Springer Theses
Recognizing Outstanding Ph.D. Research

Tim Joachim Zuehlsdorff

Computing the Optical Properties of Large Systems

Springer

Tim Joachim Zuehlsdorff

Computing the Optical Properties of Large Systems

Doctoral Thesis accepted by
the Imperial College London, UK

Springer

Author
Dr. Tim Joachim Zuehlsdorff
Cavendish Laboratory
University of Cambridge
Cambridge
UK

Supervisor
Prof. Peter D. Haynes
Department of Physics and
Department of Materials
Imperial College London
London
UK

ISSN 2190-5053 ISSN 2190-5061 (electronic)
Springer Theses
ISBN 978-3-319-19769-2 ISBN 978-3-319-19770-8 (eBook)
DOI 10.1007/978-3-319-19770-8

Library of Congress Control Number: 2015941115

Springer Cham Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

Springer Theses

Recognizing Outstanding Ph.D. Research

Aims and Scope

The series "Springer Theses" brings together a selection of the very best Ph.D. theses from around the world and across the physical sciences. Nominated and endorsed by two recognized specialists, each published volume has been selected for its scientific excellence and the high impact of its contents for the pertinent field of research. For greater accessibility to non-specialists, the published versions include an extended introduction, as well as a foreword by the student's supervisor explaining the special relevance of the work for the field. As a whole, the series will provide a valuable resource both for newcomers to the research fields described, and for other scientists seeking detailed background information on special questions. Finally, it provides an accredited documentation of the valuable contributions made by today's younger generation of scientists.

Theses are accepted into the series by invited nomination only and must fulfill all of the following criteria

- They must be written in good English.
- The topic should fall within the confines of Chemistry, Physics, Earth Sciences, Engineering and related interdisciplinary fields such as Materials, Nanoscience, Chemical Engineering, Complex Systems and Biophysics.
- The work reported in the thesis must represent a significant scientific advance.
- If the thesis includes previously published material, permission to reproduce this must be gained from the respective copyright holder.
- They must have been examined and passed during the 12 months prior to nomination.
- Each thesis should include a foreword by the supervisor outlining the significance of its content.
- The theses should have a clearly defined structure including an introduction accessible to scientists not expert in that particular field.

More information about this series at http://www.springer.com/series/8790

Supervisor's Foreword

Density-functional theory (DFT) surely ranks as one of the most extraordinary and surprising theoretical discoveries in the physical sciences. In October 2014, just after this thesis was submitted for examination, the journal *Nature* reported an analysis of the 100 most highly-cited papers of all time [*Nature* **514**, 550 (2014)]. Twelve of these relate to DFT, including two of the top ten. Back in 1998, Walter Kohn received the Nobel Prize for Chemistry for his leading role in the development of DFT. While relatively few researchers today actively work on the theory itself, tens of thousands of us—geologists, biochemists, electronic engineers and materials scientists as well as physicists and chemists—rely on it to perform computational simulations of molecules, materials and biological systems.

The remarkable achievement of DFT is that it makes the impossible possible. A key feature of quantum mechanics is that the electrons that form chemical bonds and determine the properties of materials, devices and drugs are not independent of each other. They must be described by a single wave function, and the computational effort required to determine this quantity grows exponentially with their number. A direct assault on the equations of quantum mechanics is prohibitively expensive for more than the tiniest collections of atoms. Yet DFT shows that, in principle at least, the wave function can be bypassed in favour of the electron density: the average number of electrons at each point in space. Through the development of a series of well-controlled approximations and robust, user-friendly software, the power of DFT has been fully exploited. Its dominance today relies on the fact that it balances two requirements: first it provides a sufficiently accurate description of the behaviour of electrons for most practical purposes, second the computational cost is much lower than for wave-function-based methods. This is why DFT has been so successful over the last 30 years.

