is not assumed. The Hermiticity of \hat{H} implies that $\langle \delta\Psi | \hat{H} - E | \Psi \rangle = \langle \Psi | \hat{H} - E | \delta\Psi \rangle$. Ignoring the error due to second-order terms, $\mathcal{O}(\delta\Psi^2)$, and imposing the condition of conservative systems, $\delta E [\Psi] = 0$, gives

the result:

$$\hat{H} |\Psi\rangle - E |\Psi\rangle = 0, \quad (1.19)$$

and thus the variational principle is proved.

As described in Section 1.1, any suitable wavefunction capable of describing a quantum system must be square-integrable. A Lagrangian $\mathcal{L}[\Psi]$ that takes into account this constraint can be constructed as follows:

$$\mathcal{L}[\Psi] = \langle \Psi | \hat{H} | \Psi \rangle - E (\langle \Psi | \Psi \rangle - C), \quad (1.20)$$

where C is a real-valued constant and E is an undetermined Lagrange multiplier (constant) that takes into account the constraint of square-integrability of Ψ . The wavefunction that minimises $\mathcal{L}[\Psi]$ must meet the condition of extremal:

$$\frac{\delta \mathcal{L}}{\delta \langle \Psi |} = 0 \implies \hat{H} |\Psi\rangle - E |\Psi\rangle = 0. \quad (1.21)$$

Therefore, the variational principle can be reformulated in terms of an optimisation problem where the Lagrangian $\mathcal{L}[\Psi]$ is minimised with respect to Ψ .

Now, let Ψ^{trial} be a trial wavefunction for the ground-state Ψ_0 of the system. For simplicity, the eigenspectra of \hat{H} is assumed to be discrete and non-degenerate, and its eigenvectors $\{\Psi_n\}$ mutually orthonormal. We shall assume that the energy eigenvalues are ordered in the ascending sequence $E_0 < E_1 < \dots < E_n < \dots$. Since $\{\Psi_n\}$ form a basis set of the Hilbert space \mathcal{H} , Ψ^{trial} can be expanded as:

$$|\Psi^{trial}\rangle = \sum_n |\Psi_n\rangle \lambda_n. \quad (1.22)$$

The energy of the system in the state Ψ^{trial} is:

$$E[\Psi^{trial}] = \frac{\langle \sum_n \lambda_m^* \Psi_m | \hat{H} | \sum_n \lambda_n \Psi_n \rangle}{\langle \sum_m \lambda_m^* \Psi_m | \sum_n \lambda_n \Psi_n \rangle} = \frac{\sum_n E_n |\lambda_n|^2}{\sum_n |\lambda_n|^2}. \quad (1.23)$$

Subtracting E_0 on both sides, we obtain the following inequality:

$$E[\Psi^{trial}] - E_0 = \frac{\sum_n (E_n - E_0) |\lambda_n|^2}{\sum_n |\lambda_n|^2} \geq 0, \quad (1.24)$$

since $E_n \geq E_0$. The above expression is identically equal to zero if and only if $\Psi^{trial} = C\Psi_0$. The fact that two wavefunctions that differ by a multiplicative constant yield the same energy can be used to choose the constant C without restriction, so long as it is a finite real number. A convenient choice is to assume every wavefunction to be normalised:

$$\langle \Psi | \Psi \rangle = 1, \quad (1.25)$$

which implies $C = \frac{1}{\sqrt{\langle \Psi | \Psi \rangle}}$.

The variational principle proves that solving the time-independent Schrödinger equation can be transformed into a variational problem. The solutions are the wavefunctions that minimise the Lagrangian (1.20), which includes the constraint of square-integrability of the wavefunctions. Furthermore, the variational principle yields an effective method for setting an upper boundary to ground-state energy of the system. This is a precious tool that enables practical quantum-mechanical calculations. Approximate solutions of (1.20) can be found by continuous refinements of Ψ^{trial} towards minimising the Lagrangian \mathcal{L} .

1.5 The Hellmann-Feynman theorem

The derivatives of the energy with respect to an external parameter λ can be efficiently calculated with the Hellmann-Feynman theorem [17, 18], enunciated as follows. Given a wavefunction $\Psi(\lambda)$, which is a solution of the time-independent Schrödinger equation $H(\lambda)\Psi(\lambda) = E(\lambda)\Psi(\lambda)$, such as $\langle\Psi|\Psi\rangle = 1$, then the derivative of the energy with respect to an external parameter λ is

$$\frac{dE(\lambda)}{d\lambda} = \langle\Psi|\frac{d\hat{H}}{d\lambda}|\Psi\rangle. \quad (1.26)$$

To prove this statement, we should first calculate the full derivative of the energy, which is equal to:

$$\begin{aligned} \frac{dE(\lambda)}{d\lambda} &= \langle\frac{d\Psi}{d\lambda}|\hat{H}|\Psi\rangle + \langle\Psi|\hat{H}|\frac{d\Psi}{d\lambda}\rangle + \langle\Psi|\frac{d\hat{H}}{d\lambda}|\Psi\rangle \\ &= E(\lambda) \left[\langle\frac{d\Psi}{d\lambda}|\Psi\rangle + \langle\Psi|\frac{d\Psi}{d\lambda}\rangle \right] + \langle\Psi|\frac{d\hat{H}}{d\lambda}|\Psi\rangle. \end{aligned} \quad (1.27)$$

On the other hand, the condition of square-integrability (which is inherent of the solutions of the time-independent Schrödinger equation, as shown by the variational principle), implies that:

$$\frac{d}{d\lambda} \langle\Psi|\Psi\rangle = \langle\frac{d\Psi}{d\lambda}|\Psi\rangle + \langle\Psi|\frac{d\Psi}{d\lambda}\rangle = 0. \quad (1.28)$$

Introducing this result in (1.27), it follows that:

$$\frac{dE(\lambda)}{d\lambda} = \langle\Psi|\frac{d\hat{H}}{d\lambda}|\Psi\rangle, \quad (1.29)$$

and the theorem is proved. The Hellmann-Feynman theorem establishes that the

derivative of the energy with respect to an external parameter λ can be calculated as the expectation value of the derivative of the Hamiltonian operator. The derivative of Ψ with respect to λ does not intervene in the final value.

Chapter 2

Density functional theory

The atomic models of Thomson [19], Rutherford [20] and Bohr [21–23] preceded and contributed to the rise of rigorous studies of atomic-scale interactions based on quantum mechanics. A modern realisation of the atomic model carries the idea of a nucleus formed by protons and neutrons surrounded by electrons moving in undetermined trajectories of quantised energy levels. Electrostatic interactions between atoms make them bond to one another, forming complex molecules and crystals and giving rise to the solid state of matter. The study of matter and the interactions between electrons and nuclei is the subject of solid state physics. The most complete theory to date combines quantum-mechanical wavefunction methods, many-body statistics and thermodynamics to yield remarkable accuracy in calculations on a wide range of systems. Even so, solving the equations of quantum mechanics for the many-body problem is far from being trivial, and exact solutions are only possible for a few simple academic examples.

Much progress has been made towards practical methods for describing molecular interactions at the quantum level of theory. The Hartree-Fock method [24–27] significantly reduces the complexity of the problem by introducing single-electron

molecular orbitals. The many-electron wavefunction is then constructed as an antisymmetrised Slater determinant of the single-electron orbitals [28]. Post-Hartree-Fock methods refine the formalism by including an explicit treatment of the electron correlation energy. Successful approaches include Møller-Plesset (MP) perturbation theory [29], coupled cluster [30, 31], configuration interaction (CI) [32], or multi-configurational self-consistent field (MCSCF) approaches such as CASSCF [33].

The pioneering Thomas-Fermi model [34–36] was the precursor for a theory based entirely on the electronic density. It took almost 40 years until this view materialised in the form of the Hohenberg-Kohn theorems [37], leading to the development of the density functional theory (DFT). DFT is a formal and exact reformulation of the quantum-mechanical equations in terms of the electronic density as the defining element of the system. Computational electronic structure methods based on DFT [38–42] are able to predict an extensive range of electronic properties in all sorts of materials and compounds, which have helped in the development new practical applications [43, 44]. In this chapter we describe the foundations of electronic structure calculations based on density functional theory.

2.1 The molecular many-body problem

The first step to describe an atomic system is to construct the Hamiltonian operator that defines the interactions therein. Consider a molecule in isolation, not subject to any external electric or magnetic field, formed by \mathcal{M} point-like nuclei, with position vectors $\{\mathbf{R}_A\}$, masses $\{m_A\}$ and positive charges $\{Z_A\}$, and \mathcal{N} point-like electrons with position vectors $\{\mathbf{r}_i\}$, moving in trajectories around and near the nuclei. Neglecting other fundamental interactions other than electrostatic, the

Hamiltonian corresponding to the system of $\mathcal{M} + \mathcal{N}$ mutually interacting nuclei and electrons is:

$$\hat{H} = \hat{T}_n + \hat{T}_e + \hat{V}_{ee} + \hat{V}_{ne} + \hat{V}_{nn}, \quad (2.1)$$

where \hat{T}_n and \hat{T}_e are the kinetic energy terms due to the motion of the nuclei and the electrons, respectively, and \hat{V}_{ee} , \hat{V}_{ne} and \hat{V}_{nn} are the electron-electron, nucleus-electron and nucleus-nucleus interactions, respectively. Representing both \hat{H} and the wavefunctions $|\Psi_n\rangle$ in terms of the nuclear and electronic coordinates, $\langle\{\mathbf{r}_i\}, \{\mathbf{R}_A\}|$, the time-independent molecular Schrödinger equation is ¹:

$$\begin{aligned} H(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) \Psi_n(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) = & \quad (2.2) \\ \left[-\frac{1}{2m_A} \sum_{A=1}^{\mathcal{M}} \nabla_A^2 - \frac{1}{2} \sum_{i=1}^{\mathcal{N}} \nabla_i^2 + \sum_{A=1}^{\mathcal{M}} \sum_{B>A}^{\mathcal{M}} \frac{Z_A Z_B}{|\mathbf{R}_A - \mathbf{R}_B|} \right. \\ & \left. + \sum_{i=1}^{\mathcal{N}} \sum_{j>i}^{\mathcal{N}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{\mathcal{N}} \sum_{A=1}^{\mathcal{M}} \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \right] \Psi_n(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}) = \\ E_n \Psi_n(\{\mathbf{r}_i\}, \{\mathbf{R}_A\}). \end{aligned}$$

Electrons are subatomic, fundamental particles that belong to the category of fermions. Fermions are characterised for having a half-integer spin number. The spin quantum number originates from the relativistic interpretation of the equations of quantum mechanics, which arises naturally from Dirac's equation [12]. Each individual electron has a total spin number of modulus $\frac{1}{2}$, which can take two distinct values $\sigma = \frac{1}{2}$ (spin-up, \uparrow) or $\sigma = -\frac{1}{2}$ (spin-down, \downarrow). The notation based on the spin-coordinate vectors $\mathbf{x} = (\mathbf{r}, \sigma)$ incorporates the electronic spin into

¹In atomic units: $m_e = e = \hbar = 4\pi\epsilon_0 = 1$.

the formalism. The thermodynamics of fermionic systems are governed by the Fermi-Dirac statistics [45–47], which take into account that fermions are bound to be indistinguishable and that no two fermions can be found with the same set of quantum numbers. The exclusion principle formulated by Pauli [48] can be expressed as a mathematical constraint that requires each many-body wavefunction to be antisymmetric under any permutation of the electrons coordinates:

$$\begin{aligned} \Psi_n(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_j, \dots, \mathbf{x}_N; \mathbf{R}_1, \dots, \mathbf{R}_M) = & \quad (2.3) \\ - \Psi_n(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_i, \dots, \mathbf{x}_N; \mathbf{R}_1, \dots, \mathbf{R}_M). & \end{aligned}$$

The solutions of the time-independent Schrödinger equation for the molecular systems must obey the abovementioned constraint.

The complexity of the molecular Hamiltonian makes the analytic treatment of the equations practically impossible. The Born-Oppenheimer approximation [49] addresses this issue with a method that effectively decouples the electronic and nuclear degrees of freedom. As the name suggests, the method is not exact, however the error of the approximation is inversely proportional to the sum of the nuclear masses $M_{TOT} = \sum_{A=1}^M m_A$, which is an enormous number compared to the mass of the electrons². The Born-Oppenheimer approximation is also referred to as the *adiabatic approximation*, since the terms that couple the motion of the nuclei and the electrons are neglected. We refer the reader to Ref. [50] for a complete derivation of the Born-Oppenheimer approximation. In short, an observer in the centre of mass of the system will perceive that the heavy nuclei are static compared to the lighter electrons, which are in constant motion in undetermined trajectories around

²The mass of the proton is 1836.152 times the mass of the electron.

the nuclear centres. The observer could then decide that, instantaneously, to a good degree of approximation, the dynamical state of the system is described only by the electronic degrees of freedom, while the nuclear positions are mere parameters that define the system. The many-body wavefunctions, $\{\Psi_n\}$, can then be separated into the product of an electronic wavefunction and a nuclear wavefunction, denoted by $\{\Psi_n^{elec}\}$ and $\{\Psi_n^{nuc}\}$, respectively, as:

$$\Psi_n(\{\mathbf{x}_i\}, \{\mathbf{R}_A\}) \approx \Psi_n^{BO}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}) = \Psi_n^{elec}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}) \Psi_n^{nuc}(\{\mathbf{R}_A\}). \quad (2.4)$$

The electronic states, $\{\Psi_n^{elec}\}$, are found by solving the electronic part of the molecular Schrödinger equation:

$$\begin{aligned} H^{elec}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}) \Psi_n^{elec}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}) = & \quad (2.5) \\ \left[-\frac{1}{2} \sum_{i=1}^{\mathcal{N}} \nabla_i^2 + \sum_{i=1}^{\mathcal{N}} \sum_{j>i}^{\mathcal{N}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_{i=1}^{\mathcal{N}} \sum_{A=1}^{\mathcal{M}} \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} \right] \Psi_n^{elec}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}) \\ = E_n^{elec}(\{\mathbf{R}_A\}) \Psi_n^{elec}(\{\mathbf{x}_i\}; \{\mathbf{R}_A\}). \end{aligned}$$

While the nuclear states, $\{\Psi_n^{nuc}\}$, are the solutions of the molecular equation, which now takes the form:

$$\begin{aligned}
 H(\{\mathbf{R}_A\}) \Psi_n^{nuc}(\{\mathbf{R}_A\}) &= \\
 \left[E_n^{elec}(\{\mathbf{R}_A\}) - \frac{1}{2m_A} \sum_{i=1}^{\mathcal{M}} \nabla_A^2 + \sum_{A=1}^{\mathcal{M}} \sum_{B>A}^{\mathcal{M}} \frac{Z_A Z_B}{|\mathbf{R}_A - \mathbf{R}_B|} \right] \Psi_n^{nuc}(\{\mathbf{R}_A\}) \\
 &= E_n \Psi_n^{nuc}(\{\mathbf{R}_A\}).
 \end{aligned} \tag{2.6}$$

Systematic evaluation of the total energy for a range of values of $\{\mathbf{R}_A\}$ produces the potential energy surface (PES), or adiabatic surface. Sampling the PES is the basis for studies involving nuclear-electron interactions or for relaxing the geometry of a system to its equilibrium state [51].

2.2 The Hohenberg-Kohn theorems

Under the Born-Oppenheimer approximation, the \mathcal{N} electrons move in an external potential created by the \mathcal{M} static nuclei. The kinetic energy of the nuclei is assumed to be zero, while the nucleus-nucleus repulsion term is considered a constant (for a given set of nuclear coordinates) that can be overlooked. The problem is in this way reduced to finding the eigenstates of the electronic Hamiltonian:

$$H(\{\mathbf{r}_i\}) = -\frac{1}{2} \sum_{i=1}^{\mathcal{N}} \nabla_i^2 + \sum_{i=1}^{\mathcal{N}} \sum_{j>i}^{\mathcal{N}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} + \sum_i v(\mathbf{r}_i), \tag{2.7}$$

where:

$$v(\mathbf{r}) = \sum_{A=1}^{\mathcal{M}} \frac{Z_A}{|\mathbf{r} - \mathbf{R}_A|}. \tag{2.8}$$

is the external potential. The first two terms in Eq. (2.7) depend exclusively on the electronic degrees of freedom, and therefore are universal for any \mathcal{N} -electron system, regardless of the presence of the nuclei. The external potential $v(\mathbf{r})$, on the other hand, is not universal, since it depends parametrically on $\{\mathbf{R}_A\}$ and $\{Z_A\}$. As a result, in a given atomic system, the electronic wavefunction is uniquely determined by the functional form of $v(\mathbf{r})$ and by the total number of electrons, \mathcal{N} . Therefore, the total number of variational degrees of freedom is equal to $3\mathcal{N}$, that is, the three-dimensional components of the position vector of each electron.

Hohenberg and Kohn revisited the \mathcal{N} -electron problem from a different perspective and enunciated two important theorems that laid the foundations of density functional theory (DFT) [37]. DFT is an exact reformulation of the quantum theory in terms of the electronic density, $n(\mathbf{r})$, which is given the role of the defining element of the entire electronic system. The electronic density can be calculated via the recursive integration of the system's \mathcal{N} -electron wavefunction as:

$$n(\mathbf{r}) = \mathcal{N} \int d\mathbf{r}_2, \dots, d\mathbf{r}_{\mathcal{N}} |\Psi(\mathbf{r}, \mathbf{r}_2, \dots, \mathbf{r}_{\mathcal{N}})|^2. \quad (2.9)$$

The interesting thought behind DFT is that all the information about the system is retained in $n(\mathbf{r})$, even after the integration of $3(\mathcal{N} - 1)$ electron coordinates of the \mathcal{N} -electron wavefunction. Computationally-speaking, it is much easier to represent $n(\mathbf{r})$ in a three-dimensional computational grid than $\Psi(\{\mathbf{r}_i\})$ in $3\mathcal{N}$ dimensions [52]. In their seminal work, Hohenberg and Kohn conceived elegant proofs for their two theorems, which are reproduced hereafter.

First Hohenberg-Kohn theorem: the total energy of the system is a functional of the electronic density, $n(\mathbf{r})$, with the form:

$$E[n] = F[n] + \int d\mathbf{r} n(\mathbf{r}) v(\mathbf{r}), \quad (2.10)$$

where $F[n]$ is a universal functional formed by the addition of the kinetic energy and electron-electron energy terms, and where the external potential, $v(\mathbf{r})$, is uniquely determined by $n(\mathbf{r})$. The proof of this theorem can be done by *reductio ad absurdum*. Assuming that $n(\mathbf{r})$ is positive-definite and integrable to a constant integer number \mathcal{N} , and that the ground-state Ψ_0 is non-degenerate, consider another potential $v'(\mathbf{r})$ with ground-state wavefunction Ψ'_0 , that gives rise to the same electronic density $n(\mathbf{r})$. Using the variational principle of quantum mechanics (see Section 1.4):

$$\begin{aligned} E_0 &= \langle \Psi_0 | \hat{H} | \Psi_0 \rangle < \langle \Psi'_0 | \hat{H} | \Psi'_0 \rangle = \langle \Psi'_0 | \hat{H}' | \Psi'_0 \rangle + \langle \Psi'_0 | \hat{H} - \hat{H}' | \Psi'_0 \rangle \\ &= E'_0 + \int d\mathbf{r} n(\mathbf{r}) [v(\mathbf{r}) - v'(\mathbf{r})]. \end{aligned} \quad (2.11)$$

At the same time:

$$\begin{aligned} E'_0 &= \langle \Psi'_0 | \hat{H}' | \Psi'_0 \rangle < \langle \Psi_0 | \hat{H}' | \Psi_0 \rangle = \langle \Psi_0 | \hat{H} | \Psi_0 \rangle + \langle \Psi_0 | \hat{H}' - \hat{H} | \Psi_0 \rangle \\ &= E_0 + \int d\mathbf{r} n(\mathbf{r}) [v'(\mathbf{r}) - v(\mathbf{r})]. \end{aligned} \quad (2.12)$$

Adding these two equations together leads to a result that is false by definition (*absurdum*):

$$E_0 + E'_0 < E_0 + E'_0, \quad (2.13)$$

and so the first theorem is demonstrated. The first Hohenberg-Kohn theorem establishes a bijective mapping between the external potential $v(\mathbf{r})$ and the electronic density $n(\mathbf{r})$, and proves that if one of them is determined, the other one is automatically determined too. The wavefunction is determined uniquely by $v(\mathbf{r})$, and therefore it is also uniquely determined by $n(\mathbf{r})$. The electronic density is in such way placed in the position of being the element that can completely and uniquely describe the entire system of \mathcal{N} interacting electrons.

Second Hohenberg-Kohn theorem: for any trial electronic density $n^{trial}(\mathbf{r})$:

$$E[n^{trial}] \geq E[n_0], \quad (2.14)$$

where $n_0(\mathbf{r})$ is the true ground state energy corresponding to Ψ_0 . The proof relies on the first Hohenberg-Kohn theorem, which demonstrates that $n(\mathbf{r})$ defines the external potential $v(\mathbf{r})$ and also the wavefunction Ψ . From the variational principle of quantum mechanics:

$$E[n^{trial}] = \langle \Psi^{trial} | \hat{H} | \Psi^{trial} \rangle \geq \langle \Psi_0 | \hat{H} | \Psi_0 \rangle = E[n_0]. \quad (2.15)$$

The equality is true only when $n^{trial}(\mathbf{r}) = n_0(\mathbf{r})$. The second Hohenberg-Kohn theorem is the extension of the variational principle of quantum mechanics (see Section 1.4) to the electronic density. The Lagrangian (1.20) can then be written as:

$$\mathcal{L}[n] = \langle \Psi[n] | \hat{H} | \Psi[n] \rangle - E(\langle \Psi[n] | \Psi[n] \rangle - 1). \quad (2.16)$$

Imposing the condition of $n(\mathbf{r})$ being and extremal of \mathcal{L} , gives the following opti-

minimisation scheme:

$$\frac{\delta \mathcal{L}}{\delta n} = 0 \implies \frac{\delta \mathcal{L}}{\delta \langle \Psi |} = 0 \implies \hat{H} |\Psi\rangle - E |\Psi\rangle = 0. \quad (2.17)$$

The above relation is the basis for self-consistent DFT calculations, where a trial electronic density is iteratively optimised in order to minimise the total energy functional.

The reformulation of the electronic problem proposed by the Hohenberg-Kohn theorems is exact in principle, although concerns about the v -representability of the electronic density have been raised [53, 54]. The class of v -representable electronic densities are those $n(\mathbf{r})$ such as there exist an external potential $v(\mathbf{r})$ that can generate an \mathcal{N} -electron ground-state wavefunction Ψ that produces $n(\mathbf{r})$. The issue appears by considering that a plausible electronic density $n(\mathbf{r})$ that does not correspond to any wavefunction Ψ or external potential $v(\mathbf{r})$ can be constructed. An alternative proof by Levy [53] eliminates the need for non-degenerate states and introduces the wider class of \mathcal{N} -representable electronic densities. The class of \mathcal{N} -representable densities are those that can be constructed from a \mathcal{N} -electron antisymmetric wavefunction Ψ . Given the Hamiltonian operator \hat{H} , the energy of the system is determined via the following constrained minimisation problem:

$$E[n] = \min_{\Psi \rightarrow n} \langle \Psi | \hat{H} | \Psi \rangle. \quad (2.18)$$

where Ψ are the wavefunctions that produce the electronic density $n(\mathbf{r})$.

2.3 Finite-temperature DFT

Soon after the introduction of DFT, Mermin [55] proved that the Hohenberg-Kohn theorems could be generalised to the *grand canonical ensemble*, defined by a constant temperature, \mathcal{T} , constant volume, \mathcal{V} , and constant chemical potential, μ . In what follows, we will be interested in a different thermodynamical ensemble, termed the *canonical ensemble*. The system of \mathcal{N} interacting electrons is in a mixed state, formed by a collection of microstates $\{\psi_u\}$ with energies $\{E_u\}$, constant volume \mathcal{V} and number of electrons \mathcal{N} , which are in thermal equilibrium with a reservoir at a constant temperature \mathcal{T} . The probability of the microstate ψ_u to be observed is p_u , with $0 \leq p_u \leq 1$ and $\sum_u p_u = 1$. The system is fully described by the density operator $\hat{\rho} = \sum_u p_u |\psi_u\rangle \langle\psi_u|$, and constrained to meet the conditions:

$$E = \text{tr} \left[\hat{\rho} \hat{H} \right], \quad (2.19)$$

and

$$\mathcal{N} = \text{tr} \left[\hat{\rho} \hat{N} \right], \quad (2.20)$$

where \hat{N} is the number of particles operator. There are, in principle, infinite re-organisations of the probability distribution $\{p_u\}$ that meet the conditions (2.19) and (2.20). Therefore, one must take a holistic approach, and assume that the probability distribution that best describes the current state is the one that maximises the frequency $W[\{p_u\}]$ in which a given distribution appears. Such distribution is the one that maximises the entropy associated to $\{p_u\}$, defined as:

$$S[\{p_u\}] = -k_B \text{tr} [\hat{\rho} \ln \hat{\rho}] = -k_B \sum_u p_u \ln p_u, \quad (2.21)$$

which yields the values [47]:

$$p_u = \frac{e^{-\beta E_u}}{\sum_u e^{-\beta E_u}}, \quad (2.22)$$

where $\beta = [k_B \mathcal{T}]^{-1}$ and k_B is the Boltzmann constant. The resulting density operator becomes:

$$\hat{\rho} = \frac{e^{-\beta \hat{H}}}{\text{tr} [e^{-\beta \hat{H}}]}. \quad (2.23)$$

The macroscopic state function in the canonical ensemble is the Helmholtz free energy, $A(\mathcal{N}, \mathcal{V}, \mathcal{T})$, defined as:

$$A(\mathcal{N}, \mathcal{V}, \mathcal{T}) = -k_B \mathcal{T} \ln Z(\mathcal{N}, \mathcal{V}, \mathcal{T}), \quad (2.24)$$

where $Z(\mathcal{N}, \mathcal{V}, \mathcal{T})$ is the partition function:

$$Z(\mathcal{N}, \mathcal{V}, \mathcal{T}) = \sum_u e^{-\beta E_u}. \quad (2.25)$$

The dependence of $Z(\mathcal{N}, \mathcal{V}, \mathcal{T})$ on \mathcal{N} and \mathcal{V} is contained in the energy values $\{E_u\}$ of each microstate. Using the density matrix formalism, the Helmholtz free energy becomes:

$$A_{\mathcal{T}} = \text{tr} \left[\hat{\rho} \left(\hat{H} + \frac{1}{\beta} \ln \hat{\rho} \right) \right] = E - \mathcal{T} S. \quad (2.26)$$

The dependence of \mathcal{N} and \mathcal{V} has been dropped for simplicity. Parametric dependence of $A_{\mathcal{T}}$ on \mathcal{T} will be assumed from now on.

Extrapolated to the canonical ensemble, the first Hohenberg-Kohn theorem es-

establishes that the Helmholtz free energy functional can be exactly re-written as:

$$A_{\mathcal{T}}[n] = F_{\mathcal{T}}[n] + \int d\mathbf{r} n(\mathbf{r}) v(\mathbf{r}), \quad (2.27)$$

where, using the Levy formalism:

$$F_{\mathcal{T}}[n] = \min_{\hat{\rho} \rightarrow n} \text{tr} \left[\hat{\rho} \left(\hat{T} + \hat{V}_{ee} + \frac{1}{\beta} \ln \hat{\rho} \right) \right], \quad (2.28)$$

is a universal functional of the \mathcal{N} -representable electronic density $n(\mathbf{r})$, and where the external potential $v(\mathbf{r})$ is uniquely determined by $n(\mathbf{r})$. This statement can be proven by *reductio ad absurdum* [55].

In a similar fashion, the second Hohenberg-Kohn theorem is the extension of the variational principle to the Helmholtz free energy:

$$\begin{aligned} A_{\mathcal{T}}[n^{trial}] = & \quad (2.29) \\ & \min_{\hat{\rho}^{trial} \rightarrow n^{trial}} \text{tr} \left[\hat{\rho}^{trial} \left(\hat{T} + \hat{V}_{ee} + \frac{1}{\beta} \ln \hat{\rho}^{trial} \right) \right] + \int d\mathbf{r} n^{trial}(\mathbf{r}) v(\mathbf{r}) \geq \\ & \min_{\hat{\rho}_0 \rightarrow n_0} \text{tr} \left[\hat{\rho}_0 \left(\hat{T} + \hat{V}_{ee} + \frac{1}{\beta} \ln \hat{\rho}_0 \right) \right] + \int d\mathbf{r} n_0(\mathbf{r}) v(\mathbf{r}) = A_{\mathcal{T}}[n_0]. \end{aligned}$$

The proof [55] is identical to the zero-temperature case shown in the previous section. The solution to the eigenvalue equation $\hat{H} |\Psi\rangle = E |\Psi\rangle$ can be transformed into a minimisation problem of the Lagrangian:

$$\mathcal{L}_{\mathcal{T}}[n] = \min_{\hat{\rho} \rightarrow n} \left\{ \left(\text{tr} [\hat{\rho} \hat{H}] - \mathcal{T} \text{tr} [\hat{\rho} \ln \hat{\rho}] \right) - E (\langle \Psi[\hat{\rho}] | \Psi[\hat{\rho}] \rangle - 1) \right\} \quad (2.30)$$

such as

$$\frac{\delta \mathcal{L}_{\mathcal{T}}[n]}{\delta n} = 0 \implies \begin{cases} \frac{\delta \mathcal{L}_{\mathcal{T}}}{\delta \langle \Psi_u |} = 0 \\ \frac{\partial \mathcal{L}_{\mathcal{T}}}{\partial p_u} = 0 \end{cases} \implies [\hat{\rho}, \hat{H}] = 0. \quad (2.31)$$

And hence the finite-temperature ground-state density operator $\hat{\rho}$ is the solution to the time-independent Schrödinger equation (1.17).

From Eq. (2.26), it can be inferred that at zero-temperature, $A_{\mathcal{T}=0}[n] = E[n]$. If the zero-temperature ground-state is non-degenerate, the probabilities $\{p_u\}$ become identically equal to zero except for the ground-state wavefunction, and the system is found in a pure state where the entropy term is identically zero. If, on the contrary, the zero temperature ground-state is degenerate, it remains true that $\mathcal{T} = 0 \implies A_{\mathcal{T}=0}[n] = E[n]$. However, the probabilities $\{p_u\}$ are not integer numbers, and the entropy term $S[\{p_u\}]$ is not equal to zero [47].

2.4 Kohn-Sham mapping

The Hohenberg-Kohn theorems demonstrate that, at the ground-state, the whole system of \mathcal{N} interacting electrons can be uniquely described by their electronic density, $n(\mathbf{r})$. Although this is a tremendous advance in our understanding of the \mathcal{N} -electron problem, it yields very little applicability for practical calculations, because the exact form of the universal functional $F_{\mathcal{T}}[n]$ is unknown.

The theory developed by Kohn and Sham [56] establishes an exact correspondence between the system of \mathcal{N} interacting electrons and a fictitious system of \mathcal{N} non-interacting electrons moving in an effective potential. The impressive feature of this correspondence is that the non-interacting and interacting systems are as-

sociated to the same ground-state electronic density. The original proof by Kohn and Sham was developed for the canonical ensemble at zero temperature. The extension to $\mathcal{T} > 0$ is very well described by Parr and Yang [36]. Here, we will develop the Kohn-Sham mapping for the finite-temperature case, and return to the zero-temperature scheme simply by setting $\mathcal{T} = 0$.

The form of the universal functional $F_{\mathcal{T}}[n]$ of the interacting system can be manipulated to be written exactly as:

$$\begin{aligned} F_{\mathcal{T}}[n] &= T[n] + E_{ee}[n] - \mathcal{T}S[n] \\ &= T_s[n] + E_H[n] + E_{xc}[n] - \mathcal{T}S_s[n], \end{aligned} \quad (2.32)$$

where $T_s[n]$, $E_H[n]$ and $S_s[n]$ are the kinetic, the Hartree energy and the entropy terms of the system of non-interacting electrons, respectively, and $E_{xc}[n]$ is the exchange-correlation term, which by definition is equal to:

$$E_{xc}[n] = (T[n] - T_s[n]) + (E_{ee}[n] - E_H[n]) - \mathcal{T}(S[n] - S_s[n]). \quad (2.33)$$

Assuming for now that the \mathcal{N} -electron wavefunctions are orthonormal, following the variational principle of DFT, the ground state can be found as in Eq. (1.18) by imposing $\delta\mathcal{L}_{\mathcal{T}}[n] = 0$, where $\mathcal{L}_{\mathcal{T}}$ is the Lagrangian:

$$\mathcal{L}_{\mathcal{T}}[n] = E[n] - \mathcal{T}S[n] - \mu \left(\int d\mathbf{r} n(\mathbf{r}) - N \right). \quad (2.34)$$

The parameter μ is an undetermined Lagrange multiplier that ensures conservation of the number of electrons. Once the ground-state has been found, μ is equal to the

Fermi level. The condition $\delta\mathcal{L}_{\mathcal{T}}[n] = 0$ gives:

$$\frac{\delta T_s}{\delta n} + v_{KS}(\mathbf{r}) - \mathcal{T} \frac{S_s}{\delta n} = \mu, \quad (2.35)$$

where $v_{KS}(\mathbf{r})$ is the effective potential:

$$v_{KS}(\mathbf{r}) = \int d\mathbf{r}' \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + v_{xc}(\mathbf{r}) + v(\mathbf{r}), \quad (2.36)$$

with $v_{xc}(\mathbf{r}) = \frac{\delta E_{xc}}{\delta n}$. Equation (2.35) could have been derived from a different starting point, by assuming that the \mathcal{N} electrons were indeed non-interacting but subject to an effective potential $v_{KS}(\mathbf{r})$ generated by an external source. The interpretation of this result is that the problem of \mathcal{N} interacting electrons subject to the external potential $v(\mathbf{r})$ is formally equivalent to the problem of \mathcal{N} non-interacting electrons under the effective potential $v_{KS}(\mathbf{r})$. According to the first Hohenberg-Kohn theorem, both systems are associated to the same electronic density, $n(\mathbf{r})$.

The many-electron wavefunction of the problem of non-interacting electrons can be constructed as an antisymmetrised Slater determinant that combines the products of one-electron molecular orbitals $\{\psi_i\}$ [57]. The one-electron orbitals are the eigenstates of one-electron Hamiltonians $\{\hat{h}_i\}$, such as:

$$\hat{h}_i |\psi_i\rangle = \epsilon_i |\psi_i\rangle, \quad (2.37)$$

where $\{\epsilon_i\}$ are the energy eigenvalues of the one-electron orbitals. In the real-space representation, \hat{h}_i takes the form:

$$h_i(\mathbf{r}) = -\frac{1}{2}\nabla_{\mathbf{r}}^2 + v_{KS}(\mathbf{r}). \quad (2.38)$$

The many-electron Hamiltonian operator is the sum of the one-electron Hamiltonians:

$$\hat{H} = \sum_i \hat{h}_i. \quad (2.39)$$

Each orbital ψ_i is occupied with a probability f_i , with $0 \leq f_i \leq 1$. Following the *aufbau* principle, the lower-energy states must have a higher probability to be occupied. On the other hand, the fixed number of electrons in the system imposes the constraint:

$$\sum_i f_i = \mathcal{N}. \quad (2.40)$$

In the free electron model, the occupancies of each state are given by the Fermi-Dirac distribution:

$$f_i(\epsilon_i) = \left(1 + \exp \left[\frac{\epsilon_i - \mu}{k_B \mathcal{T}} \right] \right)^{-1}. \quad (2.41)$$

The Fermi-Dirac distribution is derived from the principle that the electrons are fermions of spin $\frac{1}{2}$. As mentioned earlier, this condition implies that the electrons must be indistinguishable and that no two electrons can be found in the same state simultaneously. Associated to the Fermi-Dirac distribution there is the following entropy term:

$$S_s[\{f_i\}] = -k_B \sum_i [f_i \ln f_i + (1 - f_i) \ln (1 - f_i)]. \quad (2.42)$$

The system of non-interacting electrons is nevertheless in a mixed quantum state, formed by microstates $\{\psi_i\}$ with probabilities $\{f_i\}$. Therefore, the system can only be described in terms of a density operator. The one-particle density

operator, $\hat{\rho}$, defined as:

$$\hat{\rho} = \sum_i f_i |\psi_i\rangle \langle \psi_i|, \quad (2.43)$$

fully describes the fictitious system of \mathcal{N} non-interacting electrons, where the condition (2.40) implies:

$$\text{tr} [\hat{\rho}] = \mathcal{N}. \quad (2.44)$$

The representation of $\hat{\rho}$ in real-space is known as the one-particle density matrix, $\rho(\mathbf{r}, \mathbf{r}')$, defined as:

$$\rho(\mathbf{r}, \mathbf{r}') = \langle \mathbf{r} | \hat{\rho} | \mathbf{r}' \rangle = \sum_i f_i \langle \mathbf{r} | \psi_i \rangle \langle \psi_i | \mathbf{r}' \rangle = \sum_i f_i \psi_i(\mathbf{r}) \psi_i^*(\mathbf{r}'). \quad (2.45)$$

The collection of diagonal elements of $\rho(\mathbf{r}, \mathbf{r}')$ forms the electronic density, $n(\mathbf{r})$:

$$n(\mathbf{r}) = \sum_i f_i \psi_i(\mathbf{r}) \psi_i^*(\mathbf{r}). \quad (2.46)$$

In close-shell calculations, the one-electron orbitals of opposite spin are identical³: $\psi_i = \psi_i^\uparrow = \psi_i^\downarrow$. On the other hand, in open-shell calculations the two spin channels are considered to be different: $\psi_i^\uparrow \neq \psi_i^\downarrow$. The one-electron Hamiltonian is also assumed to be different for each spin channel: $\hat{h}_i^\uparrow \neq \hat{h}_i^\downarrow$. Different spin channels can be established for the one-particle density operator ($\hat{\rho}^\uparrow \neq \hat{\rho}^\downarrow$), density matrix ($\rho^\uparrow(\mathbf{r}, \mathbf{r}') \neq \rho^\downarrow(\mathbf{r}, \mathbf{r}')$), and electronic density ($n^\uparrow(\mathbf{r}) \neq n^\downarrow(\mathbf{r})$), such as:

³Notation: $\psi_i^\uparrow(\mathbf{r}) = \psi_i(\mathbf{r}, \uparrow)$, and $\psi_i^\downarrow(\mathbf{r}) = \psi_i(\mathbf{r}, \downarrow)$

$$\hat{\rho} = \hat{\rho}^\uparrow + \hat{\rho}^\downarrow, \quad (2.47)$$

$$\rho(\mathbf{r}, \mathbf{r}') = \rho^\uparrow(\mathbf{r}, \mathbf{r}') + \rho^\downarrow(\mathbf{r}, \mathbf{r}'), \quad (2.48)$$

$$n(\mathbf{r}) = n^\uparrow(\mathbf{r}) + n^\downarrow(\mathbf{r}). \quad (2.49)$$

Naturally, the amount of computation doubles in open-shell calculations compared to closed-shell calculations.

The hermiticity of \hat{H} implies that the Kohn-Sham states $\{\psi_i\}$ are orthogonal to each other if they belong to different eigenvalues $\{\epsilon_i\}$. If degeneracy is present, Kohn-Sham states corresponding to the same eigenvalue can be constrained to be orthonormal to each other. An unconstrained minimisation of the Lagrangian:

$$\mathcal{L}_{\mathcal{T}}[n] = E[n] - \mathcal{T}S[n] - \mu \left[\int d\mathbf{r} n(\mathbf{r}) - \mathcal{N} \right] - \sum_{ij} \lambda_{ij} [\langle \psi_i | \psi_j \rangle - \delta_{ij}], \quad (2.50)$$

where $\{\lambda_{ij}\}$ are undefined Lagrange multipliers that account for orthonormality of $\{\psi_i\}$, gives rise to orthonormal solutions of the Kohn-Sham equations. The variational principle of DFT ensures that $\mathcal{L}_{\mathcal{T}}[n^{trial}] \geq \mathcal{L}_{\mathcal{T}}[n_0]$ for any trial electronic density, $n^{trial}(\mathbf{r})$.

Although the density of the interacting and non-interacting \mathcal{N} -electron problems is identical, the energy functional is different. The energy functional that corresponds to the non-interacting system is referred to as the band-structure energy:

$$E_{BS}[n] = \text{tr} [\hat{\rho} \hat{H}] = \sum_i f_i \epsilon_i. \quad (2.51)$$

To obtain the energy of the interacting system, the band-structure energy must be corrected to take into account the double-counting of the Coulomb electron-electron repulsion and corrections to the exchange-correlation energy, resulting in the functional [58]:

$$E[n] = E_{BS}[n] - \frac{1}{2} \int_{\mathcal{V}} d\mathbf{r} d\mathbf{r}' \frac{n(\mathbf{r}) n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + \int_{\mathcal{V}} d\mathbf{r} n(\mathbf{r}) v_{xc}(\mathbf{r}) - E_{xc}[n]. \quad (2.52)$$

The Kohn-Sham mapping for zero-temperature DFT can be derived almost straightforward by setting $\mathcal{T} = 0$. Under this condition, the Lagrangian (2.34) simplifies to:

$$\mathcal{L}[n(\mathbf{r})] = E[n(\mathbf{r})] - \mu \left[\int_{\mathcal{V}} d\mathbf{r} n(\mathbf{r}) - \mathcal{N} \right], \quad (2.53)$$

and its derivative becomes:

$$\frac{\delta T_s}{\delta n} + v_{KS}(\mathbf{r}) = \mu, \quad (2.54)$$

which is the original expression obtained by Kohn and Sham. Also, at $\mathcal{T} = 0$, the Fermi-Dirac distribution collapses into the Heaviside step function (see Fig. 2.1), and $\{f_i\}$ are assigned integer numbers:

$$f_i = \begin{cases} 1 & \epsilon_i \leq \mu \\ 0 & \epsilon_i > \mu. \end{cases} \quad (2.55)$$

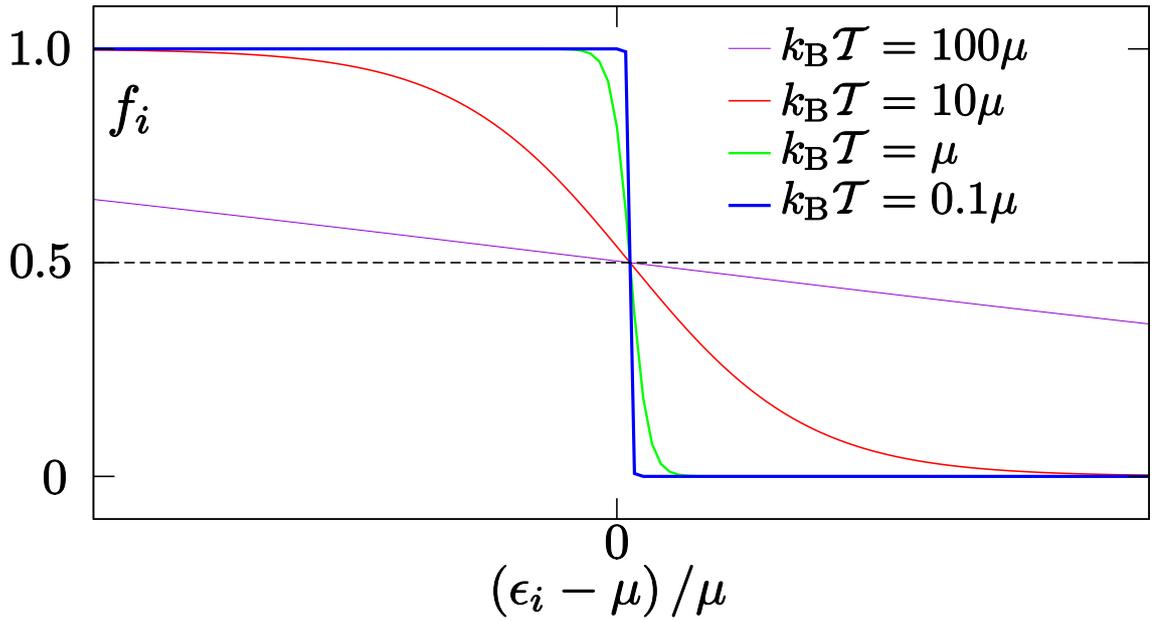


Figure 2.1: Fermi-Dirac distribution for different values of the electronic temperature.

At zero temperature, μ is an estimation of the chemical potential, and marks the barrier between occupied and unoccupied levels of the system. Since the occupancies are identically either zero or one, it follows that the one-particle density operator is idempotent:

$$\hat{\rho}^2 = \hat{\rho}. \quad (2.56)$$

According to this result, the system of \mathcal{N} non-interacting electrons is in a pure state at $\mathcal{T} = 0$.

2.5 The exchange-correlation functional

With the Kohn-Sham mapping, all the complexity inherent to the system of \mathcal{N} interacting electrons is shifted to the exchange-correlation functional, of which its exact form is unknown. The great success of the Kohn-Sham model is due to the fact that accurate approximations of $E_{xc}[n]$ can be developed for calculations on a wide range of systems. The quest towards more refined and consistent approximations of $E_{xc}[n]$ has led to a large collection of exchange-correlation functionals [59]. Eventually, the choice of the form of $E_{xc}[n]$ is motivated by the kind of system under simulation and the physical properties that are under scrutiny. A universal exchange-correlation functional is not yet available.

Numerous methods separate $E_{xc}[n]$ into the exchange and correlation parts, as in $E_{xc}[n] = E_x[n] + E_c[n]$. The Hartree-Fock exchange energy is known analytically in terms of the one-electron orbitals and their occupancies [60]:

$$E_x^{HF}[\{f_i\}, \{\psi_i\}] = -\frac{1}{2} \sum_{ij} f_i f_j \int \int d\mathbf{r} d\mathbf{r}' \frac{\psi_i^*(\mathbf{r}) \psi_j^*(\mathbf{r}') \psi_i(\mathbf{r}') \psi_j(\mathbf{r})}{|\mathbf{r} - \mathbf{r}'|}, \quad (2.57)$$

but no exact expression of $E_x[n]$ as a functional of $n(\mathbf{r})$ is known.

In the local density approximation (LDA), the total exchange-correlation energy is calculated by extrapolating the function $\varepsilon_{xc}^{LDA}(n)$, defined as the energy per electron in the uniform electron gas, to every infinitesimal volume $d\mathbf{r}$ containing $d\mathbf{r} n(\mathbf{r})$ electrons:

$$E_{xc}^{LDA}[n] = \int d\mathbf{r} n(\mathbf{r}) \varepsilon_{xc}^{LDA}(n). \quad (2.58)$$

The exchange-correlation potential calculated after differentiation of (2.58) pro-

duces:

$$v_{xc}^{LDA}(\mathbf{r}) = \frac{\delta E_{xc}^{LDA}}{\delta n} = \varepsilon_{xc}^{LDA}(n) + n(\mathbf{r}) \frac{\delta \varepsilon_{xc}^{LDA}}{\delta n}. \quad (2.59)$$

The function $\varepsilon_{xc}^{LDA}(n)$ can also be separated into the exchange and correlation parts, $\varepsilon_{xc}^{LDA}(n) = \varepsilon_x^{LDA}(n) + \varepsilon_c^{LDA}(n)$. The exchange part is calculated following the Thomas-Fermi-Dirac model [36] as:

$$\varepsilon_x^{LDA}(n) = -\frac{3}{4} \left(\frac{3}{\pi} \right)^{1/3} n^{1/3}(\mathbf{r}), \quad (2.60)$$

which, integrated as in (2.58) results in the LDA exchange energy:

$$E_x^{LDA}[n] = -\frac{3}{4} \left(\frac{3}{\pi} \right)^{1/3} \int d\mathbf{r} n^{4/3}(\mathbf{r}). \quad (2.61)$$

LDA functionals yield accurate results in calculations on systems where variations of the electronic density are smooth and uniformity is the rule. The correlation energy of the uniform electron gas has been studied analytically by Gell-Mann and Bruecker [61] and numerically by Ceperley and Alder [62]. Successful LDA functionals include the Ceperley-Alder-Perdew-Zunger (CAPZ) [63] and the Vosko-Wilk-Nusair (VWN) [64], which are different parametrisations of the LDA exchange and correlation terms.

Non-uniform systems tend to display a rapidly-varying electronic density. Non-local exchange-correlation functionals include the dependency of the exchange-correlation energy on the gradient of the electronic density, $\nabla n(\mathbf{r})$. The family of exchange-correlations produced in this way is termed the generalised gradient approximation (GGA) [65]. The generalised form of a GGA functional corresponds to the real-space integration of a function φ_{xc}^{GGA} that depends on the electronic

density and its gradient:

$$E_{xc}^{GGA}[n] = \int d\mathbf{r} \varphi_{xc}^{GGA}(n(\mathbf{r}), \nabla n(\mathbf{r})). \quad (2.62)$$

Functionals that belong to this family are the class of functionals developed by Becke (B86 [66] and B88 [67]), Lee-Yang-Parr (LYP) [68], Perdew-Wang (PW91) [69], Perdew-Burke-Ernzerhof (PBE) [70], and revised-PBE (RPBE) [71]. Another class of functionals, termed Meta-GGA functionals, include higher order derivatives of the electronic density:

$$E_{xc}^{meta-GGA}[n] = \int d\mathbf{r} \varphi_{xc}^{meta-GGA}(n(\mathbf{r}), \nabla n(\mathbf{r}), \nabla^2 n(\mathbf{r}), \dots). \quad (2.63)$$

Examples of meta-GGA parametrisations are the functionals of Perdew-Kurth-Zupan-Blaha (PKZB) [72] and Tao-Perdew-Staroverov-Scuseria (TPSS) [73].

A different class of functionals adds a fraction of the exact Hartree-Fock exchange energy to improve the accuracy of the calculation. These are called hybrid functionals, amongst which the B3LYP functional of Becke [74]:

$$E_{xc}^{B3LYP}[n] = (1 - \alpha_0 - \alpha_x) E_x^{LDA}[n] + \alpha_0 E_x^{HF}[\{f_i\}, \{\psi_i\}] \quad (2.64)$$

$$+ \alpha_x E_x^{B88}[n] + (1 - \alpha_c) E_c^{VWN}[n] + \alpha_c E_c^{LYP}[n], \quad (2.65)$$

where α_0 , α_x and α_c are chosen according to the parametrisation of Stevens *et al.* [75], must be highlighted.

According to (2.33), there is a parametric dependence of $E_{xc}[n]$ on the electronic temperature \mathcal{T} . However, to the best of our knowledge, very few approaches

take into account the effect of \mathcal{T} in the exchange-correlation term [76]. In finite-temperature Kohn-Sham DFT approaches, the parametric dependence of $E_{xc}[n]$ on \mathcal{T} is usually relegated to its effect on the Kohn-Sham orbitals and their occupancies, which gets transferred into the electronic density $n(\mathbf{r})$. The exchange-correlation term is then calculated using one of the aforementioned methods.

2.6 Basis set functions

The wavefunction space \mathcal{H} was defined in Section 1.1 as the Hilbert space of square-integrable wavefunctions capable of describing a quantum system. A subset $\mathcal{B} \in \mathcal{H}$ is said to be a basis set of \mathcal{H} if any wavefunction included in \mathcal{H} can be written as a linear combination of the elements in \mathcal{B} . The wavefunction space \mathcal{H} has infinite number of dimensions, and so \mathcal{B} must contain an infinite number of elements in order to span \mathcal{H} completely.

In Kohn-Sham DFT approaches, each one-electron molecular orbital $\psi_i \in \mathcal{H}$ can be expanded as:

$$|\psi_i\rangle = \sum_{b=1}^{\infty} |\chi_b\rangle c_{bi}, \quad (2.66)$$

where $\{\chi_b\} \in \mathcal{B}$ and $\{c_{bi}\}$ are the expansion coefficients. An expansion of the kind of (2.66) will carry an unbearable cost, since an infinite number of $\{c_{bi}\}$ must be computed and stored. A more tractable approach to DFT calculations must be found, and the answer, inevitably, must be to reduce the dimensionality of the problem by using a finite number of basis set functions.

A reasonable accuracy can be obtained in most practical calculations if only a subspace $\mathcal{H}' \subset \mathcal{H}$, is sampled, where \mathcal{H}' has a finite number of dimensions,

$0 < \varsigma < \infty$. The condition of completeness implies that $\mathcal{H}' \rightarrow \mathcal{H}$ when $\varsigma \rightarrow \infty$. The Kohn-Sham states can then be expanded as a linear combination of the basis set functions $\{\chi_b\} \in \mathcal{B}' \subset \mathcal{H}'$ as:

$$|\psi_i\rangle = \sum_{b=1}^{N_b} |\chi_b\rangle c_{bi}, \quad (2.67)$$

where N_b is the number of basis set functions in \mathcal{B}' , with $\varsigma \leq N_b < \infty$. As a result, $\{\psi_i\} \in \mathcal{H}'$. The degree of completeness of a basis set is judged based on the value of ς , which defines the dimensionality of \mathcal{H}' . A basis set capable of spanning a subspace with larger ς is expected to yield more accurate results. This is an ideal situation in which the accuracy can be systematically controlled by a single parameter, ς , which controls the dimensionality of \mathcal{H} . Very often, this task is accomplished by adding new basis set functions that are linearly-independent of all the rest. As one might imagine, adding new basis set functions is not cost-free. A larger basis set will result in a larger number of expansion coefficients $\{c_{bi}\}$, which will require more time, memory and power to be computed correctly [57, 77].

The choice of basis set is therefore crucial for the accuracy and performance of a given DFT algorithm, to the point that many available programs are defined by the type of basis set that they use. Generally speaking, the most successful choices of basis sets are those that allow systematic improvement of their degree of completeness. Basis sets formed by Gaussian-type orbitals (GTOs) [78–83] can be analytically manipulated, which makes operations such as derivatives or integrals extremely efficient. Programs such as GAUSSIAN [84] or NWCHEM [85] use multiple- ζ Gaussian basis sets for accurate DFT calculations. Basis sets of Slater-type orbitals (STOs) [86] have a different analytical form and are used in programs

such as ADF [87]. Daubechies wavelets are used in the BIGDFT program [88] with optimum results. Plane-wave basis sets [40] are the typical choice for simulating periodic systems such as crystals, surfaces and wires. Programs like ABINIT [89], CASTEP [90], GPAW [91], QUANTUM EXPRESSO [92], S/PHI/NX [93], VASP [94, 95] or Wien2K [96], use plane-wave basis sets. Plane-waves are also used in Gaussian-basis-set methods to calculate the Hartree energy in a computationally-efficient manner [97], for example in the QUICKSTEP program [98].

