LUND UNIVERSITY

Towards a multiconfigurational description of the electronic structure in solids

Larsson, Ernst Dennis

2022

Towards a multiconfigurational description of the electronic structure in solids

ERNST DENNIS LARSSON | DIVISION OF THEORETICAL CHEMISTRY | LUND UNIVERSITY

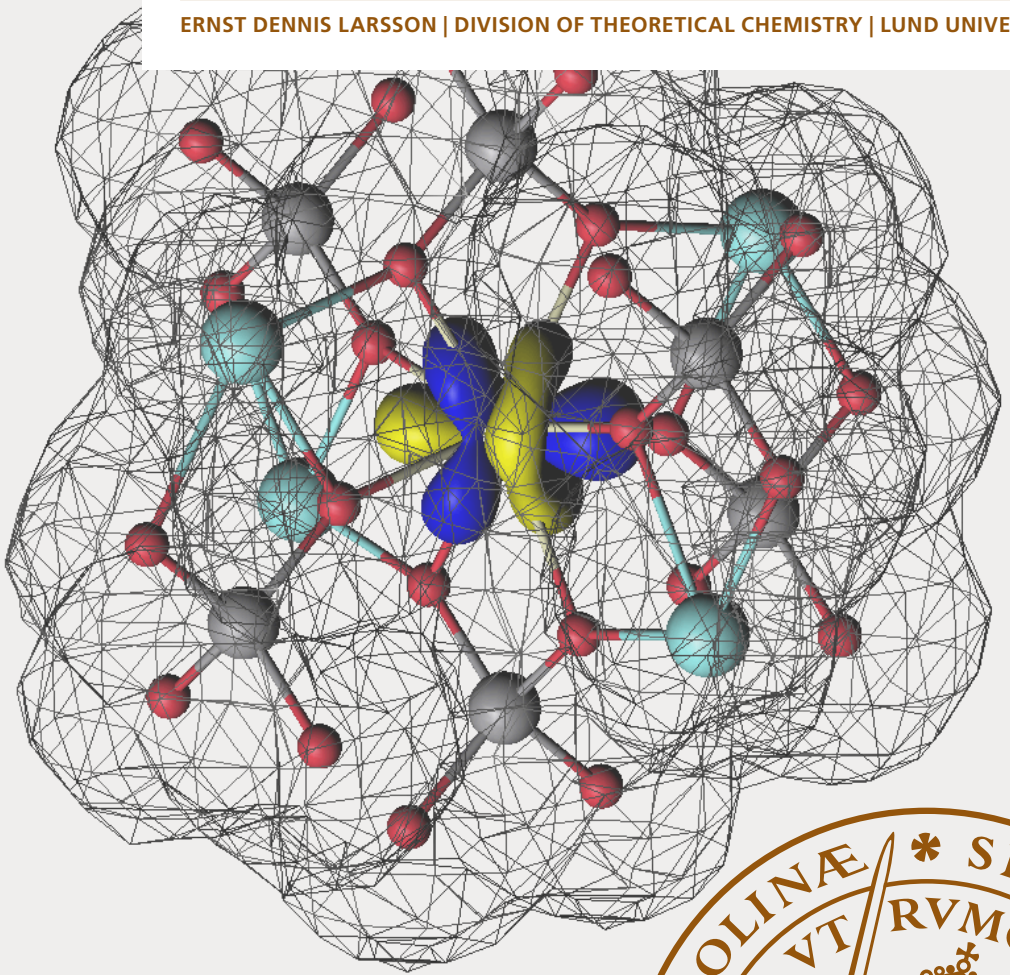

Towards a multiconfigurational description of the electronic structure in solids

Towards a multiconfigurational description of the electronic structure in solids

by Ernst Dennis Larsson

LUND
UNIVERSITY

Thesis for the degree of Doctor of Philosophy
Thesis advisors: Associate Prof. Valera Veryazov, Prof. Marie Skepö
Faculty opponent: Prof. Lev Kantorovich

To be presented, with the permission of the Division of Theoretical Chemistry of Lund University, for public criticism at Chemical Centre (Room A, floor 0, Kemicentrum) on Friday, the 9th of December 2022 at 13:00.

Organization
LUND UNIVERSITY

Department of Chemistry
Box 124
SE–221 00 LUND
Sweden

Document name
DOCTORAL THESIS

Date of disputation
2022-12-09

Sponsoring organization

Author(s)
Ernst Dennis Larsson

Title and subtitle
Towards a multiconfigurational description of the electronic structure in solids

Abstract
Materials of ionic crystals are ubiquitous in industrial chemistry. For example, materials such as cerium dioixde (CeO_2) are used in both self-cleaning ovens and to clean exhaust fumes from cars. Other materials, such as titanium dioixde (TiO_2) has been used in the solar-cell industry. So-called garnets are used in several lasers.

In common for all of these areas of application, is that they are dependant on the motion of the electrons in these materials. In order to understand how electrons behave and interact, quantum mechanics is required. A major problem that immediately arises when applying quantum mechanics to crystalline materials, is that crystals are, from a quantum mechanical perspective, enormous. One single crystal can contain as many as Avogradro's number of atoms (10^{23}). Quantum mechanical calculation are very demanding, with even the most approximate methods available today being limited to around 10 000 atoms. The type of methods used in this thesis, generally known as wavefunction theory, are roughly limited to around 100 atoms, depending a bit on what part of the periodic table that is explored and what type of property that is studied.

Methods that fall within wavefunction theory have the advantage against more approximate methods that they follow a fairly strict ladder of increasing accuracy. In other words, the predicted results can, in principle, be improved by choosing methods from higher up on the ladder. Of course, the higher up on the ladder a method is, the more computationally expensive it is. It is therefore not necessarily affordable to move enough steps on the ladder, such that the desired accuracy can be reached. For that reason, there needs to be some form om compromise when modelling crystals – in order to improve the description of the electronic structure, the atomic structure has to become more approximate. Models of that kind are usually referred to as embedding methods.

The purpose of this thesis has been to develop an embedding method for crystalline ionic materials. This was achieved by developing a computer code called SCEPIC, that generates so-called ab-inito model potentials. As a part of this thesis work, this method was evaluated in order to provide guidance to other researchers on how to best apply this method.

Key words
quantum chemistry, multiconfigurational, ionic crystals, embedding

Classification system and/or index terms (if any)

Supplementary bibliographical information

Language
English

ISSN and key title

ISBN
978-91-8039-463-5 (print)
978-91-8039-464-2 (pdf)

Recipient's notes

Number of pages
218

Price

Security classification

Signature _____ Date 2022-10-31

Towards a multiconfigurational description of the electronic structure in solids

by Ernst Dennis Larsson

LUND
UNIVERSITY

A doctoral thesis at a university in Sweden takes either the form of a single, cohesive research study (monograph) or a summary of research papers (compilation thesis), which the doctoral student has written alone or together with one or several other author(s).

In the latter case the thesis consists of two parts. An introductory text puts the research work into context and summarizes the main points of the papers. Then, the research publications themselves are reproduced, together with a description of the individual contributions of the authors. The research papers may either have been already published or are manuscripts at various stages (in press, submitted, or in draft).

Cover illustration front: Figure by V. Veryazov.

Funding information: The thesis work was financially supported by MOLCAS and the Division of Theoretical Chemistry, Lund University.

Division of Theoretical Chemistry, Department of Chemistry

ISBN: 978-91-8039-463-5 (print)
ISBN: 978-91-8039-464-2 (pdf)

Printed in Sweden by Media-Tryck, Lund University, Lund 2022

Media-Tryck is a Nordic Swan Ecolabel certified provider of printed material. Read more about our environmental work at www.mediatryck.lu.se

Printed matter 3041 0903

MADE IN SWEDEN

Dedicated to
my sister, Erica
without whom this thesis would never have been written

Contents

List of publications

This thesis is based on the following publications, referred to by their Roman numerals:

I **A program system for self-consistent embedded potentials for ionic crystals**

E. D. Larsson, M. Krośnicki, V. Veryazov
Chemical Physics, 2022, 562, art. 111549

II **Convergence of Electronic Structure Properties in Ionic Oxides Within a Fragment Approach**
E. D. Larsson, J. P. Zobel, V. Veryazov
Electronic Structure, 2022, 4, art. 014009

III **Benchmarking ANO-R basis set for multiconfigurational calculations**

E. D. Larsson, J. P. Zobel, V. Veryazov
Electronic Structure, 2022, 4, art. 014009

IV **Is density functional theory accurate for lytic polysaccharide monooxygenase enzymes?**
E. D. Larsson, G. Dong, V. Veryazov, U. Ryde, E. D. Hedegård
Dalton Transactions, 2020, 49, pp. 1501–1512

V **Modern quantum chemistry with [Open]Molcas**

F. Aquilante *et al.*
Journal of Chemical Physics, 2020, 152, art. 214117

Acknowledgements

First of all I would like to thank my supervisor Valera Veryazov, both for providing me with the opportunity to do this PhD work, as well as for many good scientific discussions during the years and support in writing this thesis. I would also like to thank Erik Hedegård, Ulf Ryde, Per-Olof Widmark and Per-Åke Malmqvist for discussions on various scientific topics. A thank you to Mickaël Delcey as well, for providing some useful comments on parts of this thesis.

A big thank you to all my fellow PhD students for many fika-room discussions on all kinds of random topics.

I would also like to thank Karin, who supported me during the majority of this work. And to Pirjo and Rolf, for being truly wonderful people.

Finally, I would like to thank my parents for all their support, ever since I started my studies in chemistry. And last, but by no means least, I would like to thank my sister, Erica, who has provided me with invaluable support over the last few months.

Abbreviations

ONV occupation-number vector

SD Slater determinant

CSF configuration state functions

BO Born-Oppenheimer

GFT group-function theory

MO molecular orbitals

EA electron affinity

LPMO lytic polysaccharide monooxygenase

PES potential energy surface

TS transition state

SOC spin-orbit coupling

IP ionisation potentials

HF Hartree-Fock

RHF restricted HF

UHF unrestricted HF

SCF self-consistent field

MCSCF multiconfigurational self-consistent field

CI configuration interaction

CC coupled cluster

MRCI multireference configuration interaction

CASSCF complete-active-space self-consistent field

DMRG density matrix renormalisation group

SAV-CASSCF state-average CASSCF

MP Møller-Plesset perturbation theory

CASPT complete-active-space perturbation theory

NEVPT n-electron valence perturbation theory

XMS-CASPT2 extended multistate CASPT2

RAS restricted-active-space

GAS generalised-active-space

DMRG density matrix renormalisation group

PBC periodic boundary conditions

DFT density functional theory

KSDFT Kohn-Sham DFT

WFT wavefunction theory

QC Quantum Chemistry

DFET density-functional embedding theory

DMET density matrix embedding theory

GPEE general purpose electrostatic embedding

MCPDFT multiconfigurational pair-density functional theory

MC-srDFT multiconfigurational short-range DFT

DKH Dougles-Kroll-Hess

X2C exact two-component

SCEI self-consistent embedded ion

PBE Perdew-Burke-Ernzerhof

GTOs Gaussian-type orbitals

pGTOs primitive Gaussian-type orbitals

cGTOs contracted Gaussian-type orbitals

PW plane-waves

ANO Atomic Natural Orbital

AIMP ab-initio model potentials

PPs pseudopotentials

PAWs projector-augmented waves

ECP effective core potentials

Populärvetenskaplig sammanfattning på svenska

Kristallina joniska material är en absolut nödvändighet i industriell kemi. Deras användningsområden är många. Exempelvis, så används materialet ceriumdioxid (CeO_2) både i självrengörande ungnar och i bilars avgasrör för att rengöra avgaser. Material som titandioxid (TiO_2) har sett användningsområden inom solcellsindustrin. Så kallade granater (ej att förväxla med det explosiva vapnet) används i flera lasrar.

Gemensamt för samtliga användningsområden, är att de är nära sammanflätade med hur elektronerna beter sig i dessa material. För att få en korrekt förståelse för hur elektroner rör sig och interagerar i ett material, så krävs kvantmekanik. Ett problem som direkt uppstår, när man vill tillämpa kvantmekanik på kristallina material, är att sådana material är, ur ett kvantmekaniskt perspektiv, enorma. En enskild kristall, kan innehålla så mycket som uppmot Avogradros tal med atomer (10^{23}). Kvantmekaniska beräkningar är väldigt krävande, även de mest approximativa metoderna som finns tillgängliga idag är begränsade till runt 10 000 atomer. För den typ av metoder som som har använts i denna avhandling, som generellt sett kallas för vågfunktionsteori, kan man i runda slängar säga att begränsningen ligger någonstans runt 100 atomer, beroende på var i det periodiska systemet man befinner sig och vad för egenskaper som undersöks.

Metoder inom vågfunktionsteorin har den attraktiva fördelen gentemont mer approximativa metoder, att de kan sägas följa en strikt noggrannhetsordning. Med andra ord, så kan man alltid förbättra sina resultat, genom att kliva upp ett steg längs noggrannhetsordningen. Självfallet följer det, att desto högre upp längs noggrannhetsordningen en metod befinner sig, desto dyrare är den. Därmed är det inte alltid givet, att man har råd att kliva upp tillräckligt många steg, för att nå önskad noggrannhet. Av den orsaken så måste man introducera någon form av kompromiss i sin beräkningmodell för kristaller – för att beskriva elektronstrukturen mer noggrannt, så krävs det att mer approximativa modeller används för att beskriva atomstrukteren. Modeller av detta slag kallas antingen, något ålderdomligt, för en inbäddning på svenska, eller för en embedding, från det motsvarande engelska ordet.

Syftet med den här avhandlingen har varit att utvecka en inbäddningsmodell för kristallina joniska material. Detta realiserades genom att utveckla en datorkod, kallad SCEPIC, som kan generera så kallade ab-initio modellpotentialer. En del av arbetet som ingått i den här avhandling har även syftat till att utvärdera denna metod och ge förslag till andra forskare på hur man bäst tillämpar den.

Popular scientific summary

Materials of ionic crystals are ubiquitous in industrial chemistry. For example, materials such as cerium dioixde (CeO_2) are used in both self-cleaning ovens and to clean exhaust fumes from cars. Other materials, such as titanium dioixde (TiO_2) has been used in the solar-cell industry. So-called garnets are used in several lasers.

In common for all of these areas of application, is that they are dependant on the motion of the electrons in these materials. In order to understand how electrons behave and interact, quantum mechanics is required. A major problem that immediately arises when applying quantum mechanics to crystalline materials, is that crystals are, from a quantum mechanical perspective, enormous. One single crystal can contain as many as Avogradro's number of atoms (10^{23}). Quantum mechanical calculation are very demanding, with even the most approximate methods available today being limited to around 10 000 atoms. The type of methods used in this thesis, generally known as wavefunction theory, are roughly limited to around 100 atoms, depending a bit on what part of the periodic table that is explored and what type of property that is studied.

Methods that fall within wavefunction theory have the advantage against more approximate methods that they follow a fairly strict ladder of increasing accuracy. In other words, the predicted results can, in principle, be improved by choosing methods from higher up on the ladder. Of course, the higher up on the ladder a method is, the more computationally expensive it is. It is therefore not necessarily affordable to move enough steps on the ladder, such that the desired accuracy can be reached. For that reason, there needs to be some form om compromise when modelling crystals – in order to improve the description of the electronic structure, the atomic structure has to become more approximate. Models of that kind are usually referred to as embedding methods.

The purpose of this thesis has been to develop an embedding method for crystalline ionic materials. This was achieved by developing a computer code called SCEPIC, that generates so-called ab-inito model potentials. As a part of this thesis work, this method was evaluated in order to provide guidance to other researchers on how to best apply this method.

Towards a multiconfigurational description of the electronic structure in solids

Chapter 1

Electronic structure theory – an overview

Chemistry is, arguably, the physics of the electrons. Most chemical phenomena are directly related to the motion of electrons in molecules or solids. For instance, in spectroscopy, we are interested in understanding the interaction of electrons with light. Chemical reactions are driven by the rearrangement of electrons. The proper description of electrons requires quantum mechanics, which lead to the development of Quantum Chemistry (QC). The goal of QC is then to describe the electronic structure in molecules and solids as accurately as possible.

Electronic structure calculations are many-body problems, such problems are difficult to solve in an accurate fashion. For this reason, many different electronic structure methods have been developed – all of them with various strengths and weaknesses. As a rule-of-thumb, the more accurate a method is, the harder it is to control and more expensive it is to compute.

This chapter will provide an overview of some of the most common electronic structure methods in use. I will try to make connections to the ultimate goal of this thesis, namely, the connection between the most powerful methods available in QC with solid materials.

1.1 Hartree-Fock theory

Before starting this chapter with a discussion on Hartree-Fock (HF) theory, some basic concepts from second quantisation will be introduced. This is due to the fact

that the language of second quantisation gives highly compact expressions for post HF methods and one can find it desirable to express HF theory in the same language as, for example, the configuration interaction (CI) method. Since the purpose of this text is not to be a complete overview of second quantisation, the interested reader is instead deferred to Refs 1,2 for more comprehensive discussions on the topic.

In second quantisation, an arbitrary wave function composed of a set of M spin-orbitals is represented by an occupation-number vector (ONV)

$$|\mathbf{k}\rangle = |k_1, k_2, \ldots, k_M\rangle \tag{1.1}$$

where k_P is one if the spin-orbital is occupied by an electron, and zero if it is not. In conventional QC, each ONV is a one-to-one map to a so-called Slater determinant (SD) [3], which additionally contains information about the spatial structure used to represent the wavefunction [1], i.e., the basis set, which will be discussed further in Chapter 2. For a set of orthonormal spin-orbitals (ψ_k), ONV's fulfil the orthonormality condition,

$$\langle \mathbf{k}|\mathbf{m}\rangle = \delta_{\mathbf{k},\mathbf{m}} \tag{1.2}$$

i.e., the overlap between two non-identical ONV's is strictly zero. Before continuing with second quantisation, I will just mention that an SD for M electrons is constructed from the spin-orbitals as

$$\Psi = \frac{1}{\sqrt{M!}} \begin{vmatrix} \psi_1(\mathbf{r}_1) & \psi_2(\mathbf{r}_1) & \ldots & \psi_M(\mathbf{r}_1) \\ \psi_1(\mathbf{r}_2) & \psi_2(\mathbf{r}_2) & \ldots & \psi_M(\mathbf{r}_2) \\ \vdots & \vdots & \vdots & \vdots \\ \psi_1(\mathbf{r}_M) & \psi_2(\mathbf{r}_M) & \ldots & \psi_M(\mathbf{r}_M) \end{vmatrix} \tag{1.3}$$

where $\mathbf{r}_i$ corresponds to the coordinates for electron i. In computer codes, the equations are generally implemented in terms of the SD's or spin-adapted linear combinations of SD's (so-called configuration state functions (CSF)) [1,2]. From the point-of-view of writing equations, however, second quantisation is more convenient, for which reason most theoretical work in QC is generally done using this formulation.

The power of second quantisation in wave function theory primarily comes from the use of creation and annihilation operators, which give compact representations of the complicated correlated wave function methods discussed in Sections 1.2.2 and 1.3. A creation operator, $\hat{a}_p^\dagger$, acts on an unoccupied orbital by placing an electron into it. In contrast, an annihilation operator, $\hat{a}_q$, acts on an occupied orbital by removing an electron from it. An excitation operator from spin-orbital q to spin-orbital p can therefore be represented by $\hat{E}_{pq} = \hat{a}_p^\dagger \hat{a}_q$.

For a molecular system, the Hamiltonian operator, $\hat{H}$, is given by

$$\hat{H} = \sum_{pq} h_{pq} \hat{a}_p^\dagger \hat{a}_q + \frac{1}{2} \sum_{pq,rs} V_{pq,rs} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s + V_{\text{nuc}} \tag{1.4}$$

where h_{pq} represent one-electron contributions to the energy and $V_{pq,rs}$ represent the Coulombic two-electron interactions. V_{nuc} represents the Coulombic interactions between the nuclei, which in the clamped-nuclei, or Born-Oppenheimer (BO), approximation, only contains a classical repulsive term and is a scalar for a given set of nuclear coordinates. In the equation above and the following text, indices p, q, r, s refer to arbitrary orbitals, while i, j refer to occupied orbitals.

In HF theory[4,5], the molecular state is represented as a single ONV, which will be denoted here as $|\Psi^{\text{HF}}\rangle$. The total energy of an HF state, E^{HF}, is obtained in the usual way in quantum mechanics, i.e., $E^{\text{HF}} = \langle \Psi^{\text{HF}} | \hat{H} | \Psi^{\text{HF}} \rangle$. The nuclear repulsion energy is, as stated above, a simple scalar in conventional calculations. The remaining part of the energy can be decomposed into two different parts, i) the one-electron energy E_1^{HF} and ii) the two-electron energy E_2^{HF}. For the one-electron energy in HF theory

$$E_1^{\text{HF}} = \sum_{pq} h_{pq} \langle \Psi^{\text{HF}} | \hat{a}_p^\dagger \hat{a}_q | \Psi^{\text{HF}} \rangle \tag{1.5}$$

which, given the orthonormality condition in Eq. 1.2, results in that only combinations where $p = q$ are non-zero. This implies that the only orbitals that contribute to the one-electron energy are the occupied orbitals, thus

$$E_1^{\text{HF}} = \sum_{i} h_{ii} \tag{1.6}$$

Similarly, when resolving the two-electron contribution to the total energy,

$$E_2^{\text{HF}} = \frac{1}{2} \sum_{pq,rs} V_{pq,rs} \langle \Psi^{\text{HF}} | \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s | \Psi^{\text{HF}} \rangle \tag{1.7}$$

the only possible solutions are that if $r = p$, then $s = q$; alternatively, if $r = q$, then $s = p$. Thus only occupied orbitals contribute to the two-electron energy and a more compact expression can be written

$$E_2^{\text{HF}} = \frac{1}{2} \sum_{i \neq j} (V_{ij,ij} - V_{ij,ji}) \tag{1.8}$$

where the first term is the classical Coulombic repulsion between two electrons and the second the exchange interaction. The (possibly) unexpected negative sign between the two terms is a consequence of commutation rules in second quantisation.