However DFT as traditionally implemented is not without its shortcomings. One of the most famous of these is that it is a ground-state theory that does not formally treat electronic excitations, and yet many important properties of materials—such as their interaction with light—depend upon this. Progress was made on this front 25 years ago with the extension of DFT to treat the influence of a time-dependent

(TD) external driving force or potential, such as the oscillating electric field of light. Developing practical approximations for TD-DFT has proven to be far more challenging than for ground-state DFT, but software tools have been developed and are in widespread use. Another limitation of DFT is that in spite of its highly competitive computational cost compared with wave-function-based methods, the scaling of the effort with the cube of the number of electrons makes it extremely expensive to treat more than a few hundreds of atoms—still small compared with relevant length-scales in materials and biological systems.

The thesis work by Tim Zuehlsdorff described in this volume aims to address both of these drawbacks by developing and applying a method to perform TD-DFT simulations with a computational cost that scales only linearly with the number of electrons. This work has been carried out within the framework of the ONETEP code, which optimises a set of local orbitals to describe the chemical environment of each atom using a basis set equivalent to plane-waves, which are a popular choice for DFT. In a previous volume (DOI:10.1007/978-3-319-00339-9), Laura Ratcliff described how this process can be extended to optimise a second set of local orbitals to describe the unoccupied or virtual electronic states—a vital prerequisite for any treatment of electronic excitations. In the work reported here, Tim has built on that foundation to implement a linear response approach to TD-DFT with a linear-scaling framework and to demonstrate its application to two systems of potential technological interest: doped organic molecular crystals for a room temperature MASER and the Fenna-Matthews-Olson pigment-protein complex. Moreover, Tim describes how the local orbital framework provides an opportunity to manipulate the description of the electronic excitations in such a way as to remedy some of the side effects of the simplest approximations commonly used in TD-DFT. It is work like this—on the development of both approximations and software—that is needed to ensure that DFT continues to serve a wide range of scientists and engineers for the next 30 years.

London
March 2015

Prof. Peter D. Haynes

Abstract

In recent years, time-dependent density-functional theory (TDDFT) has been the method of choice for calculating optical excitations in medium sized to large systems, due to its good balance between computational cost and achievable accuracy. In this thesis, TDDFT is reformulated to fit the framework of the linear-scaling density-functional theory (DFT) code `ONETEP`. The implementation relies on representing the optical response of the system using two sets of localised, atom centred, *in situ* optimised orbitals in order to ideally describe both the electron and the hole wavefunctions of the excitation. This dual representation approach requires only a minimal number of localised functions, leading to a very efficient algorithm. It is demonstrated that the method has the capability of computing low energy excitations of systems containing thousands of atoms in a computational effort that scales linearly with system size.

The localised representation of the response to a perturbation allows for the selective convergence of excitations localised in certain regions of a larger system. The excitations of the whole system can then be obtained by treating the coupling between different subsystems perturbatively. It is shown that in the limit of weakly coupled excitons, the results obtained with the coupled subsystem approach agree with a full treatment of the entire system, with a large reduction in computational cost.

The strengths of the methodology developed in this work are demonstrated on a number of realistic test systems, such as doped *p*-terphenyl molecular crystals and the exciton coupling in the Fenna-Matthews-Olson complex of bacteriochlorophyll. It is shown that the coupled subsystem TDDFT approach allows for the treatment of system sizes inaccessible by previous methods.

Acknowledgments

The research outlined in this dissertation was carried out through the EPSRC Centre for Doctoral Training on Theory and Simulation of Materials at Imperial College London.

I would like to thank my supervisors Peter Haynes, James Spencer, Nicholas Harrison and Jason Riley for their constant support, helpful advice and stimulating discussions over the past 3 years.

I am indebted to Nick Hine for furthering my understanding of the inner workings of the ONETEP code as well as the PAW method.

Niccolo Corsini and Daniel Cole provided some of the atomistic structures of the systems studied in this work.

Further thanks are due to members of my cohort at the TSM CDT for creating a fun and stimulating learning environment. Special thanks go out to Tom Poole, David Trevelyan, Nina Kearsey, Fabian Renn and Benat Gurrutxaga-Lerma for providing much needed distractions in the form of a game of catch in the corridor of the Whiteley Suite.