2.7 Periodic systems and plane waves

Systems that extend periodically in space, such as crystalline structures and surfaces, can be very well described with a basis set of plane-waves. A complete study of periodic systems and their characterisation can be found in Ref. [99]. Also, the review by Payne *et al.* [40] offers a comprehensive review of plane-wave methods applied to Kohn-Sham density functional theory and periodic systems. We refer the reader to these sources for further reading. In this section, we summarise the main features of DFT calculations on periodic systems using plane-waves.

A given atomic system is said to be periodic with the periodicity of a Bravais lattice if the Kohn-Sham potential meets the condition:

$$v_{KS}(\mathbf{r}) = v_{KS}(\mathbf{r} + \mathbf{R}), \quad (2.68)$$

where $\mathbf{R} = n_1\mathbf{a}_1 + n_2\mathbf{a}_2 + n_3\mathbf{a}_3$ is a lattice vector in real space, and n_1, n_2 and n_3 are integer numbers. The vectors $\mathbf{a}_1, \mathbf{a}_2$ and \mathbf{a}_3 are the primitive lattice vectors that define the volume of the periodic cell, $V_{cell} = \mathbf{a}_1 \cdot (\mathbf{a}_2 \times \mathbf{a}_3)$. A reciprocal Bravais lattice with primitive lattice vectors $\mathbf{b}_1, \mathbf{b}_2$ and \mathbf{b}_3 , can be defined with the rule:

$$e^{i\mathbf{G}\cdot\mathbf{R}} = 1, \quad (2.69)$$

where $\mathbf{G} = m_1\mathbf{b}_1 + m_2\mathbf{b}_2 + m_3\mathbf{b}_3$ is a reciprocal lattice vector, and m_1, m_2 and m_3 are integer numbers. Equation (2.69) implies $\mathbf{G} \cdot \mathbf{R} = 2\pi m$, with m being an integer number. Therefore:

$$\mathbf{b}_i \cdot \mathbf{a}_j = 2\pi\delta_{ij}. \quad (2.70)$$

The vectors $\mathbf{b}_1, \mathbf{b}_2$ and \mathbf{b}_3 define the volume of the first Brillouin zone:

$$V_{1BZ} = \mathbf{b}_1 \cdot (\mathbf{b}_2 \times \mathbf{b}_3) = \frac{(2\pi)^3}{V_{cell}}. \quad (2.71)$$

The first Brillouin zone corresponds to the volume that is closer to a \mathbf{k} -point than to any other \mathbf{k} -point in the reciprocal lattice. The entire reciprocal lattice can be constructed by superposing consecutive translations of the first Brillouin zone along the vectors $\mathbf{b}_1, \mathbf{b}_2$ and \mathbf{b}_3 in integer steps.

The solutions of the Kohn-Sham equations with a periodic potential, denoted by $\{\psi_{i\mathbf{k}}\}$, take the form given by Bloch's theorem [100]:

$$\psi_{i\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}}u_{i\mathbf{k}}(\mathbf{r}), \quad (2.72)$$

where $e^{i\mathbf{k}\cdot\mathbf{r}}$ is a plane-wave and $u_{i\mathbf{k}}(\mathbf{r})$ is a function with the periodicity of the Bravais lattice, i.e., $u_{i\mathbf{k}}(\mathbf{r} + \mathbf{R}) = u_{i\mathbf{k}}(\mathbf{r})$. The periodicity of $u_{i\mathbf{k}}$ implies that:

$$\psi_{i\mathbf{k}}(\mathbf{r} + \mathbf{R}) = e^{i\mathbf{k}\cdot\mathbf{R}}\psi_{i\mathbf{k}}(\mathbf{r}). \quad (2.73)$$

Each $\psi_{i\mathbf{k}}$ defines an energy band, with value $\epsilon_{i\mathbf{k}}$, and occupancy $f_{i\mathbf{k}}$ within the

range $0 \leq f_{i\mathbf{k}} \leq 1$. Expanding $u_{i\mathbf{k}}(\mathbf{r})$ as a linear combination of infinite plane-waves, the Bloch states become:

$$\psi_{i\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} c_{i,\mathbf{k}+\mathbf{G}} e^{i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}}. \quad (2.74)$$

In practice, plane-wave methods limit the sum over \mathbf{G} by establishing a spherical cut-off radius in the reciprocal space, denoted by $|\mathbf{G}_c|$. In an equivalence to the expression of the kinetic energy of an electron (in atomic units), the plane-wave kinetic energy cut-off is defined as:

$$E_c = \frac{1}{2} |\mathbf{G}_c|^2. \quad (2.75)$$

The accuracy of the plane-wave basis set is therefore well-controlled by the single parameter E_c . A higher kinetic energy cut-off automatically results in a more complete plane-wave basis set.

For each vector \mathbf{R} in the direct lattice, there is a vector \mathbf{k} in the reciprocal lattice that satisfies Eq. (2.73). If the system is periodic but finite in size, the number of periodic replica in every direction is limited, and so the number of \mathbf{k} -points is also finite. If the system is assumed to extend an infinite number of times in each direction, then the number of \mathbf{k} -points is also infinite. The range of \mathbf{k} vectors that need to be sampled can be limited by taking into account the effect of the boundary. The Bloch states can be requested to be periodic under Born-von Karman periodic boundary conditions:

$$\psi_{i\mathbf{k}}(\mathbf{r}) = \psi_{i\mathbf{k}}(\mathbf{r} + n_i \mathbf{a}_i), \quad (2.76)$$

where n_i is an integer number and \mathbf{a}_i is a primitive lattice vector. Imposing (2.73)

and (2.76) simultaneously, it follows that:

$$\psi_{i\mathbf{k}}(\mathbf{r} + n_i\mathbf{a}_i) = e^{in_i\mathbf{k}\cdot\mathbf{a}_i}\psi_{i\mathbf{k}}. \quad (2.77)$$

The \mathbf{k} vectors can be expanded in terms of the reciprocal lattice vectors as $\mathbf{k} = \sum_j x_j \mathbf{b}_j$, where $\{x_j\}$ take real values. Introducing this expansion in (2.77), and using the condition (2.70), it follows that:

$$e^{2\pi i \sum_j x_j n_j} = 1, \quad (2.78)$$

which implies that

$$x_j = \frac{m_j}{n_j}, \quad j = 1, 2, 3, \quad (2.79)$$

where $\{m_j\}$ are integers. The resulting \mathbf{k} vectors with coefficients $\{x_j\}$ span a lattice with vectors \mathbf{b}_j/n_j , of volume $(2\pi)^3/V_{cell}$. This is the volume of the first Brillouin zone, as shown in Eq. (2.71). Therefore, under periodic boundary conditions, sampling the \mathbf{k} -points within the first Brillouin zone suffices to characterise the band-structure of the entire periodic system. In an finite periodic system, there is a finite number of \mathbf{k} vectors in the first Brillouin zone. On the other hand, infinitely periodic systems show a continuous spectra of \mathbf{k} -points in the first Brillouin zone. Numerical methods such as the ones developed by Baldereschi [101] and Monkhorst and Pack [102] simplify the \mathbf{k} -point sampling by taking into consideration the symmetry groups of the lattice, and reduce the integration over a continuous range to a finite sum of certain \mathbf{k} -points.

Plane-wave approaches can also be successfully applied to non-periodic systems using the supercell approach, where the primary simulation cell in real space

is made sufficiently large as to avoid interactions between atoms in periodic replicas. The extra space is then filled with vacuum or any form of continuous medium in which the system is embedded. Enlarging the volume of the cell is not normally cost-free, as it is filled with plane-waves. As the volume of the cell increases, the first Brillouin zone becomes progressively smaller, and the reciprocal space representation of the Bravais lattice contracts, according to Eq. (2.71). As a consequence, the system can be correctly described by only one \mathbf{k} -point, ($\mathbf{k} = 0$, dubbed the Γ -point), and the formalism is greatly simplified.

Plane-wave basis sets have certain intrinsic advantages, such as the lack of the basis set superposition error (BSSE) [103] and the independence of the nuclear coordinates, which allows the calculation of the energy derivatives simply with the Hellmann-Feynman theorem. Numerically speaking, they are a convenient basis set for calculating different components of the energy functional, such as the kinetic energy or the Coulomb potential, that take a simpler form in the reciprocal space [104]. Fast Fourier transforms [105, 106] are routinely used in plane-wave methods to transform the Bloch states into the reciprocal space.

2.8 Core electrons

The electronic structure of a given atomic system is determined almost entirely by the valence electrons, which are the ones that are likely to be found at long-range distances from the nuclear centres. The core electrons, on the other hand, are tightly-bound and confined to a limited region in the vicinity of the nuclei, and are less exposed to changes in the chemical environment. Therefore, the role of the core electrons in the electronic properties of a system is typically negligible. It then becomes possible to obtain an accurate description of the electronic structure

of a system even if the treatment of the core electrons is simplified in some way [40].

From a computational perspective, there is a good reason to avoid explicit equations involving the core electrons. The combination of spatial confinement of the core electron wavefunctions and the constraint on orthonormality of the Kohn-Sham states provokes rapid oscillations of the core-electron wavefunctions in the region near the nuclear centres. An accurate plane-wave representation of the wavefunctions near the nuclei will require an extraordinarily large number of basis set functions that will make the calculation impractical.

The pseudopotential approximation [107–115] eliminates the explicit treatment of the core electrons from the problem. A pseudopotential operator, \hat{v}_{ps} , that screens the Coulomb potential in the short range and converges asymptotically to $1/r$ as $r \rightarrow \infty$, where r is the distance from the nuclear centre, is added to the system's Hamiltonian. The eigenstates of the corrected Hamiltonian, ψ_{ps} , are smoother in the short-range distances, and can now be accurately represented with a finite plane-wave basis set. The *pseudised* wavefunctions converge to the original wavefunctions in the long range. The pseudopotential approximation is depicted in Figure 2.2.

Norm-conserving pseudopotentials [111] are designed to conserve the scattering properties of the nuclei while maximising the transferability of the pseudopotential to any chemical environment. The procedure devised by Kleinman and Bylander [109] yields norm-conserving pseudopotentials composed by the sum of a local part and a non-local part:

$$\hat{v}_{ps} = \hat{v}_{local} + \hat{v}_{nl}. \quad (2.80)$$

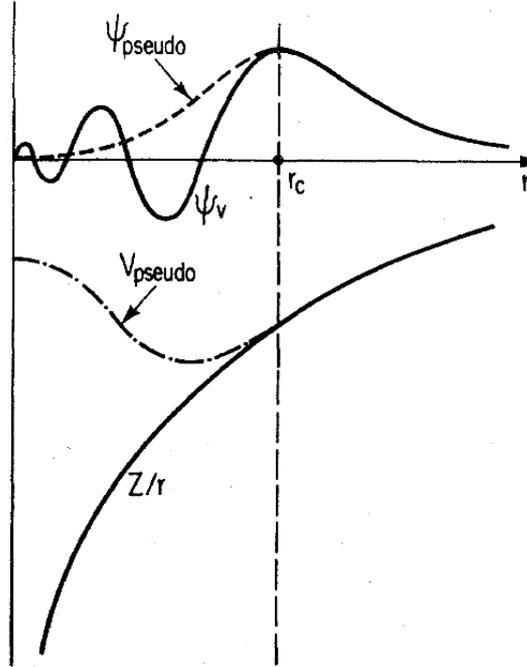


Figure 2.2: Pseudopotential approximation. The pseudopotential V_{pseudo} converges to $1/r$ after a given radial cut-off r_c . The pseudised wavefunction, ψ_{pseudo} , is smooth near the core and converges to the original one in the long-range. Image extracted from Ref. [40].

The non-local potential is constructed using the projector algebra as:

$$\hat{v}_{nl} = \sum_{l=0}^{l_{max}} \sum_{m=-l}^l \frac{|\xi_{lm}\rangle \langle \xi_{lm}|}{\langle \chi_{lm} | \hat{v}_l | \chi_{lm} \rangle}, \quad (2.81)$$

where \hat{v}_l is a different potential for each angular momentum channel, and $|\xi_{lm}\rangle = \hat{v}_l |\chi_{lm}\rangle$, where $\{\chi_{lm}\}$ are core eigenstates of each atom.

The ultrasoft pseudopotentials developed by Vanderbilt [113] allow plane-wave calculations with lower kinetic energy cut-off than the norm-conserving ones, increasing the computational speed. A completely different methodology, known as the projector-augmented wave method (PAW), developed by Blöchl [116], uses the frozen-core approximation to enable all-electron plane-wave calculations without pseudopotentials.

2.9 Linear-scaling density functional theory

The performance of a given computational algorithm can be measured in many different ways. For example, one might judge the performance of a program by the amount of computational resources that are required to execute it (memory size, number of CPUs, etc), energy efficiency (consumption of electricity in kWh in a given computer), or even in monetary terms. Performance can also be judged by the time-to-solution, that is, the lapse of time that the algorithm needs to complete all the scheduled operations when executed in a given target machine. Taking a step forward, one might ask about how the running cost would increase if the number of degrees of freedom of the problem that we wish to solve also increased. This is performance measured in terms of the scaling with the system size

Considering a problem of size N , the cost of executing a given computer program, which we denote by $\mathcal{C}(N)$, is the result of adding the cost of executing each of the individual subroutines that take place during the run-time. In more mathematical terms, a polynomial expansion of $\mathcal{C}(N)$ in terms of powers of N gives rise to the expression:

$$\mathcal{C}(N) = \sum_{s=0}^{\infty} \lambda_s N^s. \quad (2.82)$$

where $\{\lambda_i\}$ are the prefactors associated to each term of the sum. Put in this form, the condition for an algorithm to be linear-scaling is clear:

$$\{\lambda_i\}_{i=2}^{\infty} \rightarrow 0, \quad (2.83)$$

so that the polynomial expansion (2.82) can be truncated to first order, producing the straight line:

$$\mathcal{C}(N) = \lambda_0 + \lambda_1 N. \quad (2.84)$$

The term λ_0 is the prefactor, i.e., the cost of running a calculation of size $N = 0$, and the term λ_1 is the slope of the straight line, which indicates how rapidly the cost increases with N . Numerical algorithm whose cost $\mathcal{C}(N)$ scales as (2.84) are said to scale as $\mathcal{O}(N)$, and are referred to as linear-scaling methods.

Most algorithms for Kohn-Sham DFT (within the finite basis set limit) scale as $\mathcal{O}(N^3)$, and so they are said to be cubic-scaling. These implementations essentially follow a self-consistent field (SCF) optimisation using a variation of the Roothaan's algorithm [117]. Typical $\mathcal{O}(N^3)$ -scaling operations routinely performed in Kohn-Sham DFT methods are diagonalisation of the Hamiltonian matrix, orthogonalisation of $\{\psi_i\}$, and multiplication and inversion of dense matrices. In plane-wave approaches, fast Fourier transforms (FFTs) are necessary, with each FFT scaling as $\mathcal{O}(N \log N)$. With N FFTs being required, the total cost scales as $\mathcal{O}(N^2 \log N)$. If a linear-scaling DFT approach is to be devised, all non-linear-scaling steps must be completely avoided. This is not a straightforward accomplishment. Linear-scaling DFT methods rely on well controlled approximations to reduce the computational cost to the $\mathcal{O}(N)$ regime. We refer the reader to the reviews by Goedecker [118] and Bowler and Miyazaki [119] for a detailed discussion on linear-scaling DFT methods.

The complexity of the \mathcal{N} -electron problem in crystals and molecules can be greatly diminished if locality of the electronic structure can be introduced into the model. The review by Goedecker [118] has an excellent discussion on localisation in electronic structure theory. It can be shown that the one-particle density matrix, $\rho(\mathbf{r}, \mathbf{r}')$, of an insulating system decays exponentially as:

$$\rho(\mathbf{r}, \mathbf{r}') \propto \exp(-\varpi |\mathbf{r} - \mathbf{r}'|), \quad (2.85)$$

where ϖ is a constant that is proportional to the band-gap of the system [120]. Due to the large band-gap present in insulators, the effect of finite-temperature can be completely discarded, unless the value $k_B \mathcal{T}$ is abnormally large and comparable to the energy gap. In metallic systems (using the *jellium* model, or uniform electron gas), at finite temperature \mathcal{T} , $\rho(\mathbf{r}, \mathbf{r}')$ decays as an exponentially-modulated wave [120, 121]:

$$\rho(\mathbf{r}, \mathbf{r}') \propto k_F \frac{\cos(k_F |\mathbf{r} - \mathbf{r}'|)}{|\mathbf{r} - \mathbf{r}'|^2} \exp\left(-c \frac{k_B \mathcal{T}}{k_F} |\mathbf{r} - \mathbf{r}'|\right), \quad (2.86)$$

where k_F is the modulus of the wave vector associated to the Fermi surface and c is an order-1 constant. At zero temperature, the above expression yields algebraic decay:

$$\rho(\mathbf{r}, \mathbf{r}') \propto k_F \frac{\cos(k_F |\mathbf{r} - \mathbf{r}'|)}{|\mathbf{r} - \mathbf{r}'|^2}. \quad (2.87)$$

This property of localisation of $\rho(\mathbf{r}, \mathbf{r}')$ gives rise to the notion of “nearsightedness” of the electronic matter, a model put into mathematical terms by Kohn *et al.* [122, 123]. Nearsightedness cannot be exploited efficiently with a formalism based on Bloch states that are delocalised to fill all the space in the cell. Wannier [124] proposed a transformation of the Bloch states which effectively introduces a set of semi-localised orbitals, $\{w_{i\mathbf{R}}\}$, defined as:

$$w_{i\mathbf{R}}(\mathbf{r}) = \left(\frac{V}{2\pi^3}\right)^{1/2} \int_{1BZ} d\mathbf{k} \psi_{i\mathbf{k}}(\mathbf{r}) e^{i\mathbf{k}\cdot\mathbf{R}}, \quad (2.88)$$

where \mathbf{R} is a primitive lattice vector as defined in Section 2.7. Wannier functions

have been found to decay exponentially with the radial distance in insulators and finite-temperature calculations [125–128], and their spread in the real space gives a suitable measure of locality in the system [129]. The representation of the one-particle density matrix (2.45) in terms of the Wannier functions is:

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_i f_i \sum_{\mathbf{R}} w_{i\mathbf{R}}(\mathbf{r}) w_{i\mathbf{R}}^*(\mathbf{r}'). \quad (2.89)$$

One might then conclude that a small error would be introduced in those elements of $\rho(\mathbf{r}, \mathbf{r}')$ which are separated far apart, were truncated. In doing so, the numerical manipulation of the truncated elements is avoided altogether, and the computational cost is effectively reduced. After truncation, each of the Wannier functions is strictly localised within finite region of space, termed the *localisation region*. The superposition of Wannier localisation regions results in a wider molecular localisation region. In order to achieve linear-scaling, all the quantities that are related to the system are calculated only within the boundaries of the molecular localisation region, and everything outside is ignored or considered to be equal to zero. Linear-scaling methods rely on spatial truncation of the electronic structure as a well-controlled approximation to reduce the computational cost of the calculations.

A first practical implementation of linear-scaling DFT was the “divide-and-conquer” approach by Yang [130]. However, it is the approach originally proposed by Galli and Parrinello [131], based on the expansion of the Kohn-Sham states in terms of localised orbitals, that has gathered a wider attention in the electronic-structure community. There are many examples of electronic structure methods that localised orbitals in one form or another [132–136]. The large band-gap of insulators and semi-conductors allows zero-temperature Kohn-Sham DFT calculations to be performed efficiently. The occupancies $\{f_i\}$ take integer values (either

zero or one), and the entropic term vanishes. The density matrix is thus idempotent. Approaches such as the McWeeny purification transformation [137], the sign-matrices density mixing method [138], the Fermi operator expansion approach [139, 140], and the Krylov-subspace techniques [141], are methods for calculating $\rho(\mathbf{r}, \mathbf{r}')$ without explicitly diagonalising \hat{H} or orthonormalising $\{\psi_i\}$, thus avoiding some of the most important computational bottlenecks for linear-scaling DFT.

2.10 Localised orbitals

Many linear-scaling DFT approaches based on localised orbitals have been developed [131, 142–157]. The mathematical model that enables efficient algorithms based on this formalism is as follows. Consider a set of N_ϕ atom-centred orbitals $\{\phi_\alpha\}$ which are constrained to be strictly localised within a certain region of space near the atoms associated to them (typically in spherical volumes), which can be denoted generically by $\{\text{LR}(\phi_\alpha)(\mathbf{r})\}$. The Kohn-Sham states can then be expanded as a linear combination of $\{\phi_\alpha\}$ as:

$$|\psi_i\rangle = |\phi_\alpha\rangle M_i^\alpha. \quad (2.90)$$

where $\{M_i^\alpha\}$ are the expansion coefficients, which are the elements of the representation transformation matrix \mathbf{M} . In the above equation and in what follows, unless otherwise mentioned, the Einstein convention of implicit summation over repeated Greek indices is adopted. According to (2.90), the resulting Kohn-Sham states are constrained to be strictly localised within a finite region of space, formed by the union of all $\{\text{LR}(\phi_\alpha(\mathbf{r}))\}$:

$$\text{LR}(\psi_i(\mathbf{r})) = \bigcup_{\alpha=1}^{N_\phi} \text{LR}(\phi_\alpha(\mathbf{r})). \quad (2.91)$$

The Kohn-Sham states are considered to be exactly zero outside this volume. This is a well-controlled approximation, for in the limit of infinitely large localisation regions, the results with fully unconstrained (delocalised) Kohn-Sham states are recovered.

The localised functions $\{\phi_\alpha\}$ can be chosen to be orthonormal [135, 157]. However, imposing localisation and orthogonality constraints upon $\{\phi_\alpha\}$ is not always straightforward. The requirement of orthonormality can be relaxed and non-orthogonality of $\{\phi_\alpha\}$ can be assumed by default. If that is the case, however, the tensorial algebra of curved manifolds must be used [158–160]. The set of spatially-localised Kohn-Sham states $\{\psi_i\}$ form a complete basis of the complex Hilbert space $\mathcal{H}' \subset \mathcal{H}$. On the other hand, the localised functions $\{\phi_\alpha\}$, defined after the transformation \mathbf{M} (2.90), also form a complete basis set of \mathcal{H}' , albeit non-orthogonal, so long as $\{\phi_\alpha\}$ are linearly-independent functions. A set of dual functions $\{\phi^\alpha\}$ that verify the biorthonormality condition:

$$\langle \phi_\alpha | \phi^\beta \rangle = \delta_\alpha^\beta. \quad (2.92)$$

can be calculated from the duality relations:

$$|\phi^\alpha\rangle = |\phi_\beta\rangle S^{\beta\alpha}, \quad (2.93)$$

$$|\phi_\alpha\rangle = |\phi^\beta\rangle S_{\beta\alpha}. \quad (2.94)$$

where $\{S_{\alpha\beta}\}$ and $\{S^{\alpha\beta}\}$ are the elements of the overlap matrix \mathbf{S} and the inverse overlap matrix \mathbf{S}^{-1} , respectively, defined as:

$$S_{\alpha\beta} = \langle \phi_\alpha | \phi_\beta \rangle, \quad (2.95)$$

$$S^{\alpha\beta} = \langle \phi^\alpha | \phi^\beta \rangle = (S_{\alpha\beta})^{-1}. \quad (2.96)$$

The inverse overlap can be calculated at a linear-scaling cost with the Hotelling algorithm [161]. The functions $\{\phi_\alpha\}$ are rank-one covariant tensors (identified with Greek subscript), while $\{\phi^\alpha\}$ are rank-one contravariant tensors (identified with a Greek superscript). Any operator represented in terms of $\{\phi_\alpha\}$ or $\{\phi^\alpha\}$ will inherit the covariant or contravariant tensorial character, respectively. Contravariant indices can be made covariant by multiplication with \mathbf{S} , whereas covariant indices can be made contravariant by multiplication with \mathbf{S}^{-1} . Multiplication of tensors can only involve a pair of covariant and contravariant indices, and tensors of the same kind can be multiplied after mediation of the appropriate overlap matrix. The energy functional, ultimately calculated as the trace of the product of covariant and contravariant tensors, is tensor-invariant (or a rank-zero tensor). Derivatives of the total energy with respect to covariant functions $\{\phi_\alpha\}$ produce rank-one contravariant tensors:

$$\frac{\delta E}{\delta \langle \phi_\alpha |} = |g^\alpha\rangle, \quad (2.97)$$

while derivatives with respect to contravariant functions $\{\phi^\alpha\}$ produce rank-one covariant tensors:

$$\frac{\delta E}{\delta \langle \phi^\alpha |} = |g_\alpha\rangle. \quad (2.98)$$

The dual functions $\{\phi^\alpha\}$ also form a complete basis set of \mathcal{H}' . Therefore, the

Kohn-Sham orbitals $\{\psi_i\}$ can be expanded in terms of $\{\phi^\alpha\}$ as:

$$|\psi_i\rangle = |\phi^\alpha\rangle M_{\alpha i}, \quad (2.99)$$

where

$$M_{\alpha i} = S_{\alpha\beta} M_i^\beta. \quad (2.100)$$

The elements of the overlap matrices $\{S_{\alpha\beta}\}$ and $\{S^{\alpha\beta}\}$, and of the transformations $\{M_i^\alpha\}$ and $\{M_{\alpha i}\}$, are related as:

$$S_{\alpha\beta} = \sum_n M_{\alpha n} M_{n\beta}^\dagger \quad (2.101)$$

$$S^{\alpha\beta} = \sum_n M_n^\alpha M_n^{\dagger\beta}, \quad (2.102)$$

which implies that $(M_{\alpha n})^{-1} = M_n^{\dagger\alpha}$. The one-particle density matrix $\rho(\mathbf{r}, \mathbf{r}')$ can be represented in terms of $\{\phi_\alpha\}$ as:

$$\rho(\mathbf{r}, \mathbf{r}') = \phi_\alpha(\mathbf{r}) K^{\alpha\beta} \phi_\beta(\mathbf{r}'), \quad (2.103)$$

where $\{K^{\alpha\beta}\}$ are the elements of a matrix \mathbf{K} , defined as:

$$K^{\alpha\beta} = \sum_i M_i^\alpha f_i M_i^{\dagger\beta}. \quad (2.104)$$

The matrix \mathbf{K} of elements $\{K^{\alpha\beta}\}$ is the non-diagonal representation of the one-electron density operator $\hat{\rho}$ in terms of the contravariant functions $\{\phi^\alpha\}$:

$$K^{\alpha\beta} = \langle \phi^\alpha | \hat{\rho} | \phi^\beta \rangle. \quad (2.105)$$

In the $\{\phi_\alpha\}$ representation, the Hamiltonian operator becomes

$$H_{\alpha\beta} = \langle \phi_\alpha | \hat{H} | \phi_\beta \rangle. \quad (2.106)$$

The matrices \mathbf{K} , \mathbf{S} and \mathbf{H} , as well as any other matrix represented in the $\{\phi_\alpha\}$ or $\{\phi^\alpha\}$ representations, are square matrices of size N_ϕ .

One of the fundamental advantages of the formalism of localised functions is that it brings sparsity to the Hamiltonian matrix \mathbf{H} and to the overlap matrix \mathbf{S} , due to the limited overlap between localised functions centred on different atoms. Regarding \mathbf{K} , an extra approximation based on nearsightedness of $\rho(\mathbf{r}, \mathbf{r}')$ must be made: the matrix elements $K^{\alpha\beta}$ that belong to atoms separated beyond a spatial cut-off threshold, R_c , are truncated, i.e., set to zero. In such way, dense algebra, which is one of the main impediments for linear-scaling DFT calculations, is effectively substituted by sparse algebra, which is linear-scaling in the limit of sufficient sparsity.

Examples of localised functions used in linear-scaling DFT approaches are pseudoatomic orbitals (PAOs) [162, 163] (used by the SIESTA program [164–168], as well as by CONQUEST [169]), numerical atomic orbitals (NAOs) [170] (used in the FHI-AIMS program), Gaussian functions [171] and spherical waves [172]. With these approaches, multiple- ζ basis sets can be constructed in a systematic way, and the accuracy of the calculation is controlled by the number and shape of the functions included in the basis set. In some approaches, the localised functions $\{\phi_\alpha\}$ are expanded in terms of an underlying basis set of functions $\{D_p\}$, and optimised *in situ* to attain chemical accuracy. This technique has the advantage of reducing the dimensions of the matrices expressed in terms of $\{\phi_\alpha\}$ to the minimum size, usually comparable to that of a single- ζ basis set. Examples

of Kohn-Sham DFT approaches with optimised localised orbitals can be found in Refs. [149, 150, 152], and have been successfully implemented in the linear-scaling DFT packages CONQUEST [173], OPENMX [153] and ONETEP [4].

2.11 The ONETEP program

The methods that have been newly developed as part of this Ph. D. project, described in the forthcoming Chapters 3 to 6, have been implemented within the ONETEP method [4]. The ONETEP program is designed for first-principles linear-scaling Kohn-Sham DFT calculations on very large systems, using a SIMD (single input, multiple data) coding paradigm that allows calculations to be performed efficiently on modern massively-parallel computational facilities. There are a large number of publications that describe the numerical techniques used in ONETEP, cited in Refs. [4, 151, 174–187], and we refer the reader to these sources for an in-depth description. A brief summary of the most relevant features is provided in this section.

2.11.1 Calculation of the total energy

In the ONETEP method, the Kohn-Sham states, $\{\psi_i\}$, are expanded as a linear combination of *non-orthogonal generalised Wannier functions* (NGWFs) [151], which are localised within spheres of radii $\{R_\alpha\}$ centred in the atomic positions. The NGWFs are defined mathematically as:

$$\phi_\alpha(\mathbf{r}) = \frac{V}{2\pi} \int_{1BZ} d\mathbf{k} e^{-i\mathbf{k}\cdot\mathbf{R}_\alpha} \left[\sum_i \psi_{i\mathbf{k}}(\mathbf{r}) M_{i\alpha}^\dagger \right], \quad (2.107)$$

where \mathbf{R}_α is the coordinate vector of the centre of the localisation region of ϕ_α , V

is the volume of the simulation cell and $\{M_{i\alpha}^\dagger\}$ are the elements of a representation transformation matrix. An example of Kohn-Sham orbitals represented in terms of NGWFs is shown in Fig. 2.3.

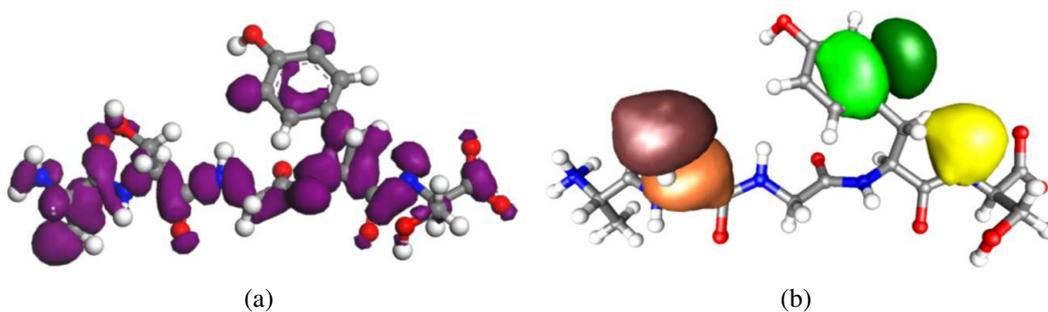


Figure 2.3: (a) Localised Kohn-Sham orbitals on a peptide. (b) Single NGWFs on the same peptide. Images taken from Ref. [4].

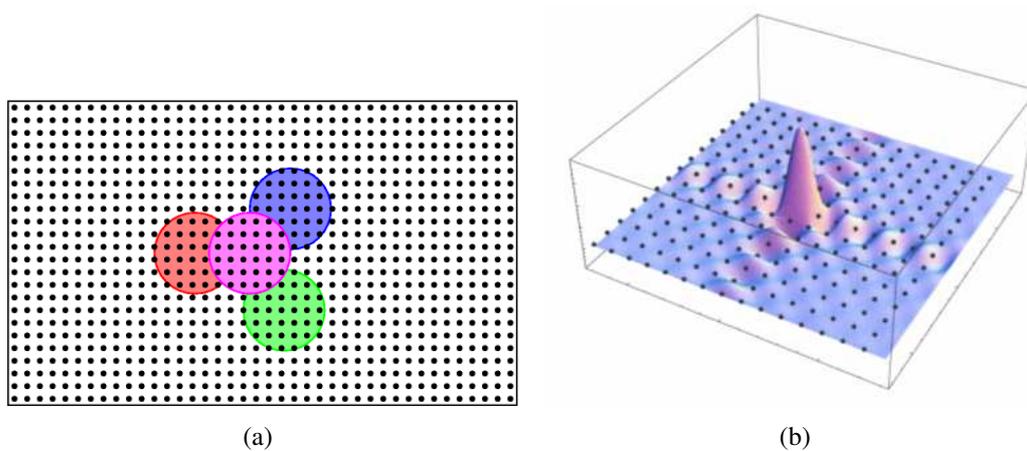


Figure 2.4: (a) Two-dimensional layout of the NGWFs in a uniform Cartesian grid, including their spherical localisation regions. (b) Three-dimensional representation of a psinc function. Images taken from Ref. [178].

In the NGWF representation, the one-particle density matrix, $\rho(\mathbf{r}, \mathbf{r}')$, takes the form (2.103), where $\{K^{\alpha\beta}\}$ are defined by Eq. (2.104) and the matrix \mathbf{K} is named the *density kernel*. The NGWFs are expanded in terms of a set of N_p periodic-sinc functions (psincs) $\{D_p\}$ [175] as:

$$\phi_\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) c_{p\alpha}, \quad (2.108)$$

where $\{c_{p\alpha}\}$ are the expansion coefficients. The psinc functions are bandwidth-limited delta functions positioned at the points of a uniform Cartesian computational grid in real space (see Figure 2.4). The psinc functions can be written as a sum of a finite set of plane-waves as:

$$D_p(\mathbf{r}) = D(\mathbf{r} - \mathbf{r}_p) = \frac{1}{N_p} \sum_{\mathbf{G}}^{\mathbf{G}_c} e^{i\mathbf{G}\cdot(\mathbf{r}-\mathbf{r}_p)}. \quad (2.109)$$

where \mathbf{r}_p is the coordinate vector of the grid point p . The quality of the psinc basis set is uniquely determined by the modulus of \mathbf{G}_c , which determines the plane-wave kinetic energy cut-off, E_c , as defined by Eq. (2.75). Standard plane-wave approaches determine a radial cut-off based on a sphere of wave vectors in the reciprocal space. The process is slightly different in ONETEP, and \mathbf{G}_c corresponds instead to the side of a cube of wave vectors in the reciprocal space [177]. Calculations performed with ONETEP can be compared to standard plane-wave methods by ensuring that the total number of basis set functions is the same. The grid spacing is selected based on the value of the Nyquist frequency, which gives the highest-frequency plane-wave that can be correctly represented in a grid of points. The magnitude of the grid spacing, Δ , is thus inversely proportional to $|\mathbf{G}_c|$. Increasing E_c reduces Δ , and augments the density of grid points per unit volume and the total number of grid points, N_p , in the calculation. Therefore, the accuracy of a ONETEP calculation is systematically controlled by a single parameter: the plane-wave kinetic energy cut-off, E_c [177]. The equivalence between psincs and plane-waves also has the advantage of providing with a well-defined framework

for performing operations in the reciprocal space representation. As explained in Section 2.9, a straight implementation of fast Fourier transforms will scale as $\mathcal{O}(N^2 \log N)$. The FFTbox technique developed within the ONETEP program [175] reduces the cost to the linear-scaling regime by performing FFTs in a box whose size is independent of the total number of atoms.

The ONETEP program is designed for running very large systems using periodic boundary conditions and the supercell approach. The simulation cell is expected to be very large, and therefore the \mathbf{k} -point sampling in ONETEP calculations is reduced to $\mathbf{k} = 0$ (Γ point). Norm-conserving pseudopotentials are employed to model the core electrons. The ion-ion interacting energy is calculated using an Ewald sum [57, 188], which includes the effect of the periodic boundary conditions, and can be regarded as a constant during the determination of the electronic states. The total Kohn-Sham energy functional of the interacting system of \mathcal{N} is calculated as:

$$E[n] = T[n] + E_{nl}[n] + E_{loc}[n] + E_H[n] + E_{xc}[n], \quad (2.110)$$

where T is the kinetic energy, E_{nl} and E_{loc} are the energy terms due to the non-local and local parts of the pseudopotential, E_H is the Hartree energy and E_{xc} is the exchange-correlation energy. The total energy calculated as in (2.110) already includes the double-counting corrections described in (2.52). An extra energy term $E_{nlcc}[n]$ must be added to account if non-linear core corrections are considered [189, 190].

In its standard functionality, ONETEP finds the Kohn-Sham ground-state at zero temperature ($\mathcal{T} = 0$), by a process of iterative constrained minimisation of $E[n]$ in two nested loops [182]. The inner loop optimises the elements of the density

kernel, $\{K^{\alpha\beta}\}$, for a fixed set of NGWFs expansion coefficients, $\{c_{p\alpha}\}$. Then, the outer loop optimises $\{c_{p\alpha}\}$ with fixed $\{K^{\alpha\beta}\}$. Mathematically, the ground-state energy is equal to:

$$\begin{aligned}
 E_0 &= \min_{\{K^{\alpha\beta}\}, \{c_{p\alpha}\}} \mathcal{L} [\{K^{\alpha\beta}\}, \{c_{p\alpha}\}] & (2.111) \\
 &= \min_{\{c_{p\alpha}\}} \min_{\{K^{\alpha\beta}\}} \mathcal{L} [\{K^{\alpha\beta}\}, \{c_{p\alpha}\}] \\
 &= \min_{\{c_{p\alpha}\}} \mathcal{L}' [\{c_{p\alpha}\}],
 \end{aligned}$$

where \mathcal{L} and \mathcal{L}' are Lagrangians that account for some of the constraints imposed upon the model. The standard methods for each of the loops are described in the next two subsections.

2.11.2 Inner loop optimisation

During the execution of the inner loop, the total energy is minimised iteratively and self-consistently with respect to the elements of the density kernel matrix [182], $\{K^{\alpha\beta}\}$, subject to the localisation constraint:

$$K^{\alpha\beta} = 0, \text{ if } |\mathbf{R}_\alpha - \mathbf{R}_\beta| \geq R_c, \quad (2.112)$$

the idempotency constraint:

$$\mathbf{KSK} = \mathbf{K}, \quad (2.113)$$

and the constraint on conservation of the number of electrons:

$$\text{tr} [\mathbf{KS}] = \mathcal{N}. \quad (2.114)$$

The idempotency constraint applies at zero temperature, and ensures that the occupancies are identically zero or one (2.55).

The method of Li, Nunes and Vanderbilt (LNV) [191, 192], based on the McWeeny purification transformation [137], is used to construct a new density kernel at a linear-scaling cost, without explicit diagonalisation of the Hamiltonian matrix. Given a so-called “auxiliary“ density kernel \mathbf{L} , subject to the same localisation constraints as \mathbf{K} , with occupancies in the range $\left[\frac{1 - \sqrt{5}}{2}, \frac{1 + \sqrt{5}}{2} \right]$, the purification transformation:

$$\mathbf{K} = 3\mathbf{L}\mathbf{S}\mathbf{L} - 2\mathbf{L}\mathbf{S}\mathbf{L}\mathbf{S}\mathbf{L}, \quad (2.115)$$

results in a “weakly-idempotent“ density kernel matrix (i.e., with occupancies in the interval $[0, 1]$). Recursive application of (2.115) yields full idempotency of \mathbf{K} [193]. Furthermore, imposing idempotency of \mathbf{K} is equivalent to imposing orthonormality of the Kohn-Sham states. The proof for this last assertion is as follows:

$$\mathbf{KSK} = \mathbf{K} \tag{2.116}$$

$$\begin{aligned} &\iff K^{\alpha\rho} S_{\rho\delta} K^{\delta\beta} = K^{\alpha\beta} \\ &\iff \sum_{rs} M_r^\alpha M_r^{\dagger\rho} S_{\rho\delta} M_s^\delta M_s^{\dagger\beta} = \sum_q M_q^\delta M_q^{\dagger\beta} \\ &\iff \sum_{rs} M_{i\alpha}^\dagger M_r^\alpha M_r^{\dagger\rho} S_{\rho\delta} M_s^\delta M_s^{\dagger\beta} M_{\beta j} = \sum_q M_{i\alpha}^\dagger M_q^\delta M_q^{\dagger\beta} M_{\beta j} \\ &\iff \sum_{rs} \delta_{ir} M_r^{\dagger\rho} S_{\rho\delta} M_s^\delta \delta_{sj} = \delta_{ij} \\ &\iff M_i^{\dagger\alpha} S_{\alpha\beta} M_j^\beta = \langle \psi_i | \psi_j \rangle = \delta_{ij}. \end{aligned}$$

This is a very important result. At $\mathcal{T} = 0$, the occupancies are perfectly known to be identically zero for the empty states and one for the occupied states. Therefore, there is no need to calculate $\{f_i\}$ from $\{\epsilon_i\}$ and the *aufbau* principle, which will require an explicit diagonalisation of the Hamiltonian matrix at a cubic-scaling cost. At the same time, the above result ensures that explicit orthonormalisation of $\{\psi_i\}$, which is also cubic-scaling (with the Gram-Schmidt [194] or Löwdin [195] methods, for instance), can be wholly avoided if the idempotency constraint holds during the energy minimisation.

The constraint of constant number of electrons, \mathcal{N} , must also be considered. The Lagrangian:

$$\mathcal{L} [\{K^{\alpha\beta}\}, \{\phi_\alpha\}] = E[\mathbf{K}, \{\phi_\alpha\}] - \mu (\text{tr}[\mathbf{KS}] - \mathcal{N}), \tag{2.117}$$

where μ is an undetermined Lagrange multiplier, ensures that the solutions that minimise (2.117) are compatible with this constraint. At any point of the itera-

tive minimisation of \mathcal{L} (except, perhaps, at convergence), the actual value of μ is unknown. To circumvent this difficulty, \mathcal{L} is reformulated in terms of a rescaled density kernel matrix, \mathbf{K}' , defined as [196]:

$$\mathbf{K}' = \frac{\mathcal{N}}{\text{tr}[\mathbf{K}\mathbf{S}]} \mathbf{K}, \quad (2.118)$$

so that, by construction and at all times, $\text{tr}[\mathbf{K}'\mathbf{S}] = \mathcal{N}$. The Lagrangian (2.117) adopts the simplified form:

$$\mathcal{L} \left[\left\{ K'^{\alpha\beta} \right\}, \left\{ \phi_\alpha \right\} \right] = E \left[\left\{ K'^{\alpha\beta} \right\}, \left\{ \phi_\alpha \right\} \right], \quad (2.119)$$

where, at convergence, $\mathbf{K} = \mathbf{K}'$. The variant of the LNV method proposed by Millam and Scuseria [197], offers a method for calculating the Lagrange multiplier μ , and makes kernel rescaling unnecessary. Both methods prove to yield optimal convergence at a linear-scaling cost. The method of kernel-rescaling is used by default in a standard ONETEP calculation.

At iteration m of the inner loop, a new auxiliary density kernel, $\mathbf{L}^{(m)}$, can be updated using the line-search step $\mathbf{L}^{(m+1)} = \mathbf{L}^{(m)} + \lambda^{(m)} \mathbf{\Delta}^{(m)}$, where $\mathbf{\Delta}^{(m)}$ is a suitable search direction compatible with the constraints. In fact, $\mathbf{\Delta}^{(m)}$ is calculated using direct energy minimisation techniques (steepest-descents or conjugate-gradients), based on the energy gradient [148]:

$$\frac{\partial E}{\partial L^{\alpha\beta}} = 6 (\mathbf{S}\mathbf{L}\mathbf{H}' + \mathbf{H}'\mathbf{L}\mathbf{S})^{\beta\alpha} - 4 (\mathbf{S}\mathbf{L}\mathbf{S}\mathbf{L}\mathbf{H}' + \mathbf{S}\mathbf{L}\mathbf{H}'\mathbf{L}\mathbf{S} + \mathbf{H}'\mathbf{L}\mathbf{S}\mathbf{L}\mathbf{S})^{\beta\alpha}, \quad (2.120)$$

where \mathbf{H}' appears as a consequence of the density kernel rescaling, and is equal to:

$$\mathbf{H}' = \mathbf{H} - \frac{\text{tr}[\mathbf{KH}]}{\text{tr}[\mathbf{KS}]} \mathbf{S}. \quad (2.121)$$

Once $\mathbf{L}^{(m+1)}$ is constructed, a new, nearly-idempotent density kernel $\mathbf{K}^{(m+1)}$ can be built with the McWeeny purification (2.115). The process can take place at a linear-scaling cost by virtue of the sparsity of \mathbf{L} and \mathbf{S} due to the localisation constraints. ONETEP incorporates a series of custom-built, parallel sparse algebra routines that allow efficient matrix multiplications using thousands of processors, expanding the scope of calculations to systems with tens of thousands of atoms [183]. At convergence of the inner loop, the Hamiltonian and the density matrix commute:

$$[\mathbf{HKS} - \mathbf{SKH}] = \mathbf{0}, \quad (2.122)$$

indicating that the time-independent Schrödinger equation has been solved.

2.11.3 Outer loop optimisation

Once the inner loop has finished, the outer loop optimises $\{c_{p\alpha}\}$ with a fixed density kernel \mathbf{K} (and auxiliary density kernel \mathbf{L}), subject to the constraints of spherical localisation of the NGWFs:

$$c_{p\alpha} = 0 \text{ if } |\mathbf{R}_\alpha - \mathbf{r}_p| \geq R_\alpha, \quad (2.123)$$

idempotency of \mathbf{K} and conservation of the number of electrons. At every iteration l of the outer loop, the NGWF expansion coefficients are updated using a line-search step $c_{p\alpha}^{(l+1)} = c_{p\alpha}^{(l)} + \lambda^{(l)} d_{p\alpha}^{(l)}$, where $\{d_{p\alpha}\}$ are the expansion coefficients of the suitable search direction vectors compatible with the constraints. The search

direction coefficients $\{d_{p\alpha}\}$ are also found with direct energy minimisation techniques, based on the gradient [148]:

$$g_p^\alpha = \frac{\partial E}{\partial c_{p\alpha}} = \left[\hat{H}\phi_\beta(\mathbf{r}) K^{\beta\alpha} - \phi_\beta(\mathbf{r}) Q_{LNV}^{\beta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (2.124)$$

where

$$\mathbf{Q}_{LNV} = 3\mathbf{LH}'\mathbf{L} - 2(\mathbf{LH}'\mathbf{LSL} + \mathbf{LSLH}'\mathbf{L}) - \frac{\text{tr}[\mathbf{KH}]}{\text{tr}[\mathbf{KS}]} \mathbf{K}. \quad (2.125)$$

If the inner loop is fully converged, $\mathbf{K} = \mathbf{LSL} = \mathbf{L}$ and $\mathbf{HKS} = \mathbf{SKH}$. Introducing these relations in (2.125), and expanding \mathbf{H}' as in (2.121), we obtain:

$$\begin{aligned} \mathbf{Q}_{LNV} &= 3\mathbf{LHL} - 2\mathbf{LHL} - 2\mathbf{LSLHL} & (2.126) \\ &+ \frac{\text{tr}[\mathbf{KH}]}{\text{tr}[\mathbf{KS}]} (3\mathbf{LSL} + 2\mathbf{LSLSL} + 2\mathbf{LSLSL} - \mathbf{K}) \\ &= -\mathbf{KHK} = -\mathbf{KHKSS}^{-1} = -\mathbf{KSKHS}^{-1} = -\mathbf{KHS}^{-1}. \end{aligned}$$

Convergence of the outer loop is judged based on the slope of the NGWF gradient along the search direction vector, calculated as:

$$\text{slope}_\phi = \sum_{p=1}^{N_p} g_p^\alpha d_{p\alpha} \quad (2.127)$$

which at convergence is equal to zero, indicating that the maximum degree of optimisation in the psinc basis set has been reached. The final NGWFs do not retain any specific angular momentum symmetry (see Fig. 2.5 for a practical example), and form a near-complete basis set for the wavefunctions that describe the valence

electrons.

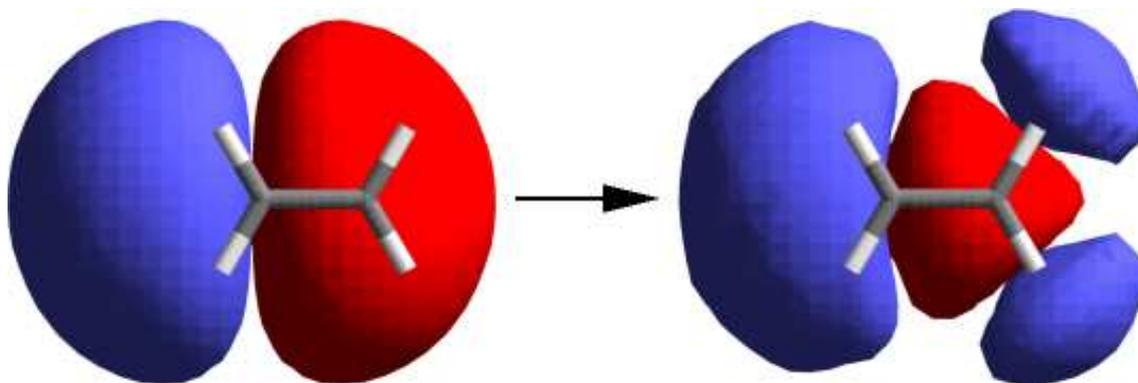


Figure 2.5: Optimisation of a single NGWF on a carbon atom in ethene. Image taken from [178].

Figure 2.6 shows a flowchart of the standard energy minimisation during a ONETEP calculation.

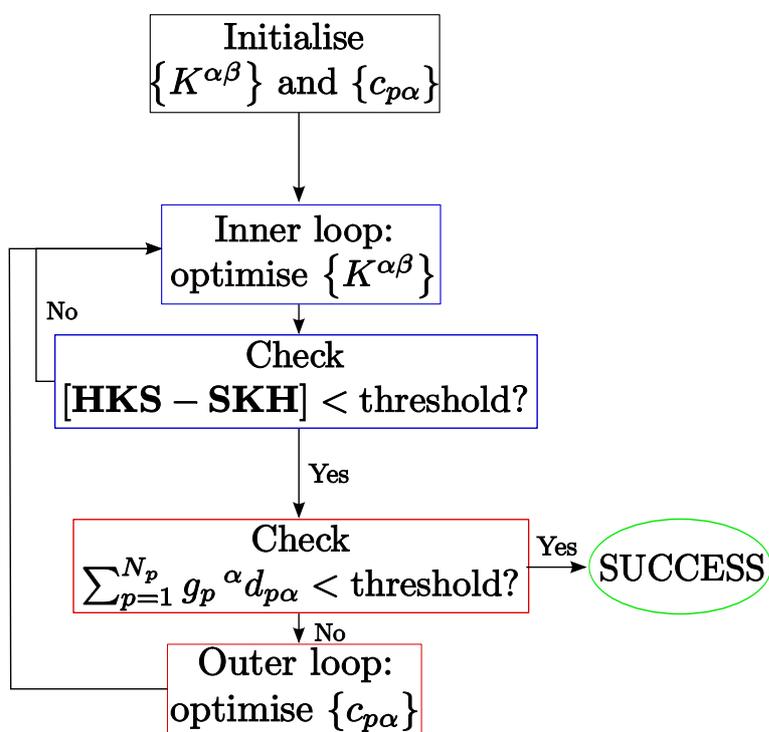


Figure 2.6: Self-consistent energy minimisation with two nested loops.

Chapter 3

Pulay forces

The total force acting on an atom is calculated from the derivative of the energy with respect to the position vector of the nucleus of that atom. The Hellmann-Feynman theorem shows that when the exact solutions of the time-independent Schrödinger equation are known, only the derivative of the Hamiltonian operator must be calculated. However, the Hellmann-Feynman theorem does not necessarily hold in variational, finite-basis-set approaches, where the wavefunctions are self-consistently improved using an iterative cycle. In some cases, a correction term to the total forces, known as Pulay forces, must be calculated for consistency between the energy and its derivatives.

It is widely accepted that Pulay forces have to be calculated if the basis set functions depend on the coordinates of the nuclei, as in PAO or Gaussian basis sets. The psinc functions that are used as a basis set in a standard ONETEP calculation are completely independent on the nuclear coordinates, since they are fixed to the points of a uniform Cartesian grid. Because of this, it was thought that the Pulay forces were null in the ONETEP method, and that the total forces could be accurately calculated simply according to the Hellmann-Feynman theorem. Some

practical tests involving geometry optimisation calculations, which rely very much on the accuracy of the atomic forces, seemed to indicate that the calculated atomic forces were not entirely correct. In particular, calculations on weakly-bound systems with hydrogen bonds often failed to relax the molecular structure and produced unreasonable results and heavily distorted geometries.

After further investigation, we discovered that the Pulay terms were in fact non-vanishing in a standard ONETEP calculation. As we explain in this chapter, the Pulay corrections are not caused by the displacement of the basis set. Instead, they appear as a direct consequence of the exact localisation constraints imposed upon the NGWF spheres. We have developed and implemented a numerical method for calculating the Pulay terms. Once these were added as a correction to the Hellmann-Feynman forces, geometry optimisation calculations consistently resulted in correct atomic structures in all sorts of systems. The addition of Pulay forces also enabled geometry optimisation calculations with fixed PAO basis sets, which have the potential to accelerate demanding calculations where the total energy has to be evaluated multiple times, such as in geometry optimisations of large DNA fragments.