Since neither E_1^{HF} nor E_2^{HF} include any contributions from the unoccupied orbitals, HF calculations are invariant to their shape. This has major implications for the design and evaluation of basis sets used to describe the spatial part of the wavefunction (Chapter 2). One downside to this approach is that we do not, in general, know the wavefunction in advance. Instead, we have to construct a first guess, that is iterated (by so-called orbital rotations) until a more satisfactory solution can be found[4,5]. Such iterations are usually called self-consistent field (SCF) iterations, for which reason both abbreviations HF and SCF are often used interchangeably for Hartree-Fock theory.

Today, the two most common ways of using HF is in the form of restricted HF (RHF) for closed-shell systems and unrestricted HF (UHF) for open-shell systems. The main difference is that in the case of RHF, the spatial part of both α and β electrons are identical, whereas in UHF these are optimised independently of each other. While UHF makes it possible to mimic open-shell states with HF theory, it is worth noting here that a single SD is, in general, not a true eigenfunction of the spin operator, which can cause problems when modelling open-shell systems, such as most metals. A closed-shell determinant is always a true spin eigenfunction, therefore the wavefunction in RHF always has proper spin-symmetry. True spin eigenfunctions can be obtained by, as stated above, making appropriate linear combinations of SD's into CSF, which is rarely done for HF theory but a very common practice for multiconfigurational self-consistent field (MCSCF) theory, which will be discussed more in Section 1.3.

As implied by the two-electron part in Eq. 1.4, the scaling of a properly implemented HF code should, for a system containing N basis functions, be maximum $O(N^4)$, which is affordable for most molecular systems on most standard computers today. Unfortunately, the restriction that only occupied orbitals contribute to the energy puts a too severe a constraint for many properties of chemical interest. The classical example for this is bond-dissociation, usually exemplified by the H_2 molecule, though other important cases relate to, for instance UV-VIS spectroscopy or conical intersections, where transitions between electronic states are of interest[1,2].

1.2 Post Hartree-Fock methods

Given that HF is insufficient to describe all of chemical space accurately, more powerful methods are required. Since the wave function obtained from an HF calculation is, for many conventional chemical systems fairly accurate, most methods in so-called wavefunction theory (WFT) generally use HF as a starting point and devising methods for adding the missing part of the exact energy, the so-called correlation energy,

E^{corr},

$$E^{\text{exact}} = E^{\text{HF}} + E^{\text{corr}} \tag{1.9}$$

Methods for computing the correlation energy generally comes in two different forms, perturbative or variational. In a nutshell, perturbative approaches are faster but more approximative than variational methods.

1.2.1 Rayleigh-Schrödinger Perturbation theory

Rayleigh-Schrödinger perturbation theory is one of the most commonly employed ways in QC to approximate solutions to many-body systems. It is based on the assumption that if a "good enough" so-called zeroth order approximation can be obtained, then adding a perturbative term to the Hamiltonian,

$$\hat{H} = \hat{H}_0 + \lambda \hat{H}_1 \tag{1.10}$$

will result in a better approximation. $\hat{H}_0$ is the zeroth-order Hamiltonian, λ is a parameter controlling the strength of the perturbation, with $0 \leq \lambda \leq 1$ and $\hat{H}_1$ is the perturbative term. While many forms of $\hat{H}_1$ can be imagined, the only widely used formulation in QC is Møller-Plesset perturbation theory (MP)[6].

The perturbed solutions are arrived at by expanding both the energy and the wavefunction in series of λ as

$$\begin{aligned} |\Psi\rangle &= |0\rangle + \lambda |1\rangle + \lambda^2 |2\rangle + \dots \\ E &= E_0 + \lambda E_1 + \lambda^2 E_2 + \dots \end{aligned} \tag{1.11}$$

In conventional QC, $|0\rangle$ should be understood to be an optimised HF state, $|\Psi^{\text{HF}}\rangle$; the alternative formulation is solely for reasons of brevity. E_1 and E_2 here should not be confused with E_1^{HF} and E_2^{HF} discussed in Section 1.1. On insertion of these expressions in the Schrödinger equation, the following set of expression are obtained up to second order

$$\begin{aligned} E_0 &= \langle 0|\hat{H}_0|0\rangle \\ E_1 &= \langle 0|\hat{H}_1|0\rangle \\ E_2 &= \langle 0|\hat{H}_1|1\rangle \end{aligned} \tag{1.12}$$

which gives the curious, although very useful, result that the first-order correction depend only on the zeroth-order wave function and that the second-order correction requires only that the first-order wavefunction is computed. To compare the last equation with Eq. 1.9, $E_0 + E_1$ is simply the HF energy, E^{HF}, (note that these terms only depend on the HF state) and the correlation energy, E^{corr}, is equal to E_2.

The convenience of these results is arguably the main reason for the prevalence of second-order perturbation theory. While not necessarily obvious from the equations above, second-order perturbation theory essentially tries to approximate single and double excitations from the reference wavefunction; third-order perturbation theory then approximates triple excitations and so-on and so-forth. Thus, unlike HF theory, this means that MPn results additionally depend on the shape of the virtual orbitals. Unfortunately, second-order perturbation theory is not always a good enough approximation and at higher orders the calculation is not much more computationally efficient than using the methods described in the following section.

1.2.2 Configuration Interaction and Coupled-cluster theory

Noting that in HF theory, only the subset of spin-orbitals that are occupied actually contribute, the most straightforward way to improve upon the results is to devise methods where unoccupied orbitals are included into the calculations. In second quantisation, the starting point is then to take an optimised HF state, $|\Psi^{\mathrm{HF}}\rangle$, and use creation and annihilation operators to generate "excited" SD's and compute the contribution from these.

Arguably, the most straightforward way is to employ a linear parametrisation of the molecular state, as is done in CI theory[7],

$$|\Psi^{\mathrm{CI}}\rangle = (1 + \hat{C})\,|\Psi^{\mathrm{HF}}\rangle \tag{1.13}$$

with

$$\hat{C} = \sum_{ia} c_i^a \hat{a}_a^\dagger a_i + \frac{1}{4}\sum_{ijab} c_{ij}^{ab} \hat{a}_a^\dagger \hat{a}_b^\dagger a_j a_i + \ldots \tag{1.14}$$

where $c_i^a, c_{ij}^{ab}, \ldots$, the so-called excitation amplitudes, become the new parameters to solve for. Alternatively, an exponential parametrisation can be used, as in coupled cluster (CC) theory[8,9],

$$|\Psi^{\mathrm{CC}}\rangle = \exp(\hat{C})\,|\Psi^{\mathrm{HF}}\rangle \tag{1.15}$$

Regardless of the parametrisation form, the obvious thing to note, is that in order to compute the exact energy, the full set of excitation amplitudes should be computed. Unfortunately, the scaling for evaluating the excitation amplitudes increases severely for each term used, with the full method usually only applicable for up to around 20 spin-orbitals, which is far too few for most chemical systems. In most practical calculations then, truncated versions has to be used with, doubles (D) being the most common truncation level, giving methods such as CISD or CCSD. If we compare this with the perturbative approach discussed in Section 1.2.1 (specifcally, MP2), the

contributions from single and double excitations are targeted in a direct fashion with CISD and CCSD, instead of approximately. At truncated levels, the CC formulation arguably outperforms the CI formulation; the most straightforward reason being that truncated CC is size-consistent, unlike truncated CI, and includes higher order corrections [1].

1.3 Multiconfigurational Quantum Chemistry

The methods described in the previous section work very well when the original HF state is a qualitatively correct description of the wavefunction. While this is true for a very large class of compounds, some systems, such as open-shell systems, which are frequently encountered among the d-block and f-block metals, are in general not satisfactorily described by a single SD [1,2]. For systems where HF is insufficient, the most straightforward solution within WFT is a class of methods called MCSCF theory.

When a single SD is insufficient to describe the wavefunction, it is common to decompose the correlation energy into two different contributions, dynamical and static. The dynamical correlation is, essentially, the same as the correlation that is missing from an HF calculation. It is primarily described by a large amount of small excitations, predominately up to doubles; usually, it is easily recovered by MP, CI or CC theories directly. The static part is associated with the insufficient reference wavefunction employed in HF theory and is recovered by allowing the reference wavefunction to be described by multiple SD's with large coefficients. Section 6.2 will devote some more effort to discuss dynamical and static correlation.

In a nutshell, MCSCF methods combine the basic CI ansatz with orbital rotations; the wavefunction is parametrised as a linear combination of some SD's, $|\Psi_i\rangle$, as

$$|\Psi^{\text{MCSCF}}\rangle = \sum_i c_i |\Psi_i\rangle \tag{1.16}$$

where both the molecular orbitals (MO) coefficients used to expand the SD's and the CI coefficients are variationally optimised. It should also be stressed that in many codes, the SD's are pre-contracted into a set of CSF's, such that the basis functions used to expand the MCSCF wavefunction are true eigenfunctions of the spin-operator [1,2].

For the same reasons that full-CI is rarely used, for MCSCF calculations to be practically feasible, there must be some way to constrain the number of SD's included in the calculation. This was realised by introducing the complete-active-space self-consistent field (CASSCF) approach [10]; instead of simply partitioning the MO space into occupied and virtual orbitals a third subspace is introduced, called the active orbitals. All

possible SD's associated with a full-CI calculation on the active orbitals are used to construct the wavefunction, with the intention that the resulting wavefunction should cover most, if not all, of the static correlation. Since the number of active space orbitals is usually only a very small subset of the complete orbital space, the selection of an appropriate active space becomes a primary focus in MCSCF calculations. The selection of appropriate active spaces is not straightforward and is the topic for a large part of the discussion on Ni:MgO in Paper I and Chapter 6 of this thesis. Realising that the CASSCF approach becomes limited, when there is a need for a very large number of orbitals to be included into the active space, as can be speculated to frequently occur for solids, alternative approaches such as the restricted-active-space (RAS) approach[11], the generalised-active-space (GAS) approach[12], stochastic CAS[13] and density matrix renormalisation group (DMRG) theory[14,15,16] have been developed. These methods are more approximate than CASSCF, but the gain from increasing the active space size often outweighs this deficiency.

As stated above, an MCSCF wavefunction only recovers the static part of the correlation energy and is therefore in general not a good enough approximation to the true solution. Dynamical correlation needs to be accounted for in some fashion, in direct analogy to how the methods in Section 1.2 improves upon the HF results. In the early days of MCSCF theory, this was done by extending CI theory to multireference configuration interaction (MRCI); while this method is powerful, it has poor scaling with basis set size and truncated versions are size-inconsistent[1,2]. Equivalently to MP theory for HF, perturbative approaches like the one discussed in Section 1.2.1 have been developed on top of MCSCF theory. This is the most common approach today to account for dynamical correlation in MCSCF calculations and is generally done in the form of complete-active-space perturbation theory (CASPT)[17] or n-electron valence perturbation theory (NEVPT)[18], usually truncated to second order[19]. CC theories for multireference systems do exist[20]; for various reasons, they have not yet become routinely employed by the QC community.

1.4 Density functional theory

An alternative formulation for QC is density functional theory (DFT), which, instead of considering an explicit wavefunction, formulates the energy directly as a functional of the spatial electron density. This is motivated on the basis of the Hohenberg-Kohn theorem[21], which states that the ground state energy of an electronic system can by computed by some universal functional of the electron density. From a computational point-of-view, this has the advantage that, unlike an N-particle wavefunction, which is a function of $3N$ spatial variables, the electron density is a function of only 3 spatial variables, a significant reduction in complexity. The major problem with DFT is that

the universal functional is not known. Modern DFT therefore generally relies on approximations, most frequently in the form of Kohn-Sham DFT (KSDFT)[22].

Since this is not a thesis on all aspects of DFT, I will not provide any derivations of the set of equations that I will discuss here, the interested reader is referred to Ref. 23. In general, we are interested in some functional, F, of the electron density, $\rho(\mathbf{r})$, that we can decompose as

$$F[\rho(\mathbf{r})] = T[\rho(\mathbf{r})] + J[\rho(\mathbf{r})] + E_{\text{ncl}}[\rho(\mathbf{r})] \tag{1.17}$$

where $T[\rho(\mathbf{r})]$ is the kinetic energy of the electrons, $J[\rho(\mathbf{r})]$ is the classical part of the Coulomb interaction and $E_{\mathbf{ncl}}[\rho(\mathbf{r})]$ all the remaining non-classical terms – the so-called self-interaction correction, exchange and Coulomb correlation. Of these, only $J[\rho(\mathbf{r})]$ is known. Kohn and Sham approached this problem by introducing a non-interacting reference system, that has the same electron density as the interacting system and is represented by an SD[22]. This choice was motivated on the basis that a single SD can be said to represent non-interacting electrons[23]. Additionally, the kinetic energy of a system represented by an SD has a known form, which we will call $T_S[\rho(\mathbf{r})]$. Using an SD therefore allows the recovery of a large portion of the kinetic energy in the interacting system (though not all of it). Much like in HF theory, the shape of the orbitals spanning the SD is not known in advance, for which reason SCF-like equations are used when working with KSDFT[22,23]. In the end, F in KSDFT can be said to take the form

$$F^{\text{KS}}[\rho(\mathbf{r})] = T_S[\rho(\mathbf{r})] + J[\rho(\mathbf{r})] + E_{XC}[\rho(\mathbf{r})] \tag{1.18}$$

where

$$E_{XC}[\rho(\mathbf{r})] = (T[\rho(\mathbf{r})] - T_S[\rho(\mathbf{r})]) + E_{\text{ncl}}[\rho(\mathbf{r})] \tag{1.19}$$

is the KSDFT exchange-correlation energy, which covers both the missing portion of the kinetic energy of the electrons, as well as the non-classical terms. Since the exact form of E_{XC} is unknown, there exists an extremely large number of different DFT functionals, as exemplified by the LibXC library – an attempt at making a standardised library of DFT functionals – which in 2017 contained 400 different functionals[24].

Further, the parametrisation of DFT functionals is not straightforward. Some groups have focused on using fitting parameters to reproduce certain properties, obtained either from experimental data or from high level WFT, while others attempt to construct functionals based on certain known constraints[25]. There have also been some attempts at defining a Jacob's ladder for improving accuracy with DFT functionals, i.e., selecting a functional from a higher level on the ladder should give better results than using one from a lower level[26]. In practice, neither selecting a more rigorously motivated functional, nor one on a higher rung of the Jacob's ladder, does in general improve the results for all properties of interest[27].

Given these issues, why is DFT so popular in both the molecular and solid-state QC communities? Simply put, KSDFT provides a fast way to include electron correlation into the calculations, with a complexity scaling is either better than of HF (non-hybrid functionals, i.e., without HF exchange) or close to it (hybrid functionals, i.e., including HF exchange) and accuracy that is usually sufficient ground state properties[25,28]. For excited states and strongly correlated systems, however, KSDFT fails much in the same way as standard HF does and a multiconfigurational approach becomes mandatory. Attempts at combining the strengths of KSDFT with those of multiconfigurational theory have been made, with two examples being multiconfigurational pair-density functional theory (MCPDFT)[29] and multiconfigurational short-range DFT (MC-srDFT)[30]. While promising in many aspects, neither approach solves the underlying issue of KSDFT that the results are still dependent on the functional of choice, for which reason the work in this thesis has refrained from using these methods.

1.5 Wavefunction theory in the solid state

Clearly, as suggested by Section 1.2, from WFT there exists several different methods which can be used to describe the electronic structure of a solid. While this might seem daunting to any one embarking on their first exploration of QC, the big advantage of WFT is that they follow a fairly strict hierarchy of increasing accuracy. As argued in the previous section, this is not true for KSDFT.

Take a crystal which we might assume is represented well by a single SD, MgO, for instance, which (in the ground state) is a closed-shell compound. In this case, we can use HF to compute a zeroth-order wavefunction and then, in principle, choose freely between PT, CI or CC in order to improve upon the results. Similarly, if the crystal is open-shell in nature, i.e., magnetite (Fe_3O_4), we can start with an MCSCF wavefunction and then use CASPT2 to recover the missing parts of the correlation energy.

Of course, the previous paragraph should be taken with some scepticism. Solids are very large systems, which in combination with the poor scaling of WFT methods generally makes it impossible to directly compute, say, the CCSD energy of MgO. For open-shell crystals, the active space demands will be far greater than those of molecules, simply due to the number of particles that needs to be modelled. Lets take a real example – the Cr_2 dimer, which has long served as a standard test case for new electronic structure methods[31,32,33,34]. Due to the complicated nature of the bonding in this molecule, as much as 22 orbitals is in general necessary to get a good description of the Cr_2 dissociation curve with multiconfigurational theory[31,32,33,34].

If we now imagine that we extend this to a chromium trimer, Cr_3, and make the assumption that we still need 11 orbitals per chromium, that would give us an active space of 33 orbitals – far beyond the capabilities of CASSCF. Instead some of the alternative approaches mentioned in Section 1.3 becomes necessary. Of these methods, only RAS and GAS are able to (as of writing this thesis) account for dynamical correlation in an efficient way (RASPT2[35], GASPT2[36]). On the other hand, these methods can be criticised by the fact that the user has to make decisions on in which active space which orbital should be placed into, as well as what type of excitations to include in the calculation. DMRG does, in principle, not have such problems and can be used with both CASPT2[37] and NEVPT2[32]. The accuracy and the computational cost of a DMRG calculation is, however, determined by a parameter known as the maximum bond dimension (the m-value) and unfortunately, both DMRG-CASPT2 and DMRG-NEVPT2 suffer from the fact that for large active spaces, a very large m-value is required for accurate perturbative corrections[38]. To the best of my knowledge, the stochastic methods has not yet been extended to include correlation beyond that captured by the active space itself. What should become apparent from this discussion – dealing with a three metal atom system such as Cr_3, is very complicated with multiconfigurational theory. Now imagine a solid such as magnetite, where a single unit-cell contains 32 metal atoms – any direct application of multiconfigurational theory to magnetite is clearly out of the question. Not to mention the fact that it is in general difficult to implement such methods with so-called periodic boundary conditions (PBC), which will be discussed further in Chapter 4.

Some compromise between the accuracy of the electronic structure method and the model of the crystal itself is therefore necessary in order to enable WFT to be successfully applied to the solid state. Enabling such a compromise, in particular for MCSCF methods and ionic crystals, has been the goal of this thesis, with Chapters 4, 5 and 6 being devoted to discussions on this topic.

Chapter 2

Atomic basis sets

Having discussed electronic structure methods in the previous chapter, I will in this chapter discuss basis sets. Basis sets are required in QC methods to approximate the spatial part of the wavefunction. Many different types of basis sets exist, with two common choices being Gaussian-type orbitals (GTOs) and plane-waves (PW). An in-depth discussion on the many different strengths and weaknesses of different types of basis sets would constitute a fairly large work on its own. For that reason, this chapter will be devoted to GTOs, which was the type of basis set used in this thesis.

GTO-type basis sets do not appear out of thin air, they have to be developed in some fashion. For this reason, there are many technical caveats that a user will need to be aware of when working with such basis sets. This chapter will discuss some of these aspects, in hope that any user of GTOs will be able to utilise such basis sets in a correct fashion after reading this work, or at the very least find appropriate sources to learn more from.

2.1 All-electron basis sets

When constructing a basis set based on GTOs, a shell-structure similar to that of atomic orbitals is usually followed[1,2]. Taking C as an example, a minimal basis set for a GTO basis set would have two s-functions describing the 1s and 2s orbitals, and one p-function describing the set of three 2p orbitals. In principle, there is nothing wrong with constructing a basis set, where each GTO appears only once, i.e., there is one GTO that describes the 1s orbital and a second that describes the 2s orbital and so-on, for all occupied shells of the atom. Such a basis set would be computationally efficient, as it would only use a minimal number of functions. It is well-

known, however, that such minimal basis sets are insufficient, as they lack the necessary mathematical flexibility to give even a qualitatively correct description of chemical bonds[1,2]. A generally applicable basis set requires far more functions, than the atomic shell structure would suggest. This rapidly leads to basis set sizes which would be computationally unfeasible, unless the GTOs are *contracted* into fewer functions, so-called contracted Gaussian-type orbitals (cGTOs). When working with cGTOs, the underlying, uncontracted, GTOs are frequently known as primitive Gaussian-type orbitals (pGTOs). In the following, ψ_i should be interpreted as cGTOs and ϕ_i as pGTOs.

The result is basis functions on the type of

$$\psi_i = \sum_k c_k \phi_k \tag{2.1}$$

where c_k are the contraction coefficients, which are conventionally precomputed and tabulated for later use. Each pGTO, is an exponential function, $\phi_k \propto \exp(-\zeta_k r^2)$, with exponents ζ_k (exact form for the radial part is given in Eq. 5.5). Therefore, when designing a new basis set, both primitive exponents ζ_k and contraction coefficients c_k have to be determined in some fashion. Obviously, the cGTO basis set can never be more accurate than the pGTO basis set from which it was contracted, for which reason coefficients c_k are, usually, chosen to minimise the truncation error with respect to some desired properties[2]. Additionally, there are more than one way to design the contraction scheme, e.g., so-called generally contracted basis sets or segmented basis sets and more. An excellent overview of some common contractions schemes is provided in Ref. 2. This work has primarily focused on generally contracted basis sets of the Atomic Natural Orbital (ANO) type[39], which will be the focus of the remainder of this text.