Finally, I thank my family, both for their love and constant support and for giving me the opportunity to embark on that journey I started 9 years ago by coming to the UK.

Contents

Chapter 1
Introduction and Theoretical Prerequisites

The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems without too much computation.

Paul Dirac [1]

1.1 Introduction

1.1.1 Practical Methods

When the theory of quantum mechanics was developed in the early 20th century, it caused the most groundbreaking revolution in scientific thinking since the publication of Newton's *Principia* over 200 years prior. It was quickly recognised that the theory did in principle allow for the calculation of all properties of matter from the atomic scale all the way up to the macroscopic scale. Or, in Dirac's perhaps slightly optimistic words, quantum mechanics provided a mathematical theory for "a large part of physics and the whole of chemistry". However, it soon became equally apparent that the equations describing the fundamental behaviour of matter were far too complicated to be soluble for more than a handful of the most trivial cases.

For this reason, ever since the development of the basic theory of quantum mechanics, a significant research effort has been directed into developing "practical methods" that would allow for calculating approximate solutions to complex systems. The earliest approximations, like the WKBJ method (named after Wentzel, Kramers, Brillouin and Jeffreys, who each independently developed it) [2–5], focused on finding so-called "semiclassical solutions" to the Schrödinger equation, the main equation governing quantum mechanics. It has been used extensively in finding the behaviour

T.J. Zuehlsdorff, *Computing the Optical Properties of Large Systems*,
Springer Theses, DOI 10.1007/978-3-319-19770-8_1

of quantum particles in simplified models of tunneling processes, most notably to derive the rate of nuclear fusion [6].

The introduction of computers helped along a further development in quantum mechanics, namely the move away from studying simplified model systems towards finding approximate numerical solutions to real systems like molecules and solids. To unlock the predictive power of quantum mechanics, it becomes necessary to make use of methods that are truly *ab initio*, in that they do not make use of any *a priori* assumptions about behaviour of the system that is studied. However, the complexity of the quantum mechanical equations that Dirac already recognised in 1926 means that a direct numerical solution requires a computational effort that grows exponentially with the size of the system that is to be solved, rendering any such direct approach impractical for anything but the most simple problems. Thus in practice, further well-controlled approximations have to be made in order to reduce the scaling from an exponential to a polynomial one.

The most successful in this new generation of "practical methods" are based on density-functional theory (DFT) [7], which after making a number of simplifying approximations to the quantum mechanical effects of exchange and correlation of electrons, yields any ground state property of a system with a computational scaling that is just the cube of the system size. Since its development in 1964 it has shown real predictive power in the study of molecules and solids and has become the standard tool in areas as diverse as chemistry, biophysics and materials science. The great impact of DFT on a wide range of scientific disciplines was recognized in 1998 by awarding the Nobel Prize of Chemistry to Walter Kohn and John Pople.

1.1.2 Spanning Lengths

While the standard formulation of DFT has been remarkably successful in the last few decades, measured in both the annual citations of the papers describing the original method and the number of different scientific areas it is now routinely used in, most of its applications are focused on two limiting cases. On one end of the scale are small, isolated molecules, while the other end of the scale contains applications to infinite crystals. The reason why most applications are limited to these two regimes can be seen in the $\mathcal{O}(N^3)$ scaling with system size that is inherent in the method, meaning that a doubling of system size leads to an 8-fold increase in computational effort. Both the case of a small, isolated molecule and an infinite (defect-free) crystal can be computed by considering only a few atoms[1] and are easily treated by standard DFT with moderate computational effort.

[1] In the case of infinite systems, this is achieved by making use of periodic boundary conditions, such that only a unit cell of the crystal has to be treated explicitly and the translational crystal symmetry is exploited.

In recent years there has been an increased interest in systems that lie between these two limiting cases and that are prohibitively large for the $\mathcal{O}(N^3)$ scaling of DFT. Examples of these systems include large biomolecules, nanocrystals and infinite crystals containing defects. The feature all of these systems have in common is that an appropriate treatment of them requires calculations containing thousands of atoms, rendering them impractical using conventional DFT even on modern supercomputers.