3.1 Origin of the Pulay forces

Following Hamilton's definition of the generalised force vector [16], the total force acting on an atom A , denoted by \mathbf{F}_A , corresponds to the negative derivative of the Lagrangian (that describes the system and some of its constraints) with respect to the position vector, \mathbf{R}_A , of the nucleus of A :

$$\mathbf{F}_A = -\frac{d\mathcal{L}}{d\mathbf{R}_A}. \quad (3.1)$$

The Hellmann-Feynman theorem (see Section 1.5) predicts that, if the system is in its exact ground-state with an \mathcal{N} -electron wavefunction Ψ , such as $\hat{H}\Psi = E\Psi$, then the force acting on atom A can simply be calculated from the derivatives of the Hamiltonian operator as:

$$\mathbf{F}_A = -\frac{d\langle\Psi|\hat{H}|\Psi\rangle}{d\mathbf{R}_A} = -\langle\Psi|\frac{d\hat{H}}{d\mathbf{R}_A}|\Psi\rangle. \quad (3.2)$$

In variational methods, however, the exact ground-state wavefunction is not known exactly, although it can be very well approximated. Instead, we must refer to a trial wavefunction, Ψ^{trial} , where, in general, $\hat{H}|\Psi^{trial}\rangle \neq E|\Psi^{trial}\rangle$. The Hellmann-Feynman theorem cannot be applied straightforwardly since an exact wavefunction is not known. This observation was first published by Pulay [198], who argued that the total force acting on atom A must also include terms containing derivatives of the trial wavefunction with respect to the nuclear coordinates.

Let us now introduce a basis set of functions $\{\Phi_b\}$, such as $|\Psi\rangle = \sum_b |\Phi_b\rangle c_b$, where $\{c_b\}$ are the expansion coefficients. Using the chain rule, the derivative of the Lagrangian becomes [199]:

$$\mathbf{F}_A = -\frac{d\mathcal{L}}{d\mathbf{R}_A} = -\left[\frac{\partial\mathcal{L}}{\partial\mathbf{R}_A} + \sum_b \frac{\partial\mathcal{L}}{\partial c_b} \frac{\partial c_b}{\partial\mathbf{R}_A} + \sum_b \int d\mathbf{r} \frac{\delta\mathcal{L}}{\delta\Phi_b(\mathbf{r})} \frac{\partial\Phi_b(\mathbf{r})}{\partial\mathbf{R}_A}\right]. \quad (3.3)$$

where $\partial\mathbf{R}_A$ indicates explicit derivation with respect to \mathbf{R}_A . The first term in (3.3) corresponds to the Hellmann-Feynman force and contains only the explicit derivatives of the Hamiltonian operator. The second and third terms are remnants of the numerical representation of the trial wavefunction. The second term in (3.3) contains the derivatives of the expansion coefficients of Ψ in terms of $\{\Phi_b\}$:

$$\mathbf{F}_A^{coeffs} = - \sum_b \frac{\partial \mathcal{L}}{\partial c_b} \frac{\partial c_b}{\partial \mathbf{R}_A}, \quad (3.4)$$

while the third terms contains the derivative of the basis set functions with respect to the atomic positions:

$$\mathbf{F}_A^{basis} = - \sum_b \int d\mathbf{r} \frac{\delta \mathcal{L}}{\delta \Phi_b(\mathbf{r})} \frac{\partial \Phi_b(\mathbf{r})}{\partial \mathbf{R}_A}. \quad (3.5)$$

The sum of these two terms is the total Pulay force:

$$\mathbf{F}_A^{Pulay} = \mathbf{F}_A^{coeffs} + \mathbf{F}_A^{basis}. \quad (3.6)$$

Looking at \mathbf{F}_A^{coeffs} , the term $\frac{\partial c_b}{\partial \mathbf{R}_A}$ is, in principle, not zero. The term $\frac{\partial \mathcal{L}}{\partial c_b}$ is also non-zero unless full self-consistent convergence to the exact ground-state has been achieved, in which case $\frac{\partial \mathcal{L}}{\partial c_b} = 0$, and \mathbf{F}_A^{coeffs} vanishes. Regarding \mathbf{F}_A^{basis} , the term $\frac{\delta \mathcal{L}}{\delta \Phi_b(\mathbf{r})}$ is always non-zero except in the limit of fully complete basis set, which will require an infinitely large basis set. If the basis set functions depend explicitly on the nuclear coordinates (such as in PAO or Gaussian basis sets), then the term $\frac{\partial \Phi_b(\mathbf{r})}{\partial \mathbf{R}_A}$ does not vanish, and \mathbf{F}_A^{basis} must be calculated. But if the basis set functions do not depend on the nuclear coordinates (such as in plane-wave or psinc basis sets), then $\frac{\partial \Phi_b(\mathbf{r})}{\partial \mathbf{R}_A} = 0$, and \mathbf{F}_A^{basis} is exactly zero.

Based on these observations, it is commonly accepted that Pulay forces do not have to be calculated in plane-wave Kohn-Sham DFT methods. This is approximately true, so long as the following condition has been met:

$$\frac{\partial \mathcal{L}}{\partial c_b} \approx 0, \quad (3.7)$$

which indicates that the calculation has converged variationally to the Kohn-Sham

ground-state within an acceptable tolerance.

3.2 Pulay forces in ONETEP

In a standard ONETEP calculation, the Lagrangian $\mathcal{L} = E \left[\left\{ K'^{\alpha\beta} \right\} \left\{ c_{p\alpha} \right\} \right]$, defined in Eq. (2.119), is minimised with respect to $\{K^{\alpha\beta}\}$ and $\{c_{p\alpha}\}$, using the LNV method and a conjugate-gradients scheme. As demonstrated in Section 2.11, this Lagrangian accounts for the constraints of orthonormality of the Kohn-Sham states (via the idempotency constraint), and for conservation of the total number of electrons. An unconstrained minimisation of \mathcal{L} will yield a solution compatible with these constraints. According to the Hohenberg-Kohn theorems, at self-consistency:

$$\frac{\delta \mathcal{L}}{\delta n} = 0 \implies \frac{\delta \mathcal{L}}{\delta \psi_i} = 0 \implies \begin{cases} \frac{\partial \mathcal{L}}{\partial K^{\beta\alpha}} = 0 \\ \frac{\partial \mathcal{L}}{\partial c_{p\alpha}} = 0, \end{cases} \quad (3.8)$$

and the exact Kohn-Sham ground-state is found.

However, so far, nothing has been said about the localisation constraints of the density kernel matrix, \mathbf{K} , and the NGWFs, $\{\phi_\alpha\}$. The fact is that \mathcal{L} does not account for the localisation constraints on either of those. Therefore, an unconstrained minimisation of \mathcal{L} will carry an unwanted consequence, which is the delocalisation of both $\{K^{\alpha\beta}\}$ and $\{\phi_\alpha\}$, and it will result in a non-linear-scaling algorithm. Instead, a *constrained* minimisation of \mathcal{L} ensures that the final density kernel and NGWFs are correctly localised. The logical path (3.8) must be altered in the following way:

$$\lim_{\{R_\alpha\} \rightarrow \infty} \lim_{R_c \rightarrow \infty} \frac{\delta \mathcal{L}}{\delta n} = 0 \implies \lim_{\{R_\alpha\} \rightarrow \infty} \lim_{R_c \rightarrow \infty} \frac{\delta \mathcal{L}}{\delta \psi_i} = 0 \implies \begin{cases} \lim_{R_c \rightarrow \infty} \frac{\partial \mathcal{L}}{\partial K^{\beta\alpha}} = 0 \\ \lim_{\{R_\alpha\} \rightarrow \infty} \frac{\partial \mathcal{L}}{\partial c_{p\alpha}} = 0, \end{cases} \quad (3.9)$$

where R_c is the radial cut-off applied to the density kernel, \mathbf{K} , and $\{R_\alpha\}$ are radii of the spherical localisation regions of the NGWFs $\{\phi_\alpha\}$. Only in the limit of infinitely large localisation regions, the constrained minimisation will produce the same result as the unconstrained minimisation. This condition is obviously equivalent to fully delocalised density kernel and NGWFs, and it never actually happens in a practical calculation. One must therefore assume that by construction, a converged ONETEP calculation with finite localisation spheres of radii $\{R_\alpha\}$ and finite radial cut-off R_c inevitably yields:

$$\frac{\partial \mathcal{L}}{\partial K^{\beta\alpha}} \neq 0, \quad (3.10)$$

and

$$\frac{\partial \mathcal{L}}{\partial c_{p\alpha}} \neq 0. \quad (3.11)$$

Another possible interpretation of this result is that the slope of the energy gradient in any direction of the multidimensional space of $\{K^{\alpha\beta}\}$ and $\{c_{p\alpha}\}$ coefficients is not generally zero. However, the slope along the search direction that is compatible with the localisation constraints becomes zero when the closest state to the fully delocalised Kohn-Sham ground-state is reached.

Now we will demonstrate how these observations are behind the non-zero Pulay forces in ONETEP. Using the chain rule, the derivative of the total energy with re-

spect to the nuclear coordinates of atom A can be split into three partial derivatives as [184]¹:

$$\mathbf{F}_A = -\frac{d\mathcal{L}}{d\mathbf{R}_A} = -\frac{\partial\mathcal{L}}{\partial\mathbf{R}_A} - \frac{\partial\mathcal{L}}{\partial K^{\alpha\beta}} \frac{\partial K^{\alpha\beta}}{\partial\mathbf{R}_A} - \sum_{p=1}^{N_p} \frac{\partial\mathcal{L}}{\partial c_{p\alpha}} \frac{\partial c_{p\alpha}}{\partial\mathbf{R}_A}, \quad (3.12)$$

where the first term is the Hellmann-Feynman force and the second and third terms are the Pulay contributions that arise from the density kernel and the NGWF psinc expansion energy gradients, respectively. The derivative $\frac{\partial K^{\alpha\beta}}{\partial\mathbf{R}_A}$ is not easy to calculate, and its exact form is not known². For this reason, in order to be able to calculate accurate forces, we must assume that the density kernel is not constrained to be localised (which is often the case in calculations with up to a few thousands of atoms), so that $\frac{\partial\mathcal{L}}{\partial K^{\alpha\beta}} = 0$ at self-consistency. Once convergence has been reached, Equation (3.12) can be simplified to:

$$\mathbf{F}_A = -\frac{d\mathcal{L}}{d\mathbf{R}_A} = -\frac{\partial\mathcal{L}}{\partial\mathbf{R}_A} - \sum_{p=1}^{N_p} \frac{\partial\mathcal{L}}{\partial c_{p\alpha}} \frac{\partial c_{p\alpha}}{\partial\mathbf{R}_A}. \quad (3.13)$$

We shall still consider the NGWFs to be localised within finite spheres.

The energy gradient $\frac{\partial\mathcal{L}}{\partial c_{p\alpha}}$, can be calculated using Eq. (2.124) for LNV methods. The derivative of the expansion coefficients with respect to \mathbf{R}_A can be calculated in an efficient manner according to the following procedure. Using the FFTbox technique [174], the NGWFs are written as a linear combination of plane waves as:

$$\phi_\alpha(\mathbf{r}) = \frac{1}{N_F} \sum_{\mathbf{G}}^{\mathbf{G}_c} \tilde{\phi}_\alpha(\mathbf{G}) e^{i\mathbf{G}(\mathbf{r}-\mathbf{R}_\alpha)}, \quad (3.14)$$

¹The psinc functions are independent of the nuclear coordinates, therefore $\frac{\partial D_p(\mathbf{r})}{\partial\mathbf{R}_A} = 0$.

²Perhaps it could be calculated using density matrix perturbation theory [200].

where N_F indicates the number of points in the FFTbox, whose size is proportional to the NGWF radii. In this representation, the coordinates of the NGWF centre appear as a translational phase factor that takes into account the displacement of the NGWF centre across the simulation cell. Direct differentiation of (3.14) provides the derivative of the NGWFs with respect to \mathbf{R}_A :

$$\frac{\partial \phi_\alpha(\mathbf{r})}{\partial \mathbf{R}_A} = \frac{1}{N_F} \sum_{\mathbf{G}}^{\mathbf{G}_c} -i\mathbf{G} \tilde{\phi}_\alpha(\mathbf{G}) e^{i\mathbf{G}(\mathbf{r}-\mathbf{R}_\alpha)} \delta_{\alpha A}. \quad (3.15)$$

where $\delta_{\alpha A}$ is a Kronecker delta, which is equal to 1 only if the centre of the NGWF ϕ_α coincides with the centre of the atom A , and is zero otherwise. Equation (3.15) is evaluated at a linear-scaling cost (i.e., $\mathcal{O}(1)$ per NGWF). The derivative of the psinc expansion coefficients is:

$$\frac{\partial c_{p\alpha}}{\partial \mathbf{R}_A} = \left[\frac{1}{N_F} \sum_{\mathbf{G}}^{\mathbf{G}_c} -i\mathbf{G} \tilde{\phi}_\alpha(\mathbf{G}) e^{i\mathbf{G}(\mathbf{r}-\mathbf{R}_\alpha)} \delta_{\alpha A} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (3.16)$$

and so the Pulay force terms are equal to:

$$\mathbf{F}_A^{Pulay} = - \sum_{p=1}^{N_p} \frac{\partial \mathcal{L}}{\partial c_{p\alpha}} \frac{\partial c_{p\alpha}}{\partial \mathbf{R}_A}. \quad (3.17)$$

The total atomic force acting on atom A is equal to the sum of the Hellmann-Feynman and the Pulay terms:

$$\mathbf{F}_A = \mathbf{F}_A^{HF} + \mathbf{F}_A^{Pulay}, \quad (3.18)$$

The terms that form \mathbf{F}_A^{HF} in a ONETEP calculation are calculated in Ref. [184] as:

$$\mathbf{F}_A^{HF} = \mathbf{F}_A^{nl} + \mathbf{F}_A^{loc} + \mathbf{F}_A^{ew} + \mathbf{F}_A^{nlcc}, \quad (3.19)$$

where \mathbf{F}_A^{nl} and $\mathbf{F}_A^{\text{loc}}$ are due to the non-local and local parts of the pseudopotential, respectively, \mathbf{F}_A^{ew} is due to the ion-ion Ewald energy and $\mathbf{F}_A^{\text{nlcc}}$ are due to the non-linear core correction term.

Regarding finite-temperature DFT calculations, many studies [201–203] prove that there is no extra term due to the fractional occupancies or the entropy term if the occupancies follow the Fermi-Dirac distribution. This is important for the calculation of the total forces in the newly-implemented method for finite-temperature calculations on metals, described in Chapter 6, where Eq. (3.18) is also valid.

3.3 Convergence of the forces

Basis set convergence of the calculations, determined by the plane-wave kinetic energy cutoff and the radii of the spherical localisation regions of the NGWFs, is vital for accuracy in the total atomic forces. This aspect has been studied in Ref. [184] for the Hellmann-Feynman forces alone. Fig. 3.1 demonstrates the effect of adding the Pulay forces to the total atomic forces. The plots refer to the force acting on an oxygen atom in CO_2 and the force acting on a hydrogen atom in a H_2O dimer. We used the CAPZ LDA exchange-correlation functional with one NGWF for hydrogen and four for carbon and oxygen atoms.

The total forces, calculated as the sum of the Hellmann-Feynman components and the Pulay corrections, converge to a single value as the kinetic energy cutoff and the NGWF radii increase, and are comparable to the results obtained with traditional plane-wave DFT methods like CASTEP. The Pulay corrections to the Hellmann-Feynman forces ensure a greater degree of consistency between ground-state energies and the total force acting on ions for any size of localisation region. The magnitude of the Pulay forces decreases as the NGWF localisation spheres

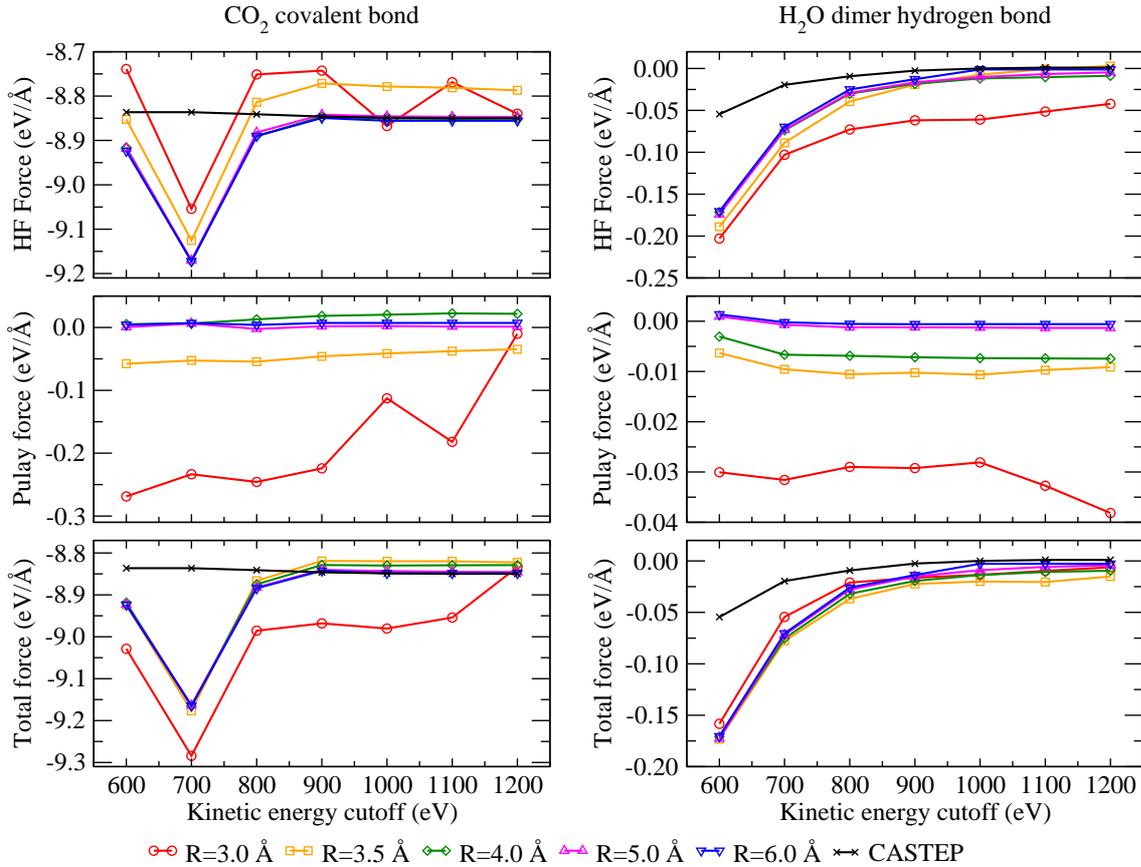


Figure 3.1: Convergence of the forces with respect to the kinetic energy cutoff and the radius, R , of the NGWF localisation spheres (assumed to be the same for all the NGWFs). Left: force acting on an oxygen atom in CO_2 along the covalent bond. Right column: force acting on a hydrogen atom in H_2O dimer along the hydrogen bond. The first row corresponds to the Hellmann-Feynman force, the second to the Pulay force and the third are the Pulay-corrected total forces (Hellmann-Feynman forces plus Pulay).

become larger, and they do not vanish as the kinetic energy cutoff increases. These results are in agreement with the theoretical observations regarding the localisation constraints, expressed in Eq. (3.9). The relative strength of the Pulay correction is greater in the case of the hydrogen bond in the water dimer (up to 15% of the Hellmann-Feynman force) than in the covalent bond of CO_2 (3%). Therefore, Pulay forces can be expected to play a more significant role in the ionic forces of weakly bonded systems, particularly when modeled with NGWFs of smaller radii.

3.4 Geometry optimisation calculations

The atomic forces play an essential role in geometry optimisation calculations, where the positions of the nuclear centres are varied until the equilibrium positions are reached. The equilibrium state is itself defined as the state where the total forces acting on the atoms are zero, leading to the structural conformation of minimum energy. Thus, the ability to calculate atomic forces accurately is critical for the success of geometry optimisation calculations. We used the pseudo Newton-Raphson method of Broyden, Fletcher, Goldfarb and Shanno (BFGS) [204] for relaxing the atomic structure. Also, we wanted to study the accuracy of the results using both optimised NGWFs (in the psinc basis set) and fixed NGWFs forming a PAO multiple- ζ basis set. The motivation is that geometry optimisation calculations can be expensive since they require a single-point energy calculation for each new set of nuclear coordinates. A fixed basis set of PAOs could potentially accelerate this type of calculations. A typical single- ζ (SZ) basis sets includes one function for hydrogen ($1s$) and four for carbon, nitrogen and oxygen ($2s2p$). Double- ζ basis sets include double the number of basis set functions to describe the same set of atomic orbitals (two for hydrogen and eight for carbon, nitrogen and oxygen). Triple- ζ basis sets include triple the number, and so on. Polarisation orbitals add atomic orbitals of higher angular momentum. A SZP basis set includes four functions for hydrogen ($1s2p$) and nine for carbon, nitrogen and oxygen ($2s2p3d$). Double and triple polarisation adds double and triple the number of polarisation orbitals. A DZP basis set includes five functions for hydrogen and thirteen for carbon, nitrogen and oxygen, while a TZDP includes nine for hydrogen and twenty-two for carbon, nitrogen and oxygen.

The quality of a geometry optimisation calculation can be affected by the so-

called “egg-box” effect [205, 206]. The egg-box effect is a numerical phenomenon associated to real-space grid methods that appears when the nuclear coordinates are displaced by a fraction of the grid-spacing in the simulation cell, generating periodic oscillations of the converged ground-state energy and forces. An estimation of the egg-box effect in the total energy and the forces is shown in Fig. 3.2. These calculations consecutively displace a CO₂ molecule along the first axis of the simulation cell (parallel to the CO₂ covalent bond) by a given step of 1/20 of the grid spacing of 0.24 Å. We use the PBE exchange-correlation functional with a kinetic energy cutoff of 1000 eV and NGWF radii of 4.23 Å to ensure convergence of the forces with the basis set. SZ, DZP and TZDP basis sets were constructed using PAOs.

The magnitude of the egg-box effect is reduced when NGWFs are used instead of PAO. A better stability is thus expected in geometry optimization calculations with optimised NGWFs. The maximum variation of the energy using NGWFs is of 1.6×10^{-3} eV, which is up to ten times smaller compared to SZ calculations. NGWFs also reduce the magnitude of the egg-box effect on the forces to 0.01 eV/Å, nearly a third of the magnitude in the case of SZ calculations.

3.4.1 Geometry optimization using NGWFs

Calculations on an adenine-thymine DNA base pair were carried out as a first test. This system is weakly bound by two hydrogen bonds, O...H and N...H, which make it sensitive to small variations of the energy and forces. We performed calculations using the PBE exchange-correlation functional with 1200 eV kinetic energy cutoff. We employed one NGWF for the hydrogen atoms and four for oxygen, carbon and nitrogen atoms, and various NGWF radii to test the convergence of

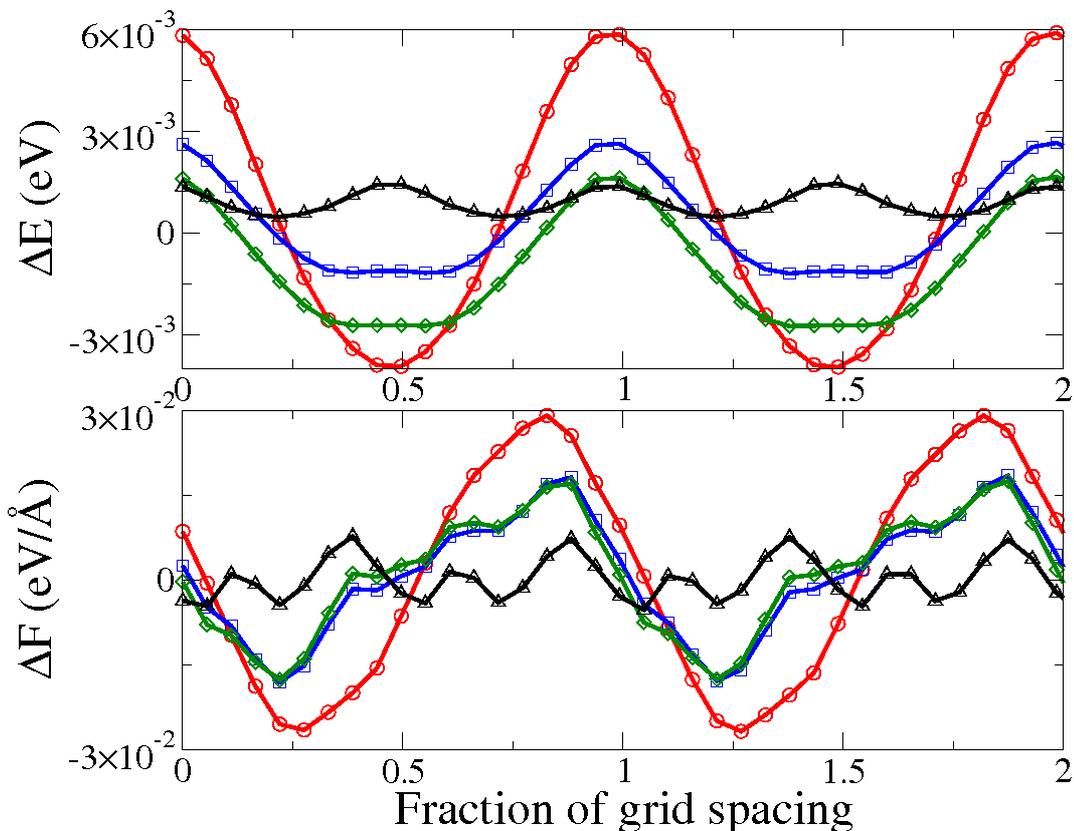


Figure 3.2: Eggbox effect on the energy (top) and force along the covalent bond (bottom) of a CO_2 molecule with SZ (red circles), DZP (blue squares) and TZDP (green diamonds) PAO basis sets and NGWFs (black triangles).

the geometry optimization calculations both with uncorrected Hellmann-Feynman forces and Pulay-corrected forces. We compared the resulting structures with those given by NWCHEM with a cc-pVTZ Gaussian basis set and CASTEP with the same kinetic energy cutoff and pseudopotentials. The results are shown in Table 3.1.

The addition of Pulay corrections to the Hellmann-Feynman forces improves the convergence of the BFGS algorithm for all the range of NGWF radii. It has a more important effect for small localisation regions, allowing relaxation of the structure to tighter tolerance thresholds. The calculations using uncorrected Hellmann-Feynman forces with small NGWF radii result in a poor description of hydrogen bonding in the adenine-thymine complex, leading to heavily distorted

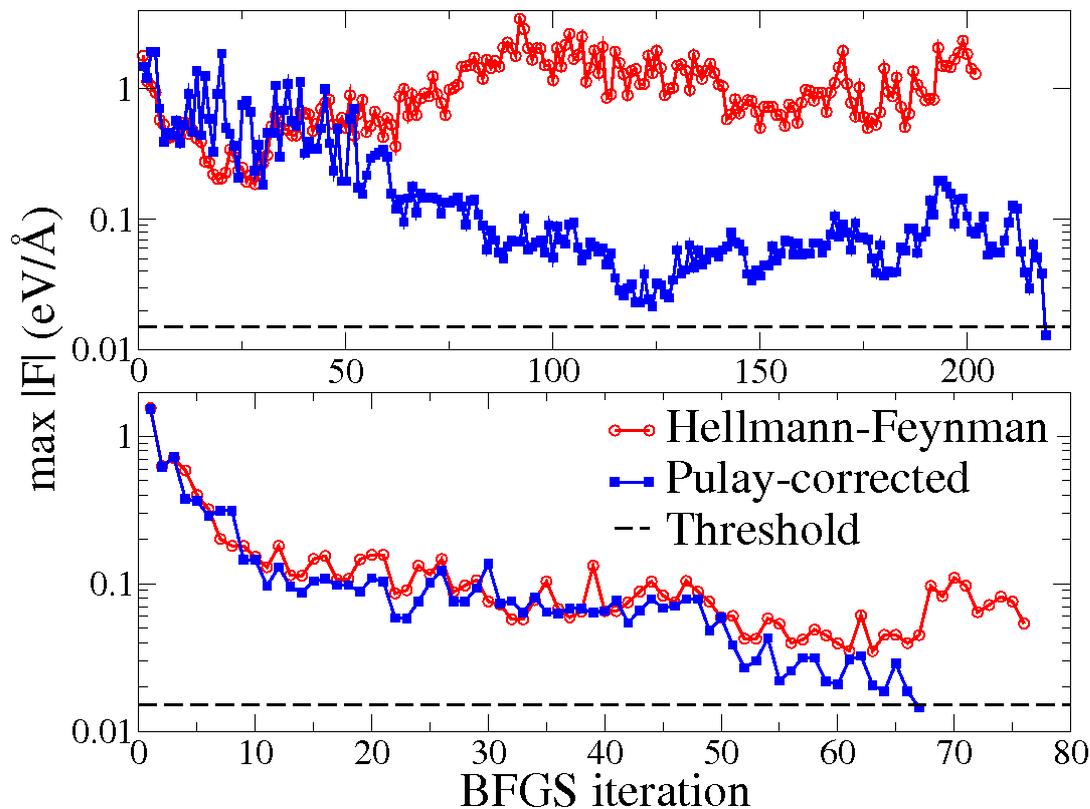


Figure 3.3: Convergence of the maximum absolute value of the force during geometry optimization of the adenine-thymine DNA base pair with NGWF radii of 3.70 Å (top) and 7.94 Å (bottom). The convergence threshold was 0.015 eV/Å.

geometry and unphysical configurations. In the uncorrected case, increasing the NGWF radii eventually results in a symmetric system that displays the correct binding properties, although the maximum value of the residual force remains higher than in the corrected case after a similar number of BFGS steps. The convergence of the maximum element of the force is plotted in Fig. 3.3, which demonstrates that accuracy in the forces is crucial to achieving efficient geometry relaxation using the BFGS method. Our results also show that the converged geometry of the adenine-thymine complex obtained with Pulay-corrected forces is still sensitive to changes in $\{R_\alpha\}$. Small NGWF localisation regions tend to overbind of the hydrogen bonds. As the NGWF radii increase, the geometry converges to

Table 3.1: Geometry optimization of adenine-thymine using ONETEP for different localisation radii R_α with Hellmann-Feynman forces (HF) and Pulay-corrected forces (PC). Results show the hydrogen bond lengths, the maximum absolute value of the force, $|F|_{\max}$, and the number of BFGS steps required.

	R_α (Å)	O...H bond (Å)		N...H bond (Å)		$ F _{\max}$ (eV/Å)		BFGS steps	
		HF	PC	HF	PC	HF	PC	HF	PC
ONETEP	3.70	9.77	1.73	10.64	1.60	1.028	0.010	203	222
	4.23	5.39	1.75	7.15	1.62	1.542	0.015	128	73
	4.86	1.90	1.80	1.65	1.65	1.542	0.015	80	73
	5.39	1.61	1.83	1.62	1.69	0.514	0.015	95	90
	5.82	1.73	1.83	1.64	1.71	1.028	0.015	100	63
	6.45	1.75	1.83	1.57	1.72	0.463	0.015	92	63
	6.98	1.85	1.84	1.62	1.72	0.463	0.015	97	65
	7.41	1.81	1.84	1.68	1.73	0.206	0.015	92	74
	7.94	1.83	1.84	1.67	1.73	0.051	0.015	78	69
NWCHEM	—	1.84		1.73		0.026		102	
CASTEP	—	1.85		1.75		0.010		88	

the optimised structures obtained with NWCHEM and CASTEP.

We have used the relaxed structures (obtained both with uncorrected and corrected forces) to calculate the binding energy of the adenine-thymine complex, defined as $E_{\text{complex}} - E_{\text{adenine}} - E_{\text{thymine}}$. The first term corresponds to the energy of the geometry-optimised complex, and the last two are the energies of the isolated adenine and thymine molecules, respectively. In each case, the individual molecules retain the same geometry as in the complex. The results, represented in Fig. 3.4, show that the calculations using uncorrected Hellmann-Feynman forces lead to extreme variations and erratic convergence with increasing NGWF radii³. In contrast, the structures obtained with Pulay-corrected forces converge systematically with respect to the size of the localisation region. In this case, the binding energies converge within 1 kcal/mol for NGWF radii larger than 5.5 Å. The convergence of these calculations is due to the simultaneous change of the basis set (as a result of increasing NGWF radii) and the change in the initial structure.

³The VMD program was used for visualisation of atoms and molecules [207].

To determine the convergence of the binding energies with respect to the basis set only, we chose the structures obtained with $R_\alpha = 3.70 \text{ \AA}$ and $R_\alpha = 7.94 \text{ \AA}$ (Pulay-corrected) and re-calculated the binding energies for different NGWF radii. In these calculations, the binding energy converges within 1 kcal/mol for NGWF radii larger than 4.2 \AA . Basis set superposition error is eliminated due to the optimization of the NGWFs in terms of the psinc basis set [180].

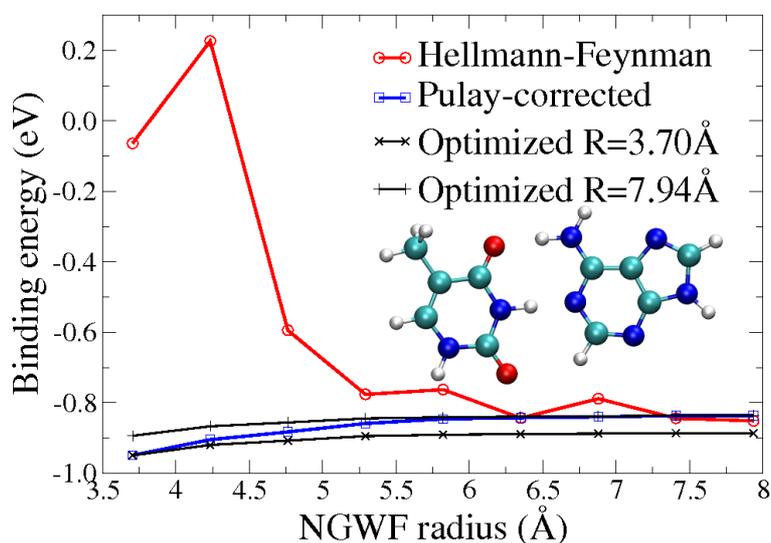


Figure 3.4: Convergence of the binding energy in the adenine-thymine complex with increasing NGWF radii. The calculations that use the structures obtained with Pulay-corrected forces show systematic convergence (blue squares) with increasing NGWF radii. In contrast, the calculations that use structures obtained with uncorrected Hellmann-Feynman forces (red circles) lack of convergence and show large variations of the binding energy. The calculations represented with black crosses show convergence of the binding energy with increasing NGWF radii, based on fixed molecular geometries obtained with Pulay-corrected forces and NGWF radii of 3.70 \AA and 7.94 \AA , respectively.

We have also performed geometry optimization calculations in the larger system of the “tennis-ball” self-assembling superstructure [208, 209], formed by 176 atoms. This system is a model for protein-ligand binding interactions, formed by two identical structures that bind each other through eight hydrogen bonds when

rotated 90 degrees. We used the PBE functional, 1200 eV kinetic energy cut-off and one NGWF per hydrogen atom and four per oxygen, nitrogen and carbon atoms with NGWF radii of 3.70 Å. The results, included in Fig. 3.5, show that calculations with uncorrected Hellmann-Feynman forces result in an asymmetric complex and in an unstable and highly oscillating minimization of the total force. On the other hand, when the Pulay corrections are taken into account, the BFGS method converges to a symmetric structure and a low threshold of force tolerance (0.015 eV/Å).

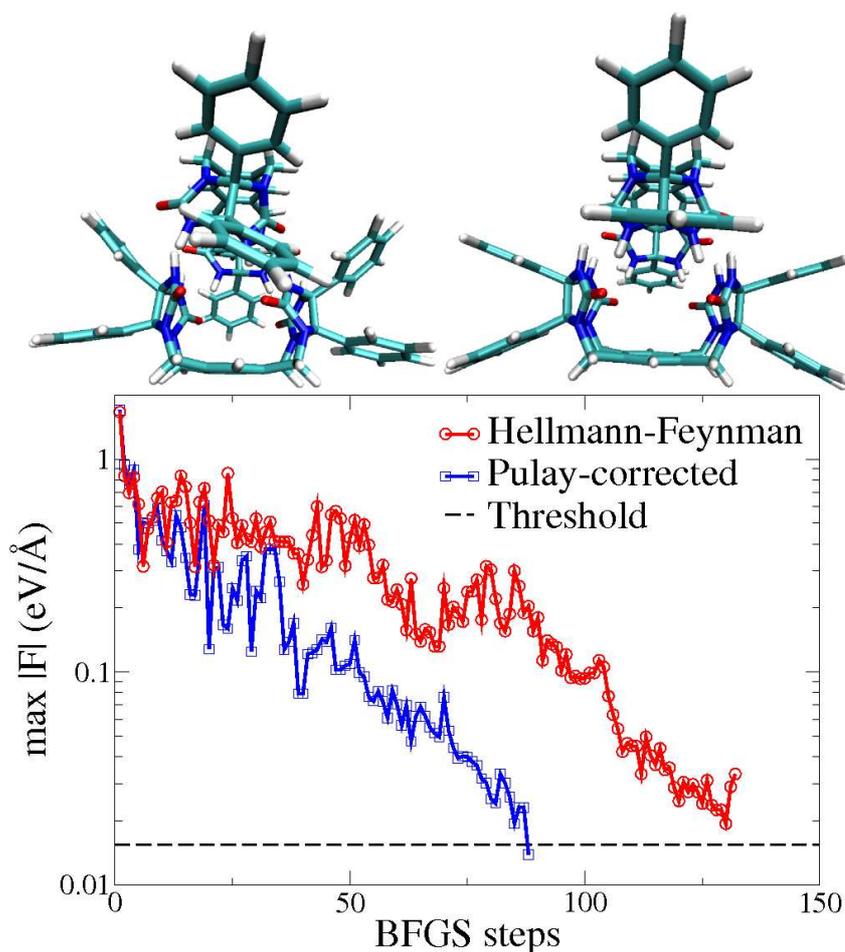


Figure 3.5: Convergence of the maximum force during the BFGS geometry optimization of the tennis-ball dimer using only the Hellmann-Feynman forces (top left molecule) and Pulay-corrected forces (top right).

3.4.2 Geometry optimization using PAOs

We performed a number of geometry optimization calculations on the adenine-thymine DNA base pair using PAO multiple- ζ basis sets, using the same exchange-correlation functional, pseudopotentials and kinetic energy cutoff than in the case of the NGWFs calculations discussed in section 3.4.1. The results, shown in Table 3.2, are consistent with those obtained with NGWFs for the same radii. The calculations with the smaller basis sets tend to overbind the hydrogen bonds. Progressively larger basis sets result in structures that converge to those obtained with NGWFs, NWCHEM and CASTEP, previously shown in Table 3.1. Basis sets of medium size such as TZP, TZDP or QZP can lead to optimum geometries.

Table 3.2: Length of the two hydrogen bonds in the adenine-thymine complex obtained with different PAO multiple- ζ basis sets of different localisation radii of size R , given in Å. The total forces converged to 0.015 eV/Å.

	O . . H bond (Å)			N . . H bond (Å)		
	R=3.70	R=4.23	R=4.86	R=3.70	R=4.23	R=4.86
SZ	1.41	1.42	1.42	1.32	1.31	1.31
SZP	1.57	1.61	1.56	1.32	1.31	1.31
DZ	1.66	1.67	1.65	1.49	1.52	1.53
DZP	1.74	1.73	1.71	1.58	1.58	1.57
TZ	1.70	1.71	1.69	1.53	1.58	1.56
TZP	1.83	1.84	1.81	1.68	1.69	1.68
TZDP	1.79	1.84	1.79	1.68	1.68	1.68
QZ	1.73	1.71	1.71	1.55	1.56	1.57
QZP	1.81	1.82	1.81	1.67	1.66	1.69

3.4.3 Geometry optimization of large systems

As an example of the application of the aforementioned methods to large-scale systems, we performed a geometry optimization calculation on a DNA fragment of 16 base pairs consisting of 1045 atoms in vacuum. The initial structure was

created using AMBER NUCGEN [210] (sequence ATCGATTGAGCTCTAG), after which we protonated the phosphate groups so that the total charge is zero. The calculations were run using the PBE exchange-correlation functional and a kinetic energy cutoff of 1200 eV. We used, for each test case, NGWFs (one for hydrogen atoms, four for oxygen, nitrogen and carbon, and nine for phosphorus), DZP and TZP PAO basis sets. The localisation radii were 3.70 Å in all cases. In such large systems, achieving the same degree of convergence of the total force as we would achieve in smaller systems can be cumbersome, due to the greater number of degrees of freedom involved in the calculation. Therefore, we consider that a higher convergence threshold of 0.16 eV/Å is suitable for the purpose of demonstrating that effective structural relaxation occurs.

Fig. 3.6 shows the convergence of the force for the different calculations, and Table 3.3 shows the optimised value of some relevant structural parameters and their change throughout the calculation. The maximum value of the force converges to the given threshold in all cases. The final structure obtained with the DZP basis set shows clear differences when compared to the structures obtained with TZP and NGWFs. It displays a shorter length, a smaller radius and much shorter N...H hydrogen bonds between the base pairs. Such observations are in agreement with the results of the adenine-thymine complex discussed in sections 3.4.1 and 3.4.2, in which the DZP basis set tends to overbind hydrogen bonds compared to TDZP and NGWFs. These last two calculations, on the other hand, produce very similar results.

Noticeably, the calculation with NGWFs requires fewer BFGS steps to complete in comparison with DZP and TZDP basis sets. This is likely to be a consequence of the more accurate description of the Kohn-Sham ground-state provided

by the NGWF optimization. Similarly, TZDP also requires fewer BFGS steps than DZP due to its better degree of completeness. For the basis sets used in this study, each BFGS step is completed faster when using PAOs, as overall fewer matrix operations are required due to the absence of the outer loop to optimise the NGWFs. We run our calculations on 120 cores, using 15 Intel Nehalem nodes with 8 cores each. On average a BFGS step took 3.2 hours with the DZP basis set, 7.6 hours with TZDP basis set, and 8.3 hours with the NGWFs optimised in the psinc basis set. As we can observe, the time saved per BFGS step through the use of a PAO basis set decreases as the number of basis set functions becomes larger. For this system in particular, using a DZP basis set is nearly twice as fast as the corresponding NGWF calculation, while the TDZP basis set is only marginally faster.

Table 3.3: Structural parameters of the optimised 1045-atom DNA fragment as optimised by ONETEP. N. . .H(4), N. . .H(8) and N. . .H(12) correspond to the hydrogen bond of the fourth, eighth and twelfth pairs, respectively, that involve a nitrogen atom.

Basis set	DZP		TZDP		NGWFs	
	Final	Change	Final	Change	Final	Change
Length (Å)	49.74	-1.18	50.62	-0.35	50.50	-0.47
Diameter (Å)	15.31	0.04	15.32	0.05	15.21	-0.06
Helix pitch (Å)	24.05	0.34	23.63	0.08	23.55	0.16
N. . .H(4)	1.66	-0.21	1.78	-0.09	1.80	-0.07
N. . .H(8)	1.48	-0.33	1.62	-0.19	1.67	-0.14
N. . .H(12)	1.54	-0.27	1.63	-0.18	1.64	-0.17
Time per BFGS step (h)	3.2		7.6		8.3	

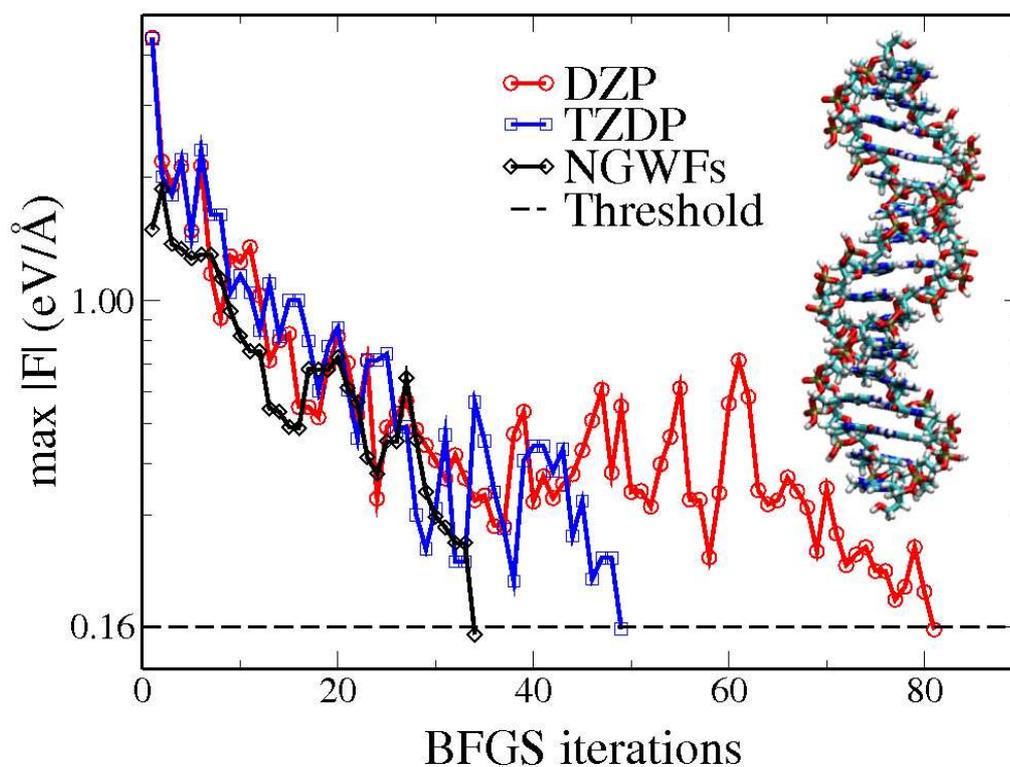


Figure 3.6: Convergence of the maximum value of the force during the geometry optimization of the DNA molecule using NGWFs and DZP and TZDP PAO basis sets.

3.5 Summary

In this chapter we have shown that the Pulay forces terms appear as a consequence of the strict localisation constraints imposed upon the density kernel and the NGWFs in a standard ONETEP calculation. We have also devised a linear-scaling numerical method for calculating the Pulay terms due to the energy gradient of the NGWF expansion coefficients. The results demonstrate that the Pulay corrections to the Hellmann-Feynman forces are essential in order to achieve accuracy in the total atomic forces. Using localised orbitals optimised *in situ* in terms of a basis set that is independent of the nuclear positions, such as psinc functions, does not eliminate the Pulay forces.

Geometry optimization calculations produce much better results when Pulay forces are taken into account. For example, Pulay forces are necessary to correctly describe weakly-bound molecular systems such as the “tennis-ball” hydrogen-bonded dimer. Our method can be used for large-scale geometry optimization calculations on systems of more than a thousand atoms. The results converge systematically with respect to the basis set.

The corrections due to the Pulay forces also allow calculations in which the localised orbitals are kept fixed to pseudoatomic orbitals (PAO) after initialisation, forming multiple- ζ basis sets of controllable accuracy. Geometry optimization calculations with medium-size PAO basis sets are faster than with optimised NGWFs.

As a future extension, a method to determine the Pulay forces due to the optimisation of the density kernel under localisation constraints might be implemented.

Chapter 4

The variationally-localised search direction (VLSD) method

Locality (or nearsightedness) of the electronic matter can be exploited to achieve linear-scaling cost in electronic-structure theory calculations. [122, 123, 130–136]. The Kohn-Sham orbitals can be expanded as a linear combination of spatially-localised functions centred on the atomic coordinates. These localised functions can additionally be optimised with an *in situ* optimisation based on their expansion in terms of a high-resolution basis set. Direct energy minimisation techniques such as steepest descent or conjugate gradients are commonly used to accomplish this task [149, 152, 156, 175]. At every iteration, the orbitals are updated by taking steps along search direction vectors calculated from the energy derivatives.

The localisation constraints must be consistently maintained during the calculation. An approach that is widely-accepted calculates the search direction vectors associated to the unconstrained problem, which are delocalised in real space, and performs an *ad-hoc* truncation step in which the values outside the assigned localisation regions are deleted [149, 151]. Hereafter, this method will be referred to

as the *truncated search direction* (TSD) method. The TSD method is compatible with the variational principle, in the sense that the energy functional converges to a single value as the size of the localisation region increases. However, at every iteration, the search direction vectors are not calculated following a fully variational approach, for it is clear that the truncation step does not guarantee that the energy will acquire a lower value along the resulting search direction vectors. The TSD method can be regarded as an approximation which implicitly assumes that the truncated search direction vectors corresponding to the unconstrained problem will resemble the search direction vectors of the problem with exact localisation constraints. Experience, however, dictates that the TSD method is capable of yielding successful convergence and of producing results which are in agreement with other electronic structure approaches [181].

The method presented in this chapter takes a fundamentally different approach. The localisation constraints are introduced into the mathematical model exactly and upfront, and propagated analytically throughout the equations that describe the system using formal tensorial algebra. Only the variational degrees of freedom associated to the problem with localisation constraints are taken into account for calculating the search direction vectors. As a result, the search direction vectors are strictly localised by construction. Truncation is therefore unnecessary and avoided altogether, and the total amount of information stored in the search direction vectors at every iteration is fully conserved. This method provides rapid convergence with any size of the localisation region in an entirely variational manner. For these reasons, this new method will be referred to hereafter as the *variationally-localised search direction* (VLSD) method.

We have implemented the new VLSD method in the ONETEP program (which

by default uses the TSD method) to work with the formalism of localised NGWFs in a non-orthogonal representation. Validation calculations demonstrate that both the TSD and VLSD methods are capable of producing correct results for the electronic and structural properties of the test molecules, compared to other computational electronic structure approaches. However, the VLSD algorithm is capable of achieving faster convergence towards the Kohn-Sham ground state. This is particularly true in atomic systems where a higher degree of natural delocalisation is expected.

4.1 Localised functions

The mathematical framework that is suitable for calculations with non-orthogonal, localised orbitals was explained in Section 2.10. With the aim of not repeating large amounts of what has already been said, we would like to ask the reader to refer to that section for a more detailed outline of the tensorial relations in curved manifolds. Refs. [158–160] are excellent sources of information regarding tensorial algebra methods applied to electronic structure theory. We shall now, briefly and for the sake of completion, refresh some of these expressions. The Einstein convention for summation over repeated Greek indices is not used in the equations presented in this chapter.

The Kohn-Sham orbitals are written in terms of a set of N_ϕ atom-centred, localised functions as:

$$|\psi_i\rangle = \sum_{\alpha=1}^{N_\phi} |\phi_\alpha\rangle M_i^\alpha. \quad (4.1)$$

where $\{\phi_\alpha\}$ are constrained to be strictly localised within a finite region of the real

space, denoted generically by $\{\text{LR}(\phi_\alpha)\}$. The set of functions $\{\phi_\alpha\}$ are covariant. A set of dual contravariant functions that is biorthogonal to $\{\phi_\alpha\}$ is defined as:

$$\langle \phi_\alpha | \phi^\beta \rangle = \delta_\alpha^\beta, \quad (4.2)$$

which implies the two duality relations:

$$|\phi^\alpha\rangle = \sum_{\beta=1}^{N_\phi} |\phi_\beta\rangle S^{\beta\alpha}, \quad (4.3)$$

$$|\phi_\alpha\rangle = \sum_{\beta=1}^{N_\phi} |\phi^\beta\rangle S_{\beta\alpha}. \quad (4.4)$$

where $\{S_{\alpha\beta}\}$ and $\{S^{\alpha\beta}\}$ are the elements of the overlap and inverse overlap matrices, respectively, defined as:

$$S_{\alpha\beta} = \langle \phi_\alpha | \phi_\beta \rangle = \sum_i M_{\alpha i} M_{i\beta}^\dagger \quad (4.5)$$

$$S^{\alpha\beta} = \langle \phi^\alpha | \phi^\beta \rangle = \sum_i M_i^\alpha M_i^{\dagger\beta}, \quad (4.6)$$

Both $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$ can be discretised in terms of an underlying, orthonormal basis set of functions $\{D_p\}$ centred in the points of a computational grid as:

$$\phi_\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) c_{p\alpha}, \quad (4.7)$$

$$\phi^\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) c_p^\alpha. \quad (4.8)$$

The expansion coefficients carry the tensorial character of the functions that they expand. The duality relations (4.3) and (4.4) apply to $\{c_{p\alpha}\}$ and $\{c_p^\alpha\}$ as:

$$c_{p\alpha} = \sum_{\beta=1}^{N_\phi} c_p^\beta S_{\beta\alpha} \quad (4.9)$$

$$c_p^\alpha = \sum_{\beta=1}^{N_\phi} c_{p\beta} S^{\beta\alpha}. \quad (4.10)$$

In the ONETEP program, $\{\phi_\alpha\}$ are NGWFs and $\{D_p\}$ are periodic-sinc (psinc) functions. In what follows, we will assume that the total Kohn-Sham energy functional at zero-temperature is minimised using the two-loop approach. We will assume that the inner loop is fully converged, and we will focus our attention to the outer loop where the NGWFs are optimised using direct energy minimisation methods.

4.2 Analysis of the constraints

4.2.1 Orthonormality of the Kohn-Sham states

The set of Kohn-Sham states $\{\psi_i\}$ forms an orthonormal basis set of the Hilbert space \mathcal{H} . At convergence, both the Hamiltonian and the one-particle density operators are diagonal in this representation, and hence their commutator $[\hat{H}, \hat{\rho}]$ is zero. Orthonormality of $\{\psi_i\}$ is a necessary physical constraint on the mathematical model that must be strictly maintained during the self-consistent optimisation. Using Eq. (4.1), orthonormality of $\{\psi_i\}$ can be written as:

$$\langle \psi_i | \psi_j \rangle = \sum_{\alpha, \beta=1}^{N_\phi} M_i^{\dagger\alpha} \langle \phi_\alpha | \phi_\beta \rangle M_j^\beta = \delta_{ij}. \quad (4.11)$$

Left-multiplying by M_{γ_i} and right-multiplying by $M_j^{\dagger\epsilon}$, and summing over α and β , the following result is obtained:

$$\sum_{ij} \sum_{\alpha, \beta=1}^{N_\phi} M_{\gamma i} M_i^{\dagger\alpha} \langle \phi_\alpha | \phi_\beta \rangle M_j^\beta M_j^{\dagger\epsilon} = \sum_{ij} M_{\gamma i} \delta_{ij} M_j^{\dagger\epsilon}. \quad (4.12)$$

Using Eq. (4.6) and the property $(M_{\alpha i})^{-1} = M_i^{\dagger\alpha}$, Eq. (4.12) yields $\langle \phi_\alpha | \phi^\beta \rangle = \delta_\alpha^\beta$, which is the biorthonormality condition as defined in Eq. (4.2). The interpretation of this result is that the transformation M_i^α converts the constraint of orthonormality of the Kohn-Sham states $\{\psi_i\}$ into a constraint of biorthonormality between the dual sets $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$, which holds if and only if Eqs. (4.3) and (4.4) are equivalent, so that duality between covariant and contravariant tensors is maintained.