Unlike other basis set types, in an ANO type basis set, each cGTO contains contribution from every pGTO for a given angular momentum. For instance, all s-type pGTOs contribute to all s-type cGTOs, all p-type pGTOs contribute to all p-type cGTOs, etc. The result is that the orbitals are very atomic like and clearly separable into core, semicore, valence and correlating orbitals[39], which is not true for all types of basis sets[40]. This is a particularly attractive feature for MCSCF methods, where a careful selection of orbitals is of high importance. The contraction coefficients c_k of an ANO basis set are determined using some correlated wave function method, e.g. MRCI or CASPT2, to describe some atomic state[2,39]. This generates a bias towards the specific atomic state used to construct the basis set, for which reason modern ANO basis sets are, in general, constructed by averaging over several electronic states[2,41]. As discussed in Paper III in this thesis, the selection of states has a large impact on the capability of the basis set to accurately describe a given property.

Before leaving this section, I would like to make a personal criticism of the name all-electron basis sets. While it is true that, when using say ANO-RCC or cc-pVxZ type basis sets, all electrons of the atoms in the system are included at the SCF level, they are rarely included in the correlation step. This simply due to the fact that post-SCF methods scale poorly with the number of electrons that are included. Core orbitals are therefore usually "frozen", meaning the the correlation step is, in fact, not all-electron. Basis sets used for WFT are usually designed with this in mind and include no correlating functions for the core orbitals[2]. For this reason, whenever using basis sets for WFT, the original publication on the basis sets should be carefully consulted, such that all orbitals that lack correlating functions are frozen when using, say, CASPT2. An example of this would be, the 1s orbital in Al with the ANO-RCC or ANO-R basis sets. Any properties which rely on an accurate treatment of core electron effects, such as core-hole spectroscopy or nuclear-magnetic resonance spectroscopy, should therefore not use conventional basis sets[40]. Rather, specially designed basis sets are required to study such properties.

2.2 A note on relativistic and nuclear effects

For heavier elements, it is well known that effects originating in Einstein's theory of special relativity must be taken into account for a proper description of the electronic structure in atoms[42]. Conventional QC was, however, formulated in a non-relativistic framework, with most computational chemistry codes being written under the assumption that relativity was not of any major importance. For that reason, most electronic structure codes treat relativistic effects as a correction, rather than direct and generally divide relativistic effects into two distinct effects, scalar relativistic and spin-orbit coupling (SOC).

Since this is not a thesis on relativistic QC, I will not dwell much longer on the specifics of these effects. Rather, I will make the comment that when modelling heavier elements with an all-electron basis set, that basis set must be designed with relativistic effects in mind[2]. If we take two ANO-type basis sets, ANO-L and ANO-RCC, one major distinction between these two basis sets is that the former is intended for non-relativistic calculations, while the latter is intended for scalar relativistic calculations (specifically, with the Dougles-Kroll-Hess (DKH) Hamiltonian[43,44,45]). Using ANO-L with scalar relativistic effects and ANO-RCC without them does, therefore, introduce potential errors in calculations based on such combinations. The same can be said for whether the basis set was designed using a point-charge model for the nucleus, or a Gaussian model[46]. In general, the Gaussian model is slightly advantageous, for which reason newer basis sets, such as the ANO-R basis set[46] (the topic of Paper III), was designed to be used with a Gaussian model for the nuclei (in ad-

dition to accounting for scalar-relativistic effects via the exact two-component (X2C) Hamiltonian[47,48,49]).

Spin-orbit coupling either has to be accounted for directly by a so-called four-component approach[50] or perturbatively via, for instance, the state interaction method of Malm-qvist[51]. The former approach is more exact and should be used for very heavy elements, such as the actinides[52], though it is very hard to pinpoint between which atoms (and for which properties) four-component approaches become mandatory. Nonetheless, in this work, SOC was modelled perturbatively (Papers I and III). From the point-of-view of perturbative SOC, there is nothing strictly wrong with using non-relativistic basis sets. Given that most elements where SOC is non-negligible are heavy elements, however, scalar relativistic basis sets are almost always more appropriate when modeling SOC.

2.3 The correct basis set for the correct electronic structure method

Let us revisit a statement in Section 1.1; the HF energy is only dependent on the occupied orbitals. On the other hand, post-HF methods also depend on the virtual orbitals (Section 1.2). Clearly, the demands on the basis sets used in a HF (or DFT calculation, for that matter) is very different from those of, say, CASPT2.

This also affects the development of the basis sets themselves. A general guideline to follow when selecting basis sets is to always go back to the original source and determine what type of method was used in the parametrisation. If it was based on some correlated WFT, then, if the basis set was appropriately designed, any correlated WFT method can be used with it[2]. On the other hand, if the basis set was designed using some DFT functional (or HF), it is ill advised to use it with correlated WFT[2]. Basis sets designed for correlated WFT, however, are in general safe to use with DFT, but such basis sets are frequently unnecessarily large and inefficient with DFT[40]. Such considerations were the underlying reason for primarily relying on correlated WFT in Paper III, where the ANO-R basis set (designed with CASPT2) was evaluated.

2.4 All-electron basis sets for solids

As stated in Section 2.1, modern ANO basis sets are typically obtained by averaging over several different atomic states[41]. This averaging generally leads to basis sets that are able to treat atomic properties derived from the included states with fairly con-

sistent accuracy. Before continuing with a discussion on basis sets, I would just like to make a comment on my nomenclature. In chemistry, it is common to use notation on either the form of M^{n+} or M(n) (with roman numerals in the latter case) for ionic states and oxidation states respectively. These concepts are highly intermixed, as a consequence, these two notations are often used interchangeably. I will make an attempt in this thesis to be strict, by using M^{n+} when I really mean that n electrons have been ionised away, while I will use M(n) to refer to a formal oxidation state. While I am under impression that using M(-n) (still with roman numerals) for anionic oxidation states is less common in chemistry than for cationic oxidation states, I will use this nomenclature in this text to be as internally consistent as possible.

Returning to a general discussion on oxidation states; in molecular systems, atoms frequently appear in oxidation states that are close to neutral. For example, in NO_2, it would be more accurate to say that the molecular state is derived from the N(0) and O(0) states, rather than N(IV) and O(-II). While it can be argued that in molecular MgO, the oxidation states are Mg(II) and O(-II), this is in fact not a particularly accurate representation. The oxidation states are probably closer to Mg(I) and O(-I), with a large degree of covalency in the bond. Thus, constructing basis set from atomic states is generally not a too bad approximation for the chemistry of the s-block and p-block elements.

In solids, particularly those where the bonds are predominantly ionic, this is not necessarily true. This is due to the Madelung potential – the electrostatic potential from the surrounding ions makes higher oxidation states more stable. Crystalline MgO is therefore fairly well represented by Mg(II) and O(-II), unlike its molecular counterpart.

So what does this mean for the ANO-type basis sets? Lets start with a discussion of cationic states, exemplified by Fe. Fe is one of many metals that have several stable oxidation states[53], with the two most common being Fe(II) and Fe(III). For instance, in magnetite (Fe_2O_3), there is a mixture of Fe(II) and Fe(III) ions. In conventional ANO type basis sets, however, only the first cationic state (Fe^+) is generally included in the averaging[46,54]. Given the fact that each successive ionisation step further contracts the ions, this is sufficient. This is clearly evidenced by Table 2.1, where the first three ionisation potentials (IP) of Fe are very well described by extended multistate CASPT2 (XMS-CASPT2). Of course, it should be noted that these oxidations states are all derived from the valence orbitals (4s3d); higher oxidation states, like a hypothetical Fe^{9+}, would start to ionise semicore orbitals, which should be assumed to be less well described by standard basis sets. We might therefore presume, that the typical construction of ANO basis sets for cations in typical oxidation states work reasonably well for solids.

Table 2.1: Ionisation potentials of Fe (First IP with respect to neutral Fe (^{5}D(s^2d^6)). The 4s3d4d orbitals were kept active, with the 1s2s2p3s orbitals frozen. ANO-RCC-VQZP was used as basis set. Experimental values from NIST 55.

	Fe^+ (^{6}D(s^1d^6)) 1st IP	Fe^{2+} (^{5}D(d^6)) 2nd IP	Fe^{3+}(^{6}S(d^5)) 3rd IP
XMS-CASPT2	7.77	16.3	30.3
Experiment	7.90	16.2	30.7

Anionic states, on the other hand, are far more complicated. Oxygen is famously usually referred to as O(-II) in ionic solids, with O(-I) being extremely rare. For this reason, we might speculate that we want to include O^{2-} in the averaging. Gaseous O^{2-}, is, however, an unbound state. In fact, there is no experimental evidence for any gas phase anions in oxidation states larger than -I. Unbound states are resonance states, such states require a time-dependent formalism, therefore typical QC methods can not describe them properly[56]. In an ionic solid, however, the O^{2-} state might be speculated to be stabilised by the external field from the surrounding system. Introducing an external field might then lead to the possibility of using MRCI or CASPT2 to generate a basis set for O(-II).

But lets make a thought experiment, O^{2-} is stable in the gas phase and we use it in the construction of an ANO-type basis set with CASPT2. The effect of this, as compared to only including O^- would, presumably, be to generate a basis set that is even more diffuse. In compact systems, such as solids, diffuse functions have long been realised to produce linear dependencies, resulting in unstable convergence patterns[57,58]. Therefore, some effort has been put into redesigning molecular basis sets for calculations in the solid state, with a typical effect being that the basis sets becomes far less diffuse[57,58]. So even if we were to include a gaseous O^{2-}, it might not actually be of any use at all.

Further, it should be stated that basis sets resulting from typical ANO design are very large. This might not matter too much for molecules, but for the solid state it quickly results in calculations that are computationally unfeasible. If we now find it desirable to use all-electron ANO-type basis sets for solids, which this thesis argues that it is given their natural connection to MCSCF theory, some redesigning will become necessary.

One could, for instance, generate ANO basis sets using similar type of fields as when designing AIMP (see Section 5.4). The downside of this approach is the same as when designing AIMPs: it is quite likely that the basis sets become material specific and possibly even specific to a given polymorph. In that case, it becomes very hard to diagnose how good the basis set is.

Chapter 3

Chemical bonds in Quantum Chemistry

In the previous chapters, the basics of QC, namely electronic structure methods and basis sets, were addressed. This chapter will address a more fundamental chemical concept, namely, chemical bonds. This is necessary as the type of chemical bonds in a system will impact the way it is approximated in so-called embedding methods, which is the primary topic of this thesis.

It could be argued that I should discuss embedding methods more thoroughly before discussing chemical bonds. From my point of view, however, some basic recap of chemical bonds becomes necessary for a discussion on embedding methods. For that reason, I defer a discussion on embedding methods to Chapter 4. Further, I will discuss some ways of diagnosing chemical bonds using QC. As I will point out in the text, QC can never provide absolute answers in this context, but such diagnostics are particularly useful when assessing if an embedding strategy is reliable or not.

3.1 Chemical bonding – considerations for embedded cluster methods

As most students of science, irrespective of field of discipline, will know, the bond between atoms is usually subdivided into different classes. Normally, chemical bonds can be discussed in terms of three extremes; ionic, covalent and metallic. As standard QC methods can describe all of these without any modifications to the Hamiltonian, the fundamental interaction in all three is electrostatic. The main divisor is where the

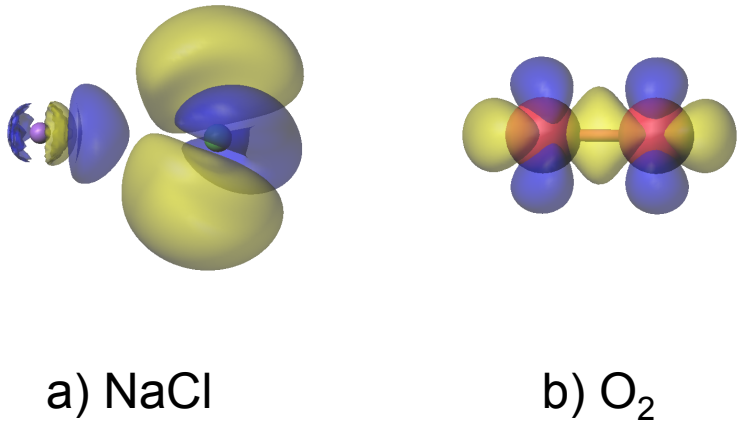

Figure 3.1: Hartree-Fock electron density differences for a) NaCl and b) O_2. The total electron density of the atoms have been subtracted from the total molecular electron densities. Yellow isosurface corresponds to increase in electron density on forming chemical bonds, with blue isosurfaces representing where the electron density decreased. Isosurface level 0.01 e/a_0^3.

electron density localises in the system.

In ionic systems, the electron density localises fully around the atoms. In covalent systems, a non-negligible portion of the density is along an axis connecting two different atoms, slightly polarised towards the more electronegative species. In metallic systems, some part of the electron density is localised in between atoms, but the notion of directionality along any axis is lost; the (valence) electron density essentially becomes delocalised. Figure 3.1 shows differential density plots of an ionic molecule (NaCl) and a covalent molecule (O_2); these show how the electron density is either localised around the anionic species in NaCl and in the middle of the bond in O_2. No visualisation is given for a metallic system, as the density is completely delocalised. Clearly, creating a single approximative method that can deal with all of these different extremes on an equal footing is a non-trivial task. Some advances towards uniformed approaches have been made, for instance in the form of density-functional embedding theory (DFET)[59] and density matrix embedding theory (DMET)[60].

For the purpose of this thesis, we will stick with only ionic and covalent materials in the remaining discussion. For ionic systems, the QM–embedding border should then be able to reproduce the charge localisation around the individual ions, while for covalent systems, the density must be allowed to polarise in a similar fashion as if a real chemical bond was in place. Obviously, ionic systems will be easier to approximate and are therefore more suitable when starting to develop new embedding strategies, as was the object of this doctoral thesis. As will be discussed in the next two chapters,

this was realised in the form the AIMP embedding strategy.

3.2 Measuring embedding quality via computational diagnostics of chemical bonding

While frequently hard to accept for chemists, in quantum mechanics, there is no operator which can directly compute the charge of an atom in a molecular or crystalline system. Since there is no operator, atomic charges are not observables. For this reason, when diagnosing the nature of the chemical bonds in a particular system, indirect routes has to be taken and no definitive answer on the atomic charges or bond orders can be given. Whenever atomic charges from QC calculations are discussed, they are either taken from features directly connected to the wavefunction, i.e. Mulliken charges[61], or from features of the spatial electron density, i.e. Bader charges[62].

While QC can not provide absolute answers in terms of atomic charges, the different approximations for atomic charges, as well as similar concepts such as bond order, can be used as diagnostic tools when evaluating embedding methods. This simply due to the fact that, if we take MgO with, say, the Perdew-Burke-Ernzerhof (PBE) DFT functional, it should not matter if we make a periodic calculation or an embedded cluster calculation. If the embedding is good enough, then the PBE derived charges and bond orders should converge to the same numbers, assuming the same basis set was used in both sets of calculations. For this reason, in Papers I and II, Mulliken charges, along with other related diagnostics, were used to diagnose the quality of the embedded clusters discussed therein. This choice was primarily due to the practicality of using Mulliken charges and that there are related diagnostics for bond indices and atomic valencies, which will be discussed below.

Most discussions and equations surrounding the used diagnostics are already provided in Paper I, for which reason I will condense the discussion here. The local properties of the electronic structure that were used to assess the quality of embedding in Papers I and II are all derived from the population matrix, $\mathbf{P}$, which is computed as $\mathbf{P} = \mathbf{DS}$ ($\mathbf{D}$, $\mathbf{S}$ are the density and overlap matrices, respectively[1]). The properties discussed in Papers I and II, namely Mulliken charges[61], Wiberg bond-indices[63], covalencies[64] and total valencies[65] are all derived from norms of the population matrix. As such, they are independent of any arbitrary rotations and are fairly robust with respect to differences in codes. They are fast to compute and therefore provide quick insight into whether the embedding is good or not, depending on how much they deviate from PBC calculations (see Papers I and II).

Additionally, as discussed in Paper II, the spatial electronic density can be compared.

While there certainly are arguments for this approach being more robust than comparing local properties[66], it is far trickier to make accurate comparisons between different codes, as will frequently be the case when comparing periodic calculations and embedded cluster calculations. For instance, one problem when comparing CP2K and Molcas, as was done in this thesis work, is that the core density is treated in an approximated fashion in CP2K[67]. Within the core region, it is impossible to get the two codes to match and it must therefore be removed. Further, when working with DFT, the spatial electron density is sensitive to the integration grid, which might additionally introduce numerical noise when comparing different codes[68,69]. This would not be a problem with, for instance, HF; unfortunately, HF becomes very expensive with periodic boundary conditions[70], for which reason PBE was used in Paper II.

Chapter 4

Periodic boundary conditions and embedded cluster models of crystals

Chapters 1–3 provide an overview of all basic concepts, necessary to understand embedding methods in electronic structure theory. In this chapter, I will start to address how these concepts relate to the solid state and the computational model that I have been developing.

This chapter will also provide an overview of the problems related to the size of crystals and discuss some problems related to the "standard" approach when studying ionic crystals with QC today. Thereafter, the two methods connected to the SCEPIC program, namely the general purpose electrostatic embedding (GPEE) method and AIMPs, will be discussed and I will argue why such a combination is sufficient, when working with ionic materials. As a part of this discussion, I will provide a general set of guidelines when modelling ionic crystals using this combination of methods.

4.1 Unit-cells and supercells

This text will primarily discuss so-called crystalline materials. Thus, it is appropriate to introduce some common nomenclature used in crystallography and computational solid state physics and chemistry, namely unit-cells and supercells, which will be the topic of this section. The introduction given here will be very brief, with more extensive discussions in this topics in, for example, Refs. 53,71,72.

In crystalline materials, the constituent atoms are regularly ordered and follow a repeating pattern (i.e,. they are translationally symmetric). An example of a two-

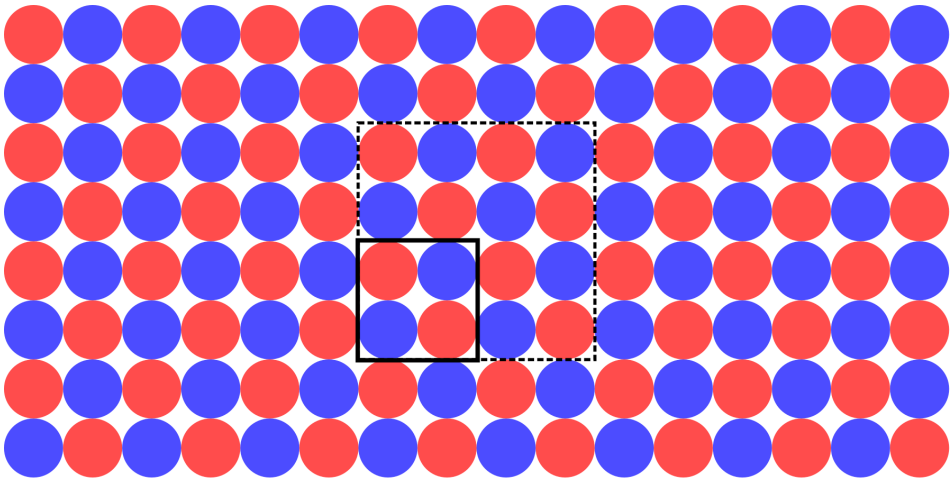

Figure 4.1: Examples of a two dimensional crystal made out of blue and red atoms. The black square marks one possible choice of unit-cell in this material, with the dashed square outlining a 2×2 supercell based on that unit-cell.

dimensional crystalline material is given in Figure 4.1.

Due to the large degree of order in a crystal, the atomic structure of a crystal is usually discussed in terms of unit-cells. A unit-cell, as marked by a block box in Figure 4.1, is a three dimensional box filled with atoms such that, if the box is repeated in space along the three vectors spanning the box, it will build up the atomic structure of the entire crystal. It is easy to get the impression that a crystal will always assume the same shape as the unit-cell, for instance, that a material with a cubic unit-cell will always appear as a cube, if put under a magnifying glass. This is a misconception; the choice of unit-cell is seldom unique for a given crystal and by making different cuts through the unit-cell, different types of surfaces can be created, giving many crystals a less regular shape, than implied by the unit-cell concept[53,72]. Not to mention the fact that defects are ubiquitous in real crystals[53,71,72], meaning that the idea of a perfectly ordered crystal is an idealised picture. Nonetheless, for the purpose of modelling bulk material where surface effects and the presence of minor defects are usually of minor importance, the unit-cell concept has proven extremely useful[71].

An arbitrary unit-cell can be taken as the combination of a set of three vectors describing the size and shape of the unit-cell, along with a set of coordinates for the atoms placed in the unit-cell. Often, these vectors are labelled as $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{c}$ (with magnitudes a, b and c). Translating any atom inside of the unit-cell by any integer multiplicative of these vectors will generate all translationally symmetric atoms. In more mathematical terms, if we let $\mathbf{r}$ be some position inside the unit-cell, and $\mathbf{R} = k_a\mathbf{a} + k_b\mathbf{b} + k_c\mathbf{c}$ (k_a, k_b and k_c, any integer), then if there is an atom posi-

tioned at $\mathbf{r}$, there will also be an atom positioned at $\mathbf{r} + \mathbf{R}$ and these two atoms are symmetry equivalent.

Using this, so-called supercells can be constructed, which are simply an extension of the unit-cell concept. In essence, a supercell is made from several unit-cells; often, they are labelled by $k_a \times k_b \times k_c$, representing the number of times a unit-cell has been repeated along a given unit-cell vector ($\mathbf{a}$, $\mathbf{b}$ or $\mathbf{c}$). In the two-dimensional case in Figure 4.1, a 2×2 supercell has been outlined in dashed lines. The purpose of supercells in computational solid-state chemistry frequently appears when it is desirable to model, for instance, dopants, as will be discussed in Section 4.3.

4.2 Scaling problem of crystals

Lets discuss what I would like to call the "scaling problem" in solid state chemistry. Take a typical crystalline material, such as the common rock-salt, NaCl, that many people enjoy seasoning their food with. Each crystal of NaCl can be viewed as a nanoparticle, with dimensions of, at the very least, 100 nm. If we approximate the crystal as a perfect cube, the number of atoms in the crystal can easily be estimated. A single (crystallographic) unit-cell of NaCl contains four Na and four Cl, with a unit-cell dimension of 5.64 Å, or 0.564 nm. Each cubic crystal of NaCl with sides of 100 nm would then corresponding to approximately a $177 \times 177 \times 177$ NaCl supercell, this leads to a total number of 44 361 864 atoms in a single crystal.