Over the last 20 years, linear-scaling DFT methods have been designed specifically to address the class of systems mentioned above and are now routinely used to treat systems containing thousands of atoms. This new generation of methods has been used with great success to treat problems that were completely inaccessible even a few years ago.

1.1.3 Light Interactions

While the development of linear-scaling techniques has opened up a wide field of new applications for DFT (see, for example [8–11]), it also comes with a new set of challenges. Since calculations of ground state properties of systems containing thousands of atoms are now possible, the problem shifts to extracting information from these calculations that can be compared directly to experimental measurements.

Consider for example a crystal with two interacting defects in it. In order to find the lowest energy arrangement of the defects in the entire structure using linear-scaling DFT one has to generate all unique configurations of the two defects in the lattice and perform a full DFT calculation on every one of these structures to compare the total energies. While each of the individual DFT calculations does scale linearly with the number of lattice sites within a finite supercell, the number of nonequivalent defect configurations grows as $\mathcal{O}(N)$ with lattice sites, leading to an overall scaling of $\mathcal{O}(N^2)$. Thus in this example the scaling problem does not originate from the DFT calculation itself but from the configurational complexity that increases once the system size is increased. To make matters worse, the above example of two interacting defects is likely to produce a large array of potential structures, all separated by very small energy differences, any of which are likely to occur in the real material whose properties need to be predicted. Due to this configurational complexity in large systems, it is often found that predicting ground state properties of the single lowest energy configuration is of little use in making predictions of the behaviour of the entire system.

A way of sidestepping the configurational complexity problem leading to unfavourable scaling when linear-scaling DFT is applied to a large system is to attempt to predict properties that are measured directly in experiments. One of the properties of interest for many practical applications is the interaction of the system with light, which can be obtained from the time-dependent extension of the DFT method (TDDFT). Accurate predictions of absorption spectra of semiconductor nanocrystals are a key ingredient to developing new generations of efficient solar

cells, while predictions of spectra of large photoactive biomolecules is of great interest in the area of biophysics.

Optical absorption spectra of nanostructures and large photoactive biomolecules that are measured in experiments are taken on timescales that are several orders of magnitude larger than the timescales of atomic vibrations around their point of equilibrium. Thus the optical spectra produced in experiments can be seen as the result of an averaging over many low energy configurational structures of the system of interest and are often well reproduced by theoretical calculations on a single, average low energy structure of that system.[2] Thus in order to make meaningful predictions on the measurable quantity of light absorption, the configurational complexity problem that troubles many potential applications of linear scaling techniques can be, to a large extent, ignored, making theoretical spectroscopy an ideal area of impact for linear-scaling methods.

The purpose of this dissertation is to extend the linear-scaling techniques that proved so successful in standard DFT over recent years to the calculation of optical spectra using TDDFT. The aim is to develop methods capable of calculating the low energy optical spectrum of systems containing thousands of atoms in an effort scaling linearly with system size. This will open up new potential areas of research in the field of theoretical spectroscopy and connect directly to experimental measurements to enable a more effective collaboration between the experimental and the theoretical community.

1.2 Theoretical Prerequisites

In this section, some of the prerequisite background knowledge that is needed for the later chapters is discussed. The main focus lies on giving a very brief overview over some important aspects of quantum mechanics and introducing the most important choices of notation that are used throughout the rest of the work.

1.2.1 The Wavefunction

The wavefunction Ψ completely describes the quantum state of a system of N particles and contains all information about the system. While the wavefunction does not have a unique representation for an arbitrary quantum system, a typical representa-

[2] It should be noted that such an average, low energy structure cannot always be readily found in large biological systems where atomic positions are derived from X-ray diffraction experiments. Furthermore, while atoms in nanostructured crystals often undergo simple oscillations around a point of equilibrium, this is not necessarily the case in biological pigment-protein complexes, where the motion of the protein happening on a much longer timescale can become important. In these systems, it is often necessary to calculate spectra of several snapshots taken from a molecular dynamics simulation in order to achieve a good representation of the relevant phase space.