4.2.2 Localisation of the non-orthogonal functions

Traditional electronic structure methods employ delocalised representations to expand the Kohn-Sham states, such as plane-waves or Gaussian functions. The constraint of localisation of $\{\phi_\alpha\}$ is a requisite for achieving linear-scaling cost with the number of atoms. In such cases, the search towards the minimum in the potential energy surface (PES) is transformed into a search along the metasurface of the PES compatible with exact localisation constraints. As a consequence, at convergence, $E_{loc} \geq E_{deloc}$, where E_{loc} and E_{deloc} are the ground state energies of the problem with and without localisation constraints, respectively [142–146]. Thus, approaches that use localised functions are variational with respect to the size of the localisation regions, and are equivalent to approaches with delocalised functions in the limit of infinitely large localisation regions. Demanding strict localisation of $\{\phi_\alpha\}$ should not result in a loss of orthonormality of $\{\psi_i\}$, or else unphysical results would be obtained. To avoid this scenario, as discussed in Sec. 4.2.1, Eqs. (4.3) and (4.4) must be verifiable and equivalent at the ground state.

A mathematically-concise manner to define the strict localisation constraints, based on the expansion of $\{\phi_\alpha\}$ in the $\{D_p\}$ basis set as in Eq. (4.7), is:

$$c_{p\alpha} = 0, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \notin \text{LR}(\phi_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi], \quad (4.13)$$

where $\text{LR}(\phi_\alpha(\mathbf{r}))$ is the localisation region of $\phi_\alpha(\mathbf{r})$. Without loss of generality, $\text{LR}(\phi_\alpha(\mathbf{r}))$ can be considered as a well-delimited, finite region of the real-space simulation cell. It is important to realise that the contravariant functions $\{\phi^\alpha\}$ are also spatially-localised in a well-delimited region of the real-space simulation cell. Figure 4.1 illustrates this property with an example. According to Eq. (4.3), the localisation region of $\phi^\alpha(\mathbf{r})$ is equal to the volume occupied by the union of all the localisation regions of $\{\phi_\beta\}$ for which $S^{\beta\alpha} \neq 0$:

$$\text{LR}(\phi^\alpha(\mathbf{r})) = \bigcup_{\substack{\beta=1 \\ S^{\beta\alpha} \neq 0}}^{N_\phi} \text{LR}(\phi_\beta(\mathbf{r})). \quad (4.14)$$

One can conclude from this result that the union of all the localisation regions of $\{\phi_\alpha\}$ and the union of all the localisation regions of $\{\phi^\alpha\}$ occupy exactly the same volume in real space. Therefore, in the $\{D_p\}$ representation (Eq. (4.8)), the localisation constraint on $\{\phi^\alpha\}$ can be defined as:

$$c_p^\alpha = 0, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \notin \text{LR}(\phi^\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]. \quad (4.15)$$

Imposing orthonormality of $\{\psi_i\}$ and localisation of $\{\phi_\alpha\}$ simultaneously does not overconstrain the Kohn-Sham equations. To prove this assertion, the number

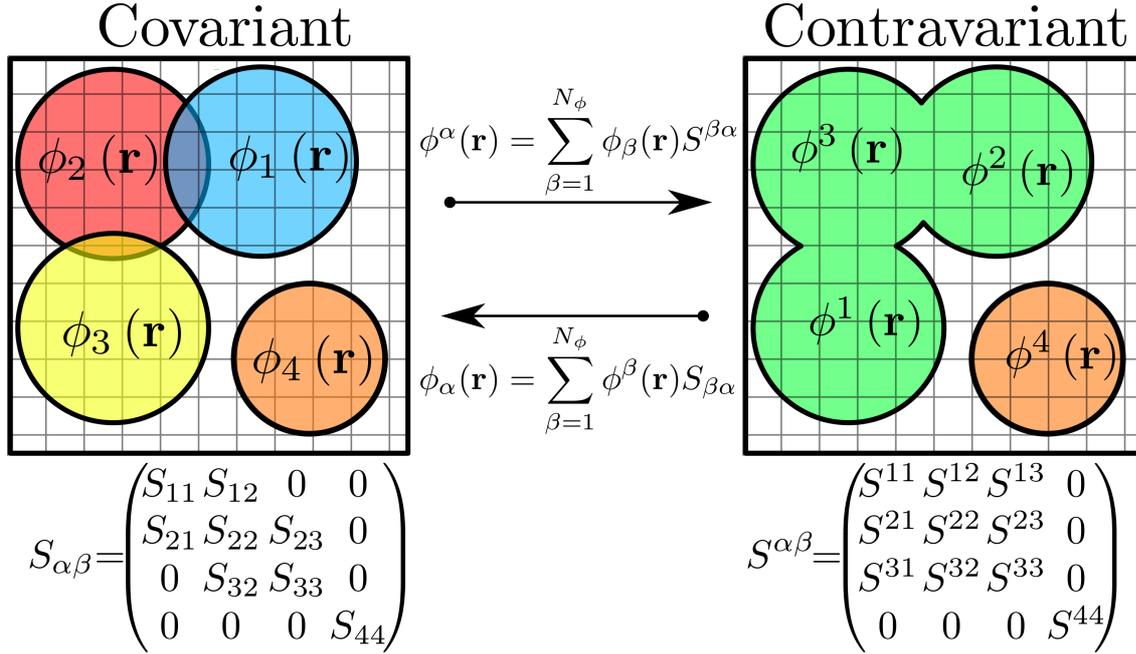


Figure 4.1: Two-dimensional depiction of the real-space localisation regions in a calculation involving four non-orthogonal localised functions. The covariant functions $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$ and $\phi_4(\mathbf{r})$ are localised within different circles, which produces a distinctive sparsity pattern in $S_{\alpha\beta}$. In contravariant space, $\phi^1(\mathbf{r})$, $\phi^2(\mathbf{r})$ and $\phi^3(\mathbf{r})$ are localised within the same localisation region, made up from the union of the localisation regions of $\phi_1(\mathbf{r})$, $\phi_2(\mathbf{r})$, $\phi_3(\mathbf{r})$, while $\phi^4(\mathbf{r})$ is localised in a circle, as a result of $\phi_4(\mathbf{r})$ not overlapping any other function. The volume covered by all the localisation regions is the same both in the covariant and contravariant representations. All the expansion coefficients associated to grid points outside of this volume are zero by definition of the localisation constraint.

irreducible degrees of freedom (DOF), defined as the minimum number of linearly-independent variational parameters capable of describing the system under these constraints, was counted. The number of DOF can be calculated as the sum of DOF of the unconstrained problem minus the number of constraints. In the unconstrained problem, each covariant function represented in real space, $\phi_\alpha(\mathbf{r})$, is expanded in terms of N_p coefficients $\{c_{p\alpha}\}$, and each contravariant function $\phi^\alpha(\mathbf{r})$ in terms of N_p coefficients $\{c_p^\alpha\}$. These add up to a total of $2 \times N_\phi \times N_p$ coefficients. Orthonormality of $\{\psi_i\}$ is imposed by ensuring that Eqs. (4.3) and (4.4) are equivalent, so only one of them needs to be taken into account when

counting the constraints. Also, they must hold point-wise, for each different grid-point \mathbf{r}_p , which amounts for a total of $N_\phi \times N_p$ constraints. The localisation constraint is imposed via Eq. (4.13). Assuming that there are N_α grid points inside $\text{LR}(\phi_\alpha(\mathbf{r}))$, the number of coefficients that are constrained to be equal to zero is $\sum_{\alpha=1}^{N_\phi} (N_p - N_\alpha) = N_\phi \times N_p - \sum_{\alpha=1}^{N_\phi} N_\alpha$. Therefore, the number of irreducible degrees of freedom, N_{DOF} , is:

$$N_{DOF} = 2 \times N_\phi \times N_p - N_\phi \times N_p - \left(N_\phi \times N_p - \sum_{\alpha=1}^{N_\phi} N_\alpha \right) = \sum_{\alpha=1}^{N_\phi} N_\alpha. \quad (4.16)$$

Perhaps the most intuitive and natural choice of N_{DOF} variational parameters to describe the system are the set of covariant coefficients associated to the points inside the localisation region of each of the covariant functions $\{\phi_\alpha\}$:

$$\Omega = \{c_{p\alpha}, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \in \text{LR}(\phi_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]\}. \quad (4.17)$$

Technically speaking, however, there are an infinite number of choices of N_{DOF} linearly-independent variational parameters that can also provide an equivalent description of the mathematical problem.

4.3 Steepest-denscent optimisation of the non-orthogonal functions

From this point onwards, it is assumed that the density kernel $K^{\alpha\beta}$ is an idempotent matrix that has been fully optimised during an inner loop. The focus will be put on

the optimisation of $\{\phi_\alpha\}$ in an outer loop, based on the steepest-descent algorithm for direct energy minimisation [40, 151]. The extension to the more efficient conjugate gradients method is straightforward. At every iteration l of the outer loop, the covariant functions $\{\phi_\alpha^{(l)}\}$ are updated as:

$$|\phi_\alpha^{(l+1)}\rangle = |\phi_\alpha^{(l)}\rangle + \lambda^{(l)} |\Delta_\alpha^{(l)}\rangle, \quad (4.18)$$

where $\{\Delta_\alpha^{(l)}\}$ are the search direction vectors compatible with the constraints applied onto $\{\phi_\alpha^{(l)}\}$, and $\lambda^{(l)}$ is the optimum step length, calculated after a line-search process. The property of tensor-invariance of the energy functional implies that a steepest-descent update of the contravariant functions $\{\phi^{\alpha(l)}\}$ of the type:

$$|\phi^{\alpha(l+1)}\rangle = |\phi^{\alpha(l)}\rangle + \lambda^{(l)} |\Delta^{\alpha(l)}\rangle, \quad (4.19)$$

must produce an equivalent energy minimisation path. The covariant and contravariant search directions are mutually dual:

$$|\Delta_\alpha^{(l)}\rangle = \sum_{\beta=1}^{N_\phi} |\Delta_\beta^{(l)}\rangle S_{\beta\alpha}^{(l)}, \quad (4.20)$$

$$|\Delta^{\alpha(l)}\rangle = \sum_{\beta=1}^{N_\phi} |\Delta_\beta^{(l)}\rangle S^{\beta\alpha(l)}. \quad (4.21)$$

In practice, the steepest-descent update takes place in the space of coefficients in terms of $\{D_p\}$:

$$c_{p\alpha}^{(l+1)} = c_{p\alpha}^{(l)} + \lambda^{(l)} d_{p\alpha}^{(l)}, \quad (4.22)$$

$$c_p^{\alpha(l+1)} = c_p^{\alpha(l)} + \lambda^{(l)} d_p^{\alpha(l)}, \quad (4.23)$$

where $\{d_{p\alpha}^{(l)}\}$ and $\{d_p^{\alpha(l)}\}$ are the expansion coefficients of $\{\Delta_\alpha^{(l)}\}$ and $\{\Delta^{\alpha(l)}\}$:

$$\Delta_\alpha^{(l)}(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) d_{p\alpha}^{(l)}, \quad (4.24)$$

$$\Delta^{\alpha(l)}(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) d_p^{\alpha(l)}, \quad (4.25)$$

with $\Delta_\alpha^{(l)}(\mathbf{r})$ and $\Delta^{\alpha(l)}(\mathbf{r})$ being the representation of $|\Delta_\alpha^{(l)}\rangle$ and $|\Delta^{\alpha(l)}\rangle$ in real space, respectively. The coefficients $\{d_{p\alpha}^{(l)}\}$ and $\{d_p^{\alpha(l)}\}$ are mutually dual:

$$d_{p\alpha}^{(l)} = \sum_{\beta=1}^{N_\phi} d_p^{\beta(l)} S_{\beta\alpha}^{(l)}, \quad (4.26)$$

$$d_p^{\alpha(l)} = \sum_{\beta=1}^{N_\phi} d_{p\beta}^{(l)} S^{\beta\alpha(l)}. \quad (4.27)$$

Orthonormality of $\{\psi_i^{(l)}\}$ must be maintained during the steepest-descent update. As discussed in Sec. 4.2.1, imposing orthonormality of $\{\psi_i\}$ is equivalent to imposing biorthonormality between the dual sets $\{\phi_\alpha^{(l)}\}$ and $\{\phi^{\alpha(l)}\}$. After updating the non-orthogonal functions as in Eqs. (4.18) and (4.19), Eq. (4.2) changes as:

$$\begin{aligned} \langle \phi_\alpha^{(l+1)} | \phi^{\beta(l+1)} \rangle &= \langle \phi_\alpha^{(l)} | \phi^{\beta(l)} \rangle \\ &+ \lambda^{(l)} \left(\langle \phi_\alpha^{(l)} | \Delta^{\beta(l)} \rangle + \langle \Delta_\alpha^{(l)} | \phi^{\beta(l)} \rangle \right) \\ &+ \left(\lambda^{(l)} \right)^2 \langle \Delta_\alpha^{(l)} | \Delta^{\beta(l)} \rangle. \end{aligned} \quad (4.28)$$

Exact biorthonormality can be imposed only if the following two equations are simultaneously true:

$$\langle \phi_\alpha^{(l)} | \Delta^{\beta(l)} \rangle = \delta_\alpha^\beta, \quad (4.29)$$

$$\langle \Delta_\alpha^{(l)} | \Delta^{\beta(l)} \rangle = \delta_\alpha^\beta. \quad (4.30)$$

The task of imposing these two conditions simultaneously is highly non-trivial. It is often more practical to impose only Eq. (4.29) and assume a small error of order $\mathcal{O}(\lambda^2)$. Thus, after the steepest descent update, biorthonormality is maintained only to first order, and orthonormality of $\{\psi_i^{(l+1)}\}$ is not exact. An immediate consequence of this is that the steepest-descent updates Eqs. (4.18) and (4.19) are not equivalent. Another manifestation of the lack of exact orthonormality of $\{\psi_i^{(l+1)}\}$ is that the total number of electrons in the system is not maintained:

$$\tilde{N}_e = \text{tr}[\hat{\rho}] = \sum_{\alpha,\beta=1}^{N_\phi} K^{\alpha\beta} S_{\beta\alpha}^{(l+1)} \neq N_e, \quad (4.31)$$

where $S_{\beta\alpha}^{(l+1)}$ is the up-to-date overlap matrix. To fix the problem of electron conservation, a rescale factor [196] can be applied to the density kernel as:

$$\tilde{K}^{\alpha\beta} = \frac{N_e}{\tilde{N}_e} K^{\alpha\beta}, \quad (4.32)$$

which guarantees that $\sum_{\alpha,\beta=1}^{N_\phi} \tilde{K}^{\alpha\beta} S_{\beta\alpha}^{(l+1)} = N_e$. This technique reduces the negative effect of the lack of orthonormality of $\{\psi_i^{(l)}\}$, which can be completely restored after re-entering the inner loop by allowing changes to the transformation matrix M_i^α .

Regarding the localisation constraint on $\{\phi_\alpha^{(l+1)}\}$, inspection of Eq. (4.18) reveals that the search directions $\{\Delta_\alpha^{(l)}\}$ must also be spatially-localised within the

same region as $\{\phi_\alpha^{(l)}\}$. This observation also applies to the contravariant functions.

Therefore:

$$\text{LR}(\Delta_\alpha(\mathbf{r})) = \text{LR}(\phi_\alpha(\mathbf{r})), \quad (4.33)$$

$$\text{LR}(\Delta^\alpha(\mathbf{r})) = \text{LR}(\phi^\alpha(\mathbf{r})). \quad (4.34)$$

In the representation in terms of $\{D_p\}$, the localisation constraints on $\{\Delta_\alpha^{(l)}\}$ and $\{\Delta^{\alpha(l)}\}$ can be written as:

$$d_{p\alpha}^{(l)} = 0, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \notin \text{LR}(\Delta_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi], \quad (4.35)$$

$$d_p^{\alpha(l)} = 0, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \notin \text{LR}(\Delta^\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]. \quad (4.36)$$

As demonstrated in Sec. 4.2.2, imposing orthonormality of $\{\psi_i\}$ and localisation of $\{\phi_\alpha\}$ does not overconstrain the Kohn-Sham equations. However, in the best-case scenario, orthonormality of $\{\psi_i\}$ can only be imposed to first order during the optimisation of $\{\phi_\alpha\}$. Methods such as kernel rescaling can stabilise the algorithm, but, ultimately, it is the task of the inner loop to completely fix and maintain orthonormality of $\{\psi_i\}$. Therefore, one can choose to introduce a larger error in the orthonormality of $\{\psi_i\}$ if in doing so, imposing localisation of $\{\phi_\alpha\}$ becomes simpler or faster. The risk of such an approach is that the inner loop might converge at a slower rate or even diverge. Thus, it is interesting to explore methods for which imposing localisation of $\{\phi_\alpha\}$ implies the lowest alteration to the biorthonormality constraint. When the self-consistent solution is reached, the search direction is perpendicular to the energy gradient, and the slope along the search direction compatible with the localisation constraints is zero:

$$\langle \Delta_\alpha^{(l)} | \frac{\delta E}{\delta \langle \phi_\alpha^{(l)} |} = 0, \quad (4.37)$$

When this occurs, a minimum on the energy metasurface that is compatible with the biorthonormality and localisation constraints has been reached.

In the following sections, the TSD and VLSD methods will be analysed. For clarity in the equations, in what follows, the iteration index l will be dropped, but it must be clear that all the steps presented hereafter have to be repeated at every new iteration of the outer loop.

4.4 Truncated search direction (TSD) method

The derivative of the energy functional with respect to $\{c_{p\alpha}\}$ is:

$$g_p^\alpha = \frac{\partial E}{\partial c_{p\alpha}} = \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) K^{\beta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (4.38)$$

where $\{g_p^\alpha\}$ are the expansion coefficients associated to the gradients $\{\Gamma^\alpha\}$:

$$\Gamma^\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) g_p^\alpha. \quad (4.39)$$

The coefficients $\{g_p^\alpha\}$ are contravariant tensors that cannot be used to update $\{c_{p\alpha}\}$ directly. Furthermore, they cannot yet be used to update $\{c_p^\alpha\}$ either, since they do not account for the constraints of orthonormality of $\{\psi_i\}$ and localisation of $\{\phi_\alpha\}$. Orthonormality of $\{\psi_i\}$ can be imposed to first order by projecting out the components of $\{\Gamma^\alpha\}$ parallel to $\{\phi_\beta\}$ [40, 211], so that Eq. (4.29) holds:

$$\tilde{g}_p^\alpha = \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) K^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_\phi} \phi_\beta(\mathbf{r}) K^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (4.40)$$

where $\{\tilde{g}_p^\alpha\}$ are the expansion coefficients of $\{\tilde{\Gamma}^\alpha\}$:

$$\tilde{\Gamma}^\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) \tilde{g}_p^\alpha. \quad (4.41)$$

LNV-type methods can avoid this projection, as accounting for idempotency of $K^{\alpha\beta}$ is equivalent to maintaining orthonormality of $\{\psi_i\}$ to first order [148]. At convergence of LNV, the derivative of the energy functional with respect to $\{c_{p\alpha}\}$ returns $\{\tilde{g}_p^\alpha\}$ directly. The covariant representation of $\{\tilde{g}_p^\alpha\}$ can be calculated by multiplying with $S_{\alpha\beta}$ as in Eq. (4.26). Kinetic energy [40, 176] and occupancy preconditioning [212] are applied, resulting in:

$$\tilde{d}_{p\alpha} = -\hat{P} \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\alpha(\mathbf{r}) - \sum_{\beta,\gamma=1}^{N_\phi} \phi_\beta(\mathbf{r}) S^{\beta\gamma} H_{\gamma\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (4.42)$$

where \hat{P} is a generic kinetic energy preconditioner. The search directions $\{\tilde{\Delta}_\alpha\}$ expanded by $\{\tilde{d}_{p\alpha}\}$ as

$$\tilde{\Delta}_\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) \tilde{d}_{p\alpha}, \quad (4.43)$$

are delocalised in real space, and therefore not yet suitable for updating $\{\phi_\alpha\}$ with the steepest-descent method. To solve the issue of delocalisation, an *ad hoc* truncation of $\{\tilde{\Delta}_\alpha\}$ can be utilised [149, 151], setting to zero all the expansion coefficients corresponding to grid points outside LR ($\Delta_\alpha(\mathbf{r})$):

$$d_{p\alpha} = \begin{cases} \tilde{d}_{p\alpha}, & \text{if } \mathbf{r}_p \in \text{LR}(\Delta_\alpha(\mathbf{r})) \\ 0, & \text{if } \mathbf{r}_p \notin \text{LR}(\Delta_\alpha(\mathbf{r})) \end{cases} \quad (4.44)$$

A steepest-descent update of $\{\phi_\alpha\}$ that is consistent with the constraint of localisation is now possible using the search directions $\{\Delta_\alpha\}$, which are represented in the $\{D_p\}$ basis set as:

$$\Delta_\alpha(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) d_{p\alpha}. \quad (4.45)$$

After truncation, the constraint of orthonormality of $\{\psi_i\}$ is not maintained, not even to first order. This can be seen by the fact that the inner product $\langle \Delta_\alpha | \phi^\beta \rangle$ is not zero for any pair $\alpha \neq \beta$. A new projection of the kind of Eq. (4.40) to re-impose biorthonormality to first order would not result in a suitable search direction, since it would delocalise $\{\Delta_\alpha\}$ again. The algorithm then continues with the updating $\{\phi_\alpha\}$ and the overlap matrices, $S_{\alpha\beta}$ and $S^{\alpha\beta}$, followed by the rescaling of the density kernel as in Eq. (4.32). The TSD algorithm relies upon the inner loop to restore orthonormality of $\{\psi_i\}$ during the next cycle.

In the TSD method, delocalisation originates when the gradient coefficients $\{g_p^\alpha\}$ are calculated in Eq. (4.38). At this step, the derivatives with respect to all the covariant coefficients $\{c_{p\alpha}\}$ are calculated, without taking into consideration that not all $c_{p\alpha}$ are variational degrees of freedom listed in Ω (Eq. (4.17)). After imposing orthonormality to first order, multiplying with the overlap matrix, and applying occupancy and kinetic energy preconditioning, the covariant search directions $\{\tilde{\Delta}_\alpha\}$ are still delocalised in real space. If they were used to update $\{\phi_\alpha\}$ in a steepest-descent method, the algorithm would fall out of the metasurface of the PES compatible with the localisation constraints. The only solution at this stage is

to truncate $\{\tilde{\Delta}_\alpha\}$, which is equivalent to eliminating all the information stored in the form of $\{\tilde{d}_{p\alpha}\}$ coefficients associated to \mathbf{r}_p outside the localisation regions.

There are two major caveats with the TSD method. First, the search direction vectors $\{\tilde{\Delta}_\alpha\}$ do not correspond to the search direction vectors of the problem with explicit localisation constraints. Instead, they are the search direction vectors corresponding to the unconstrained problem, where delocalisation is allowed. The second caveat is that the TSD method does not guarantee that $\{\Delta_\alpha\}$, after truncation, resemble the correct search direction vectors of the problem with strict localisation constraints and, therefore, the existence of a minimum along $\{\Delta_\alpha\}$ is not guaranteed either. Truncation of $\{\tilde{\Delta}_\alpha\}$ is a non-variational artifact that has to be introduced in order to obey the localisation constraints, and the resulting search direction vector is only approximate. In other words, there is always an error associated to the truncation of $\{\tilde{\Delta}_\alpha\}$. Only in the cases where the amount of information that is deleted during truncation is negligible compared to the amount of information that remains in the search direction vectors $\{\Delta_\alpha\}$, the TSD method could be considered as a valid approximation.

4.5 The variationally-localised search direction (VLSD) method

4.5.1 Mathematical principle of the VLSD method

The optimisation of the localised functions $\{\phi_\alpha\}$ is twice-constrained by the requirements of biorthonormality of $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$, and localisation of $\{\phi_\alpha\}$. While biorthonormality is a strong requirement upon the physical model, because it is equivalent to imposing orthonormality of the Kohn-Sham orbitals $\{\psi_i\}$, the locali-

sation constraints are optional. It was shown in Sec. 4.2.2 that requiring these two constraints simultaneously does not overconstrain the Kohn-Sham equations. The localisation constraints, however, reduce the variational degrees of freedom of the optimisation problem to $N_{DOF} \leq N_p$, as shown in Eq. (4.16).

The principle of the VLSD method can be enunciated as follows. Given the biorthonormality constraints Eqs. (4.9) and (4.10), and the localisation constraint Eq. (4.13), all the non-zero covariant coefficients $\{c_{p\alpha} \in \Omega\}$, where Ω was defined in Eq. (4.17) as the set of N_{DOF} irreducible degrees of freedom of the constrained optimisation problem, can be uniquely and unambiguously calculated from the subset of N_{DOF} contravariant coefficients:

$$\{c_p^\alpha, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \in \text{LR}(\phi_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]\}, \quad (4.46)$$

which correspond to grid points inside the volume of the individual localisation regions of the covariant functions, provided that the overlap matrix, $S_{\alpha\beta}$, is known.

The proof this statement begins by considering the group of all the covariant coefficients associated to a particular grid point, \mathbf{r}_p . It must be noted that there could be more than one $\{c_{p\alpha}\}$ coefficient associated to the same \mathbf{r}_p , originating from different $\{\phi_\alpha\}$. As a matter of fact, there is a finite number of $\{\phi_\alpha\}$ for which \mathbf{r}_p is inside their localisation regions, and, conversely, also a finite number of $\{\phi_\alpha\}$ for which \mathbf{r}_p is outside their localisation regions. Based on this criteria, for every single grid point \mathbf{r}_p , the following two subsets of tensorial indices can be formed:

$$\xi_p = \{\alpha \in [1, \dots, N_\phi] \mid \mathbf{r}_p \in \text{LR}(\phi_\alpha(\mathbf{r}))\} \quad (4.47)$$

$$\chi_p = \{\alpha \in [1, \dots, N_\phi] \mid \mathbf{r}_p \notin \text{LR}(\phi_\alpha(\mathbf{r}))\}. \quad (4.48)$$

This separation is unique and complete, in the sense that for every \mathbf{r}_p , a given index α always belongs to one, and only one, of these subsets. Hence, $N_\phi = N_\xi + N_\chi$, where N_ξ and N_χ are the number of indices included in each subset. The notation used in this work is as follows. Expressions such as $\{\phi_\eta \mid \forall \eta \in \xi_p\}$, refer to the subset of covariant functions ϕ_η for which $\mathbf{r}_p \in \text{LR}(\phi_\eta(\mathbf{r}))$. The notation extends to the contravariant functions. For example, $\{\phi^\tau \mid \forall \tau \in \chi_p\}$, refers to the subset of contravariant functions ϕ^τ for which $\mathbf{r}_p \notin \text{LR}(\phi_\tau(\mathbf{r}))$, with ϕ_τ being the dual of ϕ^τ , calculated as $\phi_\tau = \sum_{\alpha=1}^{N_\phi} \phi^\alpha S_{\alpha\tau}$. In the same fashion, the notation also extends to the search direction vectors and, in general, to any tensor. For example, $\{d_p^\iota \mid \forall \iota \in \xi_p\}$, refers to the subset of contravariant coefficients $\{d_p^\iota\}$ for which $\mathbf{r}_p \in \text{LR}(\phi_\iota(\mathbf{r}))$, with $d_{p\iota} = \sum_{\alpha=1}^{N_\phi} d_p^\alpha S_{\alpha\iota}$. Figure 4.2 further clarifies this notation with an example.

Based on this point-wise subset separation, three distinct overlap matrices can be defined for every single grid point \mathbf{r}_p :

$$A_{p,\eta\iota} = \langle \phi_\eta \mid \phi_\iota \rangle, \quad \forall \eta, \iota \in \xi_p \quad (4.49)$$

$$B_{p,\iota\sigma} = \langle \phi_\iota \mid \phi_\sigma \rangle, \quad \forall \iota \in \xi_p, \quad \forall \sigma \in \chi_p \quad (4.50)$$

$$C_{p,\sigma\tau} = \langle \phi_\sigma \mid \phi_\tau \rangle, \quad \forall \sigma, \tau \in \chi_p. \quad (4.51)$$

The matrices $A_{p,\eta\iota}$, $B_{p,\iota\sigma}$ and $C_{p,\sigma\tau}$ are unique different mappings of the full over-

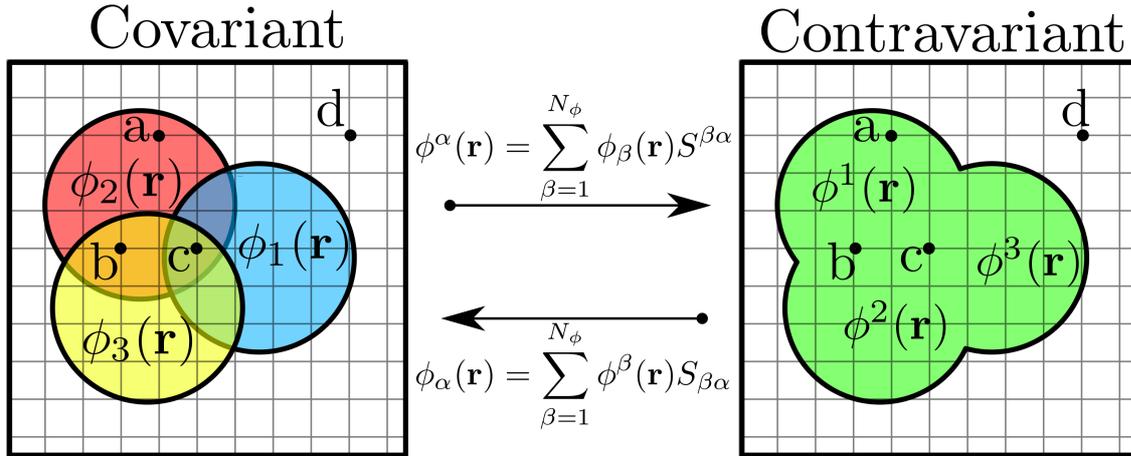


Figure 4.2: Example of point-wise subset separation in a hypothetical calculation with a two-dimensional real space simulation cell, containing three covariant functions localised within circles. The grid point \mathbf{r}_a is inside LR ($\phi_2(\mathbf{r})$), and outside LR ($\phi_1(\mathbf{r})$) and LR ($\phi_3(\mathbf{r})$). Therefore $\xi_a = \{2\}$ and $\chi_a = \{1, 3\}$. The grid point \mathbf{r}_b is inside LR ($\phi_2(\mathbf{r})$) and LR ($\phi_3(\mathbf{r})$), and outside LR ($\phi_1(\mathbf{r})$). Therefore $\xi_b = \{2, 3\}$ and $\chi_b = \{1\}$. The grid point \mathbf{r}_c is inside all the localisation regions. Therefore $\xi_c = \{1, 2, 3\}$ and χ_c is an empty subset. The grid point \mathbf{r}_d is outside all the localisation regions. Therefore, ξ_d is an empty subset and $\chi_d = \{1, 2, 3\}$.

lap matrix $S_{\alpha\beta}$. $A_{p,\eta\iota}$ is a square matrix of dimensions $N_\xi \times N_\xi$; $B_{p,\iota\sigma}$ is a rectangular matrix of dimensions $N_\xi \times N_\chi$; and $C_{p,\sigma\tau}$ is a square matrix of dimensions $N_\chi \times N_\chi$. If a the grid point \mathbf{r}_p is inside all the localisation regions, then χ_p is an empty subset, and $B_{p,\iota\sigma}$ and $C_{p,\sigma\tau}$ are matrices of rank zero. If \mathbf{r}_p is outside all the localisation regions, then ξ_p is an empty subset, and $A_{p,\eta\iota}$ and $B_{p,\iota\sigma}$ are matrices of rank zero. Equation (4.9) can now be split into two expressions:

$$c_{p\iota} = \sum_{\eta=1}^{N_\xi} c_p^\eta A_{p,\eta\iota} + \sum_{\sigma=1}^{N_\chi} c_p^\sigma B_{p,\sigma\iota}^\dagger, \quad \forall \eta, \iota \in \xi_p, \quad \forall \sigma \in \chi_p \quad (4.52)$$

$$c_{p\tau} = \sum_{\eta=1}^{N_\xi} c_p^\eta B_{p,\eta\tau} + \sum_{\sigma=1}^{N_\chi} c_p^\sigma C_{p,\sigma\tau}, \quad \forall \eta \in \xi_p, \quad \forall \sigma, \tau \in \chi_p. \quad (4.53)$$

No approximation was made in the above separation, and duality between $\{\phi_\alpha\}$

and $\{\phi^\alpha\}$ is rigorously maintained. In this notation, the localisation constraint on $\{\phi_\alpha\}$, as given in Eq. (4.13), can be written exactly as:

$$c_{p\tau} = 0, \quad \forall \tau \in \chi_p. \quad (4.54)$$

Substituting Eq. (4.54) in the left-hand side of Eq. (4.53) leads to:

$$c_p^\sigma = - \sum_{\eta=1}^{N_\xi} \sum_{\tau=1}^{N_\chi} c_p^\eta B_{p,\eta\tau} (C_p^{-1})^{\tau\sigma}, \quad \forall \eta \in \xi_p, \forall \sigma, \tau \in \chi_p. \quad (4.55)$$

Eq. (4.55) can now be introduced into (4.52) to obtain:

$$c_{p\iota} = \sum_{\eta=1}^{N_\xi} c_p^\eta \left[A_{p,\eta\iota} - \sum_{\tau,\sigma=1}^{N_\chi} B_{p,\eta\tau} (C_p^{-1})^{\tau\sigma} B_{p,\sigma\iota}^\dagger \right], \quad \forall \eta, \iota \in \xi_p, \forall \sigma, \tau \in \chi_p. \quad (4.56)$$

Equations (4.54), (4.55) and (4.56) show that, given the overlap matrix, $S_{\alpha\beta}$, the values of $\{c_p^\sigma, \forall \sigma \in \chi_p\}$ and $\{c_{p\iota}, \forall \iota \in \xi_p\}$ are uniquely determined by the values of the coefficients $\{c_p^\eta, \forall \eta \in \xi_p\}$. Therefore, the optimisation problem with biorthonormality and localisation constraints is completely and unambiguously determined by the contravariant coefficients:

$$\Omega' = \{c_p^\eta, \quad \forall \eta \in \xi_p, \quad \forall p \in [1, \dots, N_p]\}. \quad (4.57)$$

A more detailed inspection of Eq. (4.57) reveals that the $\{c_p^\eta\}$ coefficients listed in Ω' correspond exactly to all the contravariant coefficients inside the localisation regions of the covariant functions. The set Ω' can be redefined in formal terms as:

$$\Omega' = \{c_p^\alpha, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \in \text{LR}(\phi_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]\}. \quad (4.58)$$

Hence, the initial hypothesis is demonstrated. The question of whether Ω' is a valid set of N_{DOF} irreducible degrees of freedom is left unresolved in this work. To prove such condition, it would be enough to demonstrate that the overlap matrix $S_{\alpha\beta}$ can be fully built from the coefficients in Ω' .

4.5.2 The VLSD algorithm

The equations introduced in Sec. 4.5.1 can be used to design an algorithm for the constrained optimisation of $\{\phi_\alpha\}$ with the steepest descent iterative method. The search direction vectors $\{\Delta_\alpha\}$ are calculated in the current $\{\phi_\alpha\}$ representation, of which the overlap matrix $S_{\alpha\beta}$ is fully known. Also, the search direction vectors are subject to the biorthonormality constraints Eqs. (4.26) and (4.27), and the localisation constraint Eq. (4.35). The set of covariant search direction coefficients $\{d_{p\alpha}\}$ that can be used to update the variational degrees of freedom $\{c_{p\alpha} \in \Omega\}$ is denoted as:

$$\Upsilon = \{d_{p\alpha}, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \in \text{LR}(\Delta_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]\}, \quad (4.59)$$

while the set of contravariant search direction coefficients $\{d_p^\alpha\}$ that can be used to update $\{c_p^\alpha \in \Omega'\}$ is denoted as:

$$\Upsilon' = \{d_p^\alpha, \forall p \in [1, \dots, N_p] \mid \mathbf{r}_p \in \text{LR}(\Delta_\alpha(\mathbf{r})), \forall \alpha \in [1, \dots, N_\phi]\}. \quad (4.60)$$

Under these conditions, the equations in Sec. 4.5.1 demonstrate that the values of $\{d_{p\alpha} \in \Upsilon\}$ are fully and univocally determined by the values of $\{d_p^\alpha \in \Upsilon'\}$, via the recursive application of Eqs. (4.54), (4.55) and (4.56) for every grid point \mathbf{r}_p .

The VLSD algorithm begins with the calculation of the values of the contravariant search direction coefficients in Υ' . First, the energy gradients $\{g_p^\alpha\}$ are calculated as:

$$g_p^\alpha = \frac{\partial E}{\partial c_{p\alpha}} = \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) K^{\beta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad \forall c_{p\alpha} \in \Omega. \quad (4.61)$$

A central difference between of the VLSD algorithm with respect to the TSD algorithm is that the derivative of the energy is calculated exclusively with respect to $\{c_{p\alpha} \in \Omega\}$, and not with respect to all the $\{c_{p\alpha}\}$ coefficients, which is consistent with the choice of variational degrees of freedom. Orthonormality of $\{\psi_i\}$ is maintained to first order with the projection:

$$\tilde{g}_p^\alpha = \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) K^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_\phi} \phi_\beta(\mathbf{r}) K^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad \forall c_{p\alpha} \in \Omega. \quad (4.62)$$

Kinetic energy [40, 176] and occupancy preconditioning [212] are applied in the contravariant representation point-wise:

$$d_p^\alpha = -\hat{P} \left[\sum_{\beta=1}^{N_\phi} \hat{H} \phi_\beta(\mathbf{r}) S^{\beta\alpha} - \sum_{\beta,\gamma,\delta=1}^{N_\phi} \phi_\beta(\mathbf{r}) S^{\beta\gamma} H_{\gamma\delta} S^{\delta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad \forall c_{p\alpha} \in \Omega. \quad (4.63)$$

The resulting contravariant search direction coefficients correspond to the entire set $\{d_{p\alpha} \in \Upsilon\}$. At this stage, a loop runs over all the grid points \mathbf{r}_p . At each point, the subset separation (4.47) and (4.48) is established, followed by the determination of the overlap matrices $A_{p,\eta\iota}$, $B_{p,\iota\sigma}$ and $C_{p,\sigma\tau}$ with Eqs. (4.49), (4.50) and (4.51). Then, Eqs. (4.54), (4.55) and (4.56) are applied recursively to obtain the remaining non-zero covariant and contravariant search direction coefficients. The VLSD algorithm is depicted in Figure 4.3.

The localisation constraint is introduced using Eq. (4.54):

$$d_{p\sigma} = 0, \quad \forall \sigma \in \chi_p. \quad (4.64)$$

The contravariant search directions coefficients $\{d_p^\sigma, \forall \sigma \in \chi_p\}$ are calculated using Eq. (4.55):

$$d_p^\sigma = - \sum_{\eta=1}^{N_\xi} \sum_{\tau=1}^{N_\chi} d_p^\eta B_{p,\eta\tau} (C_p^{-1})^{\tau\sigma}, \quad \forall \eta \in \xi_p, \forall \sigma, \tau \in \chi_p. \quad (4.65)$$

And the non-zero covariant search direction coefficients $\{d_{p\iota}, \forall \iota \in \xi_p\}$ are calculated using Eq. (4.56):

$$d_{p\iota} = \sum_{\eta=1}^{N_\xi} d_p^\eta \left[A_{p,\eta\iota} - \sum_{\tau,\sigma=1}^{N_\chi} B_{p,\eta\tau} (C_p^{-1})^{\tau\sigma} B_{p,\sigma\iota}^\dagger \right], \quad \forall \eta, \iota \in \xi_p, \forall \sigma, \tau \in \chi_p. \quad (4.66)$$

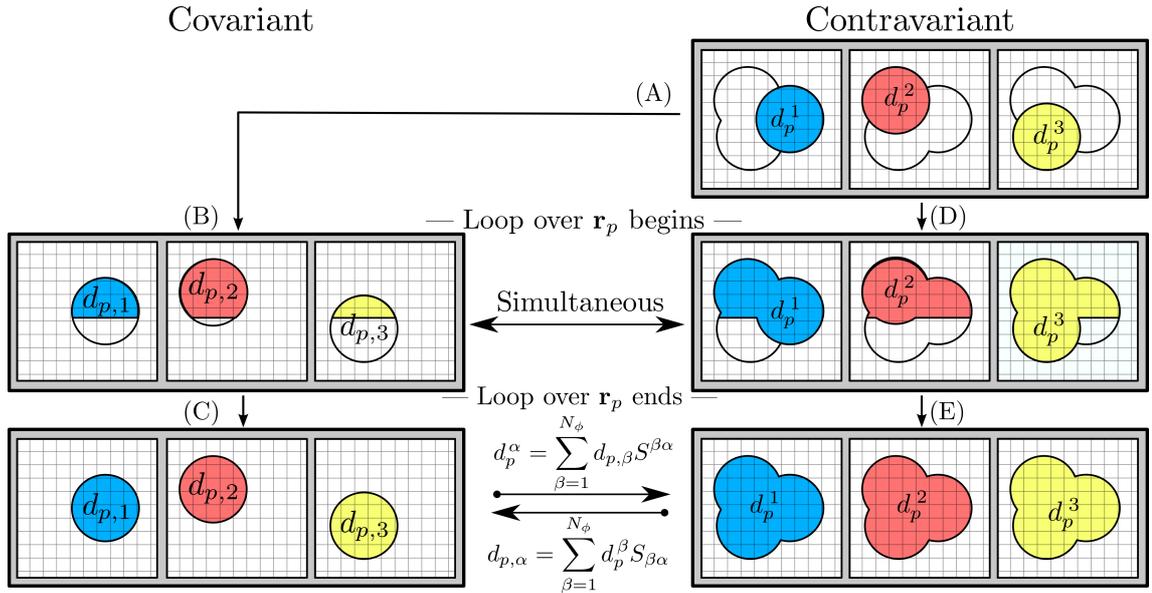


Figure 4.3: Schematic representation of the VLSD method. In this example, there are three mutually overlapping covariant functions, $\{\phi_\alpha\}$, localised within circles in a two-dimensional simulation cell. The algorithm begins in (A), calculating the derivative of the energy with respect to the variational parameters in Ω , imposing orthogonality of $\{\psi_i\}$ and applying occupancy and kinetic energy preconditioning Eq. (4.63). The result are the values of $\{d_p^\alpha \in \Upsilon'\}$, which correspond to the grid points confined within the circles. Then, a loop over all the grid points \mathbf{r}_p starts. Steps (B) and (C) represent the progress of the algorithm at a certain intermediate stage, where some, but not all, grid points have been scrutinised. In most situations, only the covariant search directions are needed to proceed with the steepest descent update. In such cases, the path through (B) and (C) calculates $\{d_{p1}\}$, $\{d_{p2}\}$ and $\{d_{p3}\}$ directly by applying Eq. (4.64) and Eq. (4.66) to the current grid point. Alternatively, the path through (D) and (E) calculates the remaining contravariant search direction coefficients using Eq. (4.65). The VLSD algorithm is capable of calculating the covariant and contravariant search direction coefficients simultaneously by running steps (B) and (D) as independent processes. At the conclusion of the look over \mathbf{r}_p , the resulting search directions are directly localised within the established localisation regions and duality between $\{d_{p\alpha}\}$ and $\{d_p^\alpha\}$ holds exactly.

At the conclusion of the loop over \mathbf{r}_p , the full covariant and contravariant search direction vectors, $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$, can be built on the grid as:

$$\Delta_\alpha(\mathbf{r}) = \sum_p^{N_p} D_p(\mathbf{r}) d_{p\alpha}, \quad (4.67)$$

$$\Delta^\alpha(\mathbf{r}) = \sum_p^{N_p} D_p(\mathbf{r}) d_p^\alpha, \quad (4.68)$$

and used to update $\{\phi_\alpha\}$ and $\{\phi^\alpha\}$, respectively, using the steepest descent method. The formalism of the VLSD method ensures that Eqs. (4.26) and (4.27) hold, and hence $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$ are mutually dual. More importantly, $\{\Delta_\alpha\}$ and $\{\Delta^\alpha\}$ are fully localised by construction. After updating $\{\phi_\alpha\}$, $S_{\alpha\beta}$ and $S^{\alpha\beta}$, the density kernel is rescaled to assure electron conservation and to reduce the impact of second-order error terms, and the total energy functional is evaluated. The inner loop is then responsible of fixing orthonormality of $\{\psi_i\}$ beyond second-order terms during the next cycle.

4.6 Results and discussion

4.6.1 Description of the calculations

The current implementation of the VLSD algorithm in ONETEP is not optimised for computational performance, and therefore execution timings are meaningless. Also, currently, calculations using the VLSD algorithm can only be run in serial (one processor), which limits the size of the molecules that can be simulated to a maximum of a few tens of atoms. The purpose of the VLSD implementation that was used to obtain the results exposed in this section is solely that of testing and benchmarking the method. In the current implementation, the covariant search direction coefficients $\{d_{p\alpha}\}$ are calculated using Eq. (4.66) directly. This is equivalent to follow the path through (A), (B) and (C) in Fig. 4.3.

The calculations with the ONETEP program used a kinetic energy cut-off of 800 eV, the PBE exchange-correlation functional and norm-conserving pseudopo-

tentials. An inner loop performed 3 LNV iterations to optimise the density kernel. The size of the localisation regions of $\{\phi_\alpha\}$ is taken into account by performing calculations with different R_α in the range of 3.0 Å to 7.0 Å. Periodic boundary conditions were required. Interactions between periodic images are largely avoided by confining the systems within a cubic simulation cell of 52.91 Å width, ensuring that a vacuum space of at least 45.46 Å exists between periodic images. The possibility of long-range ionic interactions affecting the convergence of the calculations was discarded after performing calculations with spherical Coulomb cut-off [186] and obtaining identical results. Fast Fourier transforms were calculated using cubic FFT-boxes [174, 175] of size 44.45 Å and 189 grid points on each lattice vector direction.

A set of sixteen benchmark molecular systems was chosen and organised in four groups, each of which contains molecules that share a common property. These groups will be referred to hereafter as G1, G2, G3 and G4, respectively, and are constructed as follows. G1 contains molecules with at least one chlorine atom: Cl₂ (chlorine diatomic), CH₃Cl (methyl chloride), HCl (hydrogen chloride), and NaCl (sodium chloride). The first three of these chlorinated molecules contain covalent bonds, while NaCl is an ion pair with a long ionic bond. CH₃Cl, HCl and NaCl are expected to show non-zero electric dipole. Calculations on these systems with Gaussian basis set approaches require the addition of diffuse functions in order to achieve chemical accuracy in the prediction of their electronic properties [78, 82]. Therefore, a certain degree of natural delocalisation of their molecular orbitals is expected. G2 is formed by the phosphoric acid molecule H₃PO₄ and the subsequent anions H₂PO₄⁻, HPO₄²⁻ and PO₄³⁻, of increasing total negative charge. Also, for such molecules, calculations with Gaussian basis sets need the addition

of diffuse functions to correctly describe electronic properties [78, 82]. Delocalisation of their molecular orbitals is expected to increase with the total charge. G3 encompasses a set of common small molecules that form short covalent and hydrogen bonds. These are H₂O (water), NH₃NH₃ (ammonia dimer, hydrogen bonded), C₆H₆ (benzene, aromatic) and NH₄⁺ (ammonium cation). The small size and reduced electric dipole of these test systems is expected to result in a high natural localisation of their molecular orbitals. Finally, G4 contains four amino acids that are often found in proteins and other organic complexes: lysine (lys⁺), glutamic acid (glu⁻), glycine (gly) and serine (ser). Lysine and glutamic acid are positively and negatively charged, respectively, glycine is non-polar and serine is polar. Natural localisation of the molecular orbitals is expected to be observed in these systems. The number of NGWFs per atom is as follows: one for H, four for C, N, O and Na, and nine for Cl and P.

4.6.2 Analysis of the TSD method

The effect of truncation of the search direction vectors in calculations with the TSD method was analysed. The coefficient μ , defined as

$$\mu = \sum_{\alpha=1}^{N_{\phi}} \frac{\sum_{\mathbf{r}_p \in \text{LR}(\Delta_{\alpha}(\mathbf{r}))} d_{p\alpha}^2}{\sum_{p=1}^{N_p} \tilde{d}_{p\alpha}^2}, \quad (4.69)$$

compares the sum of the squares of $\{d_{p\alpha}\}$ inside $\text{LR}(\Delta_{\alpha}(\mathbf{r}))$ and the sum of the squares of $\{\tilde{d}_{p\alpha}\}$ in the entire simulation cell prior to truncation. By definition, μ is confined to be in the range $0 \leq \mu \leq 1$, where $\mu = 0$ means that the truncation step deletes all the information available, and $\mu = 1$ means that no truncation happens at all. It must be emphasised that calculations with the VLSD method always result in $\mu = 1$, by construction. The value of μ was plotted against the

iteration counter l for all the test systems and all the different values of R_α . Figure 4.4 shows the results.

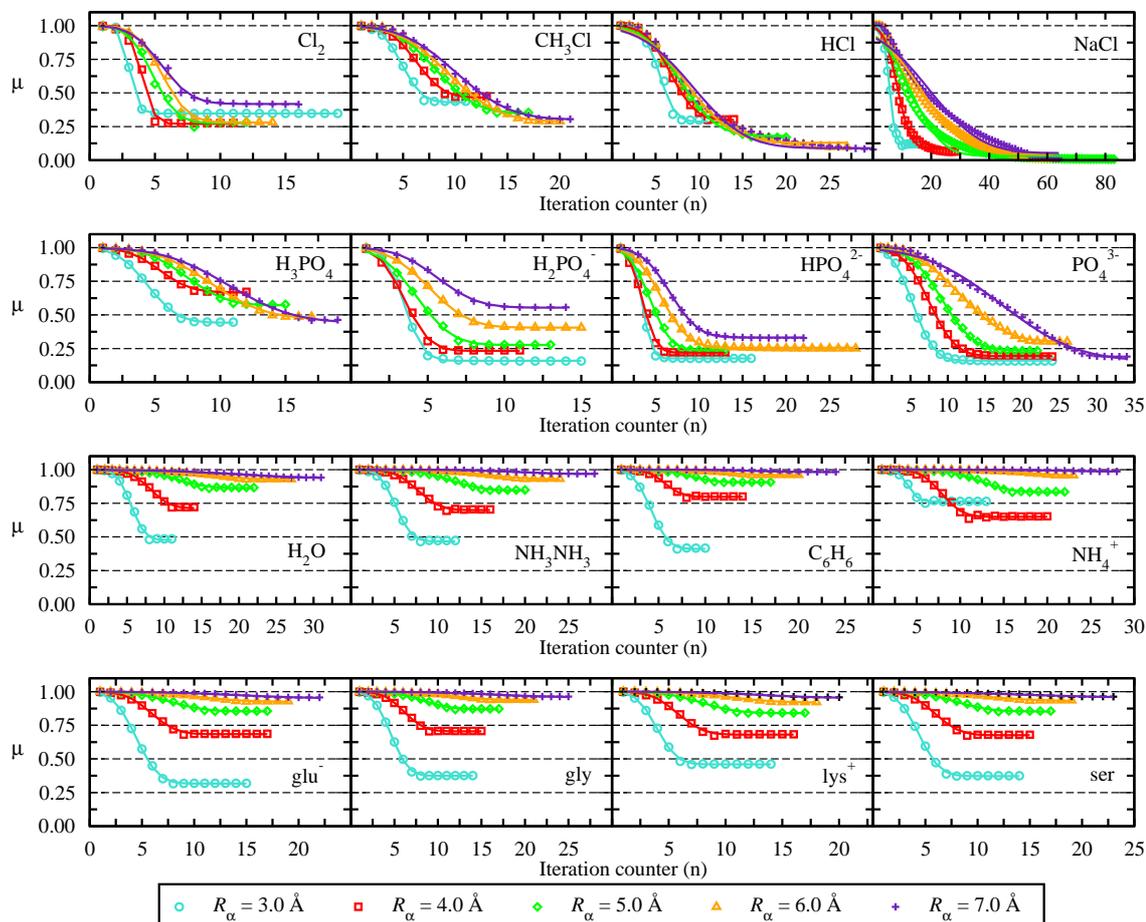


Figure 4.4: Evolution of the parameter μ with the iteration counter l for different values of R_α . The solid lines are the Gaussian broadening curves used to fit the results.

A first observation is that, as the number of iterations increases, μ saturates to a constant value, denoted by μ_0 , indicating that the ratio of the amount of information that is deleted to the amount of information that is kept does no longer change. This trend was observed in calculations on all the four groups of test systems with any value of R_α . The pace at which μ saturates to μ_0 is determined by the slope of the curve. To quantify and analyse this behaviour, a Gaussian broadening curve with equation

$$\mu(l) = \mu_0 + \frac{(1 - \mu_0)}{2} \left(1 - \operatorname{erf} \left[\frac{l - l_0}{\theta_0} \right] \right), \quad (4.70)$$

was fitted to the results, where erf is the error function, l_0 is the iteration number for which $\mu = \frac{\mu_0}{2}$, and θ_0 is the broadening factor (inversely proportional to the slope). The choice of the Gaussian broadening function was motivated after testing other alternatives, such as an exponential function, hyperbolic tangent, cumulative Gamma distribution, Gaussian distribution, and Fermi-Dirac distribution. It was found that the Gaussian broadening curve resulted in the smallest fitting error, with a correlation coefficient superior to 0.99 in all cases.

The values of μ_0 , l_0 and θ_0 for each of the test cases in every group were plotted for the different values of R_α . Figure 4.5 shows the results. The relation between μ_0 and R_α , and between l_0 , θ_0 and R_α , must be analysed separately. Calculations on molecules in G1 and G2 resulted in a very low μ_0 , typically below 0.50, and in some cases like HCl, NaCl and PO_4^{3-} , below 0.25 for most of the values of R_α . The values of μ_0 for NaCl are remarkably low, reaching a minimum of 0.03 with $R_\alpha = 5.0 \text{ \AA}$. Such low values indicate that, in calculations on molecules in G1 and G2, during the last iterations, comparatively, there is more information contained outside the localisation spheres than inside. The value of μ_0 does not increase with R_α for molecules in G1 and G2, however, it consistently increases with R_α for molecules in G3 and G4. Regarding l_0 and θ_0 , it seems plausible to think that these two variables are in some way inter-related, for a smoother decay is expected to result in a larger value of l_0 and θ_0 , and viceversa. It is clear from Figure 4.4 that l_0 and θ_0 increase with R_α in every system in every group. Therefore, the non-variational truncation of the search direction results in the deletion of less information per iteration as R_α increases, measured in relative terms.