If we now want to make a standard QC calculation on this crystal, using a minimal basis set (3s1p per Na and 3s2p per Cl), this would result in a basis set containing 332 713 980 basis functions. Noting that current limitations on most standard supercomputers are somewhere in the range of 2000 basis functions (when working with correlated wavefunction theory), modelling something as simple as a NaCl crystal in a direct fashion is clearly impossible.

One natural question to ask, is if it is truly necessary to use QC to model the full crystal, or if some alternative models or approximations can be invoked. Two different approaches to model crystals are in existence, one based on the periodicity of the atomic structure in a crystal and the latter based on dividing the interactions in the crystal into short-range and long-range and treating these at different levels of theory. The latter way of approximating a crystal is usually referred to as an embedding. Since the main part of the thesis work has been on embedding methods, a brief discussion on PBC will be given in Section 4.3, before continuing on with a discussion on the embedding method that was used here in the remainder of the chapter.

4.3 Periodic boundary conditions

As any student of solid state physics already knows; one smart way of circumventing the scaling problem for crystalline material is to utilise the translationally symmetric nature of the crystal. Rather than describing the crystal directly, PBC are imposed on the calculation. In this way, the spatial basis set required to describe NaCl with a minimal basis is reduced to just 120 basis functions. Increasing the basis set to quadruple zeta level (6s5p3d2f1g for both Na and Cl) would result in 472 basis functions, which is easily manageable on today's computers. Some increase in computational compared to vacuum calculations should be expected, due to the introduction of so-called **k**-vectors[71,73], which may be on the order of hundreds for metallic system, but only on the order of tens for ionic systems. A proper introduction to **k**-vectors (also called **k**-space) is very important for a deeper understanding of PBC, however, such a discussion would be out-of-scope for this thesis. For the remainder of this text, it is sufficient to simply accept it as a given fact that **k**-space encodes information about the electronic structure in the entire crystal.

The basis for the use of PBC in computational solid state physics and chemistry lies in Bloch's theorem. This theorem is described in detail in many standard textbooks on solid state theory[71,73]. Here, only the basic idea will be given. Basically, for an ideal crystal, the electronic potential, $V(\mathbf{r})$, is translationally symmetric, just like the atomic structure of the crystal, i.e., $V(\mathbf{r}+\mathbf{R}) = V(\mathbf{r})$. Bloch's theorem then states that, the solutions to the Schrödinger equation can be taken as $\psi(\mathbf{k}, \mathbf{r}+\mathbf{R}) = \exp(i\mathbf{kR})\psi(\mathbf{k}, \mathbf{r})$. Therefore, if the solution to the Schrödinger equation is known in one unit-cell, it can easily be obtained for the full crystal.

While PBC is a fantastic approximation in many cases and has been successfully applied for many properties[28], there are some caveats associated with it. In particular, when the unit-cell is very small. The first and most trivial is in the case of sparse dopants (atom percent of around 1 %); for instance Cu-doped NaCl. Should only a single crystallographic NaCl unit-cell be used, then one Na is substituted for one Cu; this would correspond to a doping level of 12.5 %. Therefore, a supercell approach is required. Lets take a $2 \times 2 \times 2$ repetition of NaCl; in this case, there would be 64 atoms in total in the cell and the dopant level is decreased to around 1.5 %. Using a $2 \times 2 \times 2$ supercell should supposedly then give a doping level which is acceptable (though arguably, the Cu-dopants will be too regularly ordered). Still, when working with all-electron GTO type basis set, a lot of thought has to be spent on the basis set size when using a $2 \times 2 \times 2$ supercell.

Using a quadruple zeta contraction level on all atoms (Cu, 7s6p4d3f2g1h) of the $2 \times 2 \times 2$ supercell would require 3812 basis functions, which is clearly unfeasible (for

correleated wavefunction theory) on most standard computers. A common technique that can be employed when working with GTOs basis sets is to use different contraction levels on different atoms, depending on their estimated importance. In the current example, one might speculate that a double zeta contraction is good enough for the NaCl host (4s3p1d for both Na and Cl), and that a quadruple zeta basis set is only required on Cu. Using a combination of double zeta on the host and quadruple zeta on Cu gives a total of 1229 basis functions in a $2 \times 2 \times 2$ supercell; such a calculation would be expensive, but manageable. Ideally, when using mixed contraction levels, the results should be carefully checked to be converged with respect to increasing the basis set. Such studies are unfortunately rarely done, as the original motivation for mixing the contraction level is that using larger basis sets is too expensive.

Ignoring the problem of basis sets size, there is an additional problem when it comes to using PBC. It is in general not a trivial task to take a correlated wavefunction method and adopt it to PBC, for which reason mainly DFT is used for such studies. The combination of DFT and PBC has been successful for a large range of properties[28], and it is not the intention of this text to claim that such a combination is never motivated. But for many compounds that contain either d-metals or f-metals, the accuracy of standard DFT functionals are limited. A typical example would be in the technologically important reducible oxides TiO_2 and CeO_2[74].

Oxygen vacancies are ubiquitous in TiO_2 and CeO_2[74]; as a consequence there is a mixture of M(III)/M(IV) ions in any real crystal of these materials, where the M(III) ions have localised unpaired electrons. Standard, non-hybrid, functionals can not describe this phenomenon properly, for which reason they are frequently studied with under the paradigm of DFT+U[75], where a so-called Hubbard U parameter is added to the calculation to enforce electron localisation into the Ti d-levels or the Ce f-levels. While there are some self-consistent procedures for obtaining the best value of U[76], most often, it is treated as an input parameter with a fixed value. When used in this way, there is no unique best choice of U for a given material, or DFT functional for that matter[77]. Take two examples with the PBE functional for CeO_2; the bulk electronic structure is best described with a value of around 5.5 eV[78] whereas values of around 2–3 eV are more appropriate for surface thermochemical properties[79,80]. Using hybrid functional does ensure proper localisation of the unpaired electrons, but the usual drawbacks of DFT still apply, i.e., which hybrid functional should be chosen (e.g., PBE0, HSE06, TPSSh, SCANh, etc.) with some publications noting that sometimes it is better to use non-default degrees of HF exchange in the hybrid functionals[81,82,83,84]. Moving beyond both standard and hybrid DFT functionals, into the realm of more accurate and parameter free WFT, holds the promise of eliminating such issues from consideration.

4.4 Embedding considerations

The question is then, if it is too complicated to implement PBC with most WFT methods and we at the same time require a good quality basis set to describe whatever physicochemical phenomena we are interested in, how do we deal with it? Coming back to the example of sparse dopants in a crystal, the first question we should ask ourselves is if it is truly necessary to describe the full crystal with quantum mechanics? Noting that the main two-body terms in the QC Hamiltonian are electrostatic in nature, either as a classical Coulomb interaction or as the quantum mechanical exchange interaction, there are the terms that have to be analysed.

In systems of localised charges, such as ionic systems or covalent systems, the HF exchange term has long been realised to decay fairly rapidly, leading to the development of screened hybrid DFT functionals[85]. Supposedly then, we should be able to ignore long-range exchange interaction, when studying local properties in ionic or covalent crystals. Short-range exchange interaction, however, serves as a repulsive barrier between atoms, for which reason it should not be neglected[86]. The importance of the Madelung potential in the crystal should not be understated and it should be approximated in some fashion[71,87]. One simple idea for modelling crystals is then to replace atoms far from the central region with, for instance, classical point-charges. Using classical point-charges instead of explicit atoms greatly reduces the computational effort since i) such integrals are inexpensive to compute and ii) they only effect the one-electron Hamiltoninan, thus they do not impact the cost of SCF iterations or correlated calculations. Of course, two questions still remain, one is how big do we need the quantum region to be? This question is the topic of Paper II. The second is of a more technical nature, if we expose an electron density, which by definition carries a negative charge, to a positive point-charge to model a cation, this point-charge will correspond to a singularity in the Coulomb potential, leading to an artificial "leakage" of electron density into the point-charge region. Thus, there needs to be some repulsive wall to remove this singularity.

Obviously, the repulsive potential should fulfil, at the very least, two criteria. This first and fairly obvious criteria is that the potential must prevent the electron density from spreading out into the point-charge region. The second, is that the influence from the potential should be such that atoms on the border of the quantum cluster have the same properties as if the entire material was described in an exact fashion. While the first criteria is usually fairly easy to achieve, indeed, simply replacing any positive point-charges with basis-less effective core potentials (ECP) is sufficient (the so-called capped ECP approach), the second criteria is the source of a lot of headache. The main problem being that depending on the type of chemical bonds in the system, different strategies needs to be employed. For instance, in a truly ionic system, it is most likely

sufficient that the potential prevents the electron density from leaking, as evidenced by the historic success of capped ECP[88,89] and embedded AIMP studies[90,91,92]. In more covalent systems, however, the directional chemical bonds between the quantum cluster and the embedding region must be mimicked somehow[93], which is a non-trivial task.

The history of embedding is very long, there are probably several thousands of articles on various aspects of the embedding schemes. As such, it is impossible for this thesis work to give proper citation to all relevant articles and provide a comprehensive discussion about all different methods that are in existence. Therefore, in addition to already provided references, the interested reader is referred to, e.g., Refs. 94,95,96,97,98,99,100,101 and references therein for a more thorough overview of different embedding schemes for solids.

4.5 General purpose electrostatic embedding

Before mentioning more details of the AIMP embedding strategy, the question of the Madelung potential resulting from the crystalline field must be addressed. This potential is nothing more than the result from the classical Coulomb potential created from all the ions in the crystal. In atomic units, the classical electrostatic potential at some ion j in a crystal reads as

$$V(r_j) = \sum_{i \neq j} \frac{q_i}{|r_i - r_j|} \tag{4.1}$$

where q_i and r_i are, respectively, the charge and position of ion i. As discussed in Paper I, this summation can, in principle be solved directly, but suffers from convergence issues due to the r^{-1} dependency and the fact that in an ionic crystal, the charges will be of alternating signs. Therefore, strategies such as Ewald summation[102] or Evjens method[103] are usually invoked to enforce convergence of this series.

In the work related to this thesis, a different approach was used, namely the GPEE method of P. Sushko and I. V. Abarenkov[87,104]. The basis for this method is to replace the, in principle infinite, sum in Eq. 4.1 with a summation over a discrete set of point-charges. Additionally, these point-charges are constructed in such a way, that they cancel multipole moments of arbitrary order over the unit-cell. The details of this method are discussed at length in Refs 87,104, for which reason most details will be omitted here.

The important realisation of this method was that the higher order multipole moments over a unit-cell can be placed in tetrahedra situated at the corners of the unit-

cell. The mth multipole moment of the unit-cell, Q_0 can be computed as

$$Q_0(m_1, m_2, m_3) = \sum_i^{N^{uc}} q_i x_i^{m_1} y_i^{m_2} z_i^{m_3} \tag{4.2}$$

where x_i, y_i, z_i are the internal co-ordinates of the ions in the unit-cell and N^{uc} the number of ions in the unit-cell. m_i are integers noting the multipole moments along the different unit-cell vectors, constrained such that for multipole of order $m = m_1 + m_2 + m_3$.

Using this, the multipole moments over the unit-cell can be eliminated by requiring that point-charges are placed in a tetrahedron such that the following equation is solved

$$\sum_n q_n x_n^{m_1} y_n^{m_2} z_n^{m_3} = Q_0(m_1, m_2, m_3) \tag{4.3}$$

where the summation n is over the number of positions in the tetrahedron. The positions of point-charges intended to cancel the mth multipole moment are always taken such that $|x_n|+|y_n|+|z_n| = m$ in the internal co-ordinates of the tetrahedron, under the constraint that x_n, y_n and z_n are integers.

Once a modified unit-cell has been obtained via the GPEE method, the modified unit-cell is repeated in space, along the unit-cell vectors to construct an overall electroneutral superstructure. Since the modified unit-cell is electroneutral and the resulting superstructure has a finite size, the sum in Eq. 4.1 is regularised such that it becomes absolutely convergent with increasing size of the superstructure.

The GPEE method was originally implemented in Molcas under the name EmbQ. As a part of this thesis work, the method was reimplemented in the SCEPIC program system. For most cases, the two different versions of the code should behave identical, the main difference in the SCEPIC implementation is that it was intentionally designed to work with AIMPs and has more features for automatic replacement of classical point-charges with AIMPs.

4.6 Ionic crystals with ab-initio model potentials

Once an electrostatic embedding has been obtained via the GPEE method, the embedding AIMP strategy becomes a natural extension to introduce a repulsive barrier, necessary to prevent electron leakage. This is trivially introduced by replacing point-charges close to the QM region with AIMPs. The underlying theory of AIMPs will be discussed in the next chapter, here, I will only detail how to construct a computational model, suitable for AIMP embedding. For the purpose of this discussion,

the most important point to take from the next chapter is that AIMPs introduce exchange repulsion between the QM region and the embedding environment, which then becomes the model for the repulsive barrier. Some additional details are discussed in Section 5.5, I therefore strongly advise consulting both these sections before embarking on any embedded AIMP studies.

In Papers I and II, it is shown that at distances beyond roughly 6 Å from the border of the QM region, there is no difference in the embedding potential between using AIMPs or classical point-charges. On the other hand, the cost of increasing the AIMP further is in general not that expensive, for which reason I would recommend to increase this distance to at least 10–15 Å, purely as a precaution. Such an approach was used to model MgO, CaO and CaF_2 in Papers I and II. The actual superstructure of one of the MgO clusters used is visualised in Figure 4.2, where, when going from left to right, the QM representation is first given, followed by the QM region along with surrounding AIMPs and finally the remaining classical point charges are visualised.

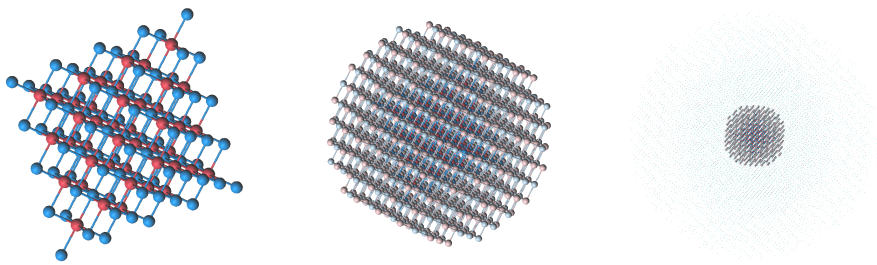

Figure 4.2: Embedding superstructure used to model MgO. From left to right, the QM representation is first given, followed by the QM region along with surrounding AIMPs and finally the remaining classical point charges are visualised.

For the materials that were studied in Papers I and II, the question of the necessary size of the QM region was addressed. The easiest way to express the recommended size is that 1) all elements of the material should be present in the QM region, i.e., no single element in a given material should be represented solely as AIMPs and 2) there should be at least one QM atom layer in between the actual ion or cluster of interest and the AIMP region. The second criterion could be expressed as, if we want to study the local properties of Ni in Ni-doped MgO, as was done in Paper I, then the smallest QM region we might be interested in would be NiO_6^{10-}, since we want the bonds between Ni and O to be present. Figure 4.3 illustrates the convergence of the local properties discussed in Section 3.2 for MgO, which clearly demonstrate the second criterion. Since the entire NiO_6^{10-} region should ideally be well described, one layer of Mg^{2+}-ions should be added, resulting in a total QM region of $NiO_6Mg_{18}^{26+}$.

Additionally, the introduction of the Mg^{2+}-layer ensures that all elements of material are present in the QM region.

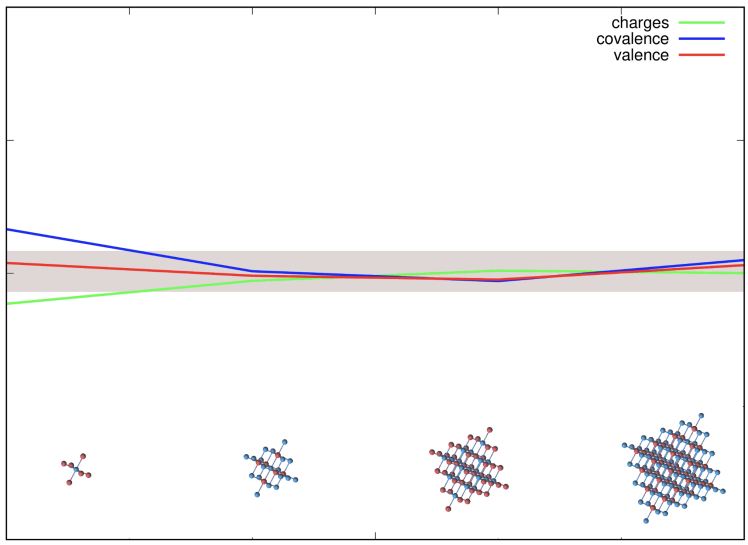

Figure 4.3: Convergence of local properties (computed with PBE) with increasing sizes of MgO clusters. Thresholds are taken as the difference in local properties from periodic calculations using $2 \times 2 \times 2$ and 3x3x3 supercells of MgO.

Chapter 5

Model potentials and the SCEPIC program

In the previous chapter, I discussed the embedding model used in this thesis to mimic ionic crystals. This was realised in the form of ab-initio model potentials. In this chapter, I will discuss the underlying theory of using ab-initio model potentials for embedding ionic crystals.

I will start with was is arguably a detour – AIMPs were first invented as a type of ECP, for which reason I will start the disscuion in this chapter on their usage as ECPs. Thereafter, I will move on to describe group-function theory (GFT) [105], which is the theoretical justification for the embedded AIMP approach. Finally, I will address the implementation of the AIMP parametrisation routine in the SCEPIC program, written as a part of this thesis work.

5.1 Model potentials as effective core potentials

In chemistry, it is a well-known phenomenon that many properties of molecules and solids are primarily described by the so-called valence electrons. Core electrons do not contribute to, for-instance, chemical bonding and valence spectroscopy [86]. The need to describe core electrons with all-electron basis sets in the wavefunction optimisation (discussed in Section 2.1), might be considered to unnecessarily increase the computational cost in certain instances, since they are anyway (usually) frozen in the correlation step. In heavier elements, which are of frequent interest in solid-state chemistry, the "inert" core electrons vastly outnumber the valence electrons. Take

Ce as an example, the neutral species has 58 electrons. Depending a bit on the purpose of the basis sets, either the 4f5s5p5d6s orbitals are kept in the correlation step (ANO-RCC, ANO-R, cc-pVxZ) [46,106,107], with the 4s4p4d orbitals sometimes added (cc-pwCVxZ) [107]. If the larger set is used, only 29 electrons are correlated, exactly half of the total number of electrons. The realisation that explicit treatment of core electrons is unnecessary in many applications, sparked the development of several different approaches for approximating the core electrons. Here, primary focus will be given to AIMPs [86], since they have been the primary topic of research in this thesis. Other, alternative approaches worth mentioning are pseudopotentials (PPs) [108] and projector-augmented waves (PAWs) [109].

To get a conceptual understanding of AIMPs, it is useful to once again start from the atomic orbital shell structure. Taking Al as an example; the electron configuration is $1s^2 2s^2 2p^6 3s^2 3p^1$. Conventional chemistry would take the 1s orbital to be the core orbital, 2s2p as the semicore orbitals (i.e., they are generally not frozen in correlated WFT) and the 3s3p as valence orbitals. Obviously, only treating the 3s3p orbitals, i.e., the valence orbitals, explicitly would be more efficient than describing the full set of orbitals. One way of realising this, as introduced by Huzinaga-Cantu [110,111], is to freeze the core orbitals (based on some atomic HF calculation) and project them out of the explicit calculation.

This can easily be achieved by simply adding a projection operator to the Hamiltonian, $\hat{p} = \sum_k \beta_k \left|\psi_k\right\rangle \left\langle\psi_k\right|$, where the summation runs over the frozen orbitals. The projection constants, β_k, need to be determined in some fashion; the conventional approach in AIMP studies is to set β_k to be $2\epsilon_k$, where ϵ_k are the one-electron canonical molecular orbital energies from HF calculations on the atomic ground state [86]. A strength with the AIMP approach is that the effect from the core-orbitals is not completely neglected and results are frequently equivalent to conventional frozen-core calculations [86]. Obviously, when using AIMPs as ECPs, the number of frozen orbitals should be equivalent to those in a conventional frozen-core calculation. Coming back to the example of Al, this would imply that only the 1s orbital can be safely frozen in correlated calculations [46,112]; clearly, significant improvements in computational efficiency is only realised for heavier elements or for materials that contain a very large number of Al atoms, such as Al_2O_3.

In addition to projecting the orbitals from the calculations, the electrostatic interaction between the core region and the electron in the system is simplified by describing the nucleus and the core electrons as a charge distribution in the form of a contracted series of s-type Gaussians as

$$\hat{V}_{coul}^{AIMP}(r) = \frac{\sum_k c_k e^{-\alpha_k r^2}}{r}, \tag{5.1}$$

where the exponents α_k and contraction coefficients c_k are fitted such that they mimic the original charge distribution of the SCF calculation.

The exchange interaction between the core electrons and valence electrons is likewise approximated as

$$\hat{V}_{exch}^{AIMP} = \sum_{l} \sum_{m=-l}^{l} \sum_{ab} |\psi_{alm}\rangle \, A_{l;ab} \, \langle\psi_{blm}| \,, \tag{5.2}$$

where $|\psi_{alm}\rangle = |\phi_a\rangle \, |\vartheta(l,m)\rangle$; $|\phi_a\rangle$ corresponds to the radial primitive basis and $|\vartheta(l,m)\rangle$ the spherical harmonics. $A_{l;ab}$ is related to the exact exchange operator, $\hat{V}_{exch}$, via

$$\mathbf{A} = \mathbf{S}^{-1}\mathbf{K}\mathbf{S}^{-1}, \tag{5.3}$$

where $\mathbf{S}$ is the overlap matrix and $\mathbf{K}$ has elements $K_{ab} = \langle\psi_{alm}| \, \hat{V}_{exch} \, |\psi_{blm}\rangle$. The fact that AIMPs directly model exchange interaction is important for the upcoming discussion on embedding AIMPs.