With the TSD method, the search direction is calculated as unconstrained and allowed to be delocalised until the final stage, where the truncation step imposes the localisation constraints. According to the results, during the first few iterations, a small fraction of the total information is deleted. Then μ decays to the saturation value μ_0 at speed that is inversely proportional to R_α . A low value of μ_0 can be interpreted as a failure of the implicit assumption of TSD: that the truncated search direction of the unconstrained problem resembles the search direction of the problem with exact localisation constraints. Truncation has a stronger effect in calculations on molecules in G1 and G2 than in calculations on molecules in G3 and G4. This is in consonance with the expectations of more diffuse molecular orbitals on systems in G1 and G2 than on G3 and G4.

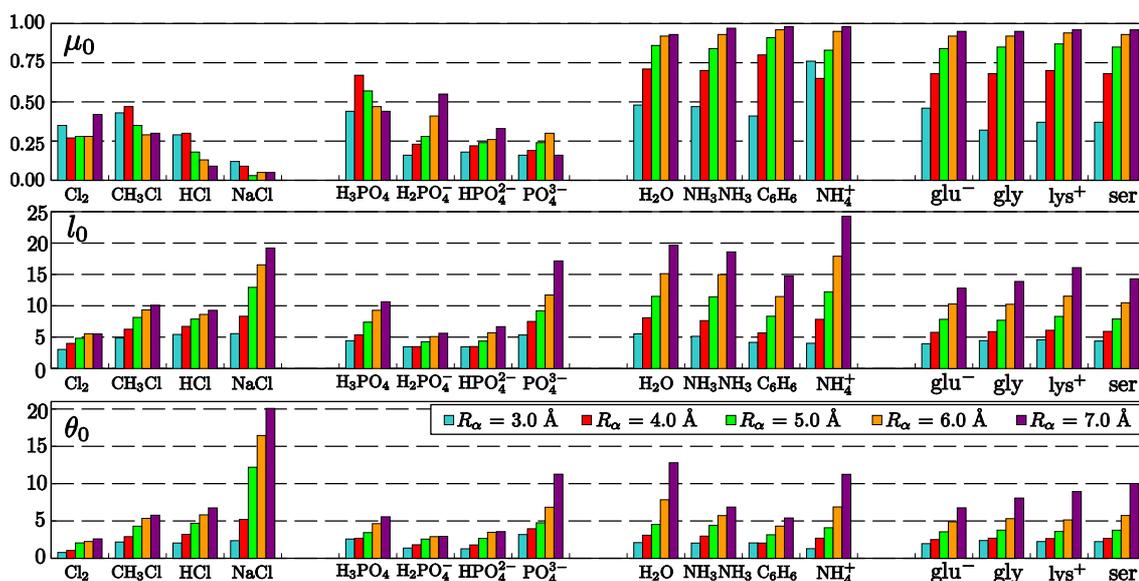


Figure 4.5: Progression of μ_0 , l_0 and θ_0 , calculated after fitting a Gaussian broadening function through the results, for different values of R_α .

4.6.3 Convergence properties

The convergence rate of the TSD and VLSD methods, defined as the number of iterations that is necessary to achieve self-consistent convergence, was compared.

Convergence is considered to be achieved when the root-mean-square (RMS) of the gradient along the search direction falls below a tolerance threshold of 3.8×10^{-5} eV, and the change in the total energy below a tolerance threshold of 1×10^{-7} eV/atom. All the TSD and VLSD calculations converged according to these criteria, with the only exception of H_2PO_4^- with $R_\alpha = 3.0 \text{ \AA}$, in which the total energy converged to the established tolerance threshold with an RMS gradient of 6.8×10^{-4} eV, both in TSD and VLSD calculations. The total number of iterations to achieve convergence with the TSD and VLSD methods, denoted as l_{TSD} and l_{VLSD} , respectively, is shown in Table 4.1. The relative speed-up coefficient $\zeta = l_{TSD}/l_{VLSD}$ gives a comparison of the relative convergence properties of the two methods: $\zeta > 1$ indicates that VLSD needs less iterations than TSD to converge, whereas $\zeta < 1$ indicates otherwise. The value of ζ in every calculation is shown in Figure 4.6.

Table 4.1: Number of iterations required to achieve self-consistent convergence in calculations with the TSD and VLSD methods. The values refer to the ratio l_{TSD}/l_{VLSD} .

R_α (Å)	Cl_2	CH_3Cl	HCl	NaCl	H_3PO_4	H_2PO_4^-	HPO_4^{2-}	PO_4^{3-}
3.0	19/12	11/9	11/10	16/13	11/10	15/14	16/16	24/11
4.0	11/8	13/11	14/12	28/17	12/10	11/14	13/14	24/11
5.0	12/12	17/14	20/18	81/25	15/13	13/11	13/14	22/13
6.0	14/12	20/15	27/32	62/28	17/15	15/12	28/12	26/15
7.0	16/9	21/18	43/22	64/31	19/16	14/12	22/12	34/20
R_α (Å)	H_2O	NH_3NH_3	C_6H_6	NH_4^+	glu^-	gly	lys^+	ser
3.0	11/11	12/10	10/8	13/11	15/15	14/14	14/14	14/13
4.0	14/15	16/14	14/15	20/14	17/16	15/14	16/16	15/14
5.0	22/20	20/17	17/17	22/24	17/18	17/16	17/16	17/16
6.0	27/21	24/21	20/18	23/26	19/19	21/21	18/17	19/17
7.0	31/29	28/23	24/22	28/27	22/20	25/22	20/19	23/22

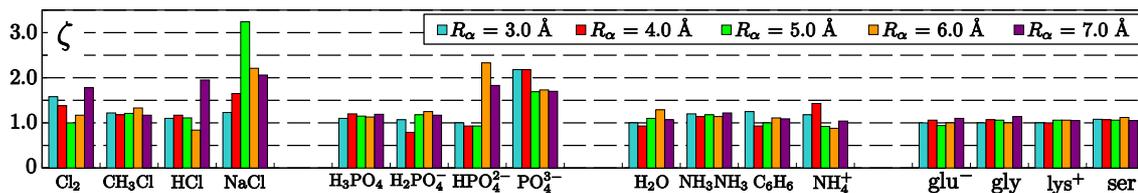


Figure 4.6: Calculated value of the relative speed-up ζ for each test system and increasing R_α .

The results show that, with some exceptions, calculations using the VLSD method require less iterations to converge than calculations with the TSD method. The analysis of the value of ζ indicates that calculations on molecules in G1 and G2 seem to benefit more of the VLSD method than calculations on molecules in G3 and G4. The largest values of ζ are obtained for NaCl, HPO_4^{2-} and PO_4^{3-} , which also show some of the lowest values of μ_0 in TSD calculations. The maximum speed-up is obtained for NaCl with $R_\alpha = 5.0 \text{ \AA}$, which has the lowest μ_0 coefficient of all in TSD calculations, and corresponds to a value of $\zeta = 3.24$, meaning that TSD needs more than three times the number of VLSD iterations to converge. On the other hand, calculations on molecules in G3 and G4 seem to require approximately the same number of iterations with TSD than with VLSD. These results seem to point out that the VLSD method is more robust compared to TSD, especially in molecules where a certain degree of delocalisation of the unconstrained molecular orbitals is expected.

4.6.4 Validation tests

A number of tests were performed to confirm that both the TSD and VLSD methods are capable of predicting structural and electronic properties correctly. According to the results in the previous section, NaCl converges very differently with the TSD and VLSD methods, whereas H_2O shows a similar convergence pattern

with both. These two molecules were chosen to perform validation tests. The structures were optimised with the BFGS geometry optimisation algorithm [204], using either the TSD or the VLSD methods to calculate the total energy and the atomic forces. The equilibrium bond length, L_0 , of the Na-Cl bond in NaCl and of the average O-H bond in H₂O were determined. Then, the magnitude of the total electric dipole moment at the equilibrium geometry, p_0 , was calculated. The results were compared to the values of L_0 and p_0 obtained with calculations using the NWCHEM program [85] and the near-complete aug-cc-pVTZ Gaussian basis set, which contains a large number of polarisation and diffuse functions [79, 80]. The values of L_0 and p_0 are shown in Table 4.2.

The equilibrium bond lengths L_0 of NaCl and H₂O calculated with the TSD and VLSD methods in ONETEP and with NWCHEM agree very well, within a < 1 % deviation margin. Convergence of L_0 to a single value with respect to R_α is fast both in NaCl and H₂O calculations.

On the other hand, the value of the total electric dipole moment seems to be more sensitive to changes in the accuracy of the basis set and the numerical method. In ONETEP calculations, p_0 converges to a single value when R_α increases, both in NaCl and H₂O calculations. The convergence of p_0 is faster in H₂O calculations than in NaCl, which might be seen as a consequence of the greater degree of delocalisation of the Kohn-Sham orbitals in NaCl. The converged value of p_0 in ONETEP calculations, with either TSD or VLSD methods, is around 8.40 Debye for NaCl and 1.86 Debye for H₂O, which differ from the values obtained with NWCHEM by ~ -3 % and 4 %, respectively. This disagreement, which is within acceptable levels, could potentially be reduced with the use of a higher kinetic energy cut-off in the ONETEP calculations.

Table 4.2: Calculated equilibrium bond length, L_0 , and total electric dipole moment at the equilibrium geometry, p_0 , of NaCl and H₂O.

	R_α (Å)	NaCl		H ₂ O	
		L_0 (Å)	p_0 (Debye)	L_0 (Å)	p_0 (Debye)
TSD	3.0	2.38	9.54	0.97	1.90
	4.0	2.37	8.95	0.97	1.89
	5.0	2.38	8.55	0.97	1.88
	6.0	2.37	8.42	0.97	1.86
	7.0	2.37	8.40	0.97	1.86
VLSD	3.0	2.38	9.51	0.97	1.90
	4.0	2.37	9.00	0.97	1.90
	5.0	2.37	8.68	0.97	1.88
	6.0	2.38	8.49	0.97	1.87
	7.0	2.38	8.42	0.97	1.86
NWCHEM		2.38	8.65	0.97	1.79

4.7 Summary

In this chapter we described a new method for imposing localisation constraints in self-consistent energy minimisation approaches for electronic structure calculations, called the variationally-localised search direction (VLSD) method. With the VLSD method, the constraints of localisation of the atom-centred functions $\{\phi_\alpha\}$ are naturally included in the mathematical formalism by means of a detailed analysis of the variational degrees of freedom and the use of formal tensorial algebra. The search direction vectors calculated with the VLSD method are fully localised by construction, and no external truncation is needed. The VLSD method is capable of achieving rapid convergence regardless of the size of the localisation region.

The conventional truncated search direction (TSD) method, in contrast, relies on the approximation that the truncated search direction vectors of the uncon-

strained problem resemble the exact search direction vectors of the constrained problem. The results of calculations on a number of molecules show that the TSD method can provoke the deletion of a large portion of information, which was pre-calculated and stored in the delocalised search direction vectors, in order to keep the localisation constraints. It was observed that the VLSD method often converges in less iterations than the TSD method. More particularly, in calculations on systems where the deletion of information by the TSD method is remarkably large, the VLSD method clearly outperforms TSD by converging in up to a third of the number of iterations. Validation tests confirm that structural and electronic properties can be accurately calculated with both methods.

In this chapter, the VLSD algorithm has been introduced as a technique to optimise non-orthogonal, localised functions in real space. However, this is only one particular application of the method. VLSD could be generalised to any optimisation problem in a non-orthogonal representation in which the search direction vector is subject to a set of analytical constraints. For example, VLSD can be used to impose that a certain subset of localised, non-orthogonal functions must remain constant while the rest are optimised, by ensuring that the relevant associated subset of search direction vectors is identically zero. The theory for developing this method is briefly outlined in Appendix B. Moreover, the VLSD method should not be limited to the case of localised functions. The optimisation of any n -th order tensor in a non-orthogonal representation can benefit from the VLSD algorithm. An example of a second-order tensor optimisation problem of special relevance in linear-scaling density functional theory approaches is the optimisation of the density matrix in a non-orthogonal representation [147, 182], which is expected to be localised, and in practice it is made localised by truncating its associated search

direction vector. The extension of the VLSD algorithm for imposing localisation constraints on the density kernel matrix is shown in [Appendix C](#).

Chapter 5

Density mixing and Hamiltonian diagonalisation methods

Density mixing is used in many DFT methods to minimise the energy functional self-consistently with respect to the electronic density. The scheme that has gathered a wider attention was proposed by Pulay [213], who baptised the new approach as “direct inversion of the iterative subspace” (DIIS). The method has since been developed by many authors [214] and applied to all sorts of systems, from insulators to metals. There are a large number of DIIS schemes, amongst which linear mixing (damping), Pulay mixing [213], C^2 -DIIS [215], ODA [216, 217], EDIIS [218, 219], ADIIS [220] or LiST [221–223], are known examples.

DIIS methods are not direct energy minimisation methods. What this means is that the objective of the iterative improvement of the electronic density is not necessarily to minimise the energy functional, but to reduce an error vector that vanishes at self-consistency. The choice of the error vector is thus fundamental, as pointed out in Ref. [224]. Generally speaking, a given DIIS method cannot be used in isolation because they tend to diverge when the initial guess of the electronic den-

sity is somewhat unoptimised. To mitigate this effect, a combination of different DIIS techniques can be used within the same calculation (such as linear-mixing and Pulay mixing, for example), or, alternatively, DIIS can be used in combination with other techniques such as Kerker preconditioning [225], level-shifting [226] or occupation smearing [202].

The implementation of DIIS methods in ONETEP is part of the current Ph. D. project as well as part of a dissertation for the M. Sc. degree on High-Performance Computing at the University of Edinburgh [227]. Two are the main motivations for including DIIS methods within ONETEP:

1. As an alternative to the LNV conjugate-gradient-based optimisation of the density kernel matrix, \mathbf{K} in the inner loop. Hamiltonian diagonalisation approaches are prototypical self-consistent field procedures for total energy minimisation. The DIIS methods that we have implemented are a very useful tool for developing and testing more complex approaches, as they do not rely on any kind of approximations to calculate the Kohn-Sham states and their occupancies.
2. To benchmark the use of parallel diagonalisers within ONETEP. The scaling of the current DIIS implementation with the system size is cubic-scaling, as a consequence of the Hamiltonian diagonalisation step. The benchmarking results turned out to be very encouraging, and opened the door for a further development: a direct-energy minimisation scheme for large-scale calculations on metallic systems, described in Chapter 6. While diagonalisation is cubic-scaling, the rest of the calculation can be executed with linear-scaling cost. This scheme has been previously proposed by Fattebert and Bernholc [150].

In this chapter, we first describe the original idea behind DIIS methods. Then, we describe the density mixing schemes that were implemented in ONETEP (linear, ODA, Pulay and LiST). Briefly, the equations for the NGWF energy gradient that is used in DIIS calculations are also outlined. The role and performance of the Hamiltonian diagonaliser is discussed afterwards, showing that large calculations of thousands of atoms are possible if efficient parallel diagonalisers (like those included in the SCALAPACK suite [228]) are used. We also show validation tests that prove that the DIIS schemes can yield a good description of the electronic structure. We conclude with some remarks on future developments for improving the convergence rate and efficiency of the DIIS methods used in ONETEP.

5.1 Direct inversion of the iterative subspace

In the original Pulay DIIS scheme [213], the total ground-state energy of the system is a functional of a set of N_{DOF} degrees of freedom, which can be interpreted as the coordinates of a parameter vector \mathbf{p} :

$$\mathbf{p} = \{p_1, p_2, \dots, p_{N_{DOF}}\}. \quad (5.1)$$

The total energy of the system becomes a functional of \mathbf{p} :

$$E = E[\mathbf{p}] = E[p_1, p_2, \dots, p_{N_{DOF}}]. \quad (5.2)$$

DIIS is inserted within an self-consistent field (SCF) iterative algorithm that minimises the total Kohn-Sham energy. At each iteration m , a new parameter vector $\mathbf{p}^{(m+1)}$, that, ideally, should be closer to the ground-state coordinates, is produced. An error vector can be associated to each new parameter vector by calculating the

distance to the previous one:

$$\Delta \mathbf{p}^{(m+1)} = \mathbf{p}^{(m+1)} - \mathbf{p}^{(m)}. \quad (5.3)$$

Convergence of the SCF cycles can be slow or non-existent for certain systems, so stable numerical algorithms to accelerate the convergence speed are needed to solve the energy minimisation problem. In this context, the DIIS algorithm improves the solution at each SCF iteration by constructing a new parameter vector from a linear combination of the solutions of previous iterations as:

$$\mathbf{p}^{(m+1)} = \sum_{j=1}^m \mathbf{p}^{(j)} \lambda_j. \quad (5.4)$$

where $\{\lambda_j\}$ are the expansion coefficients. In this way, the knowledge about the energy surface gained during a history of previous iterations is used to speed-up convergence, by generating a more refined trial vector at every iteration. The DIIS method is based on the assumption that the error vector

$$\Delta \mathbf{p}^{(m+1)} = \sum_{j=1}^m \Delta \mathbf{p}^{(j)} \lambda_j, \quad (5.5)$$

converges to zero in the mean-squares sense, provided the condition

$$\sum_{j=1}^m \lambda_j = 1, \quad (5.6)$$

to preserve the norm of \mathbf{p} at every iteration. Minimisation of (5.5) under the constraint (5.6) leads to the following system of $m + 1$ linear equations:

$$\begin{pmatrix} B_{11} & B_{12} & \dots & B_{1m} & -1 \\ B_{21} & B_{22} & \dots & B_{2m} & -1 \\ \dots & \dots & \dots & \dots & \dots \\ B_{m1} & B_{m2} & \dots & B_{mm} & -1 \\ -1 & -1 & \dots & -1 & 0 \end{pmatrix} \begin{pmatrix} \lambda_1 \\ \lambda_2 \\ \dots \\ \lambda_m \\ \nu \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \\ -1 \end{pmatrix}, \quad (5.7)$$

where $B_{jk} = \langle \Delta \mathbf{p}^{(j)} | \Delta \mathbf{p}^{(k)} \rangle$, with $\langle | \rangle$ being a suitable inner product of the error vectors, and ν being a Lagrange multiplier that imposes the condition (5.6). The solution of this system of linear equations are the $\{\lambda_j\}$ coefficients that will be used to generate the parameter vector for the next iteration, $\mathbf{p}^{(m+1)}$, according to equation (5.4).

5.2 Density kernel optimisation with DIIS methods

The DIIS algorithms implemented in ONETEP take place exclusively in the inner loop, where the Kohn-Sham energy functional is minimised with respect to the elements $\{K^{\alpha\beta}\}$ of the density kernel. The NGWFs $\{\phi_\alpha\}$ are kept constant during the process. It is important to note that the DIIS implementation presented here is specific for zero-temperature DFT calculations, although the extension to finite-temperature DFT should be possible [95]. Therefore, the density kernel is required to fulfill the constraint of idempotency (2.113), which, as demonstrated in Eq. (2.116), is equivalent to imposing orthonormality of the Kohn-Sham states.

The algorithm requires a suitable initialisation of $\{\phi_\alpha\}$ and $\{K^{\alpha\beta}\}$ to begin with. Then, at iteration m of the inner loop, the electronic density $n^{(m)}(\mathbf{r})$ is constructed as:

$$n(\mathbf{r})^{(m)} = \phi_\alpha(\mathbf{r}) K^{\alpha\beta(m)} \phi_\beta(\mathbf{r}). \quad (5.8)$$

At this point, the Hamiltonian matrix is updated with the new density, producing $\mathbf{H}^{(m)}$, of elements:

$$H_{\alpha\beta}^{(m)} = \langle \phi_\alpha | \hat{H}(\mathbf{K}^{(m)}) | \phi_\beta \rangle. \quad (5.9)$$

The total energy functional $E^{(m)} = E[\{K^{\alpha\beta(m)}\}, \{\phi_\alpha\}]$ is also evaluated. Then, the diagonalisation step:

$$H_{\alpha\beta}^{(m)} M_i^{\beta(m)} = S_{\alpha\beta} M_i^{\beta(m)} \epsilon_i^{(m)}, \quad (5.10)$$

returns $\{M_i^{\beta(m)}\}$ and $\{\epsilon_i^{(m)}\}$. A new, perfectly-idempotent density kernel $\tilde{\mathbf{K}}^{(m)}$ can now be constructed as:

$$\tilde{K}^{\alpha\beta(m)} = \sum_i^{N_{occ}} M_i^{\alpha(m)} M_i^{\dagger\beta(m)}, \quad (5.11)$$

where N_{occ} is the number of fully occupied Kohn-Sham states (i.e., $f_i = 1$), selected by applying the *aufbau* principle to $\{\epsilon_i^{(m)}\}$. The “output” density kernel $\tilde{\mathbf{K}}^{(m)}$ and the “input” Hamiltonian $\mathbf{H}^{(m)}$ share the same set of eigenvectors, and therefore the product of these matrices is commutative. However, this is not true for the input density kernel $\mathbf{K}^{(m)}$ and the input Hamiltonian, $\mathbf{H}^{(m)}$, and the commutator $[\mathbf{H}^{(m)}\mathbf{K}^{(m)} - \mathbf{K}^{(m)}\mathbf{H}^{(m)}]$ is not zero unless self-consistent convergence is achieved.

Equations (5.8) to (5.11) are common to all the DIIS schemes implemented in ONETEP. At this point, the algorithm follows different paths depending on the DIIS scheme selected for the calculation. Existing DIIS schemes can be classified

as *interpolation* methods if the mixing coefficients $\{\lambda_j\}$ are constrained to be in the range $[0, 1]$, and as *extrapolation* methods if $\{\lambda_j\}$ can take any value in the spectrum of real numbers. In both cases, the constraint (5.6) must hold. Another element that helps classifying the different DIIS schemes is the number of input and output density kernels and Hamiltonian matrices that are used to construct the input of the next iteration. The number is not only determined by how many density kernels from previous iterations are used in the mixing step (5.5) (which we will denote by N_{mix}), but also by how many matrices are necessary in order to calculate the error vector. More matrices imply that more information is available to construct a better input in the next iteration, but also that more resources are needed in order to store them on memory. A flowchart of the complete DIIS implementation is shown in Fig. 5.6 at the end of this chapter, which includes the linear, ODA, Pulay and LiST density mixing methods. The details of each DIIS method are outlined in the next subsections, and summarised in Table 5.1.

Table 5.1: Summary of the main features of the DIIS schemes implemented in ONETEP. The table shows the range of the $\{\lambda_j\}$ coefficients and the number of matrices required by each method.

Scheme	Range of $\{\lambda_j\}$	#K	# $\tilde{\mathbf{K}}$	#H	# $\tilde{\mathbf{H}}$
Linear	$[0, 1]$	1	1	1	0
ODA	$[0, 1]$	1	1	1	1
Pulay	$(-\infty, +\infty)$	1	N_{mix}	1	0
LiST	$(-\infty, +\infty)$	N_{mix}	N_{mix}	N_{mix}	N_{mix}

5.2.1 Linear mixing

Linear density mixing is the simplest of all the DIIS methods. In this scheme, the new density kernel $\mathbf{K}^{(m+1)}$ is generated solely from the input and output kernels

of the current iteration:

$$\mathbf{K}^{(m+1)} = (1 - \lambda)\mathbf{K}^{(m)} + \lambda\tilde{\mathbf{K}}^{(m)}, \quad (5.12)$$

where λ is the mixing coefficient, also called damping parameter. The peculiarity about the linear mixing is that λ is not calculated as a consequence of measuring the error vector, but instead it is selected manually and kept constant in the range $[0, 1]$. This feature makes linear mixing the DIIS method with the poorest performance regarding self-consistent convergence, but also the most stable of all, so long as λ is sufficiently small. This is an intrinsic advantage, since by manually adjusting the value of λ , an array of relatively similar and suitable density kernels can be generated during a certain number of initial iterations, which can then be used as part of a more sophisticated Pulay or LiST scheme.

5.2.2 ODA mixing

The acronym ODA stands for *optimal damping algorithm*. ODA is an improved form of linear mixing, originally developed for Hartree-Fock calculations [216] and then extended to DFT [217]. The ODA updates the density kernel by taking steps along the steepest descent direction, which is the one for which the slope:

$$s = \frac{d}{d\lambda} E \left[\mathbf{K}^{(m)} + \lambda \left(\tilde{\mathbf{K}}^{(m)} - \mathbf{K}^{(m)} \right), \{ \phi_\alpha \} \right]_{\lambda=0}, \quad (5.13)$$

is minimal. The optimal value of λ is found after minimising a third-order polynomial fitting of the energy along the steepest descent direction:

$$E(\lambda) = a\lambda^3 + b\lambda^2 + c\lambda + d. \quad (5.14)$$

In order to evaluate a , b , c and d , the electronic density must be updated with the output density kernel:

$$\tilde{n}(\mathbf{r}) = \phi_{\alpha}(\mathbf{r}) \tilde{K}^{\alpha\beta(m)} \phi_{\beta}(\mathbf{r}), \quad (5.15)$$

followed by the update of the Hamiltonian matrix:

$$\tilde{H}_{\alpha\beta}^{(m)} = \langle \phi_{\alpha} | \hat{H}(\tilde{\mathbf{K}}^{(m)}) | \phi_{\beta} \rangle, \quad (5.16)$$

and the evaluation of the total energy functional $\tilde{E}^{(m)} = E[\tilde{K}^{\alpha\beta(m)}, \{\phi_{\alpha}\}]$. Then, the values of a , b , c and d can be found analytically as:

$$d = E^{(m)} \quad (5.17)$$

$$c = 2\text{tr} \left[\mathbf{H}^{(m)} \left(\tilde{\mathbf{K}}^{(m)} - \mathbf{K}^{(m)} \right) \right] \quad (5.18)$$

$$a = 2\text{tr} \left[\tilde{\mathbf{H}}^{(m)} \left(\tilde{\mathbf{K}}^{(m)} - \mathbf{K}^{(m)} \right) \right] - 2\tilde{E}^{(m)} + c + 2d \quad (5.19)$$

$$b = \tilde{E}^{(m)} - a - c - d. \quad (5.20)$$

The two solutions of the minimisation problem $\frac{dE}{d\lambda} = 0$ are:

$$\lambda_{\pm} = \frac{-b \pm \sqrt{b^2 - 3ac}}{3a}. \quad (5.21)$$

If both values of λ_{\pm} are in the range $[0, 1]$, then one will correspond to a maximum and the other to a minimum. The λ corresponding to the minimum is chosen based on the second derivative criterion:

$$\left. \frac{d^2 E}{d\lambda^2} \right|_{\lambda=\lambda_{min}} > 0. \quad (5.22)$$

If none of the solutions are in the interval $[0, 1]$, then λ is chosen manually to a safe value (usually 0.1). The new density kernel $\mathbf{K}^{(m+1)}$ is generated as in Eq. (5.12).

5.2.3 Pulay mixing

Pulay mixing uses a history of previous N_{mix} iterations to generate the next input density kernel as:

$$\mathbf{K}^{(m+1)} = \sum_{j=m-N_{mix}}^m \tilde{\mathbf{K}}^{(j)} \lambda_j. \quad (5.23)$$

Pulay mixing is an extrapolation approach where the coefficients $\{\lambda_j\}$ can take any real values. The optimal coefficients that can be used in (5.23) are calculated by the principle of minimising the least-squares error, as described in Sec. (5.1). At iteration m , an error vector $\mathbf{R}^{(m)}$ is constructed as:

$$\mathbf{R}^{(m)} = \tilde{\mathbf{K}}^{(m)} - \mathbf{K}^{(m)}. \quad (5.24)$$

The inner product of any two error vectors can be calculated as :

$$B_{jk} = \langle \mathbf{R}^{(j)} | \mathbf{R}^{(k)} \rangle = \text{tr} \left[\mathbf{R}^{(j)} \mathbf{S} \mathbf{R}^{(k)} \mathbf{S} \right], \quad (5.25)$$

where this choice of inner product is motivated from the definition of the Frobenius norm. At every iteration, a matrix of B_{jk} elements can be formed, and the system of equations (5.7) can be solved to obtain the coefficients $\{\lambda_j\}$. The next input density kernel is calculated with Eq. (5.23).

5.2.4 LiST mixing

The acronym LiST stands for “Linear-expansion shooting technique”. There are three LiST methods, termed LiSTd (direct), LiSTi (indirect) and LiSTb (better). LiSTd is regarded as unstable and impractical [222], due to spurious linear dependencies that make the system of equations (5.7) ill-defined. On the other hand, LiSTi [222] and LiSTb [223] are essentially the same method, with the difference that the matrices associated to each method are mutually transposed, i.e., $B_{jk}^{LiSTb} = B_{kj}^{LiSTi}$. According to the authors of the LiST method, the performance of LiSTb is better than LiSTi.

In LiST methods, the corrected Hohenberg-Kohn-Sham energy functional at every iteration, $\mathcal{E}^{(m)}$, defined as [221]:

$$\mathcal{E}^{(m)} = \tilde{E}^{(m)} + \frac{1}{2} \text{tr} \left[\left(\mathbf{K}^{(m)} - \tilde{\mathbf{K}}^{(m)} \right) \left(\tilde{\mathbf{H}}^{(m)} - \mathbf{H}^{(m)} \right) \right], \quad (5.26)$$

is minimised with respect to $\mathbf{K}^{(m)}$. The idea is that when self-consistency is achieved, the second term in Eq. (5.26) vanishes, and the familiar Kohn-Sham energy functional is recovered. Therefore, $\left(\mathbf{K}^{(m)} - \tilde{\mathbf{K}}^{(m)} \right) \left(\tilde{\mathbf{H}}^{(m)} - \mathbf{H}^{(m)} \right)$ is a suitable choice of error vector per iteration, that combines not only the information stored in a history of input-output density kernels, but also in a history of input-output Hamiltonian matrices. A linear combination of the type (5.5) yields a system of equations of the type (5.7), where the elements B_{ij} are:

$$B_{jk}^{LiSTi} = \frac{1}{2} \text{tr} \left[\left(\mathbf{K}^{(k)} - \tilde{\mathbf{K}}^{(k)} \right) \left(\tilde{\mathbf{H}}^{(j)} - \mathbf{H}^{(j)} \right) \right] \quad (5.27)$$

for LiSTi, and

$$B_{jk}^{LiSTb} = \frac{1}{2} \text{tr} \left[\left(\mathbf{K}^{(j)} - \tilde{\mathbf{K}}^{(j)} \right) \left(\tilde{\mathbf{H}}^{(k)} - \mathbf{H}^{(k)} \right) \right] \quad (5.28)$$

for LiSTb. The solutions of the respective systems of equations yield the coefficients $\{\lambda_j\}$. The next input density kernel is built as in Eq. (5.23).

5.3 NGWF optimisation

The NGWFs are optimised in the outer loop using a conjugate gradients iterative scheme, with fixed density kernel \mathbf{K} . In the standard approach, the localisation constraints on the NGWFs are kept using the TSD method, described in Section 4.4. The purpose of this section is to describe how the NGWF energy gradient that is calculated when the inner loop uses a DIIS scheme for determining \mathbf{K} .

The constraints of orthogonality of the Kohn-Sham states, $\{\psi_i\}$, as well as the constraint on conservation of the number of electrons \mathcal{N} , must be taken into account during the inner loop. This task is efficiently accomplished by the Lagrangian:

$$\mathcal{L}' [\{c_{p\alpha}\}] = E' [\{c_{p\alpha}\}] - \text{tr} [\mathbf{H} (\mathbf{KSK} - \mathbf{K})] - \mu (\text{tr} [\mathbf{KS}] - \mathcal{N}). \quad (5.29)$$

The second term in \mathcal{L}' imposes idempotency of \mathbf{K} , which, we recall, is equivalent to imposing orthonormality of $\{\psi_i\}$, as shown in Eq. (2.116). The third term of \mathcal{L}' imposes conservation of the number of electrons. To eliminate the Lagrange multiplier μ , which is unknown (at least in principle), the density kernel obtained after the DIIS optimisation is rescaled as in Eq. (2.118):

$$\mathbf{K}' = \frac{\mathcal{N}}{\text{tr}[\mathbf{KS}]} \mathbf{K}, \quad (5.30)$$

and the Lagrangian \mathcal{L}' is simplified to:

$$\mathcal{L}'[\{c_{p\alpha}\}] = E'[\{c_{p\alpha}\}] - \text{tr}[\mathbf{H}(\mathbf{K}'\mathbf{S}\mathbf{K}' - \mathbf{K}')]. \quad (5.31)$$

which is now tractable with analytical techniques. The derivative of \mathcal{L}' with respect to $c_{p\alpha}$ is:

$$g_p^\alpha = \frac{\partial \mathcal{L}'}{\partial c_{p\alpha}} = \left[\hat{H} \phi_\beta(\mathbf{r}) P_{DIIS}^{\beta\alpha} - \phi_\beta(\mathbf{r}) Q_{DIIS}^{\beta\alpha} \right]_{\mathbf{r}=\mathbf{r}_p}, \quad (5.32)$$

where

$$\mathbf{P}_{DIIS} = 2\eta\mathbf{K} - \eta^2\mathbf{K}\mathbf{S}\mathbf{K}, \quad (5.33)$$

and

$$\mathbf{Q}_{DIIS} = \eta^2\mathbf{K}\mathbf{H}\mathbf{K} - 2 \left(\eta - \eta^2 \frac{\text{tr}[\mathbf{H}\mathbf{K}\mathbf{S}\mathbf{K}]}{\text{tr}[\mathbf{KS}]} \right) \mathbf{K}, \quad (5.34)$$

with $\eta = \mathcal{N}/\text{tr}[\mathbf{KS}]$. If the inner loop has fully converged, then $\mathbf{K} = \mathbf{K}\mathbf{S}\mathbf{K}$ and $\eta = 1$, and the gradient obtained with LNV methods is naturally recovered, with $\mathbf{P}_{DIIS} = \mathbf{K}$ and $\mathbf{Q}_{DIIS} = \mathbf{Q}_{LNV} = -\mathbf{K}\mathbf{H}\mathbf{S}^{-1}$, as shown in Eq. (2.126).

5.4 Numerical eigensolvers

All the DIIS schemes implemented in ONETEP require the diagonalisation of the Hamiltonian matrix as the fundamental step that ensures self-consistency. Diagonalisation is equivalent to finding the solutions to the generalised eigenvalue prob-

lem (5.10), a step that must take place at each iteration in order to construct a new guess of the idempotent density kernel. This operation is cubic-scaling, $\mathcal{O}(N^3)$, for two reasons:

1. The intrinsic cost of diagonalising \mathbf{H} is cubic-scaling.
2. The transformation matrix \mathbf{M} is dense (as opposed to sparse). Dense matrix multiplications to construct \mathbf{K} are cubic-scaling.

The size of the systems that can be simulated with DIIS methods is then expected to be limited by the cubic-scaling behaviour of the diagonalisation step. The performance of the eigensolver will therefore mark the performance of the DIIS calculations. Mathematical libraries are very important in high-performance computing for their potential to accelerate numerical algorithms. The LAPACK [229] and SCALAPACK [228] packages are a collection of optimised routines for algebraic operations involving matrices and vectors. LAPACK is optimised to be run in serial, i.e., using only one core (CPU) per problem at a time, while SCALAPACK is optimised for parallel computing, and allows to solve a single problem using many cores. For this work, we used the LAPACK eigensolver DSYGVX and its SCALAPACK parallel equivalent PDSYGVX. Both eigensolvers follow the same basic algorithm [230].

The size of the matrices in the NGWF representation is equal to N_ϕ^2 , and storing the matrix elements in memory becomes increasingly complicated as the size increases. The internal matrix handling of LAPACK implies that an entire copy of each matrix must be stored within a single core memory [231]. Therefore, LAPACK calculations are limited by the capabilities of a single core hardware. On the other hand, SCALAPACK assigns a fraction of the matrix elements to each of

the cores involved in the calculation. The memory requirements are in this manner distributed across all the processors, and the hardware limitation is effectively tackled [232]. By allocating more cores to the calculation, which might be distributed across different nodes (each having its own memory sockets), the total allocated memory to the calculation can be increased almost limitless.

Calculations with LAPACK require all-to-all communications to synchronise the results obtained with each core. This is not true parallelism, since each core is computing exactly the same problem, and should obtain identical solutions. This is an obvious bottleneck in the performance of the LAPACK eigensolver that gets worse for large systems. On the other hand, SCALAPACK calculations are communication-intensive, since the different data stored in multiple cores must be circulated correctly in order to obtain the correct result. The SCALAPACK parallel algebra engine, named PBLAS, uses a 2D block-cyclic distribution of the matrix elements within the cores that minimises the data movement between the different memory levels and optimises inter-nodal communications. The parallel communications are controlled with the BLACS interface. The overall performance of the SCALAPACK routines will be determined by the amount of computation per core and the efficiency of the communications between cores.

The ONETEP code has been interfaced to the DSYGVX and PDSYGVX eigensolvers. The choice between LAPACK and ScaLAPACK is given at compilation time, via conditional compilation flags (`#ifdef`). If no flags are invoked, LAPACK is the default solver, whereas if the flag `-DSCALAPACK` is invoked at compilation time, the ScaLAPACK solver will be enabled. The Kohn-Sham eigenvectors are required to be orthonormal to each other. Eigenvectors that correspond to different eigenvalues are orthogonal by definition, but if degeneracy is

present, each group of eigenvectors will have to be orthonormalised independently. The SCALAPACK eigensolver orthonormalises the eigenvectors using an iterative method based on convergence within a given tolerance threshold, specified by the SCALAPACK input variable `ABSTOL`. The input variable `ORFAC` tells PDSYGVX which eigenvectors must be orthonormalised (none is a possibility). The amount of workspace that must be allocated for the orthonormalisation process is directly proportional to the degree of degeneracy of the eigenstates. The variable `WORK` controls the size of the arrays allocated for the completion of the diagonalisation and orthonormalisation with PDSYGVX. For a parallel algebra eigensolver like PDSYGVX, keeping the eigenvectors orthogonal requires a large amount of all-to-all communications, which take a substantial amount of the total runtime of the diagonalisation routine.

5.5 Performance of the eigensolvers

The ONETEP program with the DIIS implementation and an interface to the LAPACK and SCALAPACK eigensolvers DSYGVX and PDSYGVX was compiled on HECToR Cray XT4, Phase 2A (3072 AMD 2.3 GHz quad-core Opteron processors and 8 GB of memory per node), and Iridis 3 (1008 Intel Nehalem nodes with two 4-core processors each and 22 GB of memory per node). Full details of the compilation of ONETEP on each of these target machines can be found in Ref. [227]. A group of benchmark systems was formed by five silicon nanorods of different sizes which should show some of the typical properties of the semiconductors, such as a medium value of the band gap. A smaller band gap means that the nearsightedness of the density matrix is not so obvious, and some linear-scaling approaches may fail during the simulation. However, recent work using this set of

silicon benchmarks [233] demonstrated that it is possible to simulate these kind of semiconductors using ONETEP without loss of accuracy. Table 5.2 summarises the set of silicon benchmarks, represented in Figure 5.1.

Table 5.2: Set of benchmarks corresponding to silicon nanorods of different size. The NGWF radii $\{R_\alpha\}$ are 7.0 Bohr for silicon atoms and 6.0 Bohr for hydrogen. The density kernel is truncated beyond a cut-off radius R_c of 25 Bohr. The table shows the number of atoms, N_{at} , number of NGWFs, N_ϕ , and plane-wave kinetic energy cut-off, E_c , of each calculation.

Test case	N_{at}	N_ϕ	E_c (eV)
$\text{Si}_{29}\text{H}_{36}$	65	297	650
$\text{Si}_{242}\text{H}_{140}$	382	2318	650
$\text{Si}_{532}\text{H}_{224}$	756	5012	600
$\text{Si}_{766}\text{H}_{402}$	1168	7296	600
$\text{Si}_{1186}\text{H}_{366}$	1552	11040	600

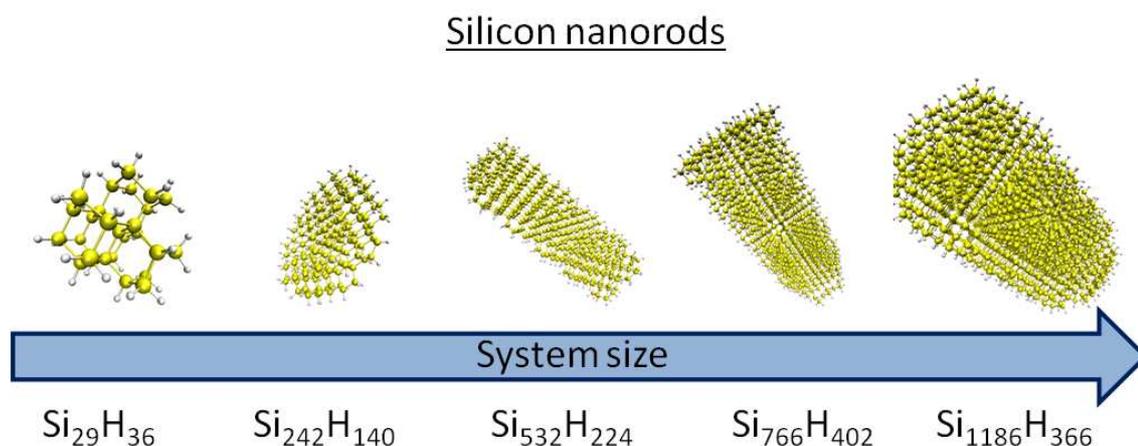


Figure 5.1: Set of silicon nanorods benchmarks.

The performance of the eigensolvers was also benchmarked in a sequence of fragments of a double-helix shaped amyloid fibril protein [234]. This biological system has a larger band gap that should make a linear-scaling approach based on the locality of the density matrix more accurate [185]. Table 5.3 and figure 5.2 summarise the computational parameters of this set of calculations.

Table 5.3: Set of benchmarks corresponding to a the amyloid protein of increasing size. The NGWF radii $\{R_\alpha\}$ are 7.0 Bohr for all species. The density kernel is truncated beyond a cut-off radius R_c of 20 Bohr. The table shows the number of atoms, N_{at} , number of NGWFs, N_ϕ , and plane-wave kinetic energy cut-off, E_c , of each calculation.

Test case	N_{at}	N_ϕ	E_c (eV)
p16_20	1712	4544	600
p32_20	3424	9088	600
p48_20	5136	13632	600
p64_20	6848	18176	600
p80_20	8560	22720	600
p96_20	10272	27264	600
p112_20	11984	31808	600
p128_20	13696	36352	600

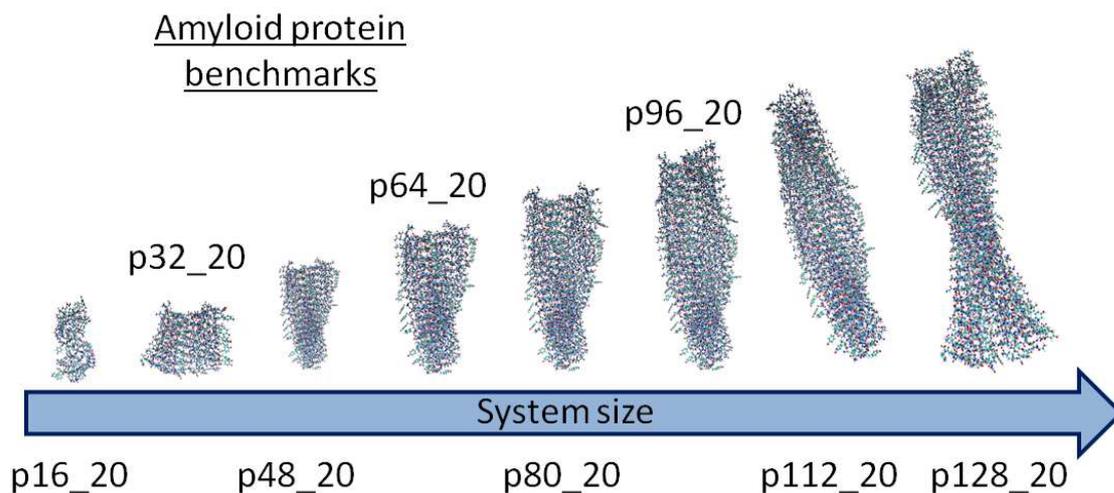


Figure 5.2: Set of benchmarks of fragments of an amyloid fibril protein.

A complete report of the benchmarking process can be found in Ref. [227]. We shall summarise here the main results which are important for further applications of parallel diagonalisers in self-consistent field algorithms.

1. Using LAPACK in a parallel ONETEP calculation yields almost no benefit. Since LAPACK is, by definition, a serial code, there is no parallel speed-up. That is, adding more cores does not make the execution of the diagonalisation

step faster.

2. Memory-wise, parallel calculations with the LAPACK eigensolver DSYGVX are hardware-limited. Each core holds an exact copy of the entire collection of dense matrices that are involved in the diagonalisation step. The memory requirements per core are thus constant regardless of the total number of cores involved in the calculation. The numbers are simple: if the matrices do not fit in a single core memory, the calculation cannot proceed. Large systems cannot be executed with LAPACK.
3. Parallel calculations with the SCALAPACK PDSYGVX eigensolver yield a noticeable parallel speed-up, however the parallel efficiency is not dramatic.
4. Parallel dense algebra packages such as SCALAPACK are essential for large-scale DFT calculations based on diagonalisation. The reason is memory efficiency: each core is in charge of storing only a certain fraction of each dense matrix, the size of which is inversely proportional to the total number of cores. Therefore, operations that involve very large matrices become feasible since the information is stored using a collaborative approach between all the cores. Should more memory be needed for a given calculation, requesting more cores will automatically increase the total memory available for executing the diagonalisation and orthonormalisation steps.

The implications of memory distribution for practical calculations are clearly visible in Table 5.4, which summarises which calculations were feasible with LAPACK and SCALAPACK on the two target machines. In the light of these performance results, one can question the performance of diagonalisation-based methods such as DIIS compared to other methods such as LNV where no explicit diagonal-

Table 5.4: Successful calculations using LAPACK and ScaLAPACK eigensolvers on HECToR and Iridis 3. In all cases, 4 cores per node have been used. The asterisk * indicates that a smaller number of nodes has not been tested and it is possible for the calculation to fit on less nodes than that. For saving computational time, the protein systems have not been simulated using LAPACK on Iridis 3.

Test case	LAPACK (HECToR)	LAPACK (Iridis 3)	SCALAPACK (HECToR)	SCALAPACK (Iridis 3)
Si ₂₉ H ₃₆	Yes	Yes	Yes (1 node)	Yes (1 node)
Si ₂₄₂ H ₁₄₀	Yes	Yes	Yes (1 node)	Yes (1 node)
Si ₅₃₂ H ₂₂₄	Yes	Yes	Yes (1 node)	Yes (1 node)
Si ₇₆₆ H ₄₀₂	No	Yes	Yes (2 nodes)	Yes (1 node)
Si ₁₁₈₆ H ₃₆₆	No	No	Yes (3 nodes)	Yes (1 node)
p16_20	Yes	-	Yes (7 nodes)	Yes (15 nodes)*
p32_20	No	-	Yes (9 nodes)	Yes (20 nodes)*
p48_20	No	-	Yes (10 nodes)	Yes (5 nodes)
p64_20	No	-	Yes (10 nodes)	Yes (5 nodes)
p80_20	No	-	Yes (11 nodes)	Yes (20 nodes)*
p96_20	No	-	Yes (13 nodes)	Yes (20 nodes)*
p112_20	No	-	Yes (17 nodes)	Yes (20 nodes)*
p128_20	No	-	Yes (19 nodes)	Yes (20 nodes)*

isation takes place. Figure 5.3 shows the scaling of one DIIS iteration and one LNV iteration with the system size. These results do not mark a definite cross-over point where the cubic-scaling diagonalisation becomes more costly than the linear-scaling approach. However, they clearly suggest that diagonalisation-based calculations of some thousands of atoms can be achieved with a computational cost that is comparable to LNV calculations. This is a very useful result that has enabled the implementation of a robust method for DFT calculations on metals with the ONETEP program, presented in Chapter 6 of this dissertation.

Scaling of LNV and DIIS with the system size / Amyloid protein benchmarks / 20 x 4 cores

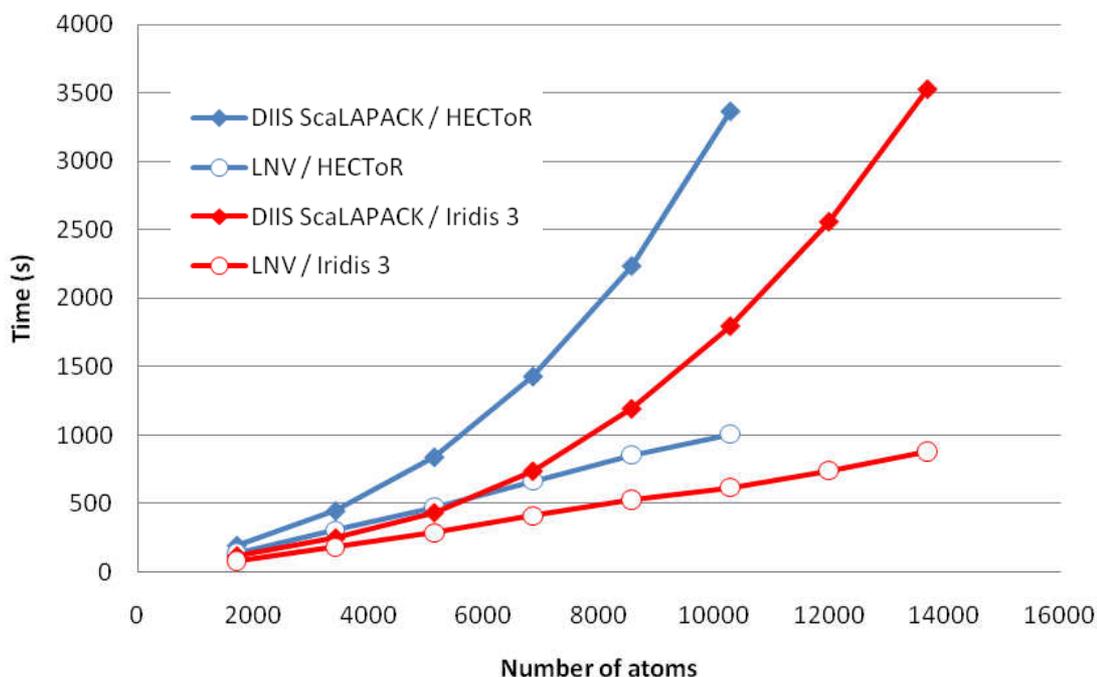
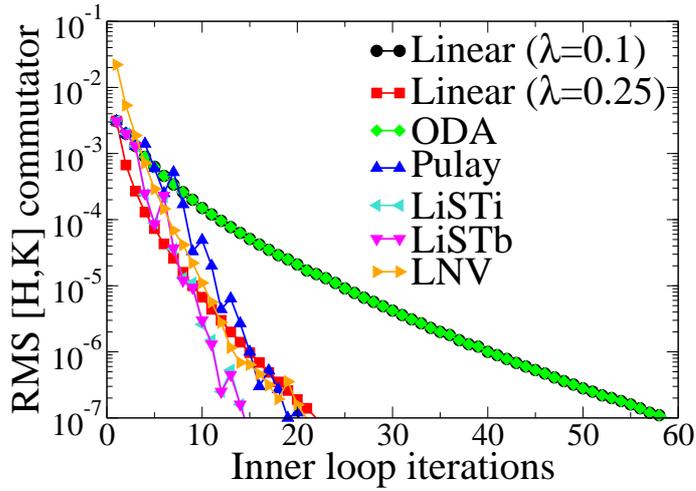


Figure 5.3: Comparison of the scaling of the LNV and DIIS algorithms with the system size for the set of amyloid benchmarks. The systems have been simulated using 20 nodes and 4 cores per node on HECToR and Iridis 3.

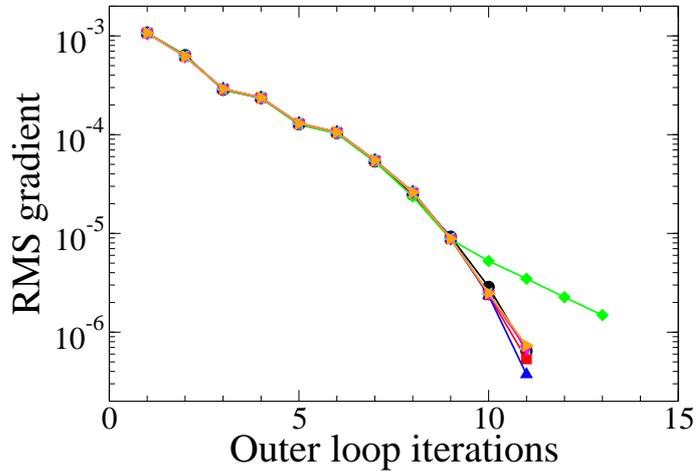
5.6 Validation tests

We have validated the implementation of DIIS in a number of test cases. Here, we show the results for the silicon nanorods $\text{Si}_{29}\text{H}_{36}$ and $\text{Si}_{242}\text{H}_{140}$ with the setup described in Table 5.2.

Figure 5.4 shows the convergence of the inner and outer loop in calculations on $\text{Si}_{29}\text{H}_{36}$. All the calculations on this system, whether with DIIS or LNV methods, converged successfully. Regarding the inner loop, it is obvious however that linear mixing with $\lambda = 0.1$ and ODA are slower than the rest, as they require almost triple number of iterations to converge to the same degree as the rest. In this relatively small system, LiST methods perform better than Pulay mixing, requiring 5 itera-



(a)



(b)

Figure 5.4: Calculations on $\text{Si}_{29}\text{H}_{36}$ silicon nanorod with DIIS. Pulay and LiST methods use a history of size $N_{mix} = 5$, and perform 2 leading linear mixing iterations with parameter $\lambda = 0.1$. (a) Convergence of the RMS of the commutator (2.122) during the execution of the inner loop. In this plot, the series corresponding to linear mixing with $\lambda = 0.1$ and ODA are overlapping. (b) Convergence of the RMS of the slope of the NGWF gradient along the search direction (2.127), with 5 iterations in the inner loop.

tions to converge (25 % less). Regarding the outer loop, ODA does not perform as good as the others and the calculation converges to a higher threshold in three more iterations. The total charge dipole moment calculated with each method is shown in Table 5.5. The value obtained with each method agrees with the others within a 1% relative error, which suggests that the electronic structure of $\text{Si}_{29}\text{H}_{36}$

is well describe with any DIIS scheme and LNV.

Table 5.5: Dipole moment of $\text{Si}_{29}\text{H}_{36}$.

Method	Dipole moment (Debye)
Linear ($\lambda = 0.1$)	-0.137
Linear ($\lambda = 0.25$)	-0.132
ODA	-0.137
Pulay	-0.130
LiSTi	-0.131
LiSTb	-0.131
LNV	-0.132

Figure 5.5 shows the convergence of the inner and outer loops in calculations on $\text{Si}_{242}\text{H}_{140}$. This system has five times more atoms and ten times more NGWFs than the $\text{Si}_{29}\text{H}_{36}$ nanorod, and therefore correct simulation of its electronic properties is more complicated. Regarding the inner loop, the linear mixing parameter has to be reduced to at least 0.05 to obtain any type of convergence. Such a small mixing parameter produces a very small change in the density kernel and convergence noticeably slower. This kind of situation is an example of charge sloshing [94], where a small change in \mathbf{K} produces a large, ill-conditioned change in the Hamiltonian \mathbf{H} . Charge-sloshing also affects the ODA mixing method, which does not perform very well and, although it eventually falls into the convergence regime, it requires a very large number of “false” iterations where the commutator oscillates. Pulay and LiST methods do converge faster, due to the large history of previous solutions that they have available for elaborating the best density kernel matrix in the following iterations. Perhaps surprisingly, convergence of the inner loop with LNV is worst than with converging DIIS methods. The commutator (2.122) is noticeably higher, which indicates that the current electronic density is not a solution to the Kohn-Sham equations. The behaviour of the inner loop is reflected in the

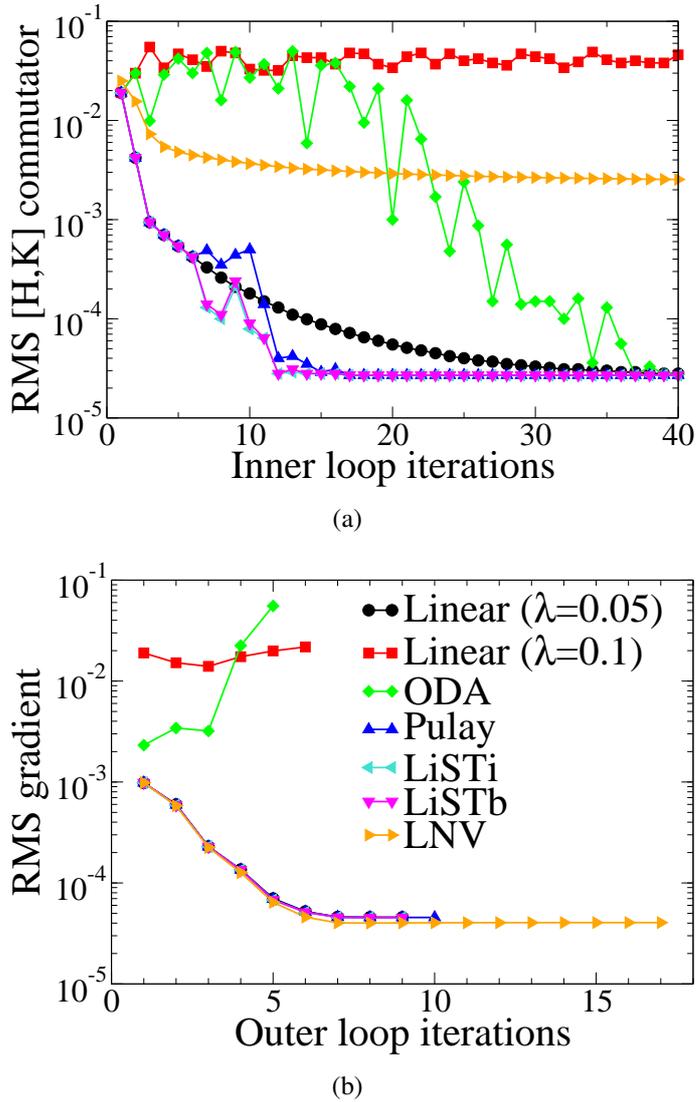


Figure 5.5: Calculations on Si₂₄₂H₁₄₀ silicon nanorod with DIIS. Pulay and LiST methods use a history of size $N_{mix} = 5$, and perform 5 leading linear mixing iterations with parameter $\lambda = 0.05$. (a) Convergence of the RMS of the commutator (2.122) during the execution of the inner loop. (b) Convergence of the RMS of the slope of the NGWF gradient along the search direction (2.127), with 10 iterations in the inner loop.

outer loop. Linear mixing with $\lambda = 0.1$ and ODA mixing fail to converge. The reason is that after ten inner loop iterations, the density kernel is not only far from convergence, but also clearly unphysical. The optimisation of the NGWFs cannot proceed without a reasonable estimate of the density kernel matrix. The outer loop does converge for all the other methods (including linear mixing with $\lambda = 0.05$).

The total energy converges to the same degree (10^{-6} eV/atom variation between two consecutive iterations) faster with the DIIS methods than with LNV. The total dipole moment of $\text{Si}_{242}\text{H}_{140}$ calculated with each method is shown in Table 5.6. Leaving aside the broken $\lambda = 0.1$ and ODA linear mixing calculations, the values obtained with the other methods match within a 10% relative error margin.

Table 5.6: Dipole moment of $\text{Si}_{242}\text{H}_{140}$.

Method	Dipole moment (Debye)
Linear ($\lambda = 0.05$)	0.088
Linear ($\lambda = 0.1$)	-25959.8
ODA	-4963.9
Pulay	0.078
LiSTi	0.083
LiSTb	0.083
LNV	0.083

The results shown in Figure 5.5 are remarkably important for the following reason. It can be seen that DIIS is able to converge the Hamiltonian-density kernel commutator to a much lower level (at least during the first NGWF iterations). To understand why the LNV loop is unable to obtain the same level of convergence, one must look at how LNV works. As opposed to DIIS methods, LNV is a direct energy minimisation method that relies on the energy gradient to construct a search direction vector (see section 2.11.2). If a finite radial cut-off R_c is applied to the elements of the density kernel, the search direction vector for LNV calculations, resulting from an unconstrained minimisation search, must be truncated in order to maintain the localisation constraints. The truncation step introduces an error in the calculation which cannot be controlled or predicted. The higher Hamiltonian-density kernel commutator in the LNV calculation is a consequence of this non-variational truncation of the search direction vector. DIIS methods, on the other

hand, perform a constrained minimisation process in which, at all stages and by construction, the localisation constraints are kept exactly. No special truncation is needed at any stage of DIIS minimisation. Instead, once the sparsity pattern of $\{K^{\alpha\beta}\}$ is decided (at the initialisation step), it is propagated throughout the equations in a fully consistent manner. This capability of DIIS makes it an ideal benchmarking code for testing the consequences of spatial truncation of the density kernel matrix.

5.7 Summary

The implementation of DIIS methods in ONETEP has served as a benchmark for testing Hamiltonian diagonalisation methods for self-consistent minimisation of the energy. The formalism of localised NGWFs makes this approach attractive despite of the inherent cubic-scaling cost since, as it has been shown in the performance results, calculations with thousands of atoms are certainly possible. There is one condition that must be met nevertheless: efficient, parallel, dense algebra routines have to be used in order to distribute the memory requirements and the computational effort between a large number of cores. In this project, the SCALAPACK package allowed us to successfully diagonalise matrices of tens of thousands of eigenvalues using a relatively modest core count. A future development can be to interface the ONETEP program to other parallel dense eigensolvers, such as the one described in Ref. [235].

DIIS approaches are conceptually simple and can produce an accurate description of the electronic structure, as have been shown with the validation results on the silicon nanorods. By definition, DIIS methods perform a constrained energy minimisation where the localisation constraints are taken into account. Therefore,

DIIS methods completely avoid errors arising from spatial truncation of the density kernel, as opposed to direct minimisation methods using the LNV technique. This property makes DIIS a perfect method for benchmarking the effects of localisation constraints imposed upon the density kernel elements.

There is, however, a drawback with DIIS methods: their reliance on an error vector that must be externally defined, rather than on the analysis of the energy functional, in the decision-making process of selecting the most suitable mixing parameters. Many times, especially in large systems, charge sloshing effects [94], where a small update of the electronic density produces an uncontrolled change to the electronic structure of the system, can cause the calculation to diverge. As a future development, Kerker preconditioning techniques [225] can be implemented with the aim of minimising charge sloshing effects and making the implementation more robust. Another further development of the DIIS methods is to substitute density kernel mixing and replace it with Hamiltonian mixing instead. Hamiltonian mixing ensures that the occupancies always follow the *aufbau* principle. In density kernel mixing, the *aufbau* principle is broken when a new input density kernel is built from the sum of previous density kernels.

Finally, it must be acknowledged that DIIS methods are commonly used for finite-temperature DFT calculations on metals [95, 219]. Also as a future development, the DIIS implementation can be expanded to consider non-integer occupancies given by the Fermi-Dirac distribution (2.41) or the Methfessel-Paxton distribution [236]. The resulting algorithm (which will need Kerker preconditioning or other similar technique to fight charge sloshing), should be able to perform efficient calculations on large metallic particles.

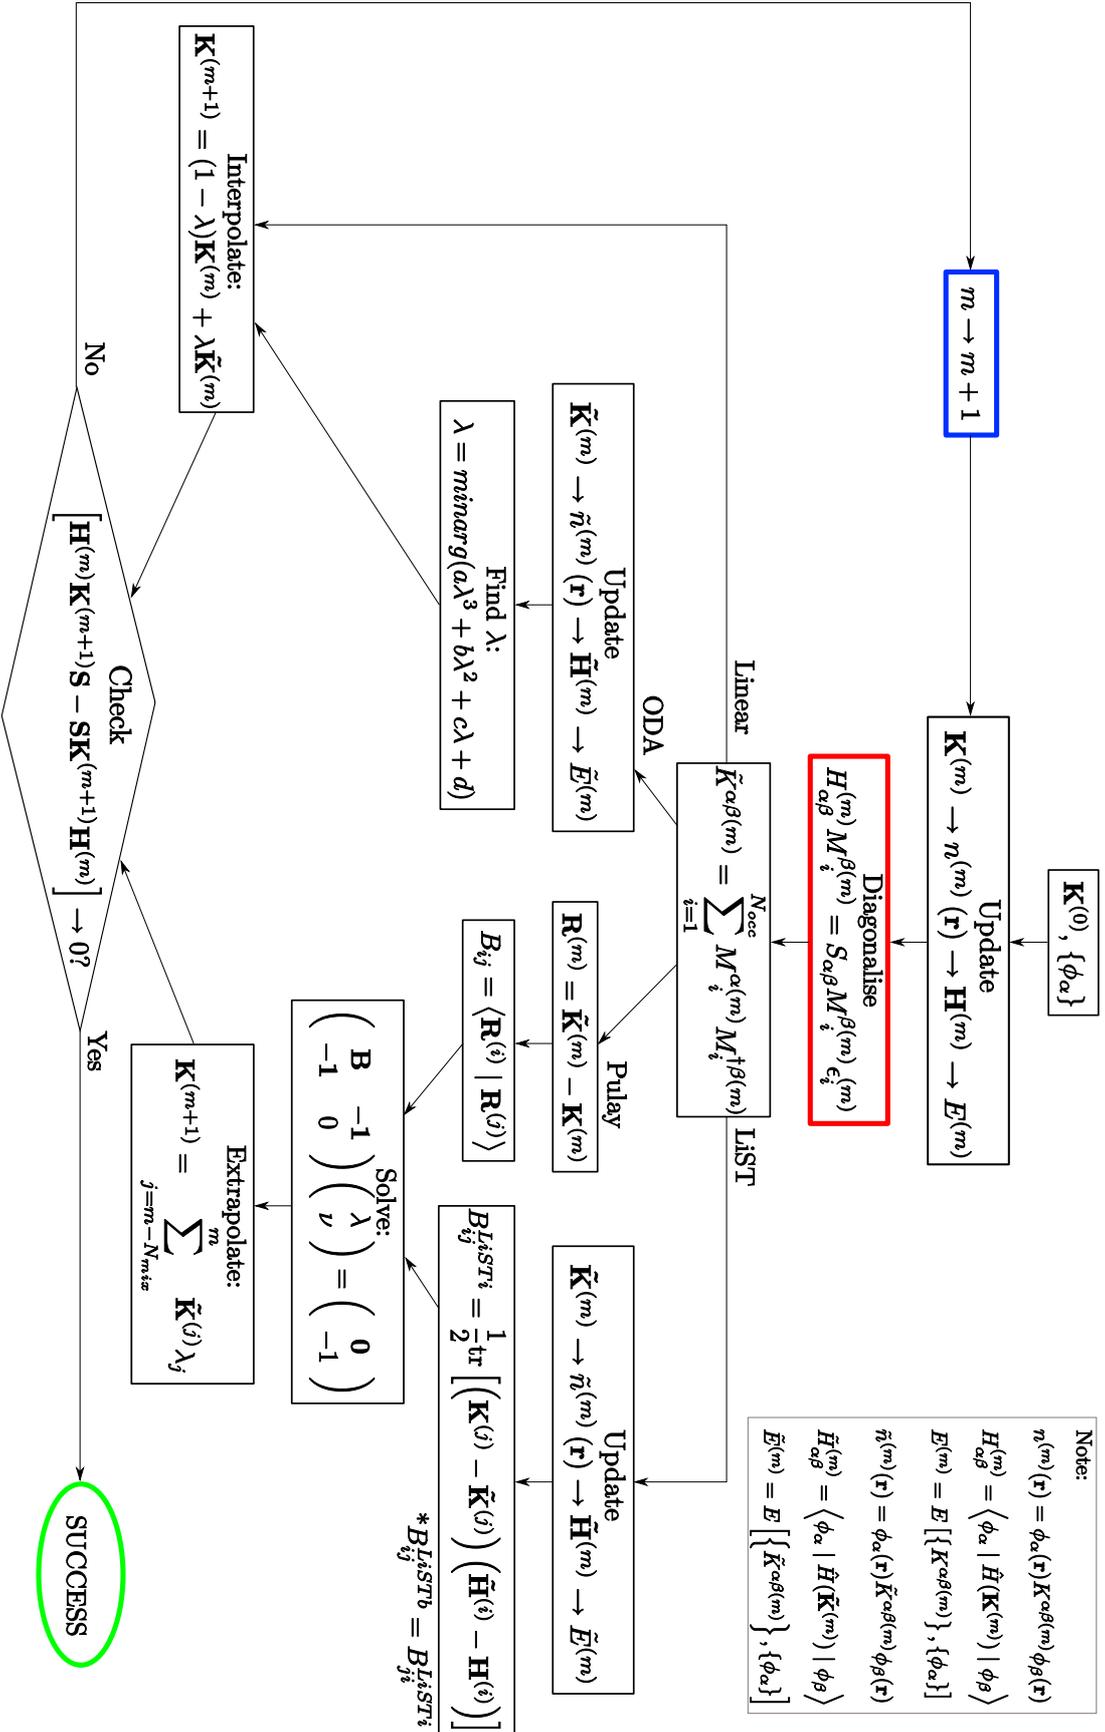


Figure 5.6: Flowchart of the DIIS implementation in ONETEP.

Chapter 6

Direct minimization method for metals

The energy levels in metallic systems form a continuum near the Fermi level that results in a zero band-gap between occupied and non-occupied states. This characteristic of metallic systems makes zero-temperature self-consistent approaches impractical, mostly due to the phenomenon of level-crossing [94], i.e., the impossibility of assigning integer occupancies (according to the *aufbau principle*) to states in the boundary of the Fermi level that are practically identical in energy. Mermin's extension of DFT to finite-temperature statistical ensembles [55], and the subsequent development of finite-temperature Kohn-Sham DFT [36], allowed calculations on metals to be performed efficiently. Within this formalism, the Kohn-Sham states have fractional occupancies determined by the Fermi-Dirac distribution (2.41), regulated by a constant electronic temperature \mathcal{T} . As described in Section 2.3, the macroscopic state function in the canonical ensemble is the Helmholtz free energy, $A_{\mathcal{T}}$, which obeys the variational principle of quantum mechanics. Many self-consistent processes to find the Kohn-Sham states that minimise the Helmholtz free energy have been envisaged, using either density mixing [95, 217, 219] or direct free energy minimisation techniques [150, 212, 237–239].

Most Kohn-Sham DFT methods are cubic-scaling because of the need for diagonalising the Hamiltonian matrix to obtain the energy eigenvalues, $\{\epsilon_i\}$, which are then used to construct the fractional occupancies of the Kohn-Sham orbitals. The size of the systems that can be simulated with finite-temperature Kohn-Sham DFT is currently limited by this scaling, and studies of large metallic systems with a potential use in new technologies are not normally viable. Examples of new metallic materials with applications in industrial processes include nanoparticles for oxygen adsorption in catalytic reactions [240–242], interactions at metallic surfaces [243–245], and characterisation of nanoalloys and amorphous materials [44, 246, 247]. Large calculations on metallic systems have nevertheless been performed successfully by the Nørskov group on gold and platinum nanoparticles [240, 241], at the expense of using tens of thousands of cores in parallel clusters.

In this chapter, we discuss the main points of a new implementation of a direct free-energy minimisation method in ONETEP. Localised orbitals have proven to be both useful and accurate in calculations on insulators and semi-conductors, but they have seldom been trialed on metallic systems in the past [168]. Our approach requires a Hamiltonian diagonalisation step which makes the overall cost scale cubically with the number of atoms in the system. However, the cost of Hamiltonian diagonalisations is reduced to the minimum as a consequence of using strictly-localised NGWFs, which are optimised in terms of the psinc basis set to achieve a high degree of accuracy. In this way, the rank of the Hamiltonian matrix is reduced to the minimum possible, and the prefactor due to diagonalisation is largely reduced. The remaining parts of the code can be computed at a linear-scaling cost due to sparsity of the matrices in the NGWF representation. This approach follows the ideas proposed by Fattebert and Bernholc [150].

The study of DIIS methods based on Hamiltonian diagonalisation, described in Chapter 5, showed that parallel dense algebra techniques are essential if calculations with thousands of atoms are to be performed. These methods serve a double purpose: they distribute both the cost and the storage of dense algebraic operations across a large number of processors. In the present implementation we used the PDSYGVX parallel eigensolver included in the SCALAPACK package [228]. The results show how our method is capable of achieving a high level of accuracy in large-scale calculations on metallic systems with thousands of atoms. The computational cost is also considerably reduced, which makes our method a powerful tool for studying very large systems which are inaccessible with conventional methods. These include some of the largest metallic nanoclusters simulated to date with the DFT level of theory, and which can have a potential application on industrial catalysts.

6.1 Implementation in the ONETEP program

The process of minimising the Helmholtz free energy functional, $A_{\mathcal{T}}[n]$, with respect to the electronic density, $n(\mathbf{r})$, is restated as a variational problem in terms of the NGWF representation of the Hamiltonian matrix, $\{H_{\alpha\beta}\}$, and the NGWFs expansion coefficients, $\{c_{p\alpha}\}$, as ¹:

$$A_{\mathcal{T}}[n] = A_{\mathcal{T}}[\{H_{\alpha\beta}\}, \{c_{p\alpha}\}] = E[\{H_{\alpha\beta}\}, \{c_{p\alpha}\}] - \mathcal{T}S[\{H_{\alpha\beta}\}]. \quad (6.1)$$

¹The entropy due to the Fermi-Dirac distribution only depends on the occupancies, $\{f_i\}$. The energy levels that give rise to the occupancies are obtained via diagonalisation of the Hamiltonian matrix in the current NGWF representation. Therefore, $\{f_i\}$ depends on $\{H_{\alpha\beta}\}$, and so the entropic term can be considered a functional of the Hamiltonian matrix elements.

The functional $A_{\mathcal{T}} [\{H_{\alpha\beta}\}, \{c_{p\alpha}\}]$ is then minimised using two nested loops, based on the following procedure:

$$\begin{aligned}
 A_{\mathcal{T}} [n_0] &= \min_{\{H_{\alpha\beta}\}, \{c_{p\alpha}\}} A_{\mathcal{T}} [\{H_{\alpha\beta}\}, \{c_{p\alpha}\}] & (6.2) \\
 &= \min_{\{c_{p\alpha}\}} \min_{\{H_{\alpha\beta}\}} A_{\mathcal{T}} [\{H_{\alpha\beta}\}, \{c_{p\alpha}\}] \\
 &= \min_{\{c_{p\alpha}\}} A'_{\mathcal{T}} [\{c_{p\alpha}\}],
 \end{aligned}$$

subject to the constraints of spatial localisation of the NGWFs (2.123), orthonormality of the Kohn-Sham states ($\langle\psi_i|\psi_j\rangle = \delta_{ij}$), and conservation of the number of electrons (2.40). Each loop uses a direct free-energy minimisation technique (steepest-descent or conjugate gradients), to achieve self-consistency.

The approach based on two nested loops was pioneered by the ensemble-DFT method of Marzari *et al.* [212], although in their method, the free energy functional was minimised in the inner loop with respect to the occupancies $\{f_i\}$ instead of $\{H_{\alpha\beta}\}$. The idea of a self-consistent procedure in terms of the Hamiltonian elements corresponds to Freysoldt *et al.* [239]. We prefer this method for two reasons:

- It ensures that the *aufbau* principle, which prescribes that the occupancies are determined by the energy levels, and not the other way around, holds at all times. The occupancies are calculated with the Fermi-Dirac distribution (2.41), which has an associated entropic term that is well known, given by Eq. (2.42). If the *aufbau* principle did not hold, then neither the Fermi-Dirac distribution or its entropic term would be valid.
- An approach based on the energy eigenvalues ensures that only the Hamiltonian matrix has to be diagonalised at every iteration. The eigenspectra of

$\{H_{\alpha\beta}\}$ (the set of energy eigenvalues $\{\epsilon_i\}$) is not compressed within a particular range of real values, and degenerate states appear in few subgroups which are orthonormal to each other. Reorthonormalising each subgroup independently does not require extensive computational resources and makes the process suitable for large calculations with parallel eigensolvers. On the other hand, an approach based on the occupancies requires the diagonalisation of the density kernel matrix, \mathbf{K} . The eigenvalues of \mathbf{K} are compressed to be in the range $[0, 1]$ and therefore degeneracy appears much more frequently and in much larger groups. Reorthonormalising large groups requires an immense effort from the parallel diagonaliser, which is very time and memory consuming. In many cases, the memory requirements were found to be so elevated that the calculations could not proceed.

6.2 Inner loop: energy eigenvalues optimisation

At iteration m of the inner loop, the Hamiltonian matrix is first diagonalised as:

$$H_{\alpha\beta}^{(m)} M_i^{\beta(m)} = S_{\alpha\beta} M_i^{\beta(m)} \epsilon_i^{(m)}, \quad (6.3)$$

after which the degenerate eigenstates are orthonormalised with the Löwdin symmetric orthonormalisation method [195]. The occupancies $\{f_i^{(m)}\}$ are calculated using the Fermi-Dirac distribution:

$$f_i^{(m)}(\epsilon_i^{(m)}) = \left(1 + \exp \left[\frac{\epsilon_i^{(m)} - \mu^{(m)}}{k_B T} \right] \right)^{-1}. \quad (6.4)$$

The Fermi level, $\mu^{(m)}$, is calculated at a negligible expense using the bisection method based on the constraint $\sum_i f_i^{(m)} = \mathcal{N}$. The entropy term $S^{(m)}$ is calculated

from $\{f_i^{(m)}\}$ as:

$$S^{(m)} \left[\{f_i^{(m)}\} \right] = -k_B \sum_i \left[f_i^{(m)} \ln f_i^{(m)} + (1 - f_i^{(m)}) \ln (1 - f_i^{(m)}) \right]. \quad (6.5)$$

The density kernel is now built from the occupancies as:

$$K^{\alpha\beta(m)} = \sum_i M_i^{\alpha(m)} f_i^{(m)} M_i^{\dagger\beta(m)}. \quad (6.6)$$

and the electronic density, $n^{(m)}(\mathbf{r})$, is updated as:

$$n^{(m)}(\mathbf{r}) = \phi_\alpha(\mathbf{r}) K^{\alpha\beta(m)} \phi_\beta^*(\mathbf{r}). \quad (6.7)$$

Then, the density-dependent components of the Hamiltonian operator $\hat{V}_H [n^{(m)}]$ and $\hat{V}_{XC} [n^{(m)}]$ are updated and the Helmholtz free energy functional $A_{\mathcal{T}}^{(m)}$, is evaluated. A new Hamiltonian matrix of elements $\{\tilde{H}_{\alpha\beta}^{(m)}\}$ is constructed as:

$$\tilde{H}_{\alpha\beta}^{(m)} = \langle \phi_\alpha | \hat{T} + \hat{V}_{ext} + \hat{V}_H [n^{(m)}] + \hat{V}_{XC} [n^{(m)}] | \phi_\beta \rangle. \quad (6.8)$$

The values of $\{\tilde{H}_{\alpha\beta}^{(m)}\}$ are used to define a search direction $\Delta^{(m)}$ of elements $\{\Delta_{\alpha\beta}^{(m)}\}$ in the multidimensional space of non-diagonal Hamiltonian matrix elements as:

$$\Delta_{\alpha\beta}^{(m)} = \tilde{H}_{\alpha\beta}^{(m)} - H_{\alpha\beta}^{(m)}. \quad (6.9)$$

The Hamiltonian matrix $H_{\alpha\beta}^{(m)}$ is updated as:

$$H_{\alpha\beta}^{(m+1)} = H_{\alpha\beta}^{(m)} + \lambda \Delta_{\alpha\beta}^{(m)}. \quad (6.10)$$

The choice of this search direction is similar to the choice of search directions in the algorithms by Marzari *et al.* [212] and Freysoldt *et al.* [239]. The above equation can be recast into $\mathbf{H}^{(m+1)} = (1 - \lambda) \mathbf{H}^{(m)} + \lambda \tilde{\mathbf{H}}^{(m)}$, where λ acts as an optimal damping parameter that has been chosen with the purpose of minimising the free energy functional. The parameter λ is calculated using a second-order polynomial fitting of the free energy functional after evaluating $A_{\mathcal{T}} [\{\epsilon_i\}, \{c_{p\alpha}\}]$ along the search direction $\Delta^{(m)}$ at $\lambda = 0$ and at two trial steps in the range $[0, 1]$. Occasionally, in the very first iterations, a third-order polynomial fitting may be required. Rarely, a fourth-order polynomial can be necessary. The algorithm quickly falls into the parabolic regime. The minimum is reached when the commutator

$$\left[\mathbf{H}^{(m+1)}, \mathbf{K}^{(m)} \right] = \mathbf{H}^{(m+1)} \mathbf{K}^{(m)} \mathbf{S} - \mathbf{S} \mathbf{K}^{(m)} \mathbf{H}^{(m+1)} \quad (6.11)$$

falls below the pre-selected tolerance threshold. Typically, 3 to 5 inner loop iterations are required per outer loop iteration.

6.3 Outer loop: NGWF optimisation

The outer loop optimises $\{c_{p\alpha}\}$, with fixed $\{H_{\alpha\beta}\}$ by minimising the projected free energy functional $A'_{\mathcal{T}} [\{c_{p\alpha}\}]$ [144, 146, 147, 150, 151]. The NGWF optimisation in the metals method follows the principles of the truncated search direction method (TSD), explained in detail in Section 4.4. The gradient $g_p^{\beta(l)}$ of the Helmholtz free-energy functional $A_{\mathcal{T}}^{(l)}$ with respect to $c_{p\beta}^{(l)}$ returns :

$$g_p^{\beta(l)} = \frac{\partial A'_{\mathcal{T}}^{(l)}}{\partial c_{p\beta}^{(l)}} = \left[\hat{H} \phi_{\gamma}^{(l)}(\mathbf{r}) K^{\gamma\beta} \right]_{\mathbf{r}=\mathbf{r}_p}. \quad (6.12)$$

Orthonormality of the KS states $\{\psi_i\}$ can be imposed to first order by projecting

out the components of $g_p^{\beta(l)}$ that are parallel to $c_{p\beta}^{(l)}$, as described in Ref. [211]. The resulting gradient is:

$$\tilde{g}_p^{\beta(l)} = \frac{\partial A_{\mathcal{I}}^{(l)}}{\partial c_{p\beta}^{(l)}} = \left[\hat{H} \phi_{\gamma}^{(l)}(\mathbf{r}) K^{\gamma\beta} - \phi_{\gamma}^{(l)}(\mathbf{r}) K^{\gamma\delta(l)} H_{\delta\nu} S^{\nu\beta(l)} \right]_{\mathbf{r}=\mathbf{r}_p}. \quad (6.13)$$

which is contravariant with respect to the index β . To obtain covariant representation, $\tilde{g}_p^{\beta(l)}$ is right-multiplied with $S_{\beta\alpha}^{(l)}$ [159]. Then, kinetic energy preconditioning [176] and occupancy preconditioning [212] are applied, resulting in the steepest-descent search direction coefficients:

$$d_{p\alpha}^{(l)} = -\hat{P} \left[\hat{H} \phi_{\alpha}^{(l)}(\mathbf{r}) - \phi_{\beta}^{(l)}(\mathbf{r}) S^{\beta\gamma(l)} H_{\gamma\alpha}^{(l)} \right]_{\mathbf{r}=\mathbf{r}_p}. \quad (6.14)$$

where \hat{P} is a generic kinetic energy preconditioning operator [40] and the leading minus is a consequence of the “downhill” update of the steepest descent method. To account for the localisation constraint of the NGWFs, those search direction coefficients that correspond to points \mathbf{r}_p outside the spherical localisation regions are set to zero. Using the psinc representation, a spherically-localised steepest-descent search direction vector $\Delta_{\alpha}^{(l)}$ is defined as:

$$\Delta_{\alpha}^{(l)}(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) d_{p\alpha}^{(l)}. \quad (6.15)$$

The conjugate gradients search direction $|\chi_{\alpha}^{(l)}\rangle$ is obtained by conjugating $|\Delta_{\alpha}^{(l)}\rangle$ to $|\chi_{\alpha}^{(l-1)}\rangle$ as:

$$|\chi_{\alpha}^{(l)}\rangle = |\Delta_{\alpha}^{(l)}\rangle + \frac{\langle \tilde{\Gamma}^{\beta(l)} | \Delta_{\beta}^{(l)} \rangle}{\langle \tilde{\Gamma}^{\gamma(l-1)} | \chi_{\gamma}^{(l-1)} \rangle} |\chi_{\alpha}^{(l-1)}\rangle, \quad (6.16)$$

where

$$\tilde{\Gamma}^{\beta(l)} = \sum_{p=1}^{N_p} D_p(\mathbf{r}) \tilde{g}_p^{\beta(l)}. \quad (6.17)$$

The conjugate gradients search direction, $\chi_\alpha^{(l)}(\mathbf{r})$, is also spatially-localised (since the localisation regions do not change), and can be expressed in the psinc basis set as:

$$\chi_\alpha^{(l)}(\mathbf{r}) = \sum_{p=1}^{N_p} D_p(\mathbf{r}) x_{p\alpha}^{(l)}, \quad (6.18)$$

where $\{x_{p\alpha}^{(l)}\}$ are its expansion coefficients. Convergence of the outer loop is judged by monitoring the slope along this search direction:

$$\text{slope}_\phi = \sum_{p=1}^{N_p} \tilde{g}_p^{\alpha(l)} x_{p\alpha}^{(l)}. \quad (6.19)$$

If self-consistency is not yet achieved, the NGWF coefficients on the grid are updated as:

$$c_{p\alpha}^{(l+1)} = c_{p\alpha}^{(l)} + \eta x_{p\alpha}^{(l)}, \quad (6.20)$$

where η is the optimal step-length obtained using a second-order polynomial fitting of $A'_T[\{c_{p\alpha}\}]$ along the search direction $|\chi_\alpha^{(l)}\rangle$ (a third-order polynomial fitting is rarely needed). The algorithm continues by updating the overlap and inverse overlap matrices $\{S_{\alpha\beta}^{(l+1)}\}$ and $\{S^{\alpha\beta(l+1)}\}$, respectively. For consistency, the transformation $\{M_i^\alpha\}$ must be represented in terms of the new NGWFs, $\{\phi_\alpha^{(l+1)}\}$.

Combining the closure relation:

$$\hat{I} = |\phi_\beta^{(l+1)}\rangle S^{\beta\gamma(l+1)} \langle\phi_\gamma^{(l+1)}| \quad (6.21)$$

and the expansion of the KS states in terms of the previous NGWF representation

$$|\psi_i^{(l)}\rangle = |\phi_\alpha^{(l)}\rangle M_i^\alpha, \quad (6.22)$$

the following expression is obtained:

$$\begin{aligned} |\psi_i^{(l)}\rangle &= |\phi_\alpha^{(l)}\rangle M_i^\alpha \\ \implies |\psi_i^{(l)}\rangle &= |\phi_\beta^{(l+1)}\rangle S^{\beta\gamma(l+1)} \langle\phi_\gamma^{(l+1)}|\phi_\alpha^{(l)}\rangle M_i^\alpha \\ \implies \langle\phi_\mu^{(l+1)}|\psi_i^{(l)}\rangle &= \langle\phi_\mu^{(l+1)}|\phi_\beta^{(l+1)}\rangle S^{\beta\gamma(l+1)} \langle\phi_\gamma^{(l+1)}|\phi_\alpha^{(l)}\rangle M_i^\alpha \\ \implies \tilde{M}_i^\mu &= \delta_\beta^\mu S^{\beta\gamma(l+1)} \langle\phi_\gamma^{(l+1)}|\phi_\alpha^{(l)}\rangle M_i^\alpha \\ \implies \tilde{M}_i^\mu &= S^{\mu\gamma(l+1)} \langle\phi_\gamma^{(l+1)}|\phi_\alpha^{(l)}\rangle M_i^\alpha. \end{aligned} \quad (6.23)$$

The density kernel is rebuilt as:

$$\tilde{K}^{\alpha\beta} = \sum_i \tilde{M}_i^\alpha f_i \tilde{M}_i^{\dagger\beta}, \quad (6.24)$$

followed by the update of the electronic density:

$$n^{(l+1)}(\mathbf{r}) = \phi_\alpha^{(l+1)}(\mathbf{r}) \tilde{K}^{\alpha\beta} \phi_\beta^{*(l+1)}(\mathbf{r}), \quad (6.25)$$

and the density-dependent terms $\hat{V}_H[n^{(l+1)}]$ and $\hat{V}_{XC}[n^{(l+1)}]$. The projected functional $A_T'^{(l+1)}$ is evaluated, followed by the update of the Hamiltonian before re-entering the inner loop. The two-nested-loop approach is repeated to self-consistency until convergence to the desired tolerance threshold is achieved. We

found that the change of representation of Eq. (6.23) is very important for the stability and robustness of the algorithm. Calculations on metallic systems do not normally converge unless this step is performed.

6.4 Scaling with the system size

A key feature of the approach presented in this work is the use of NGWFs. As a consequence of the strict spherical localisation constraints imposed on $\{\phi_\alpha\}$, most of the overlap $\{S_{\alpha\beta}\}$ and Hamiltonian $\{H_{\alpha\beta}\}$ matrix elements are zero. Therefore, both matrices are sparse, allowing the custom-built, highly parallel, optimised sparse algebra routines included in the ONETEP program [183] to be utilised. The cost of algebraic operations with sparse matrices scales as $\mathcal{O}(N)$ with the number of the atoms in the system, N . Additionally, the cost of fast Fourier transforms also scales as $\mathcal{O}(N)$ with the use of the FFT-box technique [174, 175]. As a result, many essential steps of the algorithm such as updating the electronic density on a grid, updating the Hamiltonian elements or calculating the NGWF gradient can be performed at $\mathcal{O}(N)$ cost (linear-scaling).

However, the computational cost and memory requirements to diagonalise the Hamiltonian is $\mathcal{O}(N^3)$. Therefore, the overall scaling of the algorithm is also $\mathcal{O}(N^3)$ (cubic-scaling). The strategy followed in this work is not of avoiding the diagonalisation step, but rather of reducing its impact on performance. The NGWFs are essential to this goal. The set of $\{\phi_\alpha\}$ contains a minimal number of functions which are used to represent the Hamiltonian operator. Hence, the Hamiltonian matrix to be diagonalised is reduced to have the smallest possible dimensions, which reduces the prefactor associated to diagonalisation and its impact on the total runtime of the algorithm. Chemical accuracy is achieved by means of op-

timising the NGWFs in the psinc basis set during the outer loop at $\mathcal{O}(N)$ cost. The memory requirements are often a stronger limiting factor than the computational time. This is due to the fact that $\{M_i^\alpha\}$, $\{K^{\alpha\beta}\}$ and other temporary matrices are dense, amounting for $(N_b \times N)^2$ non-zero elements per matrix that must be stored. In order to be able to simulate large systems with thousand of atoms, memory storage of matrix elements has to be shared over multiple cores. To address this difficulty, the ScaLAPACK PDSYGVX parallel eigensolver [228] was used, not only to accelerate the diagonalisation step with parallel algebra techniques, but more importantly, to divide and distribute the matrix elements over the memory of a large number of cores in an efficient way [150].

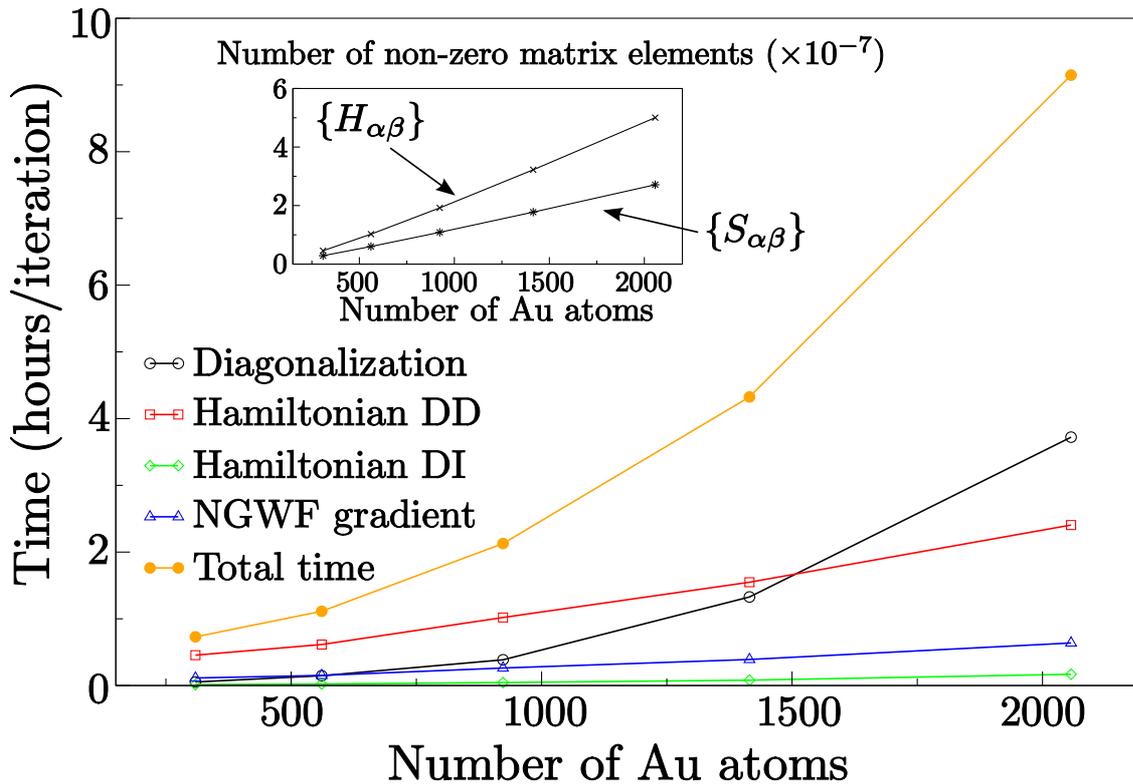


Figure 6.1: Time taken to complete one outer loop iteration with five inner loop iterations, in calculations on Au nanoparticles of increasing size. The plot shows the time taken by different parts of the algorithm. “Hamiltonian DD” and “Hamiltonian DI” must be interpreted as the density-dependent and density-independent terms of the Hamiltonian, respectively.

To demonstrate the computational demands of some of the most relevant parts of the algorithm, a set of benchmark calculations on Au cuboctahedral nanoparticles of increasing size, ranging between 309 and 2057 atoms, was performed. All the calculations were run on 120 cores distributed over 20 6-core 2.4 GHz Intel Westmere processors and 3.6 GB of memory per core. The calculations use NGWF radii of $R_\alpha = 5.0 \text{ \AA}$, a kinetic energy cut-off of 700 eV, an electronic smearing of $k_B T = 0.1 \text{ eV}$, the PBE exchange-correlation functional [70] and a norm-conserving pseudopotential with 11 valence electrons per atom. Figure 6.1 shows the time per one outer loop iteration with five inner loop iterations compared to the system size, given by the number of atoms. The inset plot shows the number of non-zero matrix elements $\{S_{\alpha\beta}\}$ and $\{H_{\alpha\beta}\}$ in each calculation. The number of non-zero $\{S_{\alpha\beta}\}$ and $\{H_{\alpha\beta}\}$ matrix elements scales linearly with the number of atoms, and both matrices remain sparse regardless of the size of the system. The fact that the Hamiltonian matrix is less sparse compared to the overlap matrix is due to the non-local part of the norm-conserving pseudopotential [175]. Algebraic operations that involve these two matrices can be carried out at a linear-scaling cost using parallel sparse algebra methods [183]. The time to execute critical steps of the algorithm such as building the density-independent and density-dependent terms of the Hamiltonian matrix (including updating the electronic density) and calculating the NGWF gradient scale linearly with the number of atoms. For small systems, building the density-dependent terms of the Hamiltonian is the most computationally-demanding part of the calculation. As the number of atoms increases, diagonalisation becomes more costly, and rapidly claims most of the computational effort at every iteration. In calculations of 1500 atoms and above, diagonalising the Hamiltonian is the most expensive part of the calculation.

Eventually, the cost of diagonalisation will limit the size of systems that can be simulated.

6.5 Results

A number of validation calculations were completed in order to test the accuracy of the method. In all the calculations shown hereafter the RMS commutator (6.11) was converged within a tolerance of at least 1×10^{-4} eV, while the RMS of the slope of the NGWF gradient along the search direction (6.19) was converged to at least 5×10^{-5} eV. With these settings, the free energy functional is converged below 10^{-6} eV/atom in all cases. First, BFGS geometry optimisation calculations [204] on a Pt₁₃ cuboctahedral nanoparticle were performed. The results were compared to those obtained with the finite-temperature Kohn-Sham DFT direct minimisation method included in CASTEP [90], which uses fully delocalised orbitals and a plane-waves basis set. The calculations with the two programs use the same settings: a kinetic energy cut-off of 1000 eV, an electronic smearing of $k_B T = 0.1$ eV, the PBE exchange-correlation functional [70], the same set of norm-conserving pseudopotentials with 10 valence electrons per atom, and a total of 117 energy bands. In the calculations with ONETEP, 9 NGWFs are assigned per atom. The optimised structures obtained with NGWF radii in the range 3.0 Å to 6.0 Å were compared. Table 6.1 shows the final value of the distance to the nearest-neighbour Pt atom. Convergence of the optimised geometry to a single structure is achieved as R_α increases. The structures obtained with ONETEP match the result obtained with CASTEP within 0.01 Å tolerance for $R_\alpha = 4.0$ Å and above. The density of states (DOS) of the optimised structures, calculated with a Gaussian smearing of 0.1 eV, is shown in Fig. 6.2. The DOS obtained with ONETEP also match the

CASTEP results for $R_\alpha = 4.0 \text{ \AA}$ and above, in the region of energies below and in the vicinity of the Fermi level μ . The ONETEP DOS is known to be inaccurate to describe the higher-energy, near-empty bands above the Fermi level, unless the conduction bands are optimised further [187].

Table 6.1: BFGS geometry optimisation of Pt_{13} with ONETEP and CASTEP. The table shows the optimised value of the distance to the nearest-neighbour Pt atom.

	Distance (\AA)
ONETEP ($R_\alpha = 3.0 \text{ \AA}$)	2.64
ONETEP ($R_\alpha = 3.5 \text{ \AA}$)	2.66
ONETEP ($R_\alpha = 4.0 \text{ \AA}$)	2.69
ONETEP ($R_\alpha = 4.5 \text{ \AA}$)	2.69
ONETEP ($R_\alpha = 5.0 \text{ \AA}$)	2.69
ONETEP ($R_\alpha = 5.5 \text{ \AA}$)	2.70
ONETEP ($R_\alpha = 6.0 \text{ \AA}$)	2.70
CASTEP	2.70

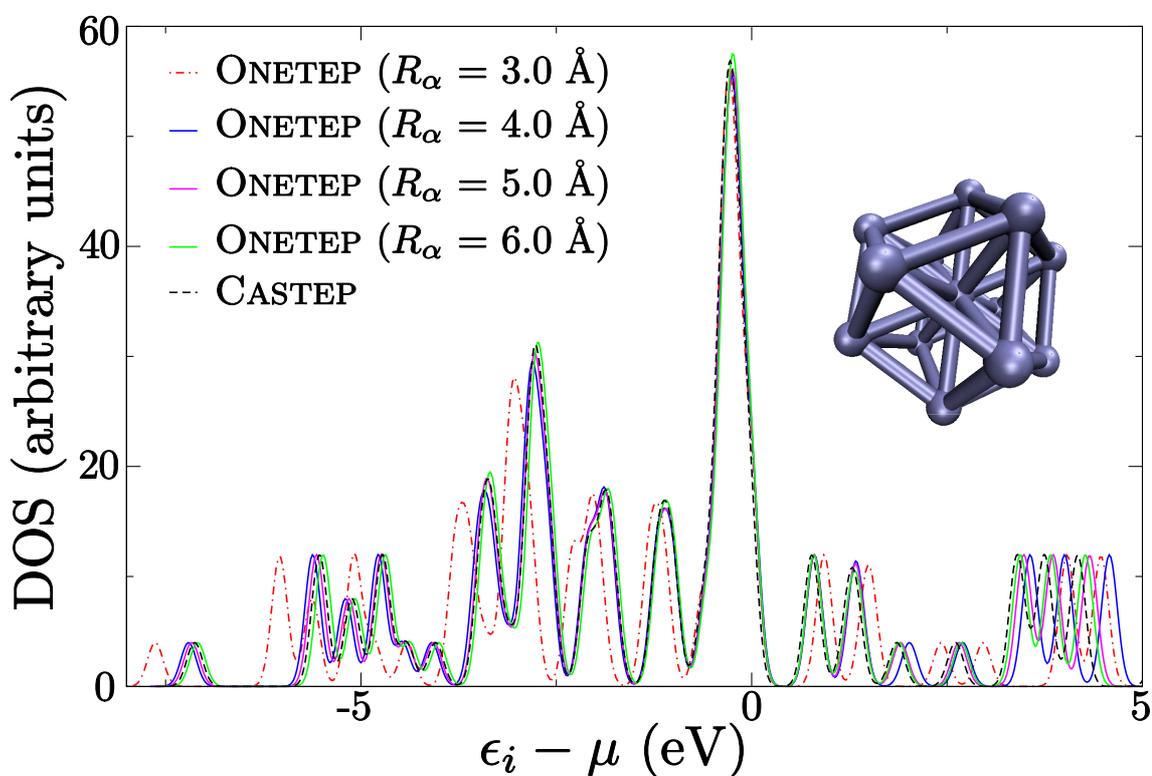


Figure 6.2: Density of states of Pt_{13} , obtained with ONETEP and CASTEP. Agreement is achieved for NGWF radii of 4.0 \AA and above.

The next validation test involved calculations on face-centred cubic (fcc) bulk Cu. The free energy potential associated with variations in the lattice parameter was obtained with ONETEP and compared against the potential obtained with CASTEP. The ONETEP calculations include only the Γ -point. In order to imitate the behaviour of the k-point grid, at least to some extent, the four-atom fcc cell of bulk Cu was replicated five times in each lattice vector direction, resulting in a $5 \times 5 \times 5$ supercell containing 500 atoms. The procedure described in Ref. [181] to vary the lattice parameter while maintaining the kinetic energy cut-off constant in the ONETEP calculations was used. The settings in both programs are the same: a kinetic energy cut-off of 898.2 eV, an electronic smearing of $k_B T = 0.1$ eV, the RPBE exchange-correlation functional [71] and the same norm-conserving pseudopotentials with 11 valence electrons per atom. In ONETEP, each Cu atom was assigned 9 NGWFs of radii $R_\alpha = 4.0 \text{ \AA}$, amounting for a total of 4500 energy bands in the system. The CASTEP calculations include one four-atom fcc cell, a Monkhorst-Pack k-point grid [102] of 5 points in each direction, and a total of 27 energy bands in the system. The results are shown in Fig. 6.3. The bulk modulus B and the equilibrium lattice parameter L_0 were calculated by fitting a curve determined by the third-order Birch-Murnaghan equation [248] to the results. These are shown in Table 6.2. The calculations with CASTEP and ONETEP agree very well in their predictions of L_0 (within 0.01 \AA) and B (within 0.06 GPa). These results on bulk Cu confirm that a similar level of accuracy in the description of crystalline systems can be achieved compared to standard Kohn-Sham DFT methods. More importantly, these calculations exemplify how calculations on periodic systems with hundreds of atoms in the simulation cell can be performed with the algorithm presented in this work. This capability can be used to study complicated

periodic structures formed by many atoms in one, two or three dimensions, or with many lattice defects. Amorphous systems could also be studied in closer detail by allowing more atoms in the simulation cell [246, 247].

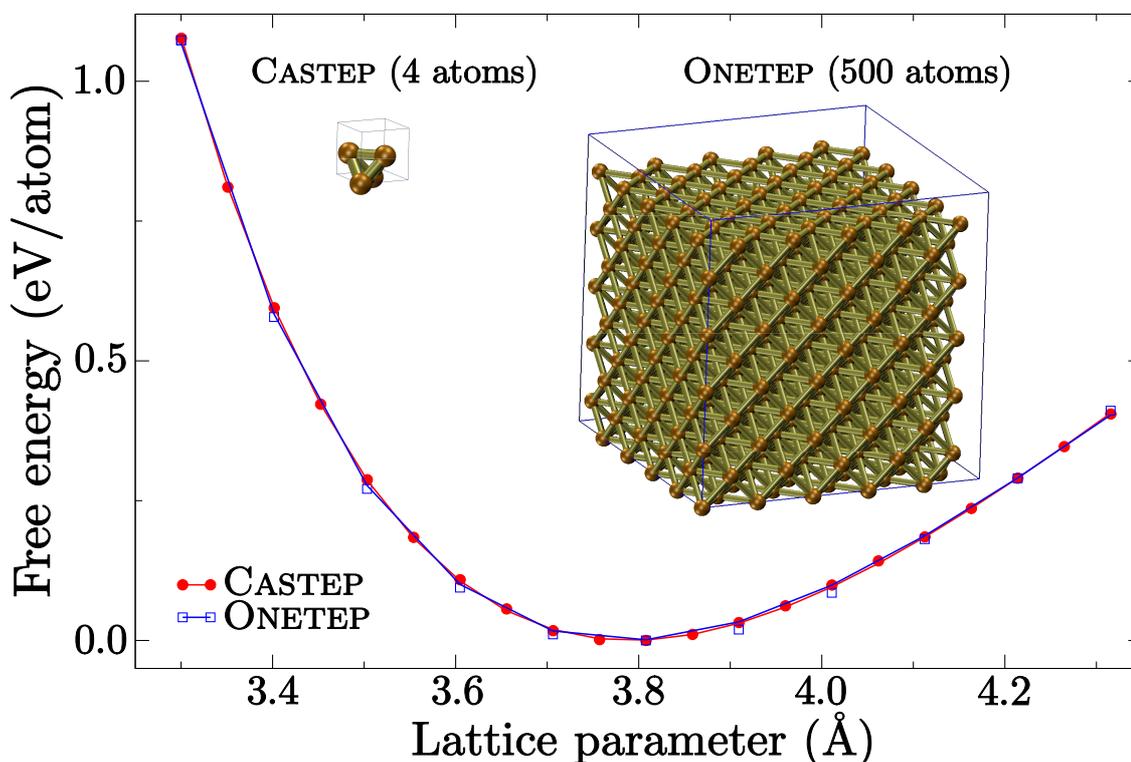


Figure 6.3: Lattice parameter stretching of bulk Cu. There are 4 atoms in the CASTEP simulation cell and 500 atoms in the ONETEP simulation cell, forming a $5 \times 5 \times 5$ super-cell.

Table 6.2: Bulk modulus, B , and equilibrium lattice parameter, L_0 , of bulk Cu, calculated with CASTEP and ONETEP. The value of χ^2 corresponding to the fitting of the results to the third-order Birch-Murnaghan equation is also shown.

	B (GPa)	L_0 (Å)	χ^2
CASTEP	93.57	3.79	6×10^{-5}
ONETEP	93.51	3.79	6×10^{-4}

The last set of tests aims to provide an insight into the capabilities of the algorithm to successfully complete calculations on metallic systems that can be of practical use on industrial processes. A series of calculations on isolated Au cuboc-

tahedral nanoparticles of different sizes, from 0.4 to 3.4 nm wide, in the range of 13 to 2057 atoms, were performed. These kind of nanoparticles are of interest for the development of novel catalysts, but their large sizes make computational studies with Kohn-Sham DFT expensive. Recent studies have reported that 32768 cores were required to perform calculations on Au_{1415} using standard plane-waves Kohn-Sham DFT methods [240].