5.2 Embedding model potentials

If only the deepest core orbitals, such as the 1s orbtial in Al, can be safely frozen in correlated wavefunction calculations, what is then the reason for discussing model potentials for embedded ionic systems, where the number of ions can number up to several hundreds? The first part of the answer to that question is of a technical nature, since AIMPs effectively are just a set of frozen orbitals, they can be used to represent a frozen SCF state by including both core and valence orbitals in the AIMP. Under the assumption that this frozen state is only weakly perturbed by the presence of other ions, AIMPs can be used to describe an effective frozen environment surrounding a much smaller region, as described by GFT[105].

Since GFT was described in a generalised way in Paper 1, I will only provide a short summary for a hypothetical system of two particles, A and B. According to GFT, if the electrons of the two particles are only weakly interacting, then the total state of the system containing both particles A and B, $|\Phi^{\text{total}}\rangle$, can be partitioned into an antisymmetric tensor product of the wavefunctions of the two particles

$$|\Phi^{\text{total}}\rangle = N\hat{A}\left(|\Phi^{\text{A}}\rangle \otimes |\Phi^{\text{B}}\rangle\right) \tag{5.4}$$

where N is a normalisation factor and $\hat{A}$ is an anti symmetrisation operator. Further, if the two states form an orthogonal set, then it is possible to optimise the wavefunctions independently of each other. Simply by expanding the two group states in their

own sets of molecular orbitals, $\{|\phi_a^A\rangle\}$, that additionally are mutually orthogonal, i.e. $\langle\phi_a^A|\phi_b^B\rangle = 0$, ensures that the states are orthogonal. In practical calculations, the orthogonality condition is ensured by introducing something called orthogonalisation functions (discussed further in Paper I), in general, they are simply the outermost functions of each angular momentum shell frozen by the AIMP. For example, in an AIMP representing Mg^{2+}, the 2s and 2p functions are used. Generally, orthogonalisation functions are only placed on the AIMPs that directly border the QM region, in order to reduce the computational cost.

The GFT theorem lies at the heart of the AIMP embedding strategy. Under the assumption that in an ionic crystal, the individual atoms only interact weakly, embedding AIMPs are optimised, for instance via the SCEPIC program written by me. The embedding AIMPs then represent a group state in GFT, which is inserted into embedded cluster calculations. Since an AIMP retains an approximate description of exchange interaction, it is superior to a pure classical point-charge embedding as it provides a natural barrier between the quantum region and the surrounding[86].

Historically, embedding AIMPs were always derived from HF calculations. In essence, this means that one assumes that a frozen HF state for the host material is a good approximation. Nothing prevents AIMPs to be based on other SCF-type methods, such as KSDFT, for which reason the implementation in SCEPIC also allows for the usage of KSDFT when optimising AIMPs. As speculated in Paper I, this is sometimes advantageous.

In principle, extensions towards methods such as CC or CASPT2 might be advantageous to make the host state explicitly correlated. Currently, this is prevented by a fairly straightforward point, the projection constants (β_k above) are, traditionally, taken from one-electron MO energies. Neither CC nor CASPT2 provide a straightforward and unique way to compute MO energies, for which reason I have not continued on such a task at the present.

5.3 The SCEPIC program

In order to design AIMPs for embedded cluster calculations, there must be some computer code to carry out the necessary parameterisation. Historically, this was achieved via the makeaimp program[86,90], by requesting the authors to perform the parametrisation. As a result, many published studies have been hard to reproduce (see, e.g., Ref. 113), since the used AIMPs were not always made publicly available. In order to parametrise an AIMP, a crystalline field, representing the Madelung potential, must be specified. At the start of this project, this field was computed by the EmbQ pro-

gram[87,104], which was ported to MOLCAS 7.0[112].

Realising that working with different codes written using different programming standards and by different authors made the AIMP parametrisation procedure fairly complicated, a major undertaking of this thesis became the rewriting of these codes into a more user friendly version. This was realised in the form of the SCEPIC program.

SCEPIC is a collection of Julia functions, released as open-source under the Academic Free Licence 3.0, written from scratch by me. It implements a reduced set of the functionality of the makeaimp program (only those directly related to AIMP parametrisation) and is a full re-implementation of EmbQ. The choice of a modern language like Julia, was due to my (potentially naive) belief, that using a modern language will facilitate future development of the AIMP embedding method. The original codes were written in FORTRAN77 (makeaimp) and Fortran90 (EmbQ); for that reason, the execution of these codes might be faster, than their corresponding implementations in SCEPIC. High performance is, however, not a main issue for SCEPIC: i) in each iteration, SCEPIC generally takes less than 0.1 s to generate a set of AIMP coefficients and exponents (see next section) and ii) the generation of a point-charge embedding takes a few minutes – but this is only done once. The next section will discuss the AIMP parametrisation in more detail and specify the necessary equations to solve.

5.4 Parametrisation of embedding AIMPs with SCEPIC

In order to parametrise a set of AIMPs for embedded cluster calculations, self-consistent embedded ion (SCEI) iterations are traditionally used. Taking MgO as an example, one embedding field for a Mg^{2+}-ion and one embedding field for a O^{2-}-ion is made; the embedding field consists of AIMPs and point-charges and should reproduce the Madelung potential of crystalline MgO. SCF calculations are then performed on the two ions (in parallel) and the resulting wavefunctions are used to derive new AIMPs, as illustrated in Figure 5.1. Two points worth mentioning is that in order to parametrise a set of embedding AIMPs, there must be some starting AIMPs provided for the SCEI iterations. In SCEPIC, these are derived from gas phase calculations on the ions, if no starting guess is provided. Further, SCEI iterations are not variational[86,90], for which reason only energy stabilisation is targeted.

For the terms required to construct an AIMP, the projection and exchange operators are taken directly from the resulting (spherically averaged) orbitals. The Coulumbic term is fitted such that the s-type Gaussian in Eq. 5.1 gives the same radial distribution as the frozen orbitals, with the charge corresponding to the total charge of the ion,

i.e., the sum of the electronic and nuclear charges of the AIMP.

Therefore, the object of the SCEPIC program itself is fairly simple i) call some wavefunction solver (in the current implementation, MOLCAS or OpenMolcas), ii) spherically average the occupied orbital shells and iii) fit a set of s-type Gaussians to mimic the radial distribution of the orbitals. Objects i) and ii) are simply a matter of "bookkeeping", while iii) requires some algebra. Note that in the following equations, ζ_k refers to the exponents of a typical pGTO, while α_k refers to the primitive Gaussians used to expand the AIMP charge density (Eqs. 2.1 and 5.1, respectively).

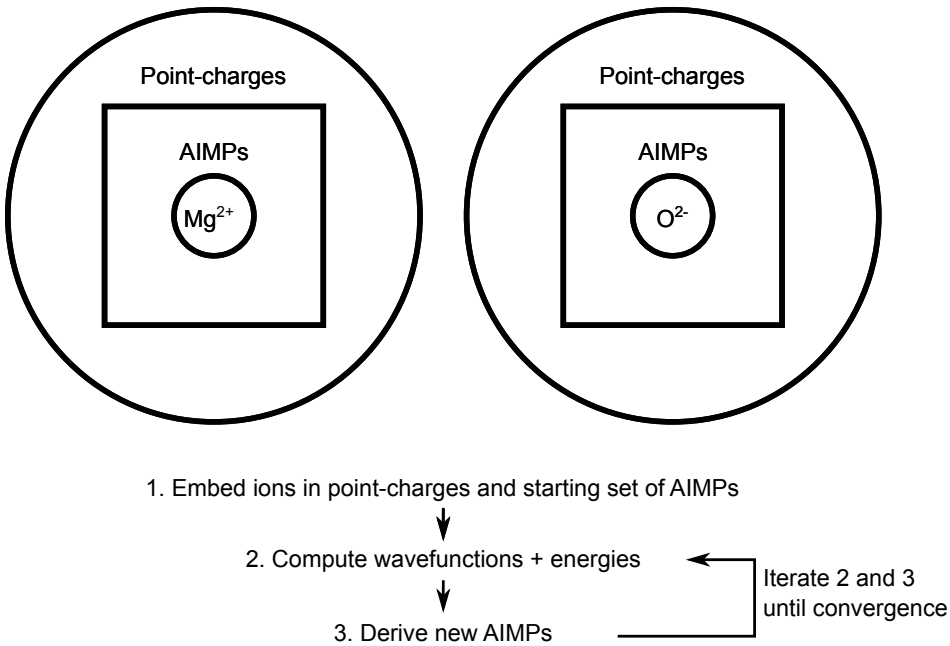

Figure 5.1: Illustration of self-consistent embedded ion iterations for MgO.

First, we take the typical radial part of a pGTO,

$$R_k(r) = \frac{2(2\zeta_k)^{3/4}}{\pi^{1/4}}\sqrt{\frac{2^l}{(2l+1)!!}}(\sqrt{2\zeta_k}r)^l \exp(-\zeta_k r^2) \tag{5.5}$$

For a set of occupied atomic orbitals represented by cGTOs, the radial distribution, or density, can then be obtained as

$$\rho(r) = r^2 \sum_i n_i \left| \sum_k c_{i,k} R_k(r) \right|^2 \tag{5.6}$$

with n_i being the occupation number and $c_{i,k}$ the contraction coefficients. Using this density, we can numerically compute some Coulomb integrals, I_m^{coul}, on a radial grid with points, r_m. Since AIMPs approximate spherical symmetry[86], r_m, is a scalar, denoting the distance of the grid point from the nucleus. Some decision has to be made on some cutoffs, in order to ensure numerical stability. In SCEPIC, the same procedure as proposed in the makeaimp program is used[86,90]. Only integrals larger than a certain threshold (10^{-4}) are retained in SCEPIC, these are then further weighted by some factor $w_m = \frac{1.0}{N^{\text{elec}} r_m}$, with N^{elec} being the number of electrons frozen by the AIMP. The weight factors are normalised by $n_m = \frac{N^{\text{points}}}{\sum_m w_m} * w_m$, with N^{points} being the number of points retained in the fitting.

Now, the object is to ensure that the AIMP Coulomb operator (Eq. 5.1) reproduces this set of integrals. This can be done using basic linear algebra, by noting that

$$\sum_k c_k \exp(-\alpha_k r_m^2) \approx Z^{\text{core}} + I_m^{\text{coul}} \tag{5.7}$$

where the left-hand side is simply Eq. 5.1 and the right-hand side the ionic point-charge potential. Both the exponents α_k and the coefficients c_k are unkowns at the start of the SCEI iterations. In SCEPIC, the exponents of the underlying s-type GTO basis set are used as a starting guess. These are then varied by making small alterations, in SCEPIC, this is generally done in step size corresponding to 1 % of the exponent size, i.e., if the exponent is 100.0, it is varied by $\pm$1.0. Each time the exponents are varied, the set of coefficients c_k are determined by recasing Eq. 5.7 into matrix form and diagonalising.

5.5 Some key points of embedding AIMPs

This chapter will end by stating some key points, important to any user who wants to model ionic solids with AIMPs. Some important points were already mentioned in Section 4.6, thus, parts of this section will be repetitious, if this thesis is read in full. First and foremost, when running SCEI iterations, some crystalline field must be specified. Take TiO_2 which has three common polymorphs; rutile, anatase and brookite. The crystalline field, or Madelung potential, in all of these three materials is different. It is therefore not guaranteed, that a set of AIMPs parametrised for rutile will be suitable for usage in anatase or brookite. Similarly, in materials where the same element appears in several Wyckoff positions, each position requires a unique set of AIMPs. To the best of my knowledge, no extensive test of the transferability of AIMPs has been made.

On the other hand, this could also be argued to be a strength of the AIMP approach. In contrast to the similar capped ECP model[88,89], where pseudopotentials are taken from some ECP-type basis set intended for molecular calculations, the embedding AIMPs are directly parametrised for a given material. Further, the capped ECP model can only be used to replace cations, whereas AIMPs can be used to represent anions.

One very important point is that the improvement from using an AIMP is fairly short ranged; around 6–8 Å surrounding the QM cluster is sufficient. The remaining part of the Madelung potential should (for the sake of efficiency) be represented by classical point-charges (for instance, from the GPEE model, Section 4.5). This leads to a potential problem with AIMPs; an embedding AIMP represent a set of frozen electrons, only integer charges are therefore natural to represent with AIMPs. This can be problematic if it expected that a reduced charge might be more representative of the electrostatic potential (for instance, ± 1.8 e in MgO[114]). From the little practical experience I have, I can only say the following, that if the guidelines for AIMP embedding advocated by me are followed (Section 4.6), the effect of only scaling the classical point-charge region is minor. Taking the electron excitation energy for MgO, which I will discuss more in-depth in Section 6.5.1, as an example. With CASSCF I obtained a value of 13.071 eV, when using HF-based AIMPs and a point-charge field derived from ± 2.0 e charges. If the classical point-charges are scaled to ± 1.8 e, the excitation energy remains virtually unchanged at 13.067 eV. Such small differences is on the same order as typical "chemical accuracy", essentially, there is no impact from scaling the classical point-charges. Clearly, when working with AIMP embedding, the proper target for improving the model is the AIMPs themselves.

Chapter 6

Practical aspects on the usage of multiconfigurational theory

In this final chapter of the thesis, I will provide an overview of the usage of multiconfigurational methods in QC. These methods are generally not usable as "black-boxes", they are multidimensional, with several potential minima that can spoil the calculations. The proper usage of methods which fall within this category therefore requires a strong understanding of both the underlying theory, as well as the particular specifics of the software that is being used. As such, this chapter will address some of the underlying, sometimes hidden, problems that might occur when working with multiconfigurational theory.

I will make an overview of different ways that has been discussed in the literature for selecting the so-called active spaces. These methods are often a good starting point for any new user of multiconfigurational methods, but have certain problems that I will address in this chapter. Further, I discuss the method I usually work with when selecting active spaces, which I hope can be of use to new users of multiconfigurational quantum chemistry.

Since this is a thesis on combining multiconfigurational quantum chemistry with ionic solids, the final discussions in this chapter is devoted to this combination. The discussion is based on my own experiences, either directly when modelling ionic crystals with AIMPs, or derived by taking into consideration related inorganic compounds. Some of it will be speculative and I make no claim that my thoughts on active spaces for ionic solids should be taken as absolutes.

The thoughts and comments in this chapter guided the results obtained in Papers I,

IV and V. Some of the text here presents slight improvements on the procedure used to obtain the results in those papers. After all, it seems as if it is only after a project is finished, that you truly understand how you should address it. The emphasis in this chapter will most likely be somewhat biased towards the implementations of these methods in the MOLCAS and OpenMolcas packages[115,116,117], as these were the electronic structure codes used in this thesis.

6.1 On Active Spaces

When reading papers on the usage of MCSCF theory, it is easy to get the impression that the selection of active space is something trivial. In reality, unless the system under study is particularly simple, choosing an appropriate active space is usually the most time consuming part. By my own estimation, I would say that more than 70 % of the total time spent on projects has been related to the selection of the active space. Sometimes due to problems related to "guiding" the convergence to the desired active space. Other times, due to the simple fact that what might look like an appropriate active space, is insufficient when used with PT2 for dynamical correlation. This chapter will discuss some of the common pitfalls when using MCSCF and provide some suggestions on how to simplify the active space selection process. In truth, however, only by either improving the dynamical correlation step, or by making methods for targeting larger active spaces more feasible, will provide an ultimate solution to all of these problems.

6.2 Static and dynamical correlation

The success of multiconfigurational theory lies in the separation of electron correlation effects into two distinct parts, so-called static correlation and dynamical correlation. But to be perfectly clear, this separation is solely for reasons of practicality, as the mathematical tool for describing both is by including more electronic configurations in the wavefunction. In general though, static correlation require a higher degree of accuracy to be appropriately accounted for, whereas dynamical correlation can be treated in a more approximate fashion[1,2]. Static correlation arises when multiple electronic configurations carry equal, or at the very least substantial, weight in a full-CI wavefunction. In such cases, the underlying assumption of HF theory, namely that a single SD is sufficient, breaks down and a multiconfigurational solution becomes mandatory. Dynamical correlation is instead an effect which describes the missing portion of electron-electron repulsion effects, which are only treated as a mean-field in HF theory. Such electron-electron repulsion effects are in correlated wavefunction

theory usually modelled as very large number of double (or higher) low weight excitations from the set of occupied orbitals into the virtual orbital space, obtainable via methods such as PT, CC or CI, for example.

Therefore, in multiconfigurational theory the usual task is first to select an appropriate reference wavefunction, with the CASSCF method being the historically most utilised method. Once an appropriate wavefunction has been selected, however, dynamical correlation has to be included. Here, CASPT2 is most likely the most commonly used method.

It is frequently a relatively trivial task to select an active space which will include the electron configurations with the largest weights, as these usually only involves quite few orbitals. At the very least for molecular systems. When treating dynamical correlation with PT2, however, it must be kept in mind that PT2 is only able to approximate up to double excitations. This is often the dominant portion of dynamical correlation, but in cases where higher excitations are of qualitative importance, all orbitals between which higher excitations are important have to be included into the active space to provide qualitatively correct results. Arguably then, using methods that also approximate triple excitations better, such as PT3, might yield better results than CASPT2 with smaller active spaces. Better methods for incorporating additional excitations than PT, such as MRCI or MRCC, could also be useful. As of writing this text, such methods are not yet readily available, at the very least when working with systems containing basis functions in the range of 1000–2000, for which reason the remainder of this chapter will focus on strategies when working with CASPT2 for dynamical correlation.

6.3 Overview of different approaches for active space selections

The art of selecting active spaces is hardly something new to this thesis work. Many different approaches can be found in the literature, ranging from approaches based on "chemical considerations" to attempts on more automatic ways, based either on MP2 natural occupation numbers or orbital entanglement calculations. This section will provide a small overview and some criticism of ways of selecting active spaces that are sometimes found in the literature.

6.3.1 Chemical considerations

Here, the selection of active orbitals is based on what is, arguably, chemical criteria. This approach is closest to the approach that I usually employ. In fact, it would be fair to say that "my" approach is derivative of the approach discussed in this section. Therefore, many details for selecting orbitals based on this approach will be reiterated in Section 6.4. Much like the approach I advocate, the active space selection process discussed here is intended to first be carried out in a small basis set, before expanding the wavefunction into a larger basis set for production calculations.

Selecting active spaces by this approach arguably evolved out of Björn Roos' group, with Refs. 2,118 providing more detailed overviews. The idea is fairly simple, preliminary calculations are made in a small basis set, usually of VDZP quality, since this generates more chemical virtual orbitals. For different groups of atoms, orbitals are selected based on a predefined rule. As an example, (taken from Ref. 118), for second row elements, the 2s and 2p orbitals should be active for elements Li–C, while for elements N–F it is sufficient to only include the 2p orbitals. Similar rules can be found for all elements all the way down to the actinides. Thus, this approach gives a fairly impressive start for any new user of multiconfigurational quantum chemistry.

So what is wrong with this approach? In Sections 6.4.3 and 6.4.4, I give some small overview of a few calculations on one excited state of Be and the electron affinity (EA) of Cl. For these two atoms, the approach discussed in Refs. 2,118 would suggest that I should only include the 3s3p orbitals of Be and only the 3p orbitals of Cl. What I find, however, is that in both cases, there is fairly strong angular correlation (to be discussed in Section 6.4.1), meaning that the 3d orbitals are important for quantitatively good results. This set of orbitals would be completely missed if the approach here is followed in absurdum.

To be fair, however, the approach in Refs. 2,118 are mainly targeting molecular systems, where my own experience indicates that such angular correlation is of less importance than in an atom, which is the target of Sections 6.4.3 and 6.4.4. Lets take an example of a crystalline calculation where this approach fails then. For lanthanides, Refs. 2,118 suggest that the most important orbitals are the 4f5d6s orbitals. Some preliminary calculations that I have done on Ce-doped YVO_4 (Ce:YVO_4), however, suggests that, at the very least for a formal Ce(III), the 5p-orbitals are non-negligible. Table 6.1 gives two set of computed spectra of the $4f^15d^0 \rightarrow 4f^05d^1$ excitations in Ce:YVO_4, modelled with a minimal QM region of only CeO_8 in an AIMP and point-charge embedding. The difference between including the 5p-orbitals or not is quite dramatic for the 2 B_2 state in the $4f^05d^1$-manifold, which is lowered from a predicted excitation energy of 8.25 eV to 3.59 eV. Noting that the lowest excited state measured experimentally is around 3 eV[119], this is a clear improvement. For this reason, my

opinion is that the guidelines provided by Refs. 2,118 should primarily be taken as a starting point, with a large degree of scepticism for elements with in high valence states (such as Ce(III) here) and for excited states.

Table 6.1: Ce:YVO_4 state energies as computed using different active spaces and ANO-RCC-VQZP basis set. Active space orbitals are taken from Ce. States ordered according to the 5p4f5d active space. Energies are in units of eV.

State	4f5d	5p4f5d
$4f^1 5d^0$-manifold		
1 B_2	0.00	0.00
1 A_1	0.02	0.00
1 E	0.31	0.13
2 E	0.31	0.26
3 E	0.43	0.27
A_2	0.15	0.38
4 E	0.49	0.42
$4f^0 5d^1$-manifold		
2 B_2	8.25	3.59
5 E	4.93	4.92
6 E	5.11	5.13
2 A_1	5.31	5.46
B_1	7.47	7.46

6.3.2 Natural orbitals

Another approach is to make a preliminary calculation using a single-reference method, e.g. MP2, and select orbitals based on the natural occupation numbers. Typically, natural orbitals with occupation number in the range of 1.98 – 0.02 should be added to the active space[2].

Using MP2 natural orbitals seemingly works very well when a sensible reference wavefunction can be obtained, in which case my manually selected orbitals for Be and Cl (discussed in the previous section and Sections 6.4.3 and 6.4.4) and the orbitals predicted by MP2 correspond to the same active space. This would be 3s3p3d for Be (excited state calculation) and 3p3d for Cl (electron affinity calculation). Note that my ground state calculation for Be (Section 6.4.3) suggest that 3s3p is good enough for the ground state and that the 3d orbitals are only important for the excited state. For the ground state of Be and Cl, MP2 natural orbitals do indeed suggest active spaces of 3s3p and 3p (Table 6.2), respectively, the same orbitals that the rules given in the previous section would suggest. With Cl, it is trivial to also make a ground state MP2 calculation of Cl^-; when doing so, MP2 does, in fact, predict that the 3d-orbitals should be added to the active space. On the other hand, making an excited state HF wavefunction for Be is non-trivial and even if it would be straightforward to do so,

MP2 is likely to diverge, making it very hard, if not impossible, to generate a reliable set of MP2 natural orbitals when studying excited states.

Table 6.2: Table of MP2 natural occupation numbers for Be and Cl.

Species	MP2 natural occ. num.		
	3s	3p	3d
Be	1.95	0.02	0.00
Cl	1.99	1.67	0.00
Cl^-	1.99	1.96	0.02

MP2 is, however, not the only possible choice of natural orbitals. Natural orbitals can be derived from either CC, CASPT2 or, in the case of open-shells, from unrestricted HF or DFT. For closed-shells with restricted HF or DFT, natural orbitals are undefined since all orbitals either have occupation number 2 or 0. Since unrestricted HF or DFT do not explicitly correlate orbitals, these methods might unfortunately miss some important contributions, for which reason I would advocate the use of wavefunction methods when computing natural orbitals. In the case of excited states, CC or CASPT2 are most likely better choices than MP2 (albeit, significantly more expensive to obtain), while DFT can be more efficient in the case of large molecular or crystalline systems. One general problem that I find with using natural orbitals, to be discussed in the next paragraph, is expected to be fairly similar for all methods.

Using natural orbitals can provide a promising alternative to the more manual selection in the previous section. But if the system starts to become large, some manual guidance is most likely going to be required anyway. Take the lytic polysaccharide monooxygenase (LPMO) enzyme studies in Paper IV, here there is both a transition metal (Cu) and several aromatic rings in the QM region. Without a doubt, MP2 natural orbitals will suggest a large number of orbitals from both parts of the system, making the calculation unfeasible. On the other hand, the targeted property (singlet–triplet energy gap) in Paper IV can be considered to be mostly located at the Cu, with the orbitals on the aromatic rings being fairly unperturbed by the spin state. In this case, there is very little reason to expect that adding more orbitals from the surrounding cluster will significantly alter the predicted singlet–triplet energy gap, despite what their MP2 occupation numbers might indicate. For this reason, using MP2 as a start might be considered to be overshooting the problem. But then again, perhaps the MP2 natural orbitals are easier to identify, than manually inspecting all orbitals? I can not provide any clear conclusion here other than stating that in future studies I am more likely to start adding MP2 (or CASPT2 for excited states) natural orbitals to my active space selection routine, simply due to the convenience of having some starting natural occupation numbers.

For generating preliminary CASPT2 natural orbitals, I would suggest following the rules discussed in Section 6.3.1, truncating the number of active orbitals to the smallest possible to get a reasonable starting wavefunction. Taking two examples from my own papers, for the LPMO enzyme in Paper IV, or for Ni-doped MgO in Paper I, my preliminary suggestion would be to select the d-orbitals for Cu and Ni, respectively. After obtaining a set of CASPT2 natural orbitals, a manual selection can be done. If it is possible to include all orbitals of importance (occupation numbers 1.98 – 0.02), then do so. Else, select orbitals that seem to be more localised in the region of the system that is of interest.

6.3.3 Automatic approaches

A fairly recent addition the active space selections is the automated approach of Stein and Reiher[120]. This approach is based on orbital entanglement measurement from a preliminary DMRG calculation. In this preliminary DMRG calculation, as many orbitals as possible are included into the active space, using a low m-value. The wavefunction is allowed to partially converge, after which an orbital entanglement measurement is performed. Based on some threshold, for which suggestions are provided in Ref. 120, the active space is reduced until it becomes computationally feasible, be it that is becomes so small that CASSCF can be used or that RASSCF or DMRG is still needed.

While I personally like the approach, it suffers from one big drawback for crystalline systems, the preliminary DMRG calculation to measure the entanglement will most likely contain too many orbitals if done in a naive way, forcing the user to anyway make a selection of some subset of orbitals. Take the LPMO system from Paper IV as an example again. Lets make a naive assumption of the LPMO backbone and ignore the active (O, O_2, OH) part. The backbone contains one Cu, one O, 14 C, five N and 19 H. We can probably ignore any orbitals originating from H in the backbone, but that still leaves us with 21 atoms to consider. From the occupied part, we can imagine that we want to include at the very least the orbitals suggested from Refs. 2,118; 10 d-orbitals for Cu (double shell) and three 2p-orbitals from each O, C and N. This gives 70 orbitals, without even starting to consider more virtual orbitals than the double shell on Cu. Clearly, this is a road to nowhere, if followed in absurdum. The automated approach will therefore still require that a preliminary selection of orbitals is made, for instance based on more stringent manual selection or more arbitrarily by some energy criterion from HF one electron energies. Even if this is done, my suspicion is that it will suffer from the same drawback as discussed in the end of the previous section; orbitals which are relatively unchanged between the singlet and triplet spin states will become included in the active space.

It should also be pointed out that, in principle, the approach of Stein and Reiher could (in some future, hypothetical, implementation of the method) be based on, say, CASPT2, instead of DMRG. Following the discussion in the last paragraph of Section 6.3.2, once the preliminary CASPT2 calculation has been done, the CASPT2 natural orbitals that are most entangled could then be added to the active space, as appropriate. Such a solution would essentially combine the three approaches discussed in this section.

6.4 Selecting an active space – How I work with multiconfigurational theory

This section will address more directly thoughts and concepts that I find useful to follow when selecting active spaces manually. Though I would like to stress that, in my future work, I am quite likely to extend on this approach, on the basis of the discussions in the previous section. Still, as I believe that for large systems, such as crystals, some manual selection will always be required, I hope that the discussion in this section can be used as guidelines for how to truncate an active space based on, say MP2 or CASPT2 natural orbitals. In the end of the section, I will demonstrate two practical examples of calculations on the Be atom and the Cl anion.

6.4.1 Correlating pairs of orbitals

The most common strategy employed by me when selecting active spaces, at least when studying ground state phenomena, is to think in terms of "correlating pairs" of orbitals. The origin for such a strategy comes from the concepts of radial and angular correlation [1]; these are two aspects of dynamical correlation. I find it simplest to think in terms of s-type orbitals when conceptualising these two phenomena.

Starting with an ns-orbital that is doubly occupied in HF theory, this orbital will usually be too contracted and the two electrons ending up too close to each other. By adding an $(n+1)$s-orbital into the active space, which contains an additional radial node compared to the ns-orbital, the probability of finding the two electrons at different distances from the atomic nucleus is increased. This is an example of radial correlation. The correlation effects between a doubly occupied bonding orbital and an unoccupied antibonding orbitals is also an example of radial correlation, such as the σ_g-orbital and the σ_u^*-orbital in the H_2 molecule at the equilibrium geometry, with the probability of finding the two electrons at different positions along the internuclear axis increasing. At the same time, it should be stressed that along the dissociation

curve of H_2, the nature of the correlation between the σ_g and the σ_u^* orbitals changes from dynamical near the equilibrium, to static as the dissociation limit is approach.

Angular correlation, on the other hand, arises when the active space is augmented by adding orbitals of higher angular momentum; usually, for an orbital of angular momentum l, a set of $(l+1)$ orbitals are added. Taking our example of an s-orbital, a set of p-orbitals would be added. This has the effect of increasing the probability of finding the two electrons in the s orbital on opposite sides of the atom. Likewise, a σ orbital can be correlated with a π orbital, increasing the probability of finding electrons on opposite sides of the intermolecular axis.

One useful guideline for correlating pairs of orbitals is that, most of the time, the sum of their natural occupation numbers is equal to two[1,2]. Thus, if some preliminary natural orbitals have been obtained, yet the shape of the orbitals is still unclear, my bet would be on checking the natural occupation numbers.

Frequently, my experience with MCSCF calculations have suggested that radial correlation is more important to include into the active space than angular correlation, when dealing with molecular systems. It would, however, be folly to suggest that including angular correlation in the active space is always negligible. This will be clearly demonstrated in Sections 6.4.3 and 6.4.4.

6.4.2 Basis set considerations – start small then go big

Textbooks often give the impression that the proper way of choosing an active space is to first make an HF calculation on the targeted system and then select the orbitals around the HOMO-LUMO level to put into the active space. While nice in theory, such an approach suffers from some problems for larger systems. The first is due to the "undefinitiveness" of the virtual orbitals. While occupied HF orbitals are usually chemically sensible, virtual orbitals will only be so due to mere chance, as they are not optimised (see Section 1.1). As most quantum chemists who have ever inspected the virtual orbitals from a calculation with a basis set that contains a large amount of diffuse function will know, the shape of virtual orbitals are usually nonsensical. All multiconfigurational results in Papers I, IV and V were obtained by the procedure described in this section.

When making a selection of virtual orbitals to include in the active space, I find that it is usually better to make preliminary calculations with a small basis set. While I argued in Paper V to use specially designed contraction levels, at least when working with ANO-type basis sets, I have slightly altered my strategy since. Today, I would make the case that one should always start by using a double zeta contraction with

polarisation functions (VDZP) when making a preliminary HF calculation and follow that by localising the virtual orbitals. This simply due to the fact that a typical VDZP basis set will contain all necessary orbitals for both radial and angular correlation, so it requires less tedious handywork with setting the basis set contraction levels.

Many schemes for localising orbitals exists and I will make no claim that any is truly superior to another. As a personal preference, however, I like the "improved virtual orbital" strategy of Widmark (IVO keyword in Molcas)[112]. With IVO, virtual orbitals from a preliminary HF, or possibly a minimal CAS, calculation are localised by diagonalising the one-electron Hamiltonian within the virtual orbital subspace, while maintaining orthogonality with the occupied orbitals. The resulting virtual orbitals are very compact and are usually chemically meaningful.

Once the targeted active space has been obtained by the means described above and converged with a small basis set, e.g. VDZP, the basis set should be expanded to obtain good results with CASPT2. In the ideal world, it would always be possible to expanded the basis set to (at the very least) VQZP quality on all atoms. For most typically studied systems, it is fairly unlikely that larger basis sets that VQZP will substantially improve the results when working with CASPT2. When dealing with crystals, this is typically not achievable and some compromise has to be made. Based on the discussions in Paper III and some practical experiences that are not published, I would make the argument that the basis set should at the very least be of VTZP quality when working with CASPT2. Using VDZP basis sets on some atoms that are very remote from the part of the system that is of interest – say a dopant in a crystal – might be permissible.

6.4.3 A case study of starting orbitals on Be

In this section, we will discuss how starting orbitals might influence the outcome of a CASSCF calculation, as well as how the active orbitals are chosen. For our test case, we will look at the Be $^1S(2s^2)$ and Be $^1P(2s^12p^1)$ states. The results are based on a minimal active space made from either the Be 2s2p3s orbitals or the Be 2s2p3d orbitals, using ANO-RCC-VDZP as a basis set. Since we are interested in four states (when accounting for the triply degenerate ^1P state), the calculations discussed in this section are based on a state-average CASSCF (SAV-CASSCF) solution over four roots. No symmetry will be used here, for the very simple reason that if symmetry were to be used, some of the errors here might be solved automatically. This section wants to focus on giving some hints at what to do when such a simple solution might not be feasible.

We will start by discussing solutions for the 2s2p3s active space. When working with

CASSCF we need a set of starting orbitals. Commonly, these are taken from an HF solution of the ground state. Alternatively, semi-empirical orbitals can be used, for instance those generated by the GuessOrb program in Molcas[112]. We will start with discussing the difference between GuessOrb and HF when making SAV-CASSCF calculation, without any manual inspection of the orbitals. In other words, we make the assumption that with both sets of starting orbitals, the orbitals will initially be ordered as typical atomic orbitals, 1s2s2p3s3p etc. In Table 6.3, it is clear that we get two very different solutions when using either GuessOrb or HF with an naive selection of the active space. Both solutions show a clear symmetry split in the excited $^1P(2s^12p^1)$-state and the differences in total energies between the two approach are on the order of 0.1 E_h, which is far beyond reasonable. So what happened? Analysing the composition of the active orbitals after convergence for the two solutions reveals the answer. In both cases, the 3s orbital was not included in the active space after convergence. Instead, a single 3d orbital was picked up during the SAV-CASSCF iterations. Since for an atom, the 3d-orbitals are part of a five-fold degenerate set, this leads to symmetry breaking.

Obviously, a naive approach to selecting active orbitals with CASSCF is prone to danger. If we now instead make a manual selection of the desired active space, by either inspecting atomic orbital coefficients or by visualising the orbitals, both the GuessOrb and HF set of starting orbitals converge to the exact same solution. Alternatively, we can take the first set of "failed" results as a guide to selecting the active space; instead of assuming that the 3s-orbital is a good addition to the active space, we attempt to naively take the full set of 3d-orbitals. As illustrated in Table 6.3, this also solved the symmetry breaking problem. Two more effects are also clearly visible from adding the 3d-orbitals, i) the energy of the ground state is nearly unchanged and ii) the energy of the excited state is drastically lowered. Therefore, it seems as if angular correlation is non-negligible and that the 3d-orbitals should be included in the active space. Importantly, results obtained with either the GuessOrb orbitals or the HF orbitals converge to the same solution, once the active orbitals had been selected in a sensible way. It is therefore up to the user to decide which approach works best for a given problem.

It should be emphasised that this example is very simple, in the sense that the symmetry of an atom allows a simple diagnostic of what went wrong with the initial solution, as well as manual inspection of orbitals to add to the active space. For molecular or crystalline systems, such a luxury is rarely present. Instead, very careful analysis of the obtained solution – convergence patterns, CI coefficients, orbital cofficients, atomic charges – has to be performed, ideally along with testing a few different active spaces.

Table 6.3: Table of total energies for the first four singlet roots of Be. Note that roots 2–4 are supposed to be degenerate in a proper solution. SE stands for semi-empirical starting orbitals, generated by GuessOrb. HF indicates that orbitals obtained from a converged HF solution on the ground-state was used as starting orbitals. Active space orbitals are given in the headers. Naive means that the results were obtained by assuming that the order of the orbitals was correct without any inspection. Manual means that the orbitals were visualised and selected by hand. Large is identical to naive, except that the number of orbitals were larger.

CASSCF root	SE 2s2p3s (naive)	SE 2s2p3s (manual)	SE 2s2p3d	HF 2s2p3s (naive)	HF 2s2p3s (manual)	HF 2s2p3d
Root 1	−14.56735989	−14.61662380	-14.61526090	-14.61522721	−14.61662380	-14.61526090
Root 2	−14.26634626	−14.36849575	-14.38588222	-14.375660679	−14.36849575	-14.38588222
Root 3	−14.26615077	-14.36849575	-14.38588222	-14.37085881	−14.36849575	-14.38588222
Root 4	−14.26610371	−14.36849575	-14.38588222	-14.37022138	−14.36849575	-14.38588222

6.4.4 A CASPT2 case study on the electron affinity of Cl

In this section, a practical example of using CASPT2 to compute the EA of Cl will be discussed, a major undertaking in Paper III. The aim is to highlight how results might change by using different settings and that the selection of active space is not necessarily obvious. When Cl binds an additional electron to form Cl^-, the electronic state changes from a $^2P(3p^5)$ state to a $^1S(3p^6)$ state. Using "conventional" arguments for active spaces, the 3s orbital will be kept outside the active space and only be correlated with CASPT2[2]. Four different active spaces will be evaluated. The first will only include the three 3p orbitals, i.e., a "valence" active space. It should be stressed that for the anionic state, this gives a singlet (6,6) active space and therefore the wavefunction contains only a single configuration and is equivalent to an HF calculation, followed by CASPT2. Ideally, this should be identical to an MP2 calculation, though due to minor differences in the reference Hamiltonian, this is rarely the case. The second active space is based on the shell structure of atomic orbitals and includes, in addition to the 3p orbitals, the 4s orbital. The third active space was designed to include radial correlation directly and therefore also correlates the 3p orbitals with the 4p orbitals, excluding the 4s orbital. The final active space instead considers angular correlation directly, thus the 3p orbitals and the 3d orbitals are included in the active space. These active spaces will be referred to as AS-I, AS-II, AS-III and AS-IV, respectively (tabulated in Table 6.4). In addition, two different settings for frozen orbitals will be used; in the first, the 1s2s2p orbitals will be frozen, as was the recommendation when the ANO-RCC basis set was developed. In the second, only the 1s2s orbitals are frozen, which is a more "modern" definition of the Cl core orbitals, used for instance in the development of the ANO-R basis set. The reference value of 3.61 eV was taken from Ref. 121.

In order to deal with the degenerate $^2P(3p^5)$ state, all energies were obtained by SAV-CASSCF followed by the XMS-CASPT2 method. No IPEA shift was used to obtain the results presented herein. ANO-RCC-VQZP was used as basis set, with the second-order Douglas-Kroll-Hess scalar-relativistic Hamiltonian (DKH2), since

ANO-RCC was developed to be used with DKH2.

Table 6.4: Table summarsing the active space nomenclature used in this Section.

Active space	Active orbitals
AS-I	3p
AS-II	3p4s
AS-III	3p4p
AS-IV	3p3d

In general, the results in Table 6.5 indicates the following; using AS-I and AS-II with 1s2s2p frozen core gives no marked changes in results, all predicting an EA of around 3.35 eV, which would correspond to an error of about 0.26 eV. Unexpectedly, with AS-III the prediction was worsened by around 0.1 eV, with respect to AS-I and AS-II. This is in contrast to the CASSCF results, from which one might expect that AS-III should outperform AS-I and AS-II. AS-IV comes closest to the experimental value, with a predicted EA of 3.48 eV. This would fit with both the EA's obtained with CASSCF as well as the CASSCF natural orbital occupation numbers given in Table 6.6, which indicate that the 3d orbitals carry a larger electron density than either the 4s or 4p orbitals. Though many might argue that the occupation number for the d-orbitals (c:a 0.02) is so small that these orbitals can safely be excluded. On the other hand, the occupation numbers suggest that the 4p orbitals should be more important than the 4s, yet the results worsened by adding these orbitals. This experience shows the importance of evaluating multiple active spaces before reaching a final conclusion.

Can we find an answer as to why the results with AS-III was worse than the others, despite that both the CASSCF EA and the occupation numbers suggest otherwise? As it turns out, the answer lies in the frozen core using with CASPT2. By additionally correlating the 2p orbitals at the CASPT2 level, the EA predicted by AS-III reaches 3.36 eV, on par with AS-I and AS-II, regardless of frozen core. No major change was observed for AS-IV by changing the frozen core. This small experiment shows then that the balance between correlation treated at the SCF level and the PT2 level needs to be considered carefully. It should also be stressed, that the choice of frozen core orbitals should always be based on what type of correlation effects where included in the design of the basis set[46]. In principle, this means that correlating the 2p-orbitals with ANO-RCC is inadvisable. Though I have frequently found that for light elements, H–Cl, correlating too many core orbitals does not seem to introduce any obvious errors when using ANO-type basis sets.

Table 6.5: Computed EAs for Cl using different active spaces and methods. All EAs are given in units of eV. Orbitals in paranthesis following CASPT2 denotes the frozen core. Values in paranthesis after EAs denoted difference with respect to reference value of 3.61 eV [121]

Active Space	CASSCF	CASPT2 (1s2s2p)	CASPT2 (1s2s)
AS-I	2.56 (-1.05)	3.36 (-0.25)	3.35 (-0.26)
AS-II	2.58 (-1.03)	3.37 (-0.24)	3.35 (-0.26)
AS-III	3.01 (-0.60)	3.24 (-0.37)	3.36 (-0.25)
AS-IV	3.19 (-0.42)	3.48 (-0.13)	3.46 (-0.15)

Table 6.6: CASSCF natural occupation numbers for the netural $^2P(3p^5)$ and anionic $^1S(3p^6)$ states of Cl. Corr. (short for correlating) is used as a collective name for the different virtual orbitals added into the active spaces, as given in Table 6.4. ** denotes that there are no correlating orbitals in AS-I.

Orbital	AS-I	AS-II	AS-III	AS-IV
		$^2P(3p^5)$		
3p	1.667	1.666	1.661	1.640
Corr.	**	0.001	0.006	0.016
		$^1S(3p^6)$		
3p	2.000	1.999	1.987	1.966
Corr.	**	0.002	0.013	0.021

On the flip side, it would be a reasonable view that all of the predicted EA's are "good enough" as they are qualitatively correct and gives errors that, for most test cases, are less than 0.3 eV. So why worry? Unfortunately, in the correct assignment of spectra, particularly for d-metals or f-metals, the density of excited states is very high, with very small differences in energy. One example of this would be Ni-doped MgO, which was studied in Paper I, where the difference between the first excited ^1E_g state and the first excited $^3T_{1g}$ states is around 0.2 eV. The ability to influence the CASPT2 results by around 0.2 eV by choosing different active spaces and frozen core (as in the case of Cl^-) will therefore impact the predicted state ordering, which, depending on the topic, may or may not be of crucial importance.