The calculations with ONETEP use the same configuration as in the calculations shown in Fig. 6.1. Convergence is achieved when the change in the free energy per atom is less than 10^{-6} eV and the RMS gradient along the NGWF search direction falls below 2×10^{-5} eV. All the calculations converged within these tolerance thresholds. Figure 6.4 shows the convergence rate with number of outer loop iterations using ONETEP. To some extent, the number of iterations required is system-dependent and increases with the number of atoms. The largest of the Au nanoparticles, containing 2057 atoms, converged within 240 hours using 300 cores distributed across 25 12-core 2.4 GHz Intel Westmere processors with 1.8 GB of memory per core. Provided there are enough computational resources, the method presented in this work could enable calculations to study even larger industrial catalysts which typically consist of thousands of metallic atoms [242].

6.6 Summary

In this work, we successfully developed an approach for large-scale calculations on metallic systems with thousands of atoms. The method is based on direct minimisation of the Helmholtz free energy functional, which makes it robust and avoids charge-sloshing and level-crossing instabilities. In its implementation within the ONETEP program, the cost of most parts of the calculation is reduced to the linear-

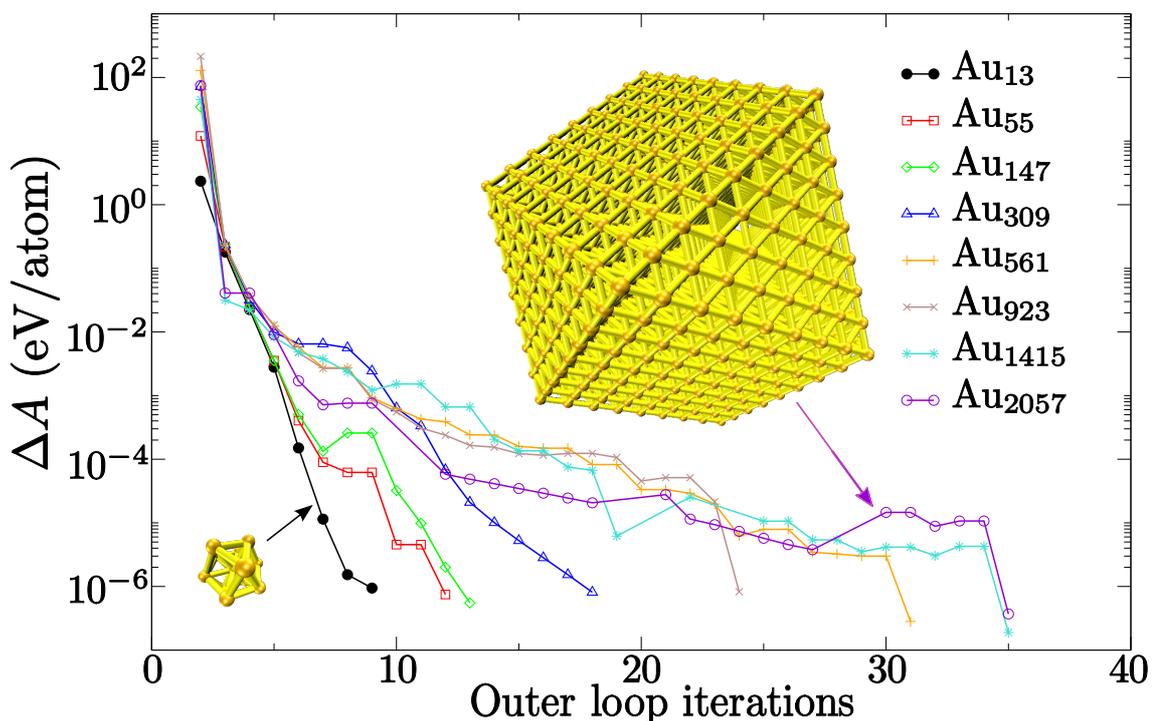


Figure 6.4: Convergence of the Helmholtz free energy functional, A_T , with the number of outer loop (NGWF optimisation) iterations, for a set of Au cuboctahedral nanoparticles of increasing sizes. The structures of Au_{13} and Au_{2057} are also shown in the plot.

scaling regime, thanks to the localisation constraints imposed upon the NGWFs. Accurate results are obtained after the optimisation of the NGWFs in the psinc basis set. The impact of Hamiltonian diagonalisation, which remains cubic-scaling, is reduced to the minimum. The memory requirements are distributed over many cores in an efficient manner using the SCALAPACK dense algebra package. Since diagonalisation is an essential part of our method, the performance of the algorithm can be improved by interfacing ONETEP to other parallel diagonalisers [235].

We validated the method with calculations on very large molecules, including a set of Au nanoparticles with up to 2057 atoms. Our approach allows this kind of calculations to take place using significantly less number of cores on modern parallel clusters. Therefore, our method has the potential to expand the range of metallic compounds that can be studied with finite-temperature Kohn-Sham DFT

and to allow simulations of larger nanoparticles or alloys with direct industrial relevance.

As a future development the code can be modified to exploit locality of the one-particle density matrix. As explained in Section 2.9, $\rho(\mathbf{r}, \mathbf{r}')$ decays exponentially (as a modulated wave) at finite temperature \mathcal{T} [120, 121]. The method of the Fermi operator expansion [139, 140] can be used to generate the Fermi-Dirac distribution without diagonalising the Hamiltonian matrix. Such functionality can be implemented in the ONETEP code to reduce the cost of metallic calculations to the quadratic or even linear scaling regimes. This functionality could potentially expand the size of metallic systems that can be simulated beyond tens of thousands of atoms.

Chapter 7

Conclusions and final remarks

Our investigations focused on new developments for computational simulations of molecules and materials based on Kohn-Sham density functional theory. More precisely, we use methods in which the single-electron molecular orbitals are expanded as a linear combination of atom-centred, spatially-localised functions to reduce the computational cost. This technique has been successfully applied to insulators and semi-conductors by many authors to produce linear-scaling DFT algorithms. We used the ONETEP program as the main platform for our development, but nevertheless we are confident that our research can reach other methods and codes that are used for electronic-structure theory calculations. Insulators show a favourable trend of localisation, which manifests itself in the form of exponential decay of the one-particle density matrix. The discrete energy levels and the broad band-gap at the Fermi level can be used to set a clear distinction between conduction and valence states. The occupancies of the Kohn-Sham states can then be assumed to be perfectly-defined integer numbers (0 for conduction bands and 1 for valence bands). The computational effort can then be focused on the optimisation of the molecular orbitals, subject to the constraints of localisation, orthonormality

and conservation of the total number of electrons.

The constraints of spatial localisation are sometimes incompatible with orthonormality of the localised functions that are used to expand the Kohn-Sham states. This is acceptable, so long as mathematical correctness is kept by using formal tensorial algebra, typical of curved manifolds. Duality appears as a consequence of non-orthogonality. Matrices and vectors acquire the tensorial character of the covariant and the contravariant spaces, and are mutually connected by multiplication with the overlap and inverse overlap matrices. Assuming that the localised functions are covariant, matrices in a covariant representation tend to display optimal localisation properties. The covariant representation of the Hamiltonian and the overlap matrices is localised and becomes sparse for large systems. Contravariant tensors, such as the density matrix (or density kernel) must be truncated in order to gain any kind of localisation.

In many self-consistent field approaches, the localised functions are optimised *in situ* using steepest-descent or conjugate-gradients approaches, which perform a line-search step along the trajectory marked by a search direction vector. Traditionally, the optimisation of the localised functions is seen as an *unconstrained* optimisation process, where the localisation constraints are not taken into account to construct the search direction vectors. With this strategy, the resulting search direction vectors, calculated from the derivative of the energy, are fully delocalised. An *ad hoc* truncation step must take place in order to restore localisation of the search direction vector and to make it suitable for updating the localised functions.

The study of spatial localisation constraints was one of our major topics of research and proved to be quite fruitful. What allowed us to make progress in this field was our realisation that the energy minimisation process was in fact a

constrained optimisation problem, where the strict spatial localisation constraints must be explicitly taken into account at every step. The *true* Kohn-Sham ground-state (the one for which the energy gradient is exactly zero), cannot be found with methods based on spatial localisation constraints, unless infinitely large localisation regions are allowed. In a practical calculation with finite localisation regions, the solution of minimum energy is an approximation to the true Kohn-Sham ground-state. Nevertheless, it is reasonable to assume, based on locality of the electronic structure of insulating systems, that the approximate ground-state and the real ground-state are not too different. The energy gradient, however, cannot normally be considered to be equal to zero.

After adopting the perspective of a constrained optimisation problem, we were able to prove, theoretically and in practice, that if localisation constraints are required, the Pulay correction terms to the total atomic forces do not vanish. The accuracy of calculations that rely on the atomic forces to progress was dramatically improved after the addition of the Pulay forces. Weakly-bound systems with hydrogen bonds, such as protein-ligand complexes, which are more susceptible to numerical inconsistencies in the forces, are now correctly described. We demonstrated this point in geometry optimisation calculations, but molecular dynamics calculations, for example, will also benefit from our developments. From a wider perspective, our research pointed out that any derivative of the energy functional with respect to any external parameter must consider terms arising from the spatial localisation of the density matrix.

The *variationally-localised search direction* (VLSD) method is one of the most important outcomes of our research. The VLSD approach is, to the best of our knowledge, a completely new point of view of the constrained optimisation prob-

lem with non-orthogonal functions and strict localisation constraints. We implemented a prototypical serial version of the VLSD algorithm in the ONETEP program for the optimisation of non-orthogonal generalised Wannier functions (NGWFs). This method is based on two key observations. First, that the set of irreducible degrees of freedom corresponds only to the subset of non-zero elements inside the spherical localisation regions. And second, that the contravariant representation (in the case of the NGWFs) is delocalised, but nevertheless constrained by the strict localisation imposed upon the covariant representation. The values of those contravariant elements that are not connected to the set of irreducible degrees of freedom are uniquely determined after the localisation constraints are imposed. The method is exact, as opposed to the standard method based on truncation of the search direction vectors, which introduces an uncontrollable error that cannot be treated variationally. As the results show, the VLSD method consistently leads to rapid convergence of the calculations.

The computational implementation of the VLSD algorithm for NGWF optimisation is not particularly straightforward. The VLSD method emphasises the connections between values at the same grid point originating from functions in different centres. An efficient implementation of the VLSD method suitable for parallel computers must take this fact into account, perhaps by promoting a domain decomposition approach to divide the computational grid. Whether the resulting code will scale linearly with the number of atoms is yet to be seen. On the positive side, the VLSD method can be extended to solve other challenging problems of constrained optimisation, such as in calculations with selective NGWF optimisation or with localisation constraints upon the density kernel matrix. These applications are discussed in Appendices [B](#) and [C](#), respectively.

Another branch of our research focused on taking advantage of localisation for designing algorithms capable of large-scale calculations on insulators and metals. We implemented two different algorithms for cubic-scaling Kohn-Sham DFT for calculations on all kinds of systems, including metals. The pillars of these methods are:

- The use of the SCALAPACK library for parallel dense algebra operations to distribute the memory allocation and the computational effort due to operations with large dense matrices. We used a SCALAPACK parallel eigensolver to diagonalise the Hamiltonian matrix and to obtain its eigenvalues and eigenvectors. This operation is cubic-scaling, and therefore, the performance of the method is strongly linked to the performance of the eigensolver.
- The use of localised functions to expand the Kohn-Sham orbitals to bring sparsity to the Hamiltonian and overlap matrices, which allows most of the operations to be executed with linear-scaling cost. Chemical accuracy is obtained via an *in situ* optimisation in terms of a high resolution basis set of psinc functions. In this representation, the size of the Hamiltonian matrix is reduced to the minimum, and as a consequence, the prefactor due to diagonalisation is also reduced to the minimum. This is a critical point that ultimately enabled calculations scale up to larger systems with thousands of atoms using relatively modest computational resources.

A self-consistent cycle based only on Hamiltonian diagonalisation is unlikely to converge. Density mixing methods, also called DIIS methods (direct inversion of the iterative subspace), can accelerate the process by gathering information from a set of previous iterations to improve the next trial electronic density. We imple-

mented two interpolation methods (linear mixing and the optimal damping algorithm, ODA), and two extrapolation methods (Pulay mixing and linear-expansion shooting techniques, LiST), as different approaches to density mixing. The validation tests on silicon nanorods show that density mixing methods can yield the same level of accuracy as the Li-Nunes-Vanderbilt approach for calculations on insulators. Typically, extrapolation methods are quicker than interpolation methods as they use a larger history of previous results.

Numerical stability of the current DIIS algorithms must be improved. The phenomenon of charge-sloshing seems to be the cause why some calculations fail to converge adequately to the state of minimum energy. As a further improvement, Kerker preconditioning must be implemented to stabilise density mixing calculations. DIIS methods can be an excellent tool for testing spatial localisation of the density kernel matrix since, by definition, density mixing is a form of constrained energy minimisation that takes into account localisation constraints. Finally, the DIIS methods can be further improved and extended to finite-temperature calculations on metals by allowing the occupancies of the Kohn-Sham states to take fractional values.

Simulations on metals are much more challenging than on insulators. This is due to the continuum of energy bands and the lack of band-gap at the Fermi level, as well as to the delocalised molecular orbitals that metals normally display. Zero-temperature DFT calculations in metals often fail due to issues with level-crossing, i.e., the impossibility of marking a clear boundary between occupied and unoccupied states which are virtually identical in energy. Finite-temperature DFT based on occupation smearing are more suitable for metallic systems. From a macroscopical point of view, the energy functional must be dropped in favour to the

Helmholtz free-energy functional, which contains the entropic term associated to the fractional occupancies. Direct minimisation techniques are remarkably robust and, in contrast to DIIS methods, do not suffer from charge sloshing.

We developed a direct Helmholtz free energy minimisation algorithm that has been included within the ONETEP program. The formalism based on parallel algebraic operations with the SCALAPACK library and *in situ* optimisation of the NGWFs allowed us to perform accurate Γ -point calculations on metallic systems with thousands of atoms. The computational cost is noticeably reduced: while standard plane-wave approaches need tens of thousands of cores to complete calculations on around one thousand atoms, we only need a few hundreds of cores. Our results are a demonstration that large-scale calculations on metals can actually be done. This new algorithm is now available to be used for practical investigation on a new range of metallic systems and compounds with direct use in engineering and industrial processes. For example, it can be used in the characterisation of metallic nanoparticles for catalytic reactions, the size of which often makes their study with traditional DFT approaches prohibitively expensive. Regarding future improvements in the code, a reformulation of the method in terms of the Fermi operator expansion would eliminate the need for explicit diagonalisation of the Hamiltonian. A Fermi-Dirac distribution could be generated by expanding the one-particle density operator as a sum of powers of the Hamiltonian matrix. The scaling with the system size could be reduced, first to quadratic-scaling, and ultimately to linear-scaling.

Appendix A

Spherical wave representation of the NGWFs

A way to store the result of a calculation on the hard drive is by decomposing the NGWFs as a linear combination of a given basis set and saving the coefficients of this expansion in a file. Spherical waves have been tested as a method for storing the NGWFs after a converged calculation. This method has two main advantages over traditional storage. First off all, it saves computational resources as less coefficients are required. Second, and most importantly, the spherical waves can be easily generated at any point of real space, which makes them a portable method for restarting calculations such as geometry optimisation or molecular dynamics without necessity of further extrapolation techniques.

A.1 Theory

Spherical waves (in what follows, referred to as SWs) are the solutions of the Helmholtz equation subject to periodic boundary conditions and confined within as spherical region of space. We refer the reader to Ref. [172] for a detailed study

of SWs and their properties. The SWs form a basis set of the Hilbert space that can be used for describing any other function in this space, such as the NGWFs. The SWs must be continuous in space, and localised within the NGWF sphere centred on each atom, which implies that the SWs are exactly zero in the border and outside of the sphere. SWs centred on the same point and localised within the same spherical region are mutually orthogonal. On the other hand, the overlap between SWs in different centres can be calculated analytically. Mathematically, the SWs, $\{\chi_{\alpha,lmn}\}$ are defined in the real space as:

$$\chi_{\alpha,lmn}(\mathbf{r}) = \begin{cases} j_l(q_{nl}r)\bar{Y}_{lm}(\Omega_{\mathbf{r}}) & r < R_{\alpha} \\ 0 & r \geq R_{\alpha} \end{cases}, \quad (\text{A.1})$$

where:

- j_l are spherical Bessel functions of the first kind, which act as radial functions that depend only on the modulus r of the coordinate vector \mathbf{r} .
- q_{nl} is a wave vector (modulus).
- \bar{Y}_{lm} are real spherical harmonics which depend upon the angular components, θ and φ (encapsulated in the single variable $\Omega_{\mathbf{r}}$).
- n indexes the Bessel zeros inside the sphere.
- l is the orbital momentum quantum number.
- m azimuthal component of the angular momentum (with values in the range $[-l, +l]$).

As the radii of the NGWFs localisation regions are known, the wave vector modulus q_{nl} is calculated using the boundary condition:

$$j_l(q_{nl}R_\alpha) = 0. \quad (\text{A.2})$$

The solutions of the above equation are the roots of the Bessel functions, which can be found numerically using a polynomial recursion formula for $j_l(r)$.

The NGWFs are expanded as a linear combination of the SWs basis set as:

$$\phi_\alpha(\mathbf{r}) = \sum_{l=0}^{l=\infty} \sum_{m=-l}^{m=+l} \sum_{n=1}^{n=\infty} \chi_{\alpha,lmn}(\mathbf{r}) C_{\alpha,lmn}, \quad (\text{A.3})$$

with the coefficients $C_{\alpha,lmn}$ calculated as:

$$C_{\alpha,lmn} = \frac{\langle \chi_{\alpha,lmn} | \phi_\alpha \rangle}{\langle \chi_{\alpha,lmn} | \chi_{\alpha,lmn} \rangle}. \quad (\text{A.4})$$

A.2 Generation of the SWS

The SWs can be generated either from their representation in real space (A.1) or in reciprocal space as:

$$\tilde{\chi}_{\alpha,lmn}(\mathbf{G}) = 4\pi i^l \bar{Y}_{lm}(\Omega_{\mathbf{G}}) e^{i\mathbf{G}\cdot\mathbf{R}_\alpha} \begin{cases} \frac{q_{nl}R_\alpha^2}{G^2 - q_{nl}^2} j_l(GR_\alpha) j_{l-1}(q_{nl}R_\alpha), & G \neq q_{nl} \\ \frac{q_{nl}R_\alpha^3}{G + q_{nl}} j_{l-1}^2(q_{nl}R_\alpha), & G = q_{nl}, \end{cases} \quad (\text{A.5})$$

where G is the modulus of the wave vector \mathbf{G} . The real-space SWs are determined with an inverse fast Fourier transform (FFT) of $\tilde{\chi}_{\alpha,lmn}(\mathbf{G})$ as:

$$\chi_{\alpha,lmn}(\mathbf{r}) = \int_0^{\mathbf{G}_{max}} \tilde{\chi}_{lmn,\alpha}(\mathbf{G}) e^{-i\mathbf{G}\cdot\mathbf{r}} d\mathbf{G}. \quad (\text{A.6})$$

In fact, what is behind this process is a decomposition of the SWs into a finite plane-wave basis set corresponding to a given kinetic energy cut-off E_c . A higher E_c will automatically enlarge the range of accessible \mathbf{G} vectors, which allows us to systematically improve the accuracy of the calculation. The grid spacing on each direction is chosen according to this energy, so a wave of energy \mathbf{G}_{max} will be correctly represented on the grid.

Initially, the SWs were generated in an FFT box [179], a box of size of 3 times the diameter of the largest localisation sphere in the simulation cell. The number of grid points that were taken into consideration was much higher than the points within the sphere, and the time to complete a SW expansion of the NGWFs was considerably long. To speed up the calculation the region to perform the inverse FFT was reduced to the *universal tightbox*. This is a box that fits tightly the largest localisation sphere in the simulation cell. This box is approximately a cube of side equal to the diameter of the sphere. The number of points in this region is near 1/27 of the points in the FFT box, so the density of \mathbf{G} vectors within the range established by the plane-wave kinetic energy cut-off was noticeably reduced. The results show that the time needed to generate one SW in the Universal Tightbox is about 1/3 of the time required with the FFT box. On the other hand, the loss of accuracy is minimal and can be regained with a more NGWF optimisation steps when the SWs representation is used to restart a calculation. These extra iterations are faster than having to generate the SWs using the FFT box. Nevertheless, the option of creating the SWs in the FFT box was made available for the cases where fine accuracy is required.

A.3 Truncation rules for the expansion of the NGWFs in terms of SWS

For a perfect, lossless description of the NGWFs in terms of the SWS, the sum in equation (A.3) has to be infinite. The issue here is to select suitable values for l_{max} and n_{max} in order to truncate this exact expansion while keeping enough accuracy of the NGWFs. This has to be done to avoid aliasing problems: the impossibility to represent a wave with smaller wave length than the Nyquist critical wave length on a computational grid [194]. The Nyquist wave length, λ_c , is related to the grid spacing, Δ , and to the plane-wave kinetic energy cut-off, E_c , as:

$$\lambda_c = 2\Delta. \quad (\text{A.7})$$

$$E_c = \frac{1}{2}|\mathbf{G}_c|^2 = \frac{1}{2} \left(\frac{2\pi}{\lambda_c} \right)^2 = \frac{1}{2} \left(\frac{\pi}{\Delta} \right)^2. \quad (\text{A.8})$$

Therefore, by increasing the plane-wave energy cut-off, both the the grid spacing and the Nyquist critical wave length are reduced, thus allowing the SWS to be represented with plane-waves of higher frequency.

The NGWFs are approximated by a linear combination of a finite set of SWS as:

$$\phi_\alpha(\mathbf{r}) \approx \sum_{l=0}^{l=l_{max}} \sum_{m=-l}^{m=l} \sum_{n=1}^{n=n_{max}} \chi_{\alpha,lmn}(\mathbf{r}) C_{\alpha,lmn} \quad (\text{A.9})$$

The SWS consist of the product of a radial function $j_l(q_{nl}r)$ times an angular part $\bar{Y}_{lm}(\Omega_{\mathbf{r}})$. The angular part is very unlikely to present problems when is represented on the grid unless the l number is extremely high. The radial part, on the other

hand, is more complicated in this sense as the spherical Bessel function, j_l shows an oscillatory behaviour whose frequency increases with the l and n numbers (see Fig. A.1). Then all truncation rules will be based on the radial part, more specifically on its two following properties:

1. The spherical Bessel functions are non-periodic, and the distance between two consecutive zeros decreases when n increases. This means that the distance between the zeros n and $(n + 1)$ will be larger than between the zeros $(n + 1)$ and $(n + 2)$, and so on. Furthermore, due to the localisation constraint, the addition of a new Bessel zero within the sphere will also affect the distance between all the rest of zeros of smaller n .
2. The distance between the same two zeros decreases when l increases. This means that for a fixed n number, the distance between the zeros labelled with n and $(n + 1)$ for a given l number is larger than the distance between the zeros n and $(n + 1)$ for $(l + 1)$.

With the current code it is possible to generate spherical Bessel functions of l number up to 8. This is very high as the vast majority of the molecules in biochemistry or nanoscience will hardly need more than $l = 3$ (f atomic orbitals). The number of SW coefficients increases with l_{max} as:

$$\text{Number of SWs coefficients} = \sum_{l=0}^{l_{max}} (2l + 1)n_{max}(l). \quad (\text{A.10})$$

The maximum number of Bessel zeros, $n_{max}(l)$, that are possible to fit within a localisation sphere without aliasing problems depends on the l number, as the distance between the same consecutive zeros is different for each of the angular momentum channels. The maximum angular momentum, l_{max} is chosen manually

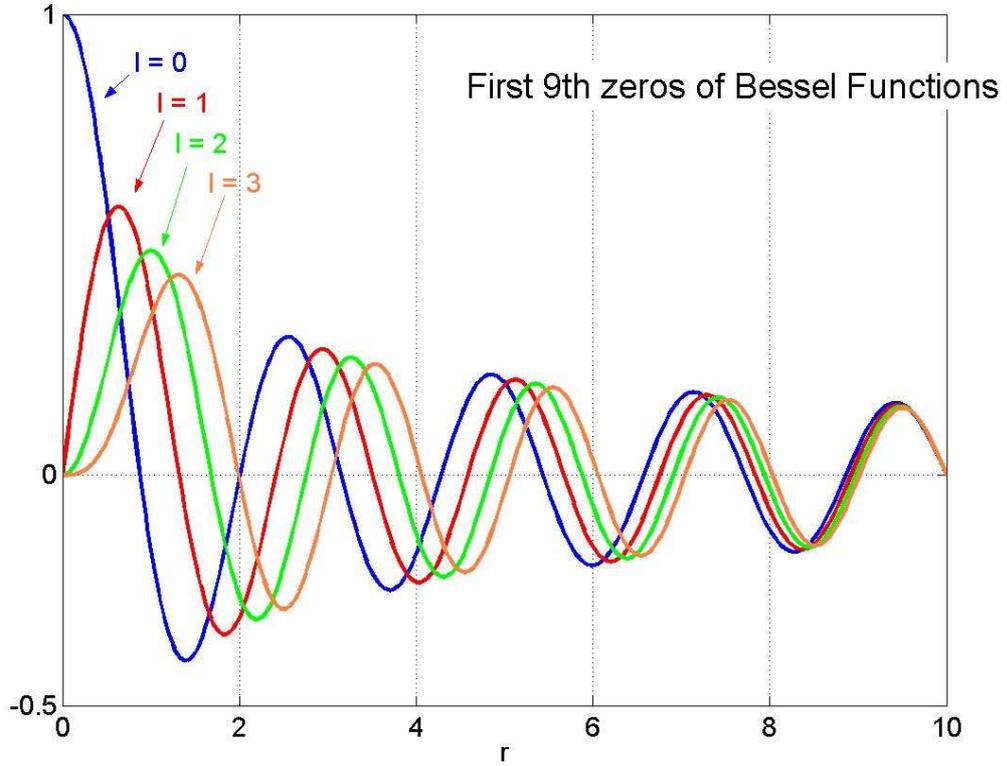


Figure A.1: Constrained spherical Bessel function

from the input file. In our implementation, $n_{max}(l)$ is calculated with the following iterative method:

1. For each l between $l = 0$ and $l = l_{max}$, compare the distance between the last two zeros of the spherical Bessel function, $n_{max}^{trial} - 1$ and n_{max}^{trial} , with the real grid spacing, Δ , for a given value of n_{max}^{trial} .
2. If this distance is larger than Δ , then the wave length of $j_l(q_{n_{max}l}r)$ is larger than the Nyquist wave length, and so there will be no aliasing problems. The algorithm then adds one unit to n_{max}^{trial} and returns to point 1 for a new iteration.
3. If that distance is shorter than Δ , then there will be aliasing problems. The

loop terminates and the value of n_{max}^{trial} at the previous iteration is accepted.

There is another important result here: the SWs are created in reciprocal space and then put in real space by doing an inverse FFT (equations (A.5) and (A.6)). The integral of the inverse FFT is limited by the cut-off energy which is related to the grid spacing. Strictly speaking, this will mean that any function that comes out as the result of performing an inverse FFT from reciprocal to real space will not suffer of aliasing problems at all. Therefore, in reciprocal space, it is possible to use a spherical Bessel function with a frequency higher than the Nyquist frequency. But it is not that easy. It is possible to show mathematically that a wave of frequency ν cannot be represented as a linear combination of a set of waves of maximum frequency smaller than ν [105]. As a result, all spherical Bessel functions with frequency above the Nyquist frequency are not correctly represented in real space after the inverse FFT. They will not have aliasing problems, but they will not be spherical Bessel functions any more, and so the SWs above the Nyquist frequency will not be exactly SWs and they will contain an error factor that will affect their properties. Because of this, those SWs above Nyquist frequency will not be orthonormal, and the expansion coefficients belonging to them will be badly affected. A expansion of this kind for the NGWFs in terms of the SWs will be:

$$\begin{aligned}
 \phi_{\alpha}(\mathbf{r}) \approx & \sum_{l=0}^{l=l_{max}} \sum_{m=-l_{max}}^{m=+l_{max}} \sum_{n=1}^{n=n_{max}} \chi_{\alpha,lmn}(\mathbf{r}) C_{\alpha,lmn} \\
 & + \sum_{l=0}^{l=l_{max}} \sum_{m=-l_{max}}^{m=+l_{max}} \sum_{n=n_{max}+1}^{n=n_{extra}} \chi_{\alpha,lmn}(\mathbf{r}) C_{\alpha,lmn}, \quad (\text{A.11})
 \end{aligned}$$

where the first summation, that goes up to the n_{max} calculated with the iterative procedure described before, contains the SWs below the Nyquist frequency, which

will be proper orthonormal SWs. On the other hand, the SWs in the second summation will not be exact SWs, so they will not be orthonormal and will introduce some error into the expansion. However, using SWs above the Nyquist frequency in reciprocal space can be beneficial in some situations as the resulting basis set will be closer to completeness.

A.4 Accuracy tests

Tests have been done for different values of l_{max} and n_{max} to check the accuracy of the SWs expansion of NGWFs. The aim is to recover the fully optimised NGWFs from the SWs representation using the minimum optimisation steps after restarting the calculation. ONETEP calculations on an ethene molecule were performed with $E_c = 400$ eV and four NGWFs per carbon atom and one hydrogen atom with $R_\alpha = 6.0$ a_0 . The number of iterations to converge after the initialisation of the NGWFs from the SWs expansion is considered as a measure of how well the final, converged NGWFs are represented in the SWs basis set. Table A.1 summarises the number of iterations required to recover the fully optimised NGWFs using each method.

We also plotted a radial section of the NGWFs in order to visualise the level of agreement between the correct NGWFs and the SW representation. The method followed in this case is to plot the values of the NGWFs at the points along the z-axis of the Universal Tightbox.

Figure A.2 shows the quality of the SWs expansion of NGWFs corresponding to an ethene molecule with different angular momenta ($l_{max} = 3, 4$ and 5), generated in reciprocal space, as well as the corresponding $l_{max} = 3$ representation in real space. It can be seen that the reciprocal space representation is much better as

Table A.1: Iterations needed to converge the NGWFs of the ethene molecule to the tolerance value.

Function	Iterations
NGWF	17
Real, $l_{max} = 3$	10
Recip, $l_{max} = 3$	8
Recip, $l_{max} = 4$	9
Recip, $l_{max} = 5$	11
Recip, $n_{extra} = 2$	10
Recip, $n_{extra} = 20$	12
Recip, $n_{extra} = 50$	45

it allows more basis set functions. Increasing l_{max} beyond 4 does not necessarily improve the accuracy of the representation.

The effect of adding extra Bessel zeros above Nyquist frequency in reciprocal space is shown in Figure A.3. Representations with $n_{extra} > 0$ can achieve good accuracy despite of the fact they contain some 'filtering' error introduced by the non-orthonormal extra basis functions. The reason for that is the larger number of basis functions and the small error that they contain as their frequency is not much higher than the frequency related to the plane-wave kinetic energy cut-off. This is the case for $n_{extra} = 2$ and $n_{extra} = 20$. Nevertheless, when the number of extra Bessel zeroes is very large, accuracy in the representation drops. This is the case of the calculations with $n_{extra} = 50$, which show deteriorating accuracy, especially in the borders of the sphere. In this case the error introduced by the extra functions exceeds the advantage of having a larger basis set, so this representation is not good at all. It is important to say that an expansion with a large number of n_{extra} needs a long time to be done. Hence, in many cases, it is more efficient to use a smaller basis set although it could result in a few more optimisation steps after restarting the calculation.

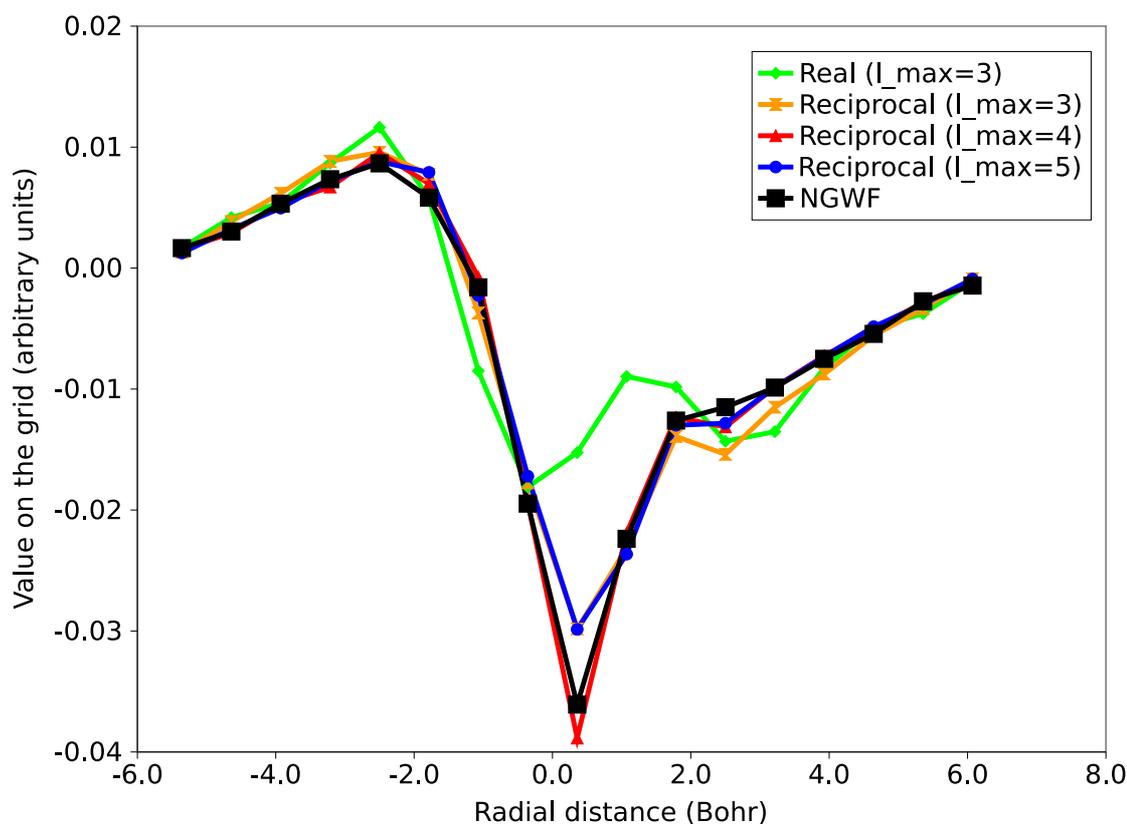


Figure A.2: NGWF for a atom of C in the ethene molecule with $l_{max} = 3, 4$ and 5 , respectively and $n_{extra} = 0$. The black line and square marks represents the fully optimised NGWF before generating the SWs expansion.

According to these results, the best expansion in terms of performance is the one with $l_{max} = 3$ and $n_{extra} = 0$, generated in the reciprocal space. This representation needs the least number of iterations to converge to the optimised NGWFs after restarting the calculation. Also, it has the advantage that the number of SWs required is much smaller than for the rest of representations, so the calculation is faster.

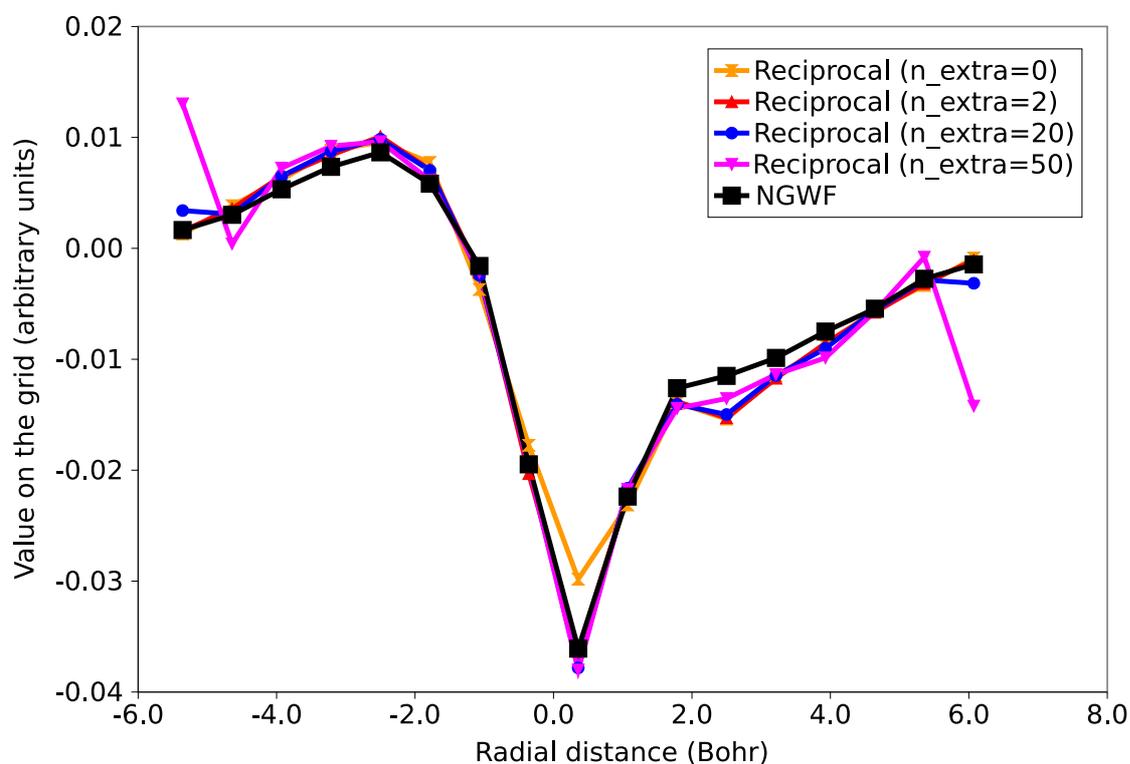


Figure A.3: NGWF for a atom of C in the ethene molecule with $l_{max} = 3$ and $n_{extra} = 2, 20$ and 50 , respectively. The black line and square marks represents the fully optimised NGWF before generating the SWs expansion.

Appendix B

A VLSD approach for QM/QM calculations with selective NGWF optimisation

In this appendix we want to lay the foundations for a new type of self-consistent calculation based on the NGWF formalism of ONETEP, but that could also be extrapolated to other approaches based on localised orbitals. We propose a method in which a selected set of NGWFs is optimised in terms of their expansion in the psinc basis set, while the rest of the NGWFs remain constant to the form they were initialised to. This approach can be seen as a “multiple-accuracy” QM/QM scheme, where a certain part of an atomic system is described at a higher level of quantum detail than the rest of the system. The scheme is fully quantum-mechanical, since every electron present in the system is associated to a single-particle molecular orbital (Kohn-Sham states in DFT approaches). The electronic interactions at the QM/QM interface are naturally taken into account by the variational optimisation of the electronic density in the entire system. This capability stands in contrast to

QM/MM methods [249], where the interactions at the interface are approximate. Large-scale calculations such as protein-ligand binding or molecules in a solvation medium can be adequately described, since a deep quantum description of the electronic structure is only necessary for a small fraction of the system.

Higher accuracy region

- Optimised density kernel
- Optimised NGWFs

QM/QM Interface

Lower accuracy region

- Optimised density kernel
- Fixed NGWFs

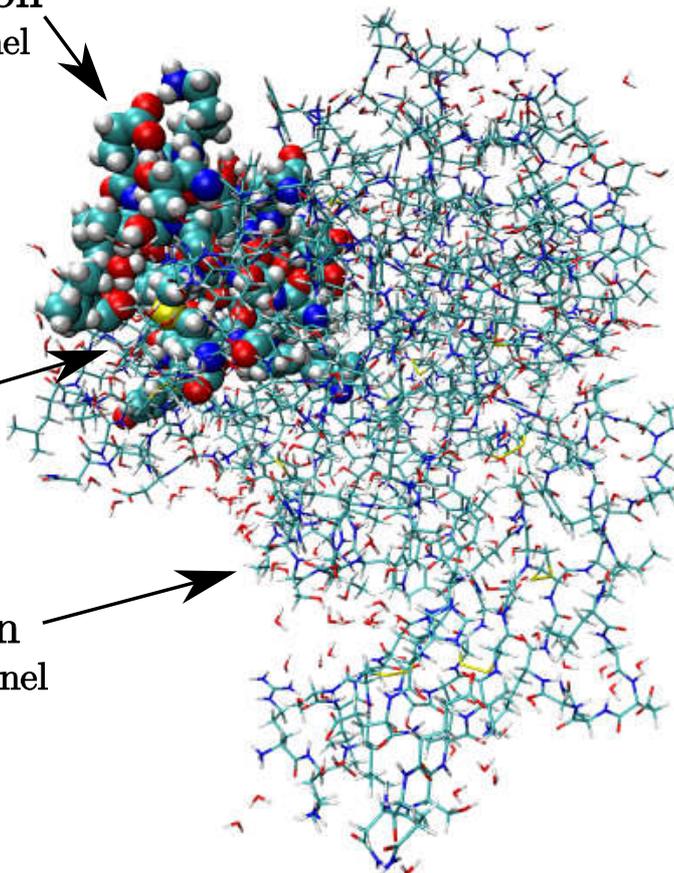


Figure B.1: Multiscale QM/QM calculation example. The top left corner of the molecule in this space representation is the high-accuracy region, where both the density kernel and the NGWFs are optimised. The rest of the molecule is the lower-accuracy region, where the NGWFs are fixed and only the density kernel is optimised

B.1 Unsuccessful attempts

A naïve attempt to the abovementioned QM/QM approach can be made simply by setting to zero the covariant search direction vector associated to the fixed NGWFs.

The search direction is previously calculated in the same way as in a standard calculation, and therefore it does not account for the new constraint that establishes that some NGWFs are fixed.

To account for this extra constraint, we can establish a VLSD approach. The set of all the covariant NGWFs, $\{\phi_{\bullet}\}$, can be split into two non-overlapping subsets as:

$$\{\phi_{\bullet}\} = \{\varphi_{\bullet}\} \cup \{\vartheta_{\bullet}\}, \quad (\text{B.1})$$

where $\{\varphi_{\bullet}\}$ is the subset that contains all the NGWFs that can be optimised in the psinc basis set, and $\{\vartheta_{\bullet}\}$ the subset that contains all the NGWFs that are constrained to be fixed. This subset separation is complete, in the sense that a given NGWF always belongs to one and only one of these subsets, but never to both at the same time. Therefore, $N_{\phi} = N_{\varphi} + N_{\vartheta}$.

We can now define three overlap matrices as:

$$A_{\alpha\beta} = \langle \varphi_{\alpha} | \varphi_{\beta} \rangle, \quad (\text{B.2})$$

$$B_{\alpha\beta} = \langle \varphi_{\alpha} | \vartheta_{\beta} \rangle, \quad (\text{B.3})$$

$$C_{\alpha\beta} = \langle \vartheta_{\alpha} | \vartheta_{\beta} \rangle, \quad (\text{B.4})$$

and prompt a VLSD optimisation where the search direction vectors of the optimised, covariant NGWFs are calculated as:

$$\Delta_{\alpha} = \Delta^{\beta} \left[A_{\beta\alpha} - B_{\beta\gamma} \left(C^{(-1)} \right)^{\gamma\delta} B_{\delta\alpha}^{\dagger} \right]. \quad (\text{B.5})$$

Figure B.2 shows the convergence of a calculation on a benzene molecule sur-

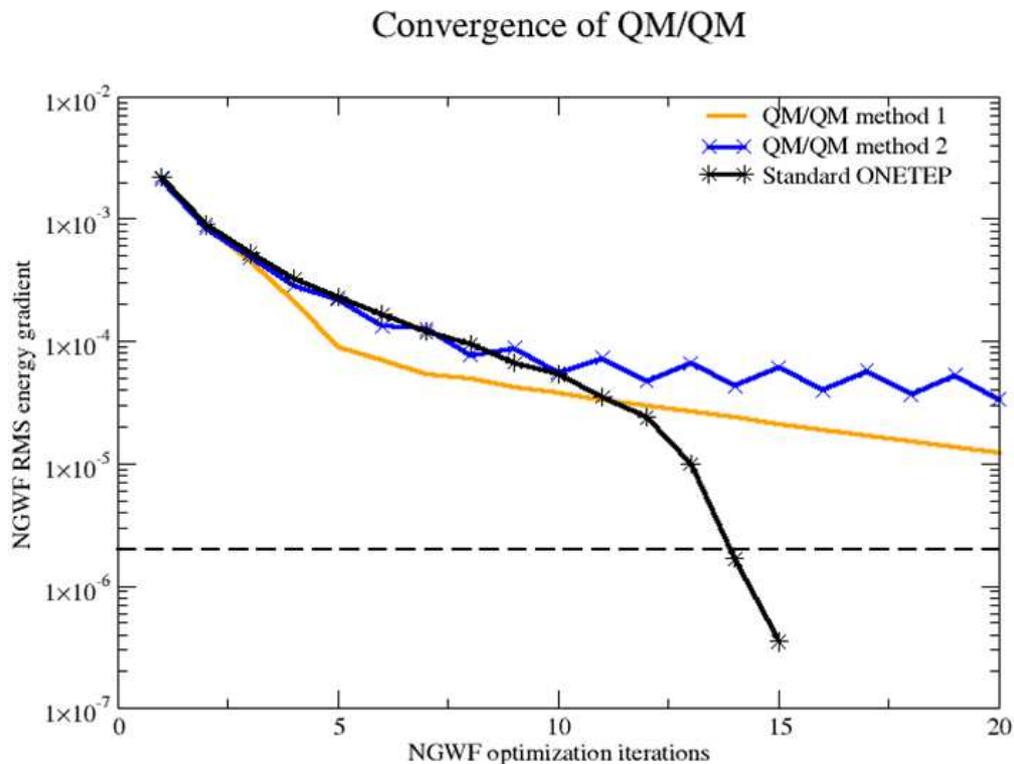


Figure B.2: QM/QM setup of benzene in solvation, with 200 water molecules in the box. The NGWFs in the benzene molecule are optimised, and the rest are fixed.

rounded by 200 water molecules. The NGWFs in the benzene are optimised, while the NGWFs in water are fixed to a SZ PAO basis set. Neither the method 1 (naïve) or method 2 (VLSD) converge very well, and neither method can be used for an efficient QM/QM approach. Most likely, the cause why these approaches failed is that the localisation constraint of the NGWFs has not been treated variationally. The two methods perform an unconstrained search, where localisation is not imposed. The naïve method is more unconstrained, so to say, since it does not consider the constraint imposed upon the fixed NGWFs either. A fully consistent and variational approach should consider localisation constraints and fixed NGWFs simultaneously.

B.2 Complete VLSD approach for QM/QM calculations

The VLSD method, originally designed for keeping strict localisation constraints on the NGWFs during a steepest-descent iterative update, can be applied to develop a multiple-accuracy QM/QM method. To understand why, one should think that the NGWF expansion coefficients corresponding to grid points that are outside the localisation spheres are effectively constant during a standard calculation with localisation constraints. In fact they are constant and *equal to zero*, but constant after all. On the other hand, a fixed NGWF expanded in the psinc basis set will yield NGWF expansion coefficients that are also constant, although not zero. To keep these psinc coefficients constant, all that is needed is a set of search direction coefficients which are constrained to be exactly zero. A robust QM/QM implementation based on variational optimisation of localised functions (like NGWFs) must take into account both the constraints of strict localisation and the new constraint imposed upon a subset of coefficients of remaining constant during the calculation. We have designed such an algorithm, but we have not yet implemented it.

The method presented in Chapter 4 only accounts for the localisation constraint. In addition to it, we can include the constraint of some NGWFs being fixed by slightly redefining the subset separation of Equations (4.47) and (4.48), which now must become:

$$\xi_p = \{\alpha \in [1, \dots, N_\phi] \mid \mathbf{r}_p \in \text{LR}(\phi_\alpha(\mathbf{r})) \text{ and } \phi_\alpha \in \{\varphi_\bullet\}\} \quad (\text{B.6})$$

$$\chi_p = \{\alpha \in [1, \dots, N_\phi] \mid \mathbf{r}_p \notin \text{LR}(\phi_\alpha(\mathbf{r})) \text{ or } \phi_\alpha \in \{\vartheta_\bullet\}\}. \quad (\text{B.7})$$

Based on this separation (point-wise, for each grid point), the subset ξ_p contains

the indices of all the covariant NGWFs for which \mathbf{r}_p is inside their covariant localisation region *and* that belong to the subset $\{\varphi_\bullet\}$ of optimisable functions. This is a much more restricted subset than the one defined in (4.47). On the other hand, the subset χ_p includes the indices of the NGWFs for which \mathbf{r}_p is outside their covariant localisation region *or* that belong to the subset $\{\vartheta_\bullet\}$ of fixed functions. Figure B.3 illustrates the subset separation with an example. Once again, it can be seen that the above separation is also complete and non-overlapping, and therefore, $N_\phi = N_\xi + N_\chi$, for each grid point with coordinate vector \mathbf{r}_p .

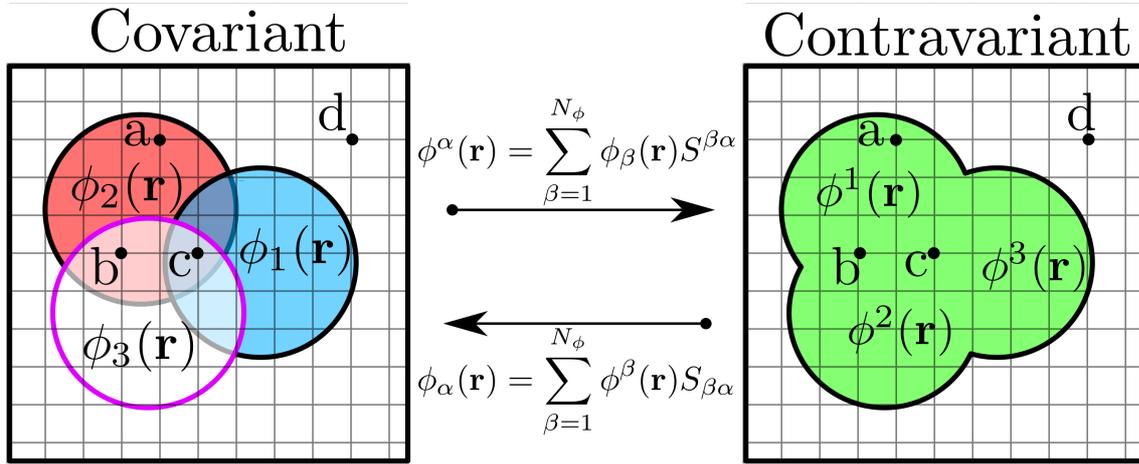


Figure B.3: Example of point-wise subset separation in a hypothetical calculation with a two-dimensional real space simulation cell, containing three covariant functions localised within circles. The covariant function ϕ_3 (magenta border and white filling) is constrained to be fixed during the calculation. The grid point \mathbf{r}_a is inside LR ($\phi_2(\mathbf{r})$), and outside LR ($\phi_1(\mathbf{r})$) and LR ($\phi_3(\mathbf{r})$). Therefore $\xi_a = \{2\}$ and $\chi_a = \{1, 3\}$. The grid point \mathbf{r}_b is inside LR ($\phi_2(\mathbf{r})$) and LR ($\phi_3(\mathbf{r})$), and outside LR ($\phi_1(\mathbf{r})$), but ϕ_3 is fixed. Therefore $\xi_b = \{2\}$ and $\chi_b = \{1, 3\}$. The grid point \mathbf{r}_c is inside all the localisation regions, but ϕ_3 is fixed. Therefore $\xi_c = \{1, 2\}$ and $\chi_c = \{3\}$. The grid point \mathbf{r}_d is outside all the localisation regions. Therefore, ξ_d is an empty subset and $\chi_d = \{1, 2, 3\}$.

The equations presented in Sec. 4.5 can now be applied, and should result in a stable algorithm for multiple-accuracy QM/QM calculations based on selective optimisation of the NGWFs. It must be noted that the number of variational degrees of freedom is reduced, since now, certain NGWFs are fixed. As a consequence,

only the derivatives with respect to the expansion coefficients of the optimisable NGWFs have to be calculated. Figure B.4 shows the QM/QM algorithm based on VLSD with localisation constraints.

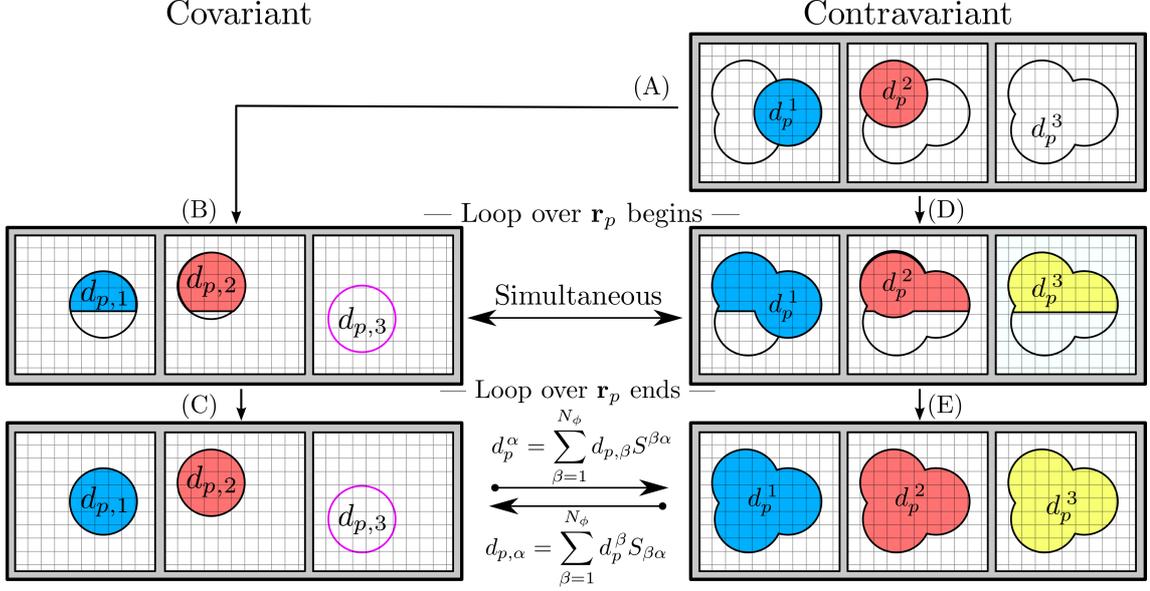


Figure B.4: Schematic representation of the VLSD method for the QM/QM scheme. In this example, there are three overlapping covariant functions, $\{\phi_\alpha\}$, localised within circles in a two-dimensional simulation cell. The covariant NGWF ϕ_3 (magenta border and white filling) is constrained to be fixed during the calculation. The algorithm begins in (A), calculating the energy gradient with respect to the variational degrees of freedom. The result are the values of $\{d_p^\alpha\}$, which correspond to the grid points confined within the circles of the optimised NGWFs only. Then, a loop over all the grid points \mathbf{r}_p starts. Steps (B) and (C) represent the progress of the algorithm at a certain intermediate stage, where some, but not all, grid points have been scrutinised. The point-wise subset separation is repeated at every grid point, as in Eq. (B.6) and Eq. (B.7). The path through (B) and (C) is the most efficient one. The covariant search directions $\{d_{p1}\}$ and $\{d_{p2}\}$ are calculated, while $\{d_{p3}\}$ becomes zero by construction. Alternatively, the path through (D) and (E) calculates the remaining contravariant search direction coefficients using Eq. (4.65). The VLSD algorithm is capable of calculating the covariant and contravariant search direction coefficients simultaneously by running steps (B) and (D) as independent processes. At the conclusion of the look over \mathbf{r}_p , the resulting search directions are directly localised within the established localisation regions and duality between $\{d_{p\alpha}\}$ and $\{d_p^\alpha\}$ holds exactly, even if $\{d_{p3}\}$ are zero.