6.5 Active Space suggestions for various kinds of metal oxides

In this section, I will discuss the selection of active spaces for various kinds of metal oxides. While I would expect similar considerations to be true of related types of ionic materials, such as other metal chalcogenides in general or metal halides, i.e., materials with anions that gives significant degrees of either σ-donation or π-donation, I will strictly speak of metal oxides, as that is the type of materials I have mostly worked with. The section will start with a discussion on "simpler" materials, i.e., where the metal is either an s-metal (alkaline metal, alkaline earth metal) or a p-metal (any metal from the p-block); the discussion will be examplified by the band gap of MgO. It

should be mentioned that CaF_2, an s-metal halide, was studied in Paper I using the methods discussed for MgO. From there, some discussions on non-doped d-metal oxides and f-metal oxides will be given. Such a discussion will primarily be speculative and advisory for anyone wanting to attempt AIMP studies of such metal oxides, as the results I myself have achieved this far in that subject are minor and not yet ready for peer-reviewing. The section will conclude by discussions on active spaces for low-concentration dopants in a host metal oxide material. The section will conclude by some suggestions for how to select active spaces surface reactions.

6.5.1 Electronic excitations in MgO

Lets start by making a revisit to electron excitations MgO, which was studied in Paper I. As argued in Paper I, this is the most straightforwards way with multiconfigurational theory to study what solid state physicists usually refer to as the band gap. It is essentially the energy required to remove on electron from the oxygen 2p-orbitals (or band) and place it into the magnesium 3s-orbitals (or band). Therefore, the band gap was modelled as a charge transfer excitation in Paper I, using a Mg-centered cluster, $[MgO_6Mg_{32}]^{26+}$. Since the second layer of Mg are border atoms in this cluster, no orbitals were included from this layer. From the oxygen surrounding the central Mg, 6 p-orbitals were included, those which could interact with the central Mg in a σ-donating fashion. On the central Mg, both the 3s orbital and 3p orbitals were included. In the end, this lead to an active space of 12 electrons in 10 orbitals, or a (12,10) active space in short. As was observed in Paper I, using HF-based AIMPs resulted in a band gap of around 10 eV with CASPT2 (13 eV with CASSCF), whereas with PBE-based AIMPs the predicted band gap was lowered to 7.5 eV, which is very close to the experimental value of around 7.8 eV. In Paper I, this result was speculated to be due to an intrinsically better description of the AIMPs by basing them on PBE, rather than HF.

One obvious criticism of the $[MgO_6Mg_{32}]^{26+}$ cluster is that it, by design, targets a localised description of the excited state. In such a picture, one Mg(II) becomes reduced to Mg(I), with the remaining hole in the occupied orbitals becoming delocalised over the surrounding oxygen ions. On the other hand, in the case of s-metal/p-metal oxides, such as MgO, it might be more appropriate with a delocalised description of the electronic structure. This is certainly the picture I get from standard textbooks on solid state chemistry and physics[71,72,73]. Therefore, the supposed failure of HF-based AIMPs to describe the band gap in MgO might have been related to the construction of the cluster, rather than the AIMPs themselves. In order to verify this, some alternative cluster should be tested, specifically designed to be more accommodating of a delocalised description of the excited state.

One fairly "straightforward" model for a more delocalised electronic structure would be a $Mg_{32}O_{32}$-cube (identical to a 2x2x2 supercell), where the active space would be localised in the inner Mg_4O_4-cube, with the remaining ions constituting the border between the central region and the embedding. Such a cluster fulfils the cluster construction criteria discussed in Papers I and II and outlined in Section 4.6. Using the same logic for the active space as in Paper I, all three p-orbitals of the four central oxygens should be included in the active space (they all point towards one central Mg each, making room for σ-donation) as well as one s-orbital and three p-orbitals per Mg. In total, the active space would then be (24,28), which is far beyond the capabilities of any modern supercomputer. Therefore, either RASSCF or DMRG would have to be used instead of CASSCF. One the other hand, a smaller active space where only one s-orbital per Mg is considered would result in a (24,16) active space; this is perfectly reasonable with CASSCF.

At this point, I will unfortunately not be able to provide a conclusion as to whether a delocalised description is better. Attempts at stabilising a (24,16) active space in a $Mg_{32}O_{32}$-cube turned out to be more challenging than I thought. For reasons of brevity, I will not discuss the (embarrassingly) long list of failed tests that lead to the (preliminary) end result I present here. Only by using a smaller (24,13) active space (only one s-type orbital, delocalised over all of the four central Mg) and making a state-average calculation over 13 roots (the ground state and twelve one-electron excitations, one from each p-orbital) could I find an active space that was somewhat reminiscent of what I was targeting. If I used (24,16), the resulting wavefunction became symmetry broken, with three of the Mg s-orbitals replaced by radially correlating orbitals from the O 3p-orbitals during the CASSCF iterations. Possibly, adding more roots to the SAV-CASSCF calculation could stabilise the inclusion of four Mg s-orbitals, but I did not test that route. Sadly, the computational burdens of MS-CASPT2 on the (24,13) active space turned out to be too severe, even for the largest computer (in terms of both disk space and memory) available to me, already at the level of ANO-RCC-VDZP as a basis set. Therefore I can give no conclusion as to the effect of constructing a more delocalised active space at the CASPT2 level-of-theory. The only preliminary insight I can give is that the CASSCF excitation energy (using HF-based AIMPs) dropped from 13.1 eV, with the cluster from Paper I, to 12.8 eV with the cubic cluster. Since this is only a marginal improvement with respect to the experimental gap of 7.8 eV and since the used active space is, in my own opinion, doubtful, more tests are necessary to decisively determine if a delocalised picture might be better. At this point, I will keep my recommendation from Paper I, i.e., that it might sometimes be advisable to test DFT-based AIMPs.

6.5.2 d-metal oxides and f-metal oxides

Oxides of d-metals and f-metals are known to be very interesting compounds, e.g. TiO_2 or CeO_2[74], which has a wide array of industrial applications. Unlike ionised s-metals and p-metals, there are many d-metal oxides and f-metal oxides where the d- or f-orbitals are partially filled. One example of this for FeO, where we would expect the Fe(II)-ions to have a d^6-configuration. In this case, all d-orbitals of any Fe included in the QM region has to be added to the active space. Otherwise, the open-shell nature of the material will be inappropriately described. Including all d-orbitals will rapidly lead to active spaces which are too large to handle, if the guidelines for AIMP embedding given in this thesis are followed. It is certainly possible to construct clusters where there are two d-metals in the QM region, but since no more than that can be included, the border with the AIMPs will be dangerously close to the transition metals.

On the other hand, in certain metal oxides, such as TiO_2 or CeO_2, the Ti(IV)/Ce(IV)-ions have a d^0/f^0-configuration, i.e., all valence electrons have been ionised away. In such materials, it is potentially safe to exclude orbitals on Ti or Ce that are far away from the cluster centre, making it easier to approach the guidelines for AIMP embedding. This statement should, however, be verified by very careful calculations before taken as an absolute. As a suggestion, either to myself in the future or to someone else reading this text, the oxygen vacancy formation energy in these materials might be a relevant property to evaluate this statement on.

It should additionally be stressed that all orbitals from the anions that provide significant σ-donation or π-donation should be included into the active space for d-metal oxides. As a rule-of-thumb, there should be one orbital from the anions per d-orbital and these orbitals should be symmetry paired[122]. For f-metal oxides, it is less straightforward how to deal with the active space, as the f-orbitals themselves do not partake in any chemical bonding[118].

6.5.3 d-metal dopants and f-metal dopants

Frequently, we are interested in studying materials that have been doped by some d-metal or f-metal, e.g., Ce-doped YVO_4, which is used as a laser. In this case, the selection of cluster model and active space is a bit more straightforward than when studying a pure material. The most sensible centre of the entire superstructure (QM region, AIMP region and point-charges) is on the dopant itself, with no real reason to test alternatives. Similarly, we should expect the active space to be primarily derived from orbitals on the dopant.

As argued in Ref 118, f-orbitals rarely partake in chemical bonding, for which reason a preliminary assumption that no orbitals from the surrounding crystals are necessary when studying doping by f-elements can be made. This is not necessarily true for d-orbitals, for which reason a sensible choice would be to include some orbitals from the crystal. Ideally, there should be at least five crystal orbitals of the same symmetry as the dopant d-orbitals[122]. Further, in the case that the host material contains open-shell species, the orbitals that are necessary to describe this has to be included in the active space. Other typical considerations for active spaces, such as a double-shell set of d-orbitals for 3d-elements (Paper I) and f-orbitals for 4f-elements[123] for elements with at least half-filled shells should be followed. Based on the discussions in Section 6.3.2, I would suggest making "minimal" CASPT2 calculations using such active spaces, compute the CASPT2 natural orbitals and based on their occupation numbers, determine if the solution is satisfactory, or increase the active space.

6.5.4 Reactions on surfaces

In heterogeneous catalysis, chemical reactions take place on the surface of a solid, often crystalline, material. Selecting active spaces for reactions is generally non-trivial, as the active space should be stable over the entire part of the potential energy surface (PES) that spans the reaction, starting from the reactant state to the final product. But in order to study the reaction, the PES has to be mapped out in some fashion, using any of the available routines for locating transition states. Since analytical gradients for methods such as CASPT2 are not generally available yet (at least, in Molcas) and are very expensive to compute, using a cheaper method such as DFT to map out the PES can be used as an alternative. Of course, some bias towards the used functional might be introduced, but there are few sensible alternatives to this approach today, in my opinion. Even if analytical gradients with CASPT2 were to be affordable, I would probably still recommend starting with a DFT generated PES, due to the active space selection process.

To construct the active space, my suggestion would be to start by selecting some orbitals for the transition state (TS). Once a satisfactory solution has been found for the TS, use the TS wavefunction as input for the remaining points along the PES. Such a starting guess is more likely to be stable over the entire PES, than for instance an active space that was initially constructed for either reactant or product. Of course, one could imagine that an appropriate active space would be a superposition of the reactant and product active spaces – that might, however, be very complicated to construct.

My recommendation to anyone attempting surface studies with AIMPs is then to first map out the PES, using periodic DFT with your functional of choice. After that, em-

bed the resulting structures in an AIMP and point-charge embedding. Construct an active space for the transition state, after which it should be verified that the solution is stable over the entire PES.

6.6 Final comments on multiconfigurational theory

In this chapter, some thoughts and comments on the usage of multiconfigurational theory – in particular, the combination of CASSCF with CASPT2 – have been outlined. The examples presented in Sections 6.4.3 and 6.4.4 are intended to provide some food for tought about simple and common mistakes that might occur when using multiconfigurational theory. In this particular case, however, the studied systems are highly symmetric (since they are atoms) and have a fairly limited set of possible active spaces. This is not true for molecular and crystalline systems, where the symmetry might be lower or completely nonexistant and the number of potential active spaces are far too many to study in a systematic way. When dealing with such systems, extra care has to be taken; the used active space(s) must always be critically scrutinised with emphasis on convergence patterns and properties of the resulting wavefunction.

Chapter 7

Summary and conclusions

This thesis work has been devoted to enabling multiconfigurational methods to be applied to problems in the solid state, in particular when the host materials is crystalline and ionic. To achieve this, the combination of two methods, GPEE and AIMPs, were used, which together gives a robust description of a crystalline environment. Some improvements on these methods were made as a part of this thesis work, with the major achievement being the creation of a user-friendly freely available program, SCEPIC, that allows the construction of AIMPs for any arbitrary ionic crystal. Further, this thesis work was (to the best of my knowledge) the first time that DFT-based AIMPs were ever used; in Paper I, this was shown to be quite useful when modelling pure materials. For very local properties, such as the spectra of dopants, is is, arguably, better to use HF-based AIMPs.

This thesis work was also the first time that systematic studies were made on the necessary size of the quantum region when using an AIMP embedding. Based on the discussions in Papers I and II, the following guidelines are proposed (stated in Chapter 4): 1) all elements of the material should be present in the QM region, i.e., no single element in a given material should be represented solely as AIMPs and 2) there should be at least one QM atom layer in between the actual ion or cluster of interest and the AIMP region. Further, the AIMP layer surrounding the QM cluster should be at the very least 6–8 Å, 10–15 Å can be used as a more conservative choice; the classical point-charges should ensure that 1) the Madelung potential is converged and 2) higher-order multipole moments are cancelled. If these criteria are fulfilled, the embedding strategy proposed in this work is robust for any ionic solid.

Paper III assessed the quality of the recently proposed ANO-R basis set. The results show that for many properties, ANO-R is of similar quality as the preceding ANO-

RCC basis set. Anionic states are, however, ill reproduced using the standard ANO-R contraction levels, for which reason the other studies in this thesis work have not used this basis set. As argued in Chapter 2, however, some time should properly be devoted to make a separate ANO-type basis set for ionic solids. This is necessary both to address the interaction between the ions and the external field of the crystal, as well as that smaller basis sets than ANO-R or ANO-RCC are desirable in order to model solids.

Papers IV and V are, arguably, a bit of a sidetrack from directly targeting ionic solids. They do, however, together, fulfil two important steps for my personal development as a user of multiconfigurational theory. In Paper IV, I began the development of how I approach active spaces, which was later discussed in Paper V and Chapter 6 of this thesis. Further, Paper IV serves as a good reminder for the fact that KSDFT is unreliable for multiconfigurational systems.

References

[1] T. Helgaker, P. Jørgensen, and J. Olsen. *Molecular Electronic Structure Theory*. John Wiley & Sons, Ltd, Chichester, England, 2000. ISBN 9780471967552.

[2] B. O. Roos, R. Lindh, P. Å. Malmqvist, V. Veryazov, and P. O. Widmark. *Multiconfigurational Quantum Chemistry*. Wiley & Sons, Hoboken, 2016. ISBN 9780470633465.

[3] J. C. Slater. The theory of complex spectra. *Physical Review*, 34(10):1293–1322, 1929.

[4] C. C. J. Roothaan. New developments in molecular orbital theory. *Reviews of Modern Physics*, 23:69–89, 1951.

[5] G. G. Hall. The molecular orbital theory of chemical valency VIII. A method of calculating ionization potentials. *Proceeding of the Royal Society A*, 205:541–552, 1951.

[6] C. Møller and M.S. Plesset. Note on an approximation treatment for many-electron systems. *Physical Review*, 46:618–622, 1934.

[7] I. Shavitt. The history and evolution of configuration interaction. *Molecular Physics*, 94(1):3–17, 1998.

[8] J. Čížek. On the correlation problem in atomic and molecular systems. Calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods. *Journal of Chemical Physics*, 45(11):4256–4266, 1966.

[9] J. Noga and W. Kutzelnigg. Coupled cluster theory that takes care of the correlation cusp by inclusion of linear terms in the interelectronic coordinates. *Journal of Chemical Physics*, 101(9):7738–7762, 1994. ISSN 00219606.

[10] B. O. Roos, P. R. Taylor, and P. E. M. Siegbahn. A Complete Active Space Super-Ci Method. *Chemical Physics*, 48(990):157–173, 1980.

[11] J. Olsen, B. O. Roos, P. Jørgensen, and H. J. Aa. Jensen. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *Journal of Chemical Physics*, 89(4):2185–2192, 1988.

[12] D. Ma, G. Li Manni, and L. Gagliardi. The generalized active space concept in multiconfigurational self-consistent field methods. *Journal of Chemical Physics*, 135(4):044128, 2011.

[13] G. Li Manni, S. D. Smart, and A. Alavi. Combining the Complete Active Space Self-Consistent Field Method and the Full Configuration Interaction Quantum Monte Carlo within a Super-CI Framework, with application to challenging metal-porphyrins. *Journal of Chemical Theory and Computation*, 12(3):1245–1258, 2016.

[14] S. R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69(19):2863–2866, 1992.

[15] S. R. White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48(14):10345–10356, 1993.

[16] G. K. L. Chan and D. Zgid. *Chapter 7 The Density Matrix Renormalization Group in Quantum Chemistry*, volume 5. Elsevier, 2009. ISBN 9780444533593.

[17] K Andersson, P.-Å. Malmqvist, and B. O. Roos. Second-order perturbation theory with a complete active space self-consistent field reference function. *Journal of Chemical Physics*, 96(2):1218–1226, 1992.

[18] C. Angeli, R. Cimiraglia, S. Evangelisti, T. Leininger, and J. P. Malrieu. Introduction of n-electron valence states for multireference perturbation theory. *Journal of Chemical Physics*, 114(23):10252, 2001.

[19] M Boggio-Pasqua, D. Jacquemin, and P. F. Loos. Benchmarking CASPT3 vertical excitation energies. *Journal of Chemical Physics*, 157(1):1–12, 2022.

[20] M. Kállay, P. G. Szalay, and P. R. Surján. A general state-selective multireference coupled-cluster algorithm. *Journal of Chemical Physics*, 117(3):980–990, 2002.

[21] P. Hohenberg and W. Kohn. Inhomogenous electron gas. *Physical Review*, 136 (3):B864–B871, 1964.

[22] W. Kohn and Sham L.J. Self-consitent equations including exchange and correlation effects. *Physical Review*, 140(4A):1133–1138, 1965.

[23] W. Koch and M. C. Holthausen. *A chemist's guide to density functional theory*. Wiley-VCH, Weinheim, Germany, 2nd edition, 2001. ISBN 9783527304226.

[24] S. Lehtola, C. Steigemann, M. J.T. Oliveira, and M. A.L. Marques. Recent developments in LIBXC — A comprehensive library of functionals for density functional theory. *SoftwareX*, 7:1–5, 2018.

[25] A. M. Teale, T. Helgaker, A. Savin, C. Adano, B. Aradi, A. V. Arbuznikov, P. Ayers, E. J. Baerends, V. Barone, P. Calaminici, E. Cances, E. A. Carter, P. K. Chattaraj, H. Chermette, I. Ciofini, T. D. Crawford, F. De Proft, J. Dobson, C. Draxl, T. Frauenheim, E. Fromager, P. Fuentealba, L. Gagliardi, G. Galli, J. Gao, P. Geerlings, N. Gidopoulos, P. M. W. Gill, P. Gori-Giorgi, A. Görling, T. Gould, S. Grimme, O. Gritsenko, H. J. Aa. Jensen, E. R. Johnson, R. O. Jones, M. Kaupp, A. Koster, L. Kronik, A. I. Krylov, S. Kvaal, A. Laestadius, M. P. Levy, M. Lewin, S. B. Liu, P.-F. Loos, N. T. Maitra, F. Neese, J. Perdew, K. Pernal, P. Pernot, P. Piecuch, E. Rebolini, L. Reining, P. Romaniello, A. Ruzsinszky, D. Salahub, M. Scheffler, P. Schwerdtfeger, V. N. Staroverov, J. Sun, E. Tellgren, D. J. Tozer, S. Trickey, C. A. Ullrich, A. Vela, G. Vignale, T. A. Wesolowski, X. Xu, and W. Yang. DFT Exchange: Sharing Perspectives on the Workhorse of Quantum Chemistry and Materials Science. *Physical Chemistry Chemical Physics*, 19:1–85, 2022.

[26] J. P. Perdew and K. Schmidt. Jacob's ladder of density functional approximations for the exchange-correlation energy. *AIP Conference Proceedings*, 577:1–20, 2001.

[27] M. G. Medvedev, I. S. Busmarinov, J. Sun, J. P. Perdew, and K. A. Lyssenko. Density functional theory is Straying from the Path toward the Exact Functional. *Science*, 355(January):49–52, 2017.

[28] P. J. Hasnip, K. Refson, M. I.J. Probert, J. R. Yates, S. J. Clark, and C. J. Pickard. Density functional theory in the solid state. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372 (2011):20130270, 2014.

[29] G. Li Manni, R. K. Carlson, S. Luo, D. Ma, J. Olsen, D. G. Truhlar, and L. Gagliardi. Multiconfiguration pair-density functional theory. *Journal of Chemical Theory and Computation*, 10(9):3669–3680, 2014.

[30] E. Fromager, J. Toulouse, and H. J. Aa. Jensen. On the universality of the long-/short-range separation in multiconfigurational density-functional theory. *Journal of Chemical Physics*, 126(7):074111, 2007.

[31] G. Li Manni, D. Ma, F. Aquilante, J. Olsen, and L. Gagliardi. SplitGAS method for strong correlation and the challenging case of Cr_2. *Journal of Chemical Theory and Computation*, 9(8):3375–3384, 2013.

[32] S. Guo, M. A. Watson, W. Hu, Q. Sun, and G. K. L. Chan. N-electron valence state perturbation theory based on a density matrix renormalization group reference function, with applications to the chromium dimer and a trimer model of poly(p-phenylenevinylene). *Journal of Chemical Theory and Computation*, 12 (4):1583–1591, 2016.

[33] S. Vancoillie, P. Å. Malmqvist, and V. Veryazov. Potential Energy Surface of the Chromium Dimer Re-re-revisited with Multiconfigurational Perturbation Theory. *Journal of Chemical Theory and Computation*, 12(4):1647–1655, 2016.

[34] C. Kollmar, K. Sivalingam, and F. Neese. An alternative choice of the zeroth-order Hamiltonian in CASPT2 theory. *Journal of Chemical Physics*, 152(21): 214110, 2020.

[35] P. Å. Malmqvist, K. Pierloot, A. R. M. Shahi, C. J. Cramer, and L. Gagliardi. The restricted active space followed by second-order perturbation theory method: Theory and application to the study of CuO_2 and Cu_2O_2 systems. *Journal of Chemical Physics*, 128(20):204109, 2008.

[36] D. Ma, G. Li Manni, J. Olsen, and L. Gagliardi. Second-order perturbation theory for generalized active space self-consistent-field wave functions. *Journal of Chemical Theory and Computation*, 12(7):3208–3213, 2016.

[37] Y. Kurashige and T. Yanai. Second-order perturbation theory with a density matrix renormalization group self-consistent field reference function: Theory and application to the study of chromium dimer. *Journal of Chemical Physics*, 135(9):094104, 2011.

[38] S. Guo, Z. Li, and G. K. L. Chan. A perturbative density matrix renormalization group algorithm for large active spaces. *Journal of Chemical Theory and Computation*, 14(8):4063–4071, 2018.

[39] J. Almlöf and P. R. Taylor. General contraction of Gaussian basis sets. I. Atomic natural orbitals for first- and second-row atoms. *Journal of Chemical Physics*, 86 (7):4070–4077, 1987.

[40] Y. J. Franzke, L. Spiske, P. Pollak, and F. Weigend. Segmented contracted error-consistent basis sets of quadruple-ζ valence quality for one- and two-component relativistic all-electron calculations. *Journal of Chemical Theory and Computation*, 16(9):5658–5674, 2020.

[41] P. O. Widmark, P.-Å. Malmqvist, and B. O. Roos. Density matrix averaged atomic natural orbital (ANO) basis sets for correlated molecular wave functions - I. First row atoms. *Theoretica Chimica Acta*, 77(5):291–306, 1990.

[42] P. Pyykkö. Relativistic effects in structural chemistry. *Chemical Reviews*, 88(3): 563–594, 1988.

[43] M. Douglas and N. M. Kroll. Quantum electrodynamical corrections to the fine structure of helium. *Annals of Physics*, 82(1):89–155, 1974.

[44] B. A. Hess. Applicability of the no-pair equation with free-particle projection operators to atomic and molecular structure calculations. *Physical Review A*, 32: 756–763, 1985.

[45] B A. Hess. Relativistic electronic-structure calculations employing a two-component no-pair formalism with external-field projection operators. *Physical Review A*, 33:3742–3748, 1986.

[46] J. P. Zobel, P. O. Widmark, and V. Veryazov. The ANO-R Basis Set. *Journal of Chemical Theory and Computation*, 16(1):278–294, 2020.

[47] W. Kutzelnigg and W. Liu. Quasirelativistic theory equivalent to fully relativistic theory. *Journal of Chemical Physics*, 123(13):241102, 2005.

[48] W. Liu. Ideas of relativistic quantum chemistry. *Molecular Physics*, 108(13): 1679–1706, 2010.

[49] D. Peng and M. Reiher. Exact decoupling of the relativistic Fock operator. *Theoretical Chemistry Accounts*, 131:1081, 2012.

[50] T. Saue. Relativistic hamiltonians for chemistry: A primer. *ChemPhysChem*, 12 (17):3077–3094, 2011.

[51] P. Å. Malmqvist, B. O. Roos, and B. Schimmelpfennig. The restricted active space (RAS) state interaction approach with spin-orbit coupling. *Chemical Physics Letters*, 357(3-4):230–240, 2002.

[52] S. Knecht, H. J. Aa Jensen, and T. Saue. Relativistic quantum chemical calculations show that the uranium molecule U_2 has a quadruple bond. *Nature Chemistry*, 11:40–44, 2018.

[53] P. Atkins, T. Overton, J. Rourke, M. Weller, and F. Armstrong. *Shriver & Atkins' Inorganic Chemistry*. Oxford University Press, Oxford, England, 5th edition, 2010. ISBN 9780199236176.

[54] B. O. Roos, R. Lindh, P. Å. Malmqvist, V. Veryazov, and P. O. Widmark. New relativistic ANO basis sets for transition metal atoms. *Journal of Physical Chemistry A*, 109(29):6575–6579, 2005.

[55] D. E. Kelleher, W. C. Martin, W. L. Wiese, J. Sugar, J. R. Fuhr, K. Olsen, A. Musgrove, P. J. Mohr, J. Reader, and G. R. Dalton. The new NIST atomic spectra database. *Physica Scripta*, T83:158–161, 1999.

[56] Sakurai J. J. and J. Napolitano. *Modern Quantum Mechanics*. Cambridge University Press, 2nd edition, 2017. ISBN 9781108422413.

[57] M. F. Peintinger, D. V. Oliveira, and T. Bredow. Consistent Gaussian basis sets of triple-zeta valence with polarization quality for solid-state calculations. *Journal of Computational Chemistry*, 34(6):451–459, 2013.

[58] H. Z. Ye and T. C. Berkelbach. Correlation-consistent gaussian basis sets for solids made simple. *Journal of Chemical Theory and Computation*, 18(3):1595–1606, 2022.

[59] F. Libisch, C. Huang, and E. A. Carter. Embedded correlated wavefunction schemes: Theory and applications. *Accounts of Chemical Research*, 47(9):2768–2775, 2014.

[60] G. Knizia and G. K. L. Chan. Density matrix embedding: A simple alternative to dynamical mean-field theory. *Physical Review Letters*, 109(18):1–5, 2012.

[61] R. S. Mulliken. Electronic population analysis on LCAO-MO molecular wave functions. I. *Journal of Chemical Physics*, 23:1833–1840, 1955.

[62] R. F.W. Bader. *Atoms in Molecules: A Quantum Theory*. New York: Oxford University Press, 1990. ISBN 9780198558651.

[63] K.B. Wiberg. Application of the Pople-Santry-Segal CNDO method to the cyclopropylcarbinyl and cyclobutyl cation and to bicyclobutane. *Tetrahedron*, 24(3):1083–1096, 1968.

[64] D.R. Armstrong, P.G. Perkins, and J.J.P. Stewart. Bond indices and valency. *Journal Chemical Society Dalton Transactions*, pages 838–840, 1973.

[65] R. A. Evarestov and V. A. Veryazov. Quantum-chemical definition of the atomic valence in molecules and crystals. *Theoretica Chimica Acta*, 81(1-2):95–103, 1991.

[66] C. Müller and K. Hermansson. Assessment methods for embedding schemes - Ceria as an example. *Surface Science*, 603(23):3329–3338, 2009.

[67] J. Hutter, M. Parrinello, and G. Lippert. The Gaussian and augmented-plane-wave density functional method for ab initio molecular dynamics simulations. *Theoretical Chemistry Accounts*, 103:124–140, 1999.

[68] J. M.L. Martin, C. W. Bauschlicher, and A. Ricca. On the integration accuracy in molecular density functional theory calculations using Gaussian basis sets. *Computer Physics Communications*, 133(2-3):189–201, 2001.

[69] R. Lindh, P. Å. Malmqvist, and L. Gagliardi. Molecular integrals by numerical quadrature. I. Radial integration. *Theoretical Chemistry Accounts*, 106(3):178–187, 2001.

[70] M. Guidon, J. Hutter, and J. Vandevondele. Auxiliary density matrix methods for Hartree-Fock exchange calculations. *Journal of Chemical Theory and Computation*, 6(8):2348–2364, 2010.

[71] R. Dronskowski. *Computational Chemistry of Solid State Materials*. Wiley-VCH, Hoboken, 2005. ISBN 9783527314102.

[72] A. R. West. *Solid State Chemistry*. Wiley & Sons, Chichester, England, 2nd edition, 2014. ISBN 9781119942948.

[73] D. Snoke. *Solid State Physics*. Cambridge University Press, Cambridge, United Kingdom, 2020. ISBN 9781107191983.

[74] T. Montini, M. Melchionna, M. Monai, and P. Fornasiero. Fundamentals and catalytic applications of CeO_2-based materials. *Chemical Reviews*, 116(10):5987–6041, 2016.

[75] S.L. Dudarev, G.A. Botton, S.Y. Savrasov, C.J. Humphreys, and A.P. Sutton. Electron-energy-loss spectra and the structural stability of nickel oxide: An LSDA+U study. *Physical Review B - Condensed Matter and Materials Physics*, 57(3):1505–1509, 1998.

[76] I. Timrov, N. Marzari, and M. Cococcioni. Self-consistent Hubbard parameters from density-functional perturbation theory in the ultrasoft and projector-augmented wave formulations. *Physical Review B*, 103(4):45141, 2021.

[77] C. Loschen, J. Carrasco, K. M. Neyman, and F. Illas. First-principles LDA+U and GGA+U study of cerium oxides: Dependence on the effective U parameter. *Physical Review B - Condensed Matter and Materials Physics*, 75(3):1–8, 2007.

[78] C. W.M. Castleton, J. Kullgren, and K. Hermansson. Tuning LDA+U for electron localization and structure at oxygen vacancies in ceria. *Journal of Chemical Physics*, 127(24):1–12, 2007.

[79] G. Pacchioni. First principles calculations on oxide-based heterogeneous catalysts and photocatalysts: problems and advances. *Catalysis Letters*, 145(1):80–94, 2015.

[80] M. J. Wolf, E. D. Larsson, and K. Hermansson. Oxygen chemistry of halogen-doped CeO2(111). *Physical Chemistry Chemical Physics*, 23(35):19375–19385, 2021.

[81] J. Graciani, A. M. Márquez, J. J. Plata, Y. Ortega, N. C. Hernández, A. Meyer, C. M. Zicovich-Wilson, and J. F. Sanz. Comparative study on the performance of hybrid DFT functionals in highly correlated oxides: The case of CeO2 and Ce2O3. *Journal of Chemical Theory and Computation*, 7(1):56–65, 2011.

[82] D. H. Seo, A. Urban, and G. Ceder. Calibrating transition-metal energy levels and oxygen bands in first-principles calculations: Accurate prediction of redox potentials and charge transfer in lithium transition-metal oxides. *Physical Review B - Condensed Matter and Materials Physics*, 92(11):1–11, 2015.

[83] D. Du, M. J. Wolf, K. Hermansson, and P. Broqvist. Screened hybrid functionals applied to ceria: Effect of Fock exchange. *Physical Review B*, 97(23):1–14, 2018.

[84] D. Du, J. Kullgren, K. Hermansson, and P. Broqvist. From ceria clusters to nanoparticles: Superoxides and supercharging. *Journal of Physical Chemistry C*, 123(3):1742–1750, 2019.

[85] J. Heyd and G. E. Scuseria. Efficient hybrid density functional calculations in solids: Assessment of the Heyd-Scuseria-Ernzerhof screened Coulomb hybrid functional. *Journal of Chemical Physics*, 121(3):1187–1192, 2004.

[86] L. Seijo and Z. Barandiarán. The Ab Initio Model Potential method: A common strategy for effective core potential and embedded cluster calculations. *Computational Chemistry: Reviews of Current Trends*, 4:55–152, 1999.

[87] P. V. Sushko and I. V. Abarenkov. General purpose electrostatic embedding potential. *Journal of Chemical Theory and Computation*, 6(4):1323–1333, 2010.

[88] A. Dittmer, R. Izsák, F. Neese, and D. Maganas. Accurate band gap predictions of semiconductors in the framework of the similarity transformed equation of motion coupled cluster theory. *Inorganic Chemistry*, 58(14):9303–9315, 2019.

[89] B. X. Shi, V. Kapil, A. Zen, J. Chen, A. Alavi, and A. Michaelides. General embedded cluster protocol for accurate modeling of oxygen vacancies in metal-oxides. *Journal of Chemical Physics*, 156(12):124704, 2022.

[90] Z. Barandiarán and L. Seijo. The ab initio model potential representation of the crystalline environment. Theoretical study of the local distortion on NaCl:Cu^+. *Journal of Chemical Physics*, 89(9):5739–5746, 1988.

[91] M. A. Nygren, L. G.M. Pettersson, Z. Barandiarán, and L. Seijo. Bonding between CO and the MgO(001) surface: A modified picture. *Journal of Chemical Physics*, 100(3):2010–2018, 1994.

[92] C. De Graaf, C. Sousa, and R. Broer. Ionization and excitation energies in CuCl and NiO within different embedding schemes. *Journal of Molecular Structure (Theochem)*, 458(1-2):53–60, 1998.

[93] V. Kairys and J. H. Jensen. QM/MM boundaries across covalent bonds: A frozen localized molecular orbital-based approach for the effective fragment potential method. *Journal of Physical Chemistry A*, 104(28):6656–6665, 2000.

[94] L. N. Kantorovich. An embedded-molecular-cluster method for calculating the electronic structure of point defects in non-metallic crystals. 1. General theory. *Journal of Physics C: Solid State Physics*, 21(29):5041–5056, 1988.

[95] D. Bakowies and W. Thiel. Hybrid models for combined quantum mechanical and molecular mechanical approaches. *Journal of Physical Chemistry*, 100(25): 10580–10594, 1996.

[96] V. B. Sulimov, P. V. Sushko, A. H. Edwards, A. L. Shluger, and A. M. Stoneham. Asymmetry and long-range character of lattice deformation by neutral oxygen vacancy in α-quartz. *Physical Review B*, 66(2):024108, 2002.

[97] A. A. Sokol, S. T. Bromley, S. A. French, C. R. A. Catlow, and P. Sherwood. Hybrid QM/MM embedding approach for the treatment of localized surface states in ionic materials. *International Journal of Quantum Chemistry*, 99(5 SPEC. ISS.):695–712, 2004.

[98] O. Danyliv and L. Kantorovich. Comparison of localization procedures for applications in crystal embedding. *Physical Review B*, 70(7):075113, 2004.

[99] P. Huang and E. A. Carter. Advances in correlated electronic structure methods for solids, surfaces, and nanostructures. *Annual Review of Physical Chemistry*, 59:261–290, 2008.

[100] B. Paulus and H. Stoll. The method of increments - a wavefunction-based correlation method for extended systems. In F. Manby, editor, *Accurate Condensed-Phase Quantum Chemistry*, pages 57–84. CRC Press, 2010.

[101] C. Müller and B. Paulus. Wavefunction-based electron correlation methods for solids. *Physical Chemistry Chemical Physics*, 14(21):7605–7614, 2012.

[102] P. P. Ewald. Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Annalen der Physik*, 64:253, 1921.

[103] H.M. Evjen. On the stability of certain heteropolar crystals. *Physical Review*, 39:675–687, 1932.

[104] I. V. Abarenkov. Unit cell for a lattice electrostatic potential. *Physical Review B - Condensed Matter and Materials Physics*, 76:165127, 2007.

[105] R. McWeeny. *Methods of Molecular Quantum Mechanics*. Academic Press, London, United Kingdom, 1992. ISBN 9780124865518.

[106] B. O. Roos, R. Lindh, P. Å. Malmqvist, V. Veryazov, P. O. Widmark, and A. C. Borin. New relativistic atomic natural orbital basis sets for lanthanide atoms with applications to the Ce diatom and LuF_3. *Journal of Physical Chemistry A*, 112(45):11431–11435, 2008.

[107] Q. Lu and K. A. Peterson. Correlation consistent basis sets for lanthanides: The atoms La-Lu. *Journal of Chemical Physics*, 145(5):054111, 2016.

[108] M. Dolg, H. Stoll, H. Preuss, and R. M. Pitzer. Relativistic and correlation effects for element 105 (hahnium, Ha). A comparative study of M and MO (M = Nb, Ta, Ha) using energy-adjusted ab initio pseudopotentials. *Journal of Physical Chemistry*, 97(22):5852–5859, 1993.

[109] P. E. Blöch. Projector augmented-wave method. *Physical Review B*, 50(24): 17953–17979, 1994.

[110] S. Huzinaga and A. A. Cantu. Theory of separability of many-electron systems. *Journal of Chemical Physics*, 55(12):5543–5549, 1971.

[111] S. Huzinaga, D. McWilliams, and A. A. Cantu. Projection operators in Hartree-Fock theory. *Advances in Quantum Chemistry*, 7(C):187–220, 1973.

[112] F. Aquilante, L. De Vico, N. Ferré, G. Ghigo, P.-Å. Malmqvist, P. Neogrády, T. B. Pedersen, M. Pitoňák, M. Reiher, B. O. Roos, L. Serrano-Andrés, M. Urban, V. Veryazov, and R. Lindh. MOLCAS 7: The Next Generation. *Journal of Computational Chemistry*, 31(1):225–247, 2010.

[113] N. A. Bogdanov, G. Li Manni, S. Sharma, O. Gunnarsson, and A. Alavi. Enhancement of superexchange due to synergetic breathing and hopping in corner-sharing cuprates. *Nature Physics*, 18(2):190–195, 2022.

[114] D. K. Kanan, S. Sharifzadeh, and E. A. Carter. Quantum mechanical modeling of electronic excitations in metal oxides: Magnesia as a prototype. *Chemical Physics Letters*, 519-520:18–24, 2012.

[115] F. Aquilante, J. Autschbach, R. K. Carlson, L. F. Chibotaru, M. G. Delcey, L. De Vico, I. Fdez. Galván, N. Ferré, L. M. Frutos, L. Gagliardi, M. Garavelli, A. Giussani, C. E. Hoyer, G. Li Manni, H. Lischka, D. Ma, P. Å. Malmqvist, T. Müller, A. Nenov, M. Olivucci, T. B. Pedersen, D. Peng, F. Plasser, B. Pritchard, M. Reiher, I. Rivalta, I. Schapiro, Javier Segarra-Martí, Michael Stenrup, Donald G. Truhlar, Liviu Ungur, Alessio Valentini, Steven Vancoillie, Valera Veryazov, V. P. Vysotskiy, O. Weingart, F. Zapata, and R. Lindh. Molcas 8: New capabilities for multiconfigurational quantum chemical calculations across the periodic table. *Journal of Computational Chemistry*, 2016.

[116] I. Fdez. Galván, N. Vacher, A. Alavi, C. Angeli, F. Aquilante, J. Autschbach, J. J. Bao, S. I. Bokarev, N. A. Bogdanov, R. K. Carlson, L. F. Chibotaru, J. Creutzberg, N. Dattani, M. G. Delcey, S. S. Dong, A. Dreuw, L. Freitag, L. M. Frutos, L. Gagliardi, F. Gendron, A. Giussani, L. González, G. Grell, M. Guo, C. E. Hoyer, M. Johansson, S. Keller, S. Knecht, G. Kovačević, E. Källman, G. Li Manni, M. Lundberg, Y. Ma, S. Mai, J. P. Malhado, P. Å. Malmqvist, P. Marquetand, S. A. Mewes, Jesper Norell, Massimo Olivucci, Markus Oppel, Quan Manh Phung, Kristine Pierloot, F. Plasser, M. Reiher, A. M. Sand, I. Schapiro, P. Sharma, C. J. Stein, L. K. Sørensen, D. G. Truhlar, M. Ugandi, L. Ungur, A. Valentini, S. Vancoillie, V. Veryazov, O. Weser, T. A. Wesołowski, P.-O. Widmark, S. Wouters, A. Zech, J. P. Zobel, and R. Lindh. OpenMolcas: From source code to insight. *Journal of Chemical Theory and Computation*, 2019.

[117] F. Aquilante, J. Autschbach, A. Baiardi, S. Battaglia, V. A. Borin, L. F. Chibotaru, I. Conti, L. De Vico, M. Delcey, I. Fdez. Galván, N. Ferré, L. Freitag, M. Garavelli, X. Gong, S. Knecht, E. D. Larsson, R. Lindh, M. Lundberg, P. Å. Malmqvist, A. Nenov, J. Norell, M. Odelius, M. Olivucci, T. B. Pedersen, L. Pedraza-González, Q. M. Phung, K. Pierloot, M. Reiher, I. Schapiro, J. Segarra-Martí, F. Segatta, L. Seijo, S. Sen, D. C. Sergentu, C. J. Stein, L. Ungur, M. Vacher, A. Valentini, and V. Veryazov. Modern quantum chemistry with [Open]Molcas. *Journal of Chemical Physics*, 152(21), 2020.

[118] B. O. Roos V. Veryazov, P. Å. Malmqvist. How to select active space for multiconfigurational quantum chemistry? *International Journal of Quantum Chemistry*, 111:3329–3338, 2011.

[119] R. Yongfeng, Z. Shouchao, L. Shuhua, L. Guanghui, L. Wenrun, and L. Jingwang. Growth and spectrum properties of Ce:YVO_4 Single Crystal. *Journal of Rare Earths*, 25(SUPPL. 1):122–124, 2007.

[120] C. J. Stein and M. Reiher. Automated selection of active orbital spaces. *Journal of Chemical Theory and Computation*, 12(4):1760–1771, 2016.

[121] T. Andersen, H. K. Haugen, and H. Hotop. Binding energies in atomic negative ions: II. *Journal of Physical and Chemical Reference Data*, 28(6):1511–1533, 1999.

[122] K. Pierloot and L. G. Vanquickenborne. The ligand field spectrum of the hexafluorochromate (III) anion: An ab initio study including correlation effects. *Journal of Chemical Physics*, 93(6):4154–4163, 1990.

[123] Z. Barandiarán and L. Seijo. Intervalence charge transfer luminescence: Interplay between anomalous and 5d - 4f emissions in Yb-doped fluorite-type crystals. *Journal of Chemical Physics*, 141(23):234704, 2014.

Scientific publications

Author contributions

Paper I: A program system for self-consistent embedded potentials for ionic crystals

I participated in the planning of the research. I wrote the computer code evaluated from scratch (using older codes as templates/references). I wrote the majority of the text. I chose the systems that were studied as a part of the multiconfigurational test calculations. I performed all calculations. I made the majority of the analysis of results.

Paper II: Convergence of Electronic Structure Properties in Ionic Oxides Within a Fragment Approach

I participated in the planning of the research. I performed all calculations. I wrote the computer code to compute electronic structure properties and electron density comparison from scratch. I contributed to analysis of results and writing of 'Results and discussion' section.

Paper III: Benchmarking ANO-R basis set for multiconfigurational calculations

I made the majority of the calculations. I contributed to analysis of results and wrote a substantial part of the 'Results and discussion' and 'Conclusions' sections.

Paper IV: Is density functional theory accurate for lytic polysaccharide monooxygenase enzymes?

I performed all multiconfigurational calculations. I contributed to analysis of results and writing of 'Results and discussion' and 'Conclusions' sections.

Paper V: Modern quantum chemistry with [Open]Molcas

I was the lone author of the section on the use of ExpBas (Section 2F in the main article and Section S7 in the supporting information) and made all test calculations related to it.

Printed by Media-Tryck, Lund 2022 NORDIC SWAN ECOLABEL 3041 0903

ISBN: 978-91-8039-463-5

Theoretical Chemistry
Department of Chemistry
Faculty of Science
Lund University