Appendix C

VLSD method for spatial localisation of the density kernel

The density kernel matrix, \mathbf{K} , of elements $\{K^{\alpha\beta}\}$, is, by definition, the contravariant NGWF representation of the one-particle density operator, $\hat{\rho}$:

$$K^{\alpha\beta} = \langle \phi^\alpha | \hat{\rho} | \phi^\beta \rangle. \quad (\text{C.1})$$

In a calculation of N_ϕ NGWFs, there are $N_\phi \times N_\phi$ density kernel matrix elements. One can argue that \mathbf{K} is a symmetric matrix (provided that $\{\phi^\alpha\}$ are real, which they are in a standard calculation), but for the sake of generality we will not assume this condition. Therefore, there are $N_\phi \times N_\phi$ distinct pairs of NGWFs. The matrix formed by the elements $\{K^{\alpha\beta}\}$ is a second-order contravariant tensor. The covariant density kernel matrix can be found by left- and right-multiplication with the NGWF overlap matrix:

$$K_{\alpha\beta} = \sum_{\gamma=1}^{N_\phi} \sum_{\delta=1}^{N_\phi} S_{\alpha\gamma} K^{\gamma\delta} S_{\delta\beta}, \quad (\text{C.2})$$

$$K^{\alpha\beta} = \sum_{\gamma=1}^{N_\phi} \sum_{\delta=1}^{N_\phi} S^{\alpha\gamma} K_{\gamma\delta} S^{\delta\beta}. \quad (\text{C.3})$$

Let us consider now the localisation constraint imposed upon $\{K^{\alpha\beta}\}$, defined in (2.112) as:

$$K^{\alpha\beta} = 0, \text{ if } |\mathbf{R}_\alpha - \mathbf{R}_\beta| \geq R_c, \quad (\text{C.4})$$

where \mathbf{R}_α and \mathbf{R}_β are the centres of the covariant NGWFs ϕ_α and ϕ_β , respectively, and R_c is a radial cut-off. It is worth noticing the *pair-wise* nature of the expressions that relate to the density kernel elements, which is a consequence of \mathbf{K} being a second-order tensor. Let us now label the set of all possible pairs of NGWF indices as:

$$\Pi = \{(\alpha, \beta) \mid \forall \alpha, \beta \in [1, \dots, N_\phi]\}. \quad (\text{C.5})$$

One can now establish a subset separation of Π according to the localisation constraint (C.4). The subsets $\Lambda \in \Pi$ and $\Xi \in \Pi$ can be defined as:

$$\Lambda = \{(\alpha, \beta) \mid |\mathbf{R}_\alpha - \mathbf{R}_\beta| < R_c\}, \quad (\text{C.6})$$

$$\Xi = \{(\alpha, \beta) \mid |\mathbf{R}_\alpha - \mathbf{R}_\beta| \geq R_c\}. \quad (\text{C.7})$$

Once again, the two subsets Λ and Ξ do not overlap, and the separation is complete, since a given element $K^{\alpha\beta}$ always belongs to one, and only one, of these subsets. Therefore, $N_\phi \times N_\phi = N_\Lambda + N_\Xi$. This is the kind of situation that can prompt a VLSD algorithm for imposing the localisation constraint on the density kernel

matrix. Figure C.1 shows an example of this subset separation, where the set of all the matrix elements is divided into two non-overlapping subsets, defined as:

$$\Lambda = \{(1, 1), (1, 2), (2, 1), (2, 2), (3, 3), (4, 4)\}, \quad (\text{C.8})$$

$$\Xi = \{(1, 3), (1, 4), (2, 3), (2, 4), (3, 1), (3, 2), (3, 4), (4, 1), (4, 2), (4, 3)\}. \quad (\text{C.9})$$

$$\begin{pmatrix} K^{11} & K^{12} & 0 & 0 \\ K^{21} & K^{22} & 0 & 0 \\ 0 & 0 & K^{33} & 0 \\ 0 & 0 & 0 & K^{44} \end{pmatrix}$$

Figure C.1: Example of pair-wise subset separation in a hypothetical calculation with four NGWFs that give rise to a 4×4 density kernel matrix. The localisation constraints with a finite radial cut-off R_c impose some off-diagonal elements being zero. The elements over the coloured background belong to the subset Λ , while the elements over the white background belong to the subset Ξ .

Equation (C.3) can be separated into two terms, one containing covariant elements with pairs of indices in the Λ subset, and another one with the pairs of indices in the subset Ξ , as:

$$K^{\alpha\beta} = \sum_{\substack{\eta,\iota=1 \\ (\eta,\iota) \in \Lambda}} S^{\alpha\eta} K_{\eta\iota} S^{\iota\beta} + \sum_{\substack{\sigma,\tau=1 \\ (\sigma,\tau) \in \Xi}} S^{\alpha\sigma} K_{\sigma\tau} S^{\tau\beta}, \quad \forall (\alpha, \beta) \in \Lambda \quad (\text{C.10})$$

$$K^{\alpha\beta} = \sum_{\substack{\eta,\iota=1 \\ (\eta,\iota) \in \Lambda}} S^{\alpha\eta} K_{\eta\iota} S^{\iota\beta} + \sum_{\substack{\sigma,\tau=1 \\ (\sigma,\tau) \in \Xi}} S^{\alpha\sigma} K_{\sigma\tau} S^{\tau\beta}, \quad \forall (\alpha, \beta) \in \Xi. \quad (\text{C.11})$$

Equations (C.10) and (C.11) are rather complicated and require a further explanation. They rely on *combinatorics*, i.e., on the detailed study of all the possible combinations of indices such as $(\eta, \iota) \in \Lambda$ and $(\sigma, \tau) \in \Xi$, to establish the accessible values of η, ι, σ and τ that will determine the entries of the overlap and inverse overlap matrices that must be used in the summations. This is not a straightforward task, and it is obvious that the effort to complete it increases very rapidly with the number of NGWFs in the system.

The localisation constraint (C.4) is imposed exactly by equating (C.11) to zero, resulting in:

$$K_{\sigma,\tau} = - \sum_{\substack{\mu,\nu=1 \\ (\mu,\nu) \in \Xi}} \sum_{\substack{\eta,\iota=1 \\ (\eta,\iota) \in \Lambda}} S_{\sigma\mu} S^{\mu\eta} K_{\eta\iota} S^{\iota\nu} S_{\nu\tau}, \quad \forall (\sigma, \tau) \in \Xi. \quad (\text{C.12})$$

Introducing (C.12) into (C.10) we obtain:

$$K^{\alpha\beta} = \sum_{\substack{\eta,\iota=1 \\ (\eta,\iota) \in \Lambda}} \left[S^{\alpha\eta} K_{\eta\iota} S^{\iota\beta} - \sum_{\substack{\sigma,\tau=1 \\ (\sigma,\tau) \in \Xi}} \sum_{\substack{\mu,\nu=1 \\ (\mu,\nu) \in \Xi}} S^{\alpha\sigma} S_{\sigma\mu} S^{\mu\eta} K_{\eta\iota} S^{\iota\nu} S_{\nu\tau} S^{\tau\beta} \right], \quad (\text{C.13})$$

$\forall (\alpha\beta) \in \Lambda.$

Therefore, the all the contravariant density kernel elements, $\{K^{\alpha\beta}\} \in \Lambda$, which are the variational degrees of freedom of the problem, are uniquely determined by the covariant density kernel elements $\{K_{\eta\iota}\}$ such as $(\eta, \iota) \in \Lambda$. This result is a consequence of imposing strict localisation constraints and consistent tensorial formalism. Equations (C.12) and (C.13) are exact. Once the VLSD method has

been established for the density kernel, the same logic can be applied to the search direction vectors, since they are also required to be subject to the localisation constraint.

It is not obvious how a practical version of this algorithm could be implemented. Subset separation of order-two tensors implies that the resulting equations are very complex. Nevertheless, what this algorithm is telling us is that the elements corresponding to terms in Ξ do not have to be calculated, neither in the covariant or contravariant representations, in order to proceed with the energy minimisation.

Bibliography

- [1] A. K. Geim and K. S. Novoselov. The rise of graphene. *Nat. Mater.* **6**(3), 183, 2007.
- [2] Q. A. Pankhurst, J. Connolly, S. K. Jones, and J. Dobson. Applications of magnetic nanoparticles in biomedicine. *J. Phys. D: Appl. Phys.* **36**(13), R167, 2003.
- [3] W. Kohn. Electronic structure of matter – wave functions and density functionals. *Nobel Lecture*, 1999.
- [4] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne. Introducing ONETEP: Linear-scaling density functional simulations on parallel computers. *J. Chem. Phys.* **122**(8), 084119, 2005.
- [5] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum mechanics*. John Wiley & Sons, 1977.
- [6] L. D. Landau and E. M Lifshitz. *Quantum mechanics (non-relativistic theory)*. Reed, third edition, 1999.
- [7] A. Messiah. *Quantum mechanics*. Dover, 1999.
- [8] E. Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Phys. Rev.* **28**(6), 1049, 1926.

BIBLIOGRAPHY

- [9] W. Heisenberg. Über quantentheoretische Umdeutung kinematischer und mechanischer Beziehungen. *ZS. f. Phys.* **33**, 879, 1925.
- [10] M. Born and P. Jordan. Zur Quantenmechanik. *ZS. f. Phys.* **34**(1), 858, 1925.
- [11] M. Born, W. Heisenberg, and P. Jordan. Zur Quantenmechanik II. *ZS. f. Phys.* **35**(8-9), 557, 1926.
- [12] P. A. M. Dirac. *The principles of quantum mechanics*. Clarendon Press, Oxford, fourth edition, 1958.
- [13] J. Von Neumann. *Mathematische Grundlagen der Quantenmechanik*. Springer, Berlin 1932.
- [14] P. A. M. Dirac. A new notation for quantum mechanics. *Math. Proc. Cambridge Philos. Soc.* **35**(3), 416, 1939.
- [15] W. Heisenberg. Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik. *ZS. f. Phys.* **43**(3-4), 172, 1927.
- [16] H. Goldstein, C. Poole, and J. Safko. *Classical mechanics*. Pearson - Addison Wesley, third edition, 2002.
- [17] H. Hellmann. *Einführung in die Quantenchemie*. Deuticke, Leipzig, 1937.
- [18] R. P. Feynman. Forces in molecules. *Phys. Rev.* **56**(4), 340, 1939.
- [19] J. J. Thomson. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *Philos. Mag.* **7**(39), 237, 1904.

- [20] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philos. Mag.* **21**(125), 669, 1911.
- [21] N. Bohr. On the constitution of atoms and molecules I. *Philos. Mag.* **26**(151), 1, 1913.
- [22] N. Bohr. On the constitution of atoms and molecules II. *Philos. Mag.* **26**(153), 476, 1913.
- [23] N. Bohr. On the constitution of atoms and molecules III. *Philos. Mag.* **26**(155), 857, 1913.
- [24] D. R. Hartree. The wave mechanics of an atom with a non-Coulomb central field. Part I. Theory and methods. *Math. Proc. Cambridge Philos. Soc.* **24**(1), 89, 1928.
- [25] D. R. Hartree. The wave mechanics of an atom with a non-Coulomb central field. Part II. Some results and discussion. *Math. Proc. Cambridge Philos. Soc.* **24**(1), 111, 1928.
- [26] D. R. Hartree. The wave mechanics of an atom with a non-Coulomb central field. Part III. Term values and intensities in series in optical spectra. *Math. Proc. Cambridge Philos. Soc.* **24**(3), 426, 1928.
- [27] V. Fock. Näherungsmethode zur Lösung des quantenmechanischen Mehrkörperproblems. *ZS. f. Phys.* **61**(1-2), 126, 1930.
- [28] J. C. Slater. Note on Hartree's method. *Phys. Rev.* **35**(2), 210, 1930.
- [29] Chr. Møller and M. S. Plesset. Note on an approximation treatment for many-electron systems. *Phys. Rev.* **46**(7), 618, 1934.

- [30] J. Cizek. On the correlation problem in atomic and molecular systems. calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods. *J. Chem. Phys.* **45**(11), 4256, 1966.
- [31] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon. A fifth-order perturbation comparison of electron correlation theories. *Chem. Phys. Lett.* **157**(6), 479, 1989.
- [32] B. Roos. A new method for large-scale CI calculations. *Chem. Phys. Lett.* **15**(2), 153, 1972.
- [33] B. O. Roos, P. R. Taylor, and P. E. M. Siegbahn. A complete active space SCF method (CASSCF) using a density-matrix formulated super-CI approach. *Chem. Phys.* **48**(2), 157, 1980.
- [34] L. H. Thomas. The calculation of atomic fields. *Math. Proc. Cambridge Philos. Soc.* **23**(5), 542, 1927.
- [35] E. Fermi. Un metodo statistico per la determinazione di alcune proprietà dell'atomo“. *Rend. Lincei* **6**, 602, 1927.
- [36] R. G. Parr and W. Yang. *Density-functional theory of atoms and molecules*. Oxford University Press, 1989.
- [37] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.* **136**(3B), B864, 1964.
- [38] E. Runge and E. K. U. Gross. Density-functional theory for time-dependent systems. *Phys. Rev. Lett.* **52**(12), 997, 1984.
- [39] R. Car and M. Parrinello. Unified approach for molecular dynamics and density-functional theory. *Phys. Rev. Lett.* **55**(22), 2471, 1985.

- [40] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos. Iterative minimization techniques for abinitio total-energy calculations: molecular-dynamics and conjugate gradients. *Rev. Mod. Phys.* **64**(4), 1045, 1992.
- [41] F. Aryasetiawan and O. Gunnarsson. The GW method. *Rep. Prog. Phys.* **61**(3), 237, 1998.
- [42] S. Baroni, S. de Gironcoli, A. Dal Corso, and P. Giannozzi. Phonons and related crystal properties from density-functional perturbation theory. *Rev. Mod. Phys.* **73**(2), 515, 2001.
- [43] M. D. Segall, P. J. D. Lindan, M. J. Probert, C. J. Pickard, P. J. Hasnip, S. J. Clark, and M. C. Payne. First-principles simulation: ideas, illustrations and the CASTEP code. *J. Phys. Condens. Matter* **14**(11), 2717, 2002.
- [44] J. Hafner. Ab-initio simulations of materials using VASP: Density-functional theory and beyond. *J. Comput. Chem.* **29**(13), 2044, 2008.
- [45] E. Fermi. Sulla quantizzazione del gas perfetto monoatomico. *Rend. Lincei* **3**, 145, 1926.
- [46] P. A. M. Dirac. On the theory of quantum mechanics. *Proc. R. Soc. London, Ser. A* **112**(762), 661, 1926.
- [47] R. K. Pathria. *Statistical mechanics*. Pergamon Press, first edition, 1972.
- [48] W. Pauli. Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren. *ZS. f. Phys.* **31**(1), 765, 1925.

- [49] M. Born and R. Oppenheimer. Zur Quantentheorie der Molekeln. *Ann. Phys.* **389**(20), 457, 1927.
- [50] F. Jensen. *Introduction to computational chemistry*. John Wiley & Sons, second edition, 2007.
- [51] R. M. Martin. *Electronic structure: basic theory and practical methods*. Cambridge University Press, first edition, 2004.
- [52] K. Capelle. A bird's-eye view of density-functional theory. *Braz. J. Phys.* **36**(4A), 1318, 2006.
- [53] M. Levy. Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v -representability problem. *PNAS* **76**(12), 6062, 1979.
- [54] M. Levy. Electron densities in search of Hamiltonians. *Phys. Rev. A* **26**(3), 1200, 1982.
- [55] N. D. Mermin. Thermal properties of the inhomogeneous electron gas. *Phys. Rev.* **137**(5A), A1441, 1965.
- [56] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**(4A), A1133, 1965.
- [57] C. J. Cramer. *Essentials of computational chemistry: Theories and models*. John Wiley & Sons, second edition, 2006.
- [58] G. Galli. Linear scaling methods for electronic structure calculations and quantum molecular dynamics simulations. *Curr. Opin. Solid State Mater. Sci.* **1**(6), 864, 1996.

- [59] M. A. L. Marques, M. J. T. Oliveira, and T. Burnus. Libxc: A library of exchange and correlation functionals for density functional theory. *Comput. Phys. Commun.* **183**(10), 2272, 2012.
- [60] J. C. Slater. A simplification of the Hartree-Fock method. *Phys. Rev.* **81**(3), 385, 1951.
- [61] M. Gell-Mann and K. A. Brueckner. Correlation energy of an electron gas at high density. *Phys. Rev.* **106**(2), 364, 1957.
- [62] D. M. Ceperley and B. J. Alder. Ground-state of the electron-gas by a stochastic method. *Phys. Rev. Lett.* **45**(7), 566, 1980.
- [63] J. P. Perdew and A. Zunger. Self-interaction correction to density-functional approximations for many-electron systems. *Phys. Rev. B* **23**(10), 5048, 1981.
- [64] S. H. Vosko, L. Wilk, and M. Nusair. Accurate spin-dependent electron liquid correlation energies for local spin density calculations: a critical analysis. *Can. J. Phys.* **58**(8), 1200, 1980.
- [65] J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, and C. Fiolhais. Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation. *Phys. Rev. B* **46**(11), 6671, 1992.
- [66] A. D. Becke. On the large-gradient behavior of the density functional exchange energy. *J. Chem. Phys.* **85**(12), 7184, 1986.
- [67] A. D. Becke. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A* **38**(6), 3098, 1988.

- [68] C. Lee, W. Yang, and R. G. Parr. Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B* **37**(2), 785, 1988.
- [69] J. P. Perdew and Y. Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B* **45**(23), 13244, 1992.
- [70] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.* **77**(18), 3865, 1996.
- [71] B. Hammer, L. B. Hansen, and J. K. Norskov. Improved adsorption energetics within density-functional theory using revised Perdew-Burke-Ernzerhof functionals. *Phys. Rev. B* **59**(11), 7413, 1999.
- [72] J. P. Perdew, S. Kurth, A. Zupan, and P. Blaha. Accurate density functional with correct formal properties: A step beyond the generalized gradient approximation. *Phys. Rev. Lett.* **82**(12), 2544, 1999.
- [73] J. Tao, J. P. Perdew, V. N. Staroverov, and G. E. Scuseria. Climbing the density functional ladder: Nonempirical meta-generalized gradient approximation designed for molecules and solids. *Phys. Rev. Lett.* **91**(14), 146401, 2003.
- [74] A. D. Becke. Density-functional thermochemistry. III. The role of exact exchange. *J. Chem. Phys.* **98**(7), 5648, 1993.
- [75] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch. Ab-initio calculation of vibrational absorption and circular-dichroism spectra using density-functional force-fields. *J. Phys. Chem.* **98**(45), 11623, 1994.

- [76] B. Farid, N. H. March, and A. K. Theophilou. Many-body partition function and thermal Hartree-Fock approximations. *Phys. Rev. E* **62**(1), 134, 2000.
- [77] J. A. Pople, P. M. W. Gill, and B. G. Johnson. Kohn-Sham density-functional theory within a finite basis set. *Chem. Phys. Lett.* **199**(6), 557, 1992.
- [78] E. R. Davidson and D. Feller. Basis set selection for molecular calculations. *Chem. Rev.* **86**(4), 681, 1986.
- [79] T. H. Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen. *J. Chem. Phys.* **90**(2), 1007, 1989.
- [80] D. E. Woon and T. H. Dunning Jr. Gaussian basis sets for use in correlated molecular calculations. III. The atoms aluminum through argon. *J. Chem. Phys.* **98**(2), 1358, 1993.
- [81] P.R. Briddon and R. Jones. LDA calculations using a basis of Gaussian orbitals. *Phys. Status Solidi B* **217**(1), 131, 2000.
- [82] J. G. Hill. Gaussian basis sets for molecular applications. *Int. J. Quantum Chem.* **113**(1), 21, 2013.
- [83] K. L. Schuchardt, B. T. Didier, T. Elsethagen, L. Sun, V. Gurumoorthi, J. Chase, J. Li, and T. L. Windus. Basis Set Exchange: A community database for computational sciences. *J. Chem. Inf. Model.* **47**(3), 1045, 2007.
- [84] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, B. Mennucci, G. A. Petersson, H. Nakatsuji, M. Caricato, X. Li, H. P. Hratchian, A. F. Izmaylov, J. Bloino, G. Zheng, J. L. Sonnenberg, M. Hada, M. Ehara, K. Toyota, R. Fukuda,

- J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, Jr. J. A. Montgomery, J. E. Peralta, F. Ogliaro, M. Bearpark, J. J. Heyd, E. Brothers, K. N. Kudin, V. N. Staroverov, R. Kobayashi, J. Normand, K. Raghavachari, A. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, N. Rega, J. M. Millam, M. Klene, J. E. Knox, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, R. L. Martin, K. Morokuma, V. G. Zakrzewski, G. A. Voth, P. Salvador, J. J. Dannenberg, S. Dapprich, A. D. Daniels, Ö. Farkas, J. B. Foresman, J. V. Ortiz, J. Cioslowski, and D. J. Fox. Gaussian 09. Gaussian Inc. Wallingford CT, 2009.
- [85] M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. J. van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, and W. A. de Jong. Nwchem: A comprehensive and scalable open-source solution for large scale molecular simulations. *Comput. Phys. Commun.* **181**(9), 1477, 2010.
- [86] D. P. Chong, E. Van Lenthe, S. Van Gisbergen, and E. J. Baerends. Even-tempered slater-type orbitals revisited: From hydrogen to krypton. *J. Comput. Chem.* **25**(8), 1030, 2004.
- [87] G. te Velde, F. M. Bickelhaupt, E. J. Baerends, C. Fonseca Guerra, S. J. A. van Gisbergen, J. G. Snijders, and T. Ziegler. Chemistry with ADF. *J. Comput. Chem.* **22**(9), 931, 2001.
- [88] L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, and R. Schneider.

- Daubechies wavelets as a basis set for density functional pseudopotential calculations. *J. Chem. Phys.* **129**(1), 014109, 2008.
- [89] X. Gonze, B. Amadon, P. M. Anglade, J. M. Beuken, F. Bottin, P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Cote, T. Deutsch, L. Genovese, Ph. Ghosez, M. Giantomassi, S. Goedecker, D. R. Hamann, P. Hermet, F. Jollet, G. Jomard, S. Leroux, M. Mancini, S. Mazevet, M. J. T. Oliveira, G. Onida, Y. Pouillon, T. Rangel, G. M. Rignanese, D. Sangalli, R. Shaltaf, M. Torrent, M. J. Verstraete, G. Zerah, and J. W. Zwanziger. ABINIT: First-principles approach to material and nanosystem properties. *Comput. Phys. Commun.* **180**(12), 2582, 2009.
- [90] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. J. Probert, K. Refson, and M. C. Payne. First principles methods using CASTEP. *Z. Kristallogr. Kristallgeom. Kristallphys. Kristallphys.* **220**(5-6), 567, 2005.
- [91] J. Enkovaara, C. Rostgaard, J. J. Mortensen, J. Chen, M. Dułak, L. Ferrighi, J. Gavnholt, C. Glinsvad, V. Haikola, H. A. Hansen, H. H. Kristoffersen, M. Kuisma, A. H. Larsen, L. Lehtovaara, M. Ljungberg, O. Lopez-Acevedo, P. G. Moses, J. Ojanen, T. Olsen, V. Petzold, N. A. Romero, J. Stausholm-Møller, M. Strange, G. A. Tritsarlis, M. Vanin, M. Walter, B. Hammer, H. Häkkinen, G. K. H. Madsen, R. M. Nieminen, J. K. Nørskov, M. Puska, T. T. Rantala, J. Schiøtz, K. S. Thygesen, and K. W. Jacobsen. Electronic structure calculations with GPAW: a real-space implementation of the projector augmented-wave method. *J. Phys. Condens. Matter* **22**(25), 253202, 2010.
- [92] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni,
-

- D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch. Quantum espresso: a modular and open-source software project for quantum simulations of materials. *J. Phys. Condens. Matter* **21**(39), 395502, 2009.
- [93] S. Boeck, C. Freysoldt, A. Dick, L. Ismer, and J. Neugebauer. The object-oriented DFT program library S/PHI/nX. *Comput. Phys. Commun.* **182**(3), 543, 2011.
- [94] G. Kresse and J. Furthmüller. Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**(16), 11169, 1996.
- [95] G. Kresse and J. Furthmüller. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **6**(1), 15, 1996.
- [96] K. Schwarz and P. Blaha. Solid state calculations using WIEN2k. *Comput. Mater. Sci.* **28**(2), 259, 2003.
- [97] G. Lippert, J. Hutter, and M. Parrinello. A hybrid Gaussian and plane wave density functional scheme. *Mol. Phys.* **92**(3), 477, 1997.
- [98] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter. QUICKSTEP: Fast and accurate density functional calculations

- using a mixed Gaussian and plane waves approach. *Comput. Phys. Commun.* **167**(2), 103, 2005.
- [99] N. W. Ashcroft and N. D. Mermin. *Solid state physics*. Saunders College Publishing, 1976.
- [100] F. Bloch. Über die Quantenmechanik der Elektronen in Kristallgittern. *ZS. f. Phys.* **52**(7-8), 555, 1929.
- [101] A. Baldereschi. Mean-value point in the Brillouin zone. *Phys. Rev. B* **7**(12), 5212, 1973.
- [102] H. J. Monkhorst and J. D. Pack. Special points for Brillouin-zone integrations. *Phys. Rev. B* **13**(12), 5188, 1976.
- [103] J. P. Bowen, J. B. Sorensen, and K. N. Kirschner. Calculating interaction energies using first principle theories: Consideration of basis set superposition error and fragment relaxation. *J. Chem. Edu.* **84**(7), 1225, 2007.
- [104] J. Ihm, A. Zunger, and M. L. Cohen. Momentum-space formalism for the total energy of solids. *J. Phys. C: Solid State Phys.* **12**(21), 4409, 1979.
- [105] G. B. Arfken and H. J. Weber. *Mathematical methods for physicists*. Elsevier Academic Press, sixth edition, 2005.
- [106] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proc. IEEE* **93**(2), 216, 2005.
- [107] A. Zunger and M. L. Cohen. First-principles nonlocal-pseudopotential approach in the density-functional formalism. II. Application to electronic and structural properties of solids. *Phys. Rev. B* **20**(10), 4082, 1979.

- [108] D. R. Hamann, M. Schlüter, and C. Chiang. Norm-conserving pseudopotentials. *Phys. Rev. Lett.* **43**(20), 1494, 1979.
- [109] L. Kleinman and D. M. Bylander. Efficacious form for model pseudopotentials. *Phys. Rev. Lett.* **48**(20), 1425, 1982.
- [110] P. J. H. Denteneer and W. van Haeringen. The pseudopotential-density-functional method in momentum space: details and test cases. *J. Phys. C: Solid State Phys.* **18**(21), 4127, 1985.
- [111] D. R. Hamann. Generalized norm-conserving pseudopotentials. *Phys. Rev. B* **40**(5), 2980, 1989.
- [112] A. M. Rappe, K. M. Rabe, E. Kaxiras, and J. D. Joannopoulos. Optimized pseudopotentials. *Phys. Rev. B* **41**(2), 1227, 1990.
- [113] D. Vanderbilt. Soft self-consistent pseudopotentials in a generalized eigenvalue formalism. *Phys. Rev. B* **41**(11), 7892, 1990.
- [114] N. Troullier and J. L. Martins. Efficient pseudopotentials for plane-wave calculations. *Phys. Rev. B* **43**(3), 1993, 1991.
- [115] J. S. Lin, A. Qteish, M. C. Payne, and V. Heine. Optimized and transferable nonlocal separable *ab initio* pseudopotentials. *Phys. Rev. B* **47**(8), 4174, 1993.
- [116] P. E. Blöchl. Projector augmented-wave method. *Phys. Rev. B* **50**(24), 17953, 1994.
- [117] C. C. J. Roothaan. New developments in molecular orbital theory. *Rev. Mod. Phys.* **23**(2), 69, 1951.

- [118] S. Goedecker. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **71**(4), 1085, 1999.
- [119] D. R. Bowler and T. Miyazaki. $\mathcal{O}(N)$ methods in electronic structure calculations. *Rep. Prog. Phys.* **75**(3), 036503, 2012.
- [120] S. Ismail-Beigi and T. A. Arias. Locality of the density matrix in metals, semiconductors, and insulators. *Phys. Rev. Lett.* **82**(10), 2127, 1999.
- [121] S. Goedecker. Decay properties of the finite-temperature density matrix in metals. *Phys. Rev. B* **58**(7), 3501, 1998.
- [122] W. Kohn. Density functional and density matrix method scaling linearly with the number of atoms. *Phys. Rev. Letters* **76**(17), 3168, 1996.
- [123] E. Prodan and W. Kohn. Nearsightedness of electronic matter. *PNAS* **102**(33), 11635, 2005.
- [124] G. H. Wannier. The structure of electronic excitation levels in insulating crystals. *Phys. Rev.* **52**(3), 191, 1937.
- [125] W. Kohn. Analytic properties of bloch waves and wannier functions. *Phys. Rev.* **115**(4), 809, 1959.
- [126] W. Kohn. Theory of the insulating state. *Phys. Rev.* **133**(1A), A171, 1964.
- [127] J. des Cloizeaux. Energy bands and projection operators in a crystal: Analytic and asymptotic properties. *Phys. Rev.* **135**(3A), A685, 1964.
- [128] J. des Cloizeaux. Analytical properties of n -dimensional energy bands and Wannier functions. *Phys. Rev.* **135**(3A), A698, 1964.

- [129] N. Marzari and D. Vanderbilt. Maximally localized generalized Wannier functions for composite energy bands. *Phys. Rev. B* **56**(20), 12847, 1997.
- [130] W. Yang. Direct calculation of electron density in density-functional theory. *Phys. Rev. Lett.* **66**(11), 1438, 1991.
- [131] G. Galli and M. Parrinello. Large scale electronic structure calculations. *Phys. Rev. Lett.* **69**(24), 3547, 1992.
- [132] S. Baroni and P. Giannozzi. Towards very large-scale electronic-structure calculations. *Europhys. Lett.* **17**(6), 547, 1992.
- [133] E. Tsuchida and M. Tsukada. Large-scale electronic-structure calculations based on the adaptive finite-element method. *J. Phys. Soc. Jpn.* **67**(11), 3844, 1998.
- [134] D. Alfè and M. J. Gillan. Efficient localized basis set for quantum Monte Carlo calculations on condensed matter. *Phys. Rev. B* **70**(16), 161101, 2004.
- [135] J. E. Subotnik, A. D. Dutoi, and M. Head-Gordon. Fast localized orthonormal virtual orbitals which depend smoothly on nuclear coordinates. *J. Chem. Phys.* **123**(11), 114108, 2005.
- [136] J. E. Subotnik, A. Sodt, and M. Head-Gordon. A near linear-scaling smooth local coupled cluster algorithm for electronic structure. *J. Chem. Phys.* **125**(7), 074116, 2006.
- [137] R. McWeeny. Some recent advances in density matrix theory. *Rev. Mod. Phys.* **32**(2), 335, 1960.
- [138] K. Nemeth and G. E. Scuseria. Linear scaling density matrix search based on sign matrices. *J. Chem. Phys.* **113**(15), 6035, 2000.

- [139] S. Goedecker and L. Colombo. Efficient linear scaling algorithm for tight-binding molecular dynamics. *Phys. Rev. Lett.* **73**(1), 122, 1994.
- [140] W. Z. Liang, C. Saravanan, Y. H. Shao, R. Baer, A. T. Bell, and M. Head-Gordon. Improved Fermi operator expansion methods for fast electronic structure calculations. *J. Chem. Phys.* **119**(8), 4117, 2003.
- [141] T. Ozaki. O(N) Krylov-subspace method for large-scale ab initio electronic structure calculations. *Phys. Rev. B* **74**(24), 245101, 2006.
- [142] F. Mauri, G. Galli, and R. Car. Orbital formulation for electronic-structure calculations with linear system-size scaling. *Phys. Rev. B* **47**(15), 9973, 1993.
- [143] P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin. Unconstrained minimization approach for electronic computations that scales linearly with system size. *Phys. Rev. B* **48**(19), 14646, 1993.
- [144] F. Mauri and G. Galli. Electronic-structure calculations and molecular-dynamics simulations with linear system-size scaling. *Phys. Rev. B* **50**(7), 4316, 1994.
- [145] J. Kim, F. Mauri, and G. Galli. Total-energy global optimizations using nonorthogonal localized orbitals. *Phys. Rev. B* **52**(3), 1640, 1995.
- [146] P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach. Linear system-size scaling methods for electronic-structure calculations. *Phys. Rev. B* **51**(3), 1456, 1995.
- [147] E. Hernández and M. J. Gillan. Self-consistent first-principles technique with linear scaling. *Phys. Rev. B* **51**(15), 10157, 1995.

- [148] E. Hernández, M. J. Gillan, and C. M. Goringe. Linear-scaling density-functional-theory technique: The density-matrix approach. *Phys. Rev. B* **53**(11), 7147, 1996.
- [149] E. Hernández, M. J. Gillan, and C. M. Goringe. Basis functions for linear-scaling first-principles calculations. *Phys. Rev. B* **55**(20), 13485, 1997.
- [150] J.-L. Fattebert and J. Bernholc. Towards grid-based $o(n)$ density-functional theory methods: Optimized nonorthogonal orbitals and multigrid acceleration. *Phys. Rev. B* **62**(3), 1713, 2000.
- [151] C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, O. Dieguez, and M. C. Payne. Nonorthogonal generalized Wannier function pseudopotential plane-wave method. *Phys. Rev. B* **66**(3), 035119, 2002.
- [152] T. Ozaki. Variationally optimized atomic orbitals for large-scale electronic structures. *Phys. Rev. B* **67**(15), 155108, 2003.
- [153] T. Ozaki and H. Kino. Efficient projector expansion for the *ab initio* LCAO method. *Phys. Rev. B* **72**(4), 045121, 2005.
- [154] J.-L. Fattebert and F. Gygi. Linear scaling first-principles molecular dynamics with controlled accuracy. *Comput. Phys. Commun.* **162**(1), 24, 2004.
- [155] J.-L. Fattebert and F. Gygi. Linear-scaling first-principles molecular dynamics with plane-waves accuracy. *Phys. Rev. B* **73**(11), 115124, 2006.
- [156] J.-L. Fattebert. Adaptive localization regions for $O(N)$ density functional theory calculations. *J. Phys. Condens. Matter* **20**(29), 294210, 2008.
- [157] E. Tsuchida. Augmented orbital minimization method for linear scaling electronic structure calculations. *J. Phys. Soc. Jpn.* **76**(3), 034708, 2007.

- [158] E. Artacho and L. Miláns del Bosch. Nonorthogonal basis sets in quantum mechanics: Representations and second quantization. *Phys. Rev. A* **43**(11), 5770, 1991.
- [159] C. A. White, P. Maslen, M. S. Lee, and M. Head-Gordon. The tensor properties of energy gradients within a non-orthogonal basis. *Chem. Phys. Lett.* **276**(1-2), 133, 1997.
- [160] D. D. O'Regan. *Optimised Projections for the Ab Initio Simulation of Large and Strongly Correlated Systems*. PhD thesis University of Cambridge, 2011.
- [161] T. Ozaki. Efficient recursion method for inverting an overlap matrix. *Phys. Rev. B* **64**(19), 195110, 2001.
- [162] O. F. Sankey and D. J. Niklewski. Abinitio multicenter tight-binding model for molecular-dynamics simulations and other applications in covalent systems. *Phys. Rev. B* **40**(6), 3979, 1989.
- [163] M. Chen, G.-C. Guo, and L. He. Systematically improvable optimized atomic basis sets for ab initio calculations. *J. Phys. Condens. Matter* **22**(44), 445501, 2010.
- [164] J. M. Soler, E. Artacho, J. D. Gale, A. Garcia, J. Junquera, P. Ordejon, and D. Sanchez-Portal. The SIESTA method for ab initio order-N materials simulation. *J. Phys. Condens. Matter* **14**(11), 2745, 2002.
- [165] E. Artacho, D. Sánchez-Portal, P. Ordejón, A. García, and J. M. Soler. Linear-scaling ab-initio calculations for large and complex systems. *Phys. Status Solidi B* **215**(1), 809, 1999.

- [166] J. Junquera, Ó. Paz, D. Sánchez-Portal, and E. Artacho. Numerical atomic orbitals for linear-scaling calculations. *Phys. Rev. B* **64**(23), 235111, 2001.
- [167] E. Anglada, J. M. Soler, J. Junquera, and E. Artacho. Systematic generation of finite-range atomic basis sets for linear-scaling calculations. *Phys. Rev. B* **66**(20), 205101, 2002.
- [168] S. García-Gil, A. García, N. Lorente, and P. Ordejón. Optimal strictly localized basis sets for noble metal surfaces. *Phys. Rev. B* **79**(7), 075441, 2009.
- [169] A. S. Torralba, M. Todorovic, V. Brazdova, R. Choudhury, T. Miyazaki, M. J. Gillan, and D. R. Bowler. Pseudo-atomic orbitals as basis sets for the O(N) DFT code CONQUEST. *J. Phys. Condens. Matter* **20**(29), 294206, 2008.
- [170] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* **180**(11), 2175, 2009.
- [171] G. E. Scuseria. Linear Scaling Density Functional Calculations with Gaussian Orbitals. *J. Phys. Chem. A* **103**(25), 4782, 1999.
- [172] P. D. Haynes and M. C. Payne. Localised spherical-wave basis set for O(N) total-energy pseudopotential calculations. *Comput. Phys. Commun.* **102**(1–3), 17, 1997.
- [173] M. J. Gillan, D. R. Bowler, A. S. Torralba, and T. Miyazaki. Order-N first-principles calculations with the CONQUEST code. *Comput. Phys. Commun.* **177**(1-2), 14, 2007.

- [174] C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, C. J. Pickard, and M. C. Payne. Accurate kinetic energy evaluation in electronic structure calculations with localized functions on real space grids. *Comput. Phys. Commun.* **140**(3), 315, 2001.
- [175] A. A. Mostofi, C.-K. Skylaris, P. D. Haynes, and M. C. Payne. Total-energy calculations on a real space grid with localized functions and a plane-wave basis. *Comput. Phys. Commun.* **147**(3), 788, 2002.
- [176] A. A. Mostofi, P. D. Haynes, C.-K. Skylaris, and M. C. Payne. Preconditioned iterative minimization for linear-scaling electronic structure calculations. *J. Chem. Phys.* **119**(17), 8842, 2003.
- [177] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne. Using ONETEP for accurate and efficient $O(N)$ density functional calculations. *J. Phys. Condens. Matter* **17**(37), 5757, 2005.
- [178] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne. ONETEP: linear-scaling density-functional theory with local orbitals and plane waves. *Phys. Status Solidi B* **243**(11), 2489, 2006.
- [179] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne. Implementation of linear-scaling plane wave density functional theory on parallel computers. *Phys. Status Solidi B* **243**(5), 973, 2006.
- [180] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne. Elimination of basis set superposition error in linear-scaling density-functional calculations with local orbitals optimised in situ. *Chem. Phys. Lett.* **422**(4-6), 345, 2006.

- [181] C.-K. Skylaris and P. D. Haynes. Achieving plane wave accuracy in linear-scaling density functional theory applied to periodic systems: A case study on crystalline silicon. *J. Chem. Phys.* **127**(16), 164712, 2007.
- [182] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne. Density kernel optimization in the ONETEP code. *J. Phys.: Condens. Matter* **20**(29), 294207, 2008.
- [183] N. D. M. Hine, P. D. Haynes, A. A. Mostofi, C.-K. Skylaris, and M. C. Payne. Linear-scaling density-functional theory with tens of thousands of atoms: Expanding the scope and scale of calculations with ONETEP. *Comput. Phys. Commun.* **180**(7), 1041, 2009.
- [184] N. D. M. Hine, M. Robinson, P. D. Haynes, C.-K. Skylaris, M. C. Payne, and A. A. Mostofi. Accurate ionic forces and geometry optimization in linear-scaling density-functional theory with local orbitals. *Phys. Rev. B* **83**(19), 195102, 2011.
- [185] S. J. Fox, C. Pittock, T. Fox, C. S. Tautermann, N. Malcolm, and C.-K. Skylaris. Electrostatic embedding in large-scale first principles quantum mechanical calculations on biomolecules. *J. Chem. Phys.* **135**(22), 224107, 2011.
- [186] N. D. M. Hine, J. Dziedzic, P. D. Haynes, and C.-K. Skylaris. Electrostatic interactions in finite systems treated with periodic boundary conditions: Application to linear-scaling density functional theory. *J. Chem. Phys.* **135**(20), 204103, 2011.
- [187] L. E. Ratcliff, N. D. M. Hine, and P. D. Haynes. Calculating optical absorp-

- tion spectra for large systems using linear-scaling density functional theory. *Phys. Rev. B* **84**(16), 165131, 2011.
- [188] P. P. Ewald. Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Ann. Phys.* **369**(3), 253, 1921.
- [189] S. G. Louie, S. Froyen, and M. L. Cohen. Nonlinear ionic pseudopotentials in spin-density-functional calculations. *Phys. Rev. B* **26**(4), 1738, 1982.
- [190] A. Dal Corso, S. Baroni, R. Resta, and S. de Gironcoli. *Ab initio* calculation of phonon dispersions in II-VI semiconductors. *Phys. Rev. B* **47**(7), 3588, 1993.
- [191] X. P. Li, R. W. Nunes, and D. Vanderbilt. Density-matrix electronic-structure method with linear system-size scaling. *Phys. Rev. B* **47**(16), 10891, 1993.
- [192] R. W. Nunes and D. Vanderbilt. Generalization of the density-matrix method to a nonorthogonal basis. *Phys. Rev. B* **50**(23), 17611, 1994.
- [193] A. H. R. Palser and D. E. Manolopoulos. Canonical purification of the density matrix in electronic-structure theory. *Phys. Rev. B* **58**(19), 12704, 1998.
- [194] W. H. Press, W. T. Vetterling, S. A. Teulosky, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1994.
- [195] P. O. Lowdin. On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals. *J. Chem. Phys.* **18**(3), 365, 1950.
- [196] A. D. Rabuck and G. E. Scuseria. Improving self-consistent field convergence by varying occupation numbers. *J. Chem. Phys.* **110**(2), 695, 1999.

- [197] J. M. Millam and G. E. Scuseria. Linear scaling conjugate gradient density matrix search as an alternative to diagonalization for first principles electronic structure calculations. *J. Chem. Phys.* **106**(13), 5569, 1997.
- [198] P. Pulay. Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. *Mol. Phys.* **17**(2), 197, 1969.
- [199] P. Pulay. Second and third derivatives of variational energy expressions: Application to multiconfigurational self-consistent field wave functions. *J. Chem. Phys.* **78**(8), 5043, 1983.
- [200] A. M. N. Niklasson, V. Weber, and M. Challacombe. Nonorthogonal density-matrix perturbation theory. *J. Chem. Phys.* **123**(4), 044107, 2005.
- [201] M. Weinert and J. W. Davenport. Fractional occupations and density-functional energies and forces. *Phys. Rev. B* **45**(23), 13709, 1992.
- [202] R. W. Warren and B. I. Dunlap. Fractional occupation numbers and density functional energy gradients within the linear combination of Gaussian-type orbitals approach. *Chem. Phys. Lett.* **262**(3–4), 384, 1996.
- [203] A. M. N. Niklasson. A note on the Pulay force at finite electronic temperatures. *J. Chem. Phys.* **129**(24), 244107, 2008.
- [204] B. G. Pfrommer, M. Cote, S. G. Louie, and M. L. Cohen. Relaxation of crystals with the quasi-Newton method. *J. Comput. Phys.* **131**(1), 233, 1997.
- [205] M. Tafipolsky and R. Schmid. A general and efficient pseudopotential Fourier filtering scheme for real space methods using mask functions. *J. Chem. Phys.* **124**(17), 174102, 2006.

- [206] E. Artacho, E. Anglada, O. Dièguez, J. D. Gale, A. García, J. Junquera, R. M. Martin, P. Ordejón, J. M. Pruneda, D. Sánchez-Portal, and J. M. Soler. The SIESTA method; developments and applicability. *J. Phys. Condens. Matter* **20**(6), 064208, 2008.
- [207] W. Humphrey, A. Dalke, and K. Schulten. VMD: Visual molecular dynamics. *J. Mol. Graphics* **14**(1), 33, 1996.
- [208] S. Fox, H. G. Wallnoefer, T. Fox, C. S. Tautermann, and C.-K. Skylaris. First principles-based calculations of free energy of binding: Application to ligand binding in a self-assembling superstructure. *J. Chem. Theory Comput.* **7**(4), 1102, 2011.
- [209] N. Branda, R. Wyler, and J. Rebek. Encapsulation of methane and other small molecules in a self-assembling superstructure. *Science* **263**(5151), 1267, 1994.
- [210] D. A. Case, T. A. Darden, III T. E. Cheatham, C. L. Simmerling, J. Wang, R. E. Duke, R. Luo, R. C. Walker, W. Zhang, K. M. Merz, B. P. Roberts, B. Wang, S. Hayik, A. Roitberg, G. Seabra, I. Kolossváry, K. F. Wong, F. Paesani, J. Vanicek, J. Liu, X. Wu, S. R. Brozell, T. Steinbrecher, H. Gohlke, Q. Cai, X. Ye, J. Wang, M.-J. Hsieh, G. Cui, D. R. Roe, D. H. Mathews, M. G. Seetin, C. Sagui, V. Babin, T. Luchko, S. Gusarov, A. Kovalenko, and P. A. Kollman. AMBER 11 2010. University of California, San Francisco.
- [211] C. K. Gan, P. D. Haynes, and M. C. Payne. Preconditioned conjugate gradient method for the sparse generalized eigenvalue problem in electronic structure calculations. *Comput. Phys. Commun.* **134**(1), 33, 2001.

- [212] N. Marzari, D. Vanderbilt, and M. C. Payne. Ensemble density-functional theory for ab initio molecular dynamics of metals and finite-temperature insulators. *Phys. Rev. Lett.* **79**(7), 1337, 1997.
- [213] P. Pulay. Convergence acceleration of iterative sequences - the case of SCF iteration. *Chem. Phys. Lett.* **73**(2), 393, 1980.
- [214] V. Eyert. A comparative study on methods for convergence acceleration of iterative vector sequences. *J. Comput. Phys.* **124**(2), 271, 1996.
- [215] H. Sellers. The C²-DIIS convergence acceleration algorithm. *Int. J. Quantum Chem.* **45**(1), 31, 1993.
- [216] E. Cancès and C. Le Bris. Can we outperform the diis approach for electronic structure calculations? *Int. J. Quantum Chem.* **79**(2), 82, 2000.
- [217] E. Cancès. Self-consistent field algorithms for Kohn–Sham models with fractional occupation numbers. *J. Chem. Phys.* **114**(24), 10616, 2001.
- [218] K. N. Kudin, G. E. Scuseria, and E. Cancès. A black-box self-consistent field convergence algorithm: One step closer. *J. Chem. Phys.* **116**(19), 8255, 2002.
- [219] E. Cancès, K. N. Kudin, G. E. Scuseria, and G. Turinici. Quadratically convergent algorithm for fractional occupation numbers in density functional theory. *J. Chem. Phys.* **118**(12), 5364, 2003.
- [220] X. Hu and W. Yang. Accelerating self-consistent field convergence with the augmented Roothaan–Hall energy function. *J. Chem. Phys.* **132**(5), 054109, 2010.

- [221] Y. A. Zhang and Y. A. Wang. Perturbative total energy evaluation in self-consistent field iterations: Tests on molecular systems. *J. Chem. Phys.* **130**(14), 144116, 2009.
- [222] Y. A. Wang, C. Y. Yam, Y. K. Chen, and G. Chen. Linear-expansion shooting techniques for accelerating self-consistent field convergence. *J. Chem. Phys.* **134**(24), 241103, 2011.
- [223] Y. K. Chen and Y. A. Wang. LISTb: a better direct approach to LIST. *J. Chem. Theory Comput.* **7**(10), 3045, 2011.
- [224] I. V. Ionova and E. A. Carter. Error vector choice in direct inversion in the iterative subspace method. *J. Comput. Chem.* **17**(16), 1836, 1996.
- [225] G. P. Kerker. Efficient iteration scheme for self-consistent pseudopotential calculations. *Phys. Rev. B* **23**(6), 3082, 1981.
- [226] V. R. Saunders and I. H. Hillier. Level-shifting method for converging closed shell Hartree-Fock wave-functions. *Int. J. Quantum Chem.* **7**(4), 699, 1973.
- [227] Á. Ruiz-Serrano. DIIS and Hamiltonian diagonalisation for total-energy minimisation in the ONETEP program using the ScaLAPACK eigensolver. Master's thesis, University of Edinburgh, 2010.
- [228] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.
- [229] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen.

- LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [230] J. Choi, J. J. Dongarra, L. S. Ostrouchov, A. P. Petitet, D. W. Walker, and R. C. Whaley. Design and implementation of the ScaLAPACK LU, QR, and Cholesky factorization routines. *Scientific Programming* **5**(3), 173, 1996.
- [231] E. C. Anderson and J. Dongarra. Performance of LAPACK - a portable library of numerical linear algebra routines. *Proc. IEEE* **81**(8), 1094, 1993.
- [232] J. Choi, J. Demmel, I. Dhillon, J. Dongarra, S. Ostrouchov, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. ScaLAPACK: A portable linear algebra library for distributed memory computers - Design issues and performance. *Comput. Phys. Commun.* **97**(1-2), 1, 1996.
- [233] N. Zonias, P. Lagoudakis, and C.-K. Skylaris. Large-scale first principles and tight-binding density functional theory calculations on hydrogen-passivated silicon nanorods. *J. Phys.: Condens. Matter* **22**(2), 025303, 2010.
- [234] J. T. Berryman, S. E. Radford, and S. A. Harris. Thermodynamic description of polymorphism in Q- and N-rich peptide aggregates revealed by atomistic simulation. *Biophys. J.* **97**(1), 1, 2009.
- [235] T. Auckenthaler, V. Blum, H. J. Bungartz, T. Huckle, R. Johanni, L. Kraemer, B. Lang, H. Lederer, and P. R. Willems. Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations. *Parallel Comput.* **37**(12, SI), 783, 2011.
- [236] M. Methfessel and A. T. Paxton. High-precision sampling for Brillouin-zone integration in metals. *Phys. Rev. B* **40**(6), 3616, 1989.

- [237] M. J. Gillan. Calculation of the vacancy formation energy in aluminium. *J. Phys. Condens. Matter* **1**(4), 689, 1989.
- [238] M. P. Grumbach, D. Hohl, R. M. Martin, and R. Car. Ab initio molecular dynamics with a finite-temperature density functional. *J. Phys. Condens. Matter* **6**(10), 1999, 1994.
- [239] C. Freysoldt, S. Boeck, and J. Neugebauer. Direct minimization technique for metals in density functional theory. *Phys. Rev. B* **79**(24), 241103, 2009.
- [240] J. Kleis, J. Greeley, N. A. Romero, V. A. Morozov, H. Falsig, A. H. Larsen, J. Lu, J. J. Mortensen, M. Dułak, K. S. Thygesen, J. K. Nørskov, and K. W. Jacobsen. Finite size effects in chemical bonding: From small clusters to solids. *Catal. Lett.* **141**(8), 1067, 2011.
- [241] L. Li, A. H. Larsen, N. A. Romero, V. A. Morozov, C. Glinsvad, F. Abild-Pedersen, J. Greeley, K. W. Jacobsen, and J. K. Nørskov. Investigation of catalytic finite-size-effects of platinum metal clusters. *J. Phys. Chem. Lett.* **4**(1), 222, 2013.
- [242] F. Baletto and R. Ferrando. Structural properties of nanoclusters: Energetic, thermodynamic, and kinetic effects. *Rev. Mod. Phys.* **77**(1), 371, 2005.
- [243] B. G. Walker, N. Marzari, and C. Molteni. Layering at liquid metal surfaces: Friedel oscillations and confinement effects. *J. Phys. Condens. Matter* **18**(19), L269, 2006.
- [244] R. Kavanagh, X.-M. Cao, W.-F. Lin, C. Hardacre, and P. Hu. Origin of low CO₂ selectivity on platinum in the direct ethanol fuel cell. *Angew. Chem. Int. Ed.* **51**(7), 1572, 2012.

- [245] Z. P. Liu and P. Hu. General rules for predicting where a catalytic reaction should occur on metal surfaces: A density functional theory study of C-H and C-O bond breaking/making on flat, stepped, and kinked metal surfaces. *JACS* **125**(7), 1958, 2003.
- [246] R. Ferrando, J. Jellinek, and R. L. Johnston. Nanoalloys: From theory to applications of alloy clusters and nanoparticles. *Chem. Rev.* **108**(3), 845, 2008.
- [247] J. Oviedo and R. E. Palmer. Amorphous structures of Cu, Ag, and Au nanoclusters from first principles calculations. *J. Chem. Phys.* **117**(21), 9548, 2002.
- [248] F. Birch. Finite elastic strain of cubic crystals. *Phys. Rev.* **71**(11), 809, 1947.
- [249] D. Bakowies and W. Thiel. Hybrid models for combined quantum mechanical and molecular mechanical approaches. *J. Chem. Phys.* **100**(25), 10580, 1996.