DENSITY FUNCTIONAL THEORY

A Practical Introduction

DAVID S. SHOLL Georgia Institute of Technology

JANICE A. STECKEL National Energy Technology Laboratory



DENSITY FUNCTIONAL THEORY

DENSITY FUNCTIONAL THEORY

A Practical Introduction

DAVID S. SHOLL Georgia Institute of Technology

JANICE A. STECKEL National Energy Technology Laboratory



Copyright © 2009 by John Wiley & Sons, Inc. All rights reserved. Prepared in part with support by the National Energy Technology Laboratory

Published by John Wiley & Sons, Inc., Hoboken, New Jersey Published simultaneously in Canada

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permission.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www. wiley.com.

Library of Congress Cataloging-in-Publication Data:

Sholl, David S.
Density functional theory : a practical introduction / David S. Sholl and Jan Steckel. p. cm.
Includes index.
ISBN 978-0-470-37317-0 (cloth)
1. Density functionals. 2. Mathematical physics. 3. Quantum chemistry. I. Steckel, Jan.
II. Title.
QC20.7.D43S55 2009
530.14'4—dc22

2008038603

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

CONTENTS

Preface

1 What Is Density Functional Theory?

1

xi

- 1.1 How to Approach This Book, 1
- 1.2 Examples of DFT in Action, 2
 - 1.2.1 Ammonia Synthesis by Heterogeneous Catalysis, 2
 - 1.2.2 Embrittlement of Metals by Trace Impurities, 4
 - 1.2.3 Materials Properties for Modeling Planetary Formation, 6
- 1.3 The Schrödinger Equation, 7
- 1.4 Density Functional Theory—From Wave Functions to Electron Density, 10
- 1.5 Exchange–Correlation Functional, 14
- 1.6 The Quantum Chemistry Tourist, 16
 - 1.6.1 Localized and Spatially Extended Functions, 16
 - 1.6.2 Wave-Function-Based Methods, 18
 - 1.6.3 Hartree-Fock Method, 19
 - 1.6.4 Beyond Hartree–Fock, 23
- 1.7 What Can DFT Not Do?, 28
- 1.8 Density Functional Theory in Other Fields, 30
- 1.9 How to Approach This Book (Revisited), 30

References, 31

Further Reading, 32

2 DFT Calculations for Simple Solids

- 2.1 Periodic Structures, Supercells, and Lattice Parameters, 35
- 2.2 Face-Centered Cubic Materials, 39
- 2.3 Hexagonal Close-Packed Materials, 41
- 2.4 Crystal Structure Prediction, 43
- 2.5 Phase Transformations, 44

Exercises, 46

Further Reading, 47

Appendix Calculation Details, 47

3 Nuts and Bolts of DFT Calculations

- 3.1 Reciprocal Space and k Points, 50
 - 3.1.1 Plane Waves and the Brillouin Zone, 50
 - 3.1.2 Integrals in k Space, 53
 - 3.1.3 Choosing *k* Points in the Brillouin Zone, 55
 - 3.1.4 Metals—Special Cases in k Space, 59
 - 3.1.5 Summary of *k* Space, 60
- 3.2 Energy Cutoffs, 61
 - 3.2.1 Pseudopotentials, 63
- 3.3 Numerical Optimization, 65
 - 3.3.1 Optimization in One Dimension, 65
 - 3.3.2 Optimization in More than One Dimension, 69
 - 3.3.3 What Do I Really Need to Know about Optimization?, 73
- 3.4 DFT Total Energies—An Iterative Optimization Problem, 73
- 3.5 Geometry Optimization, 75
 - 3.5.1 Internal Degrees of Freedom, 75
 - 3.5.2 Geometry Optimization with Constrained Atoms, 78
 - 3.5.3 Optimizing Supercell Volume and Shape, 78

Exercises, 79

References, 80

Further Reading, 80

Appendix Calculation Details, 81

4 DFT Calculations for Surfaces of Solids

- 4.1 Importance of Surfaces, 83
- 4.2 Periodic Boundary Conditions and Slab Models, 84
- 4.3 Choosing k Points for Surface Calculations, 87
- 4.4 Classification of Surfaces by Miller Indices, 88
- 4.5 Surface Relaxation, 94
- 4.6 Calculation of Surface Energies, 96

49

83

4.7 Symmetric and Asymmetric Slab Models, 98 4.8 Surface Reconstruction, 100 4.9 Adsorbates on Surfaces, 103 4.9.1 Accuracy of Adsorption Energies, 106 4.10 Effects of Surface Coverage, 107 Exercises, 110 References, 111 Further Reading, 111 Appendix Calculation Details, 112 **DFT Calculations of Vibrational Frequencies** 5 113 5.1 Isolated Molecules, 114 Vibrations of a Collection of Atoms, 117 5.2 5.3 Molecules on Surfaces, 120 5.4 Zero-Point Energies, 122 5.5 Phonons and Delocalized Modes, 127 Exercises. 128 Reference, 128 Further Reading, 128 Appendix Calculation Details, 129 6 Calculating Rates of Chemical Processes Using **Transition State Theory** 131 6.1 **One-Dimensional Example**, 132 6.2 Multidimensional Transition State Theory, 139 6.3 Finding Transition States, 142 6.3.1 Elastic Band Method, 144 6.3.2 Nudged Elastic Band Method, 145 6.3.3 Initializing NEB Calculations, 147 6.4 Finding the Right Transition States, 150 Connecting Individual Rates to Overall Dynamics, 153 6.5 6.6 Quantum Effects and Other Complications, 156 High Temperatures/Low Barriers, 156 6.6.1 6.6.2 Quantum Tunneling, 157 6.6.3 Zero-Point Energies, 157 Exercises, 158 Reference, 159 Further Reading, 159 Appendix Calculation Details, 160

7	Equilibrium Phase Diagrams from <i>Ab Initio</i> Thermodynamics		163
	7.1	Stability of Bulk Metal Oxides, 1647.1.1 Examples Including Disorder—Configurational Entropy, 169	
	7.2 7.3	Stability of Metal and Metal Oxide Surfaces, 172 Multiple Chemical Potentials and Coupled Chemical Reactions, 174	
	Exe	prcises, 175	
	Refe	erences, 176	
	App	pendix Calculation Details, 177	
8	Elec	ctronic Structure and Magnetic Properties	179
	8.1 8.2 8.3 Exer	Electronic Density of States, 179 Local Density of States and Atomic Charges, 186 Magnetism, 188 precises, 190	
	Further Reading, 191		
	Арр	bendix Calculation Details, 192	
9	Ab Initio Molecular Dynamics		
	9.1	Classical Molecular Dynamics, 193 9.1.1 Molecular Dynamics with Constant Energy, 193	
		9.1.2 Molecular Dynamics in the Canonical Ensemble, 196	
		9.1.3 Practical Aspects of Classical Molecular Dynamics, 197	
	9.2	Ab Initio Molecular Dynamics, 198	
	9.3	 Applications of <i>Ab Initio</i> Molecular Dynamics, 201 9.3.1 Exploring Structurally Complex Materials: Liquids and Amorphous Phases, 201 	
		9.3.2 Exploring Complex Energy Surfaces, 204	
	Exercises, 207		
	Reference, 207		
	Furt	ther Reading, 207	
	Арр	bendix Calculation Details, 208	

10 Accuracy and Methods beyond "Standard" Calculations 209

- 10.1 How Accurate Are DFT Calculations?, 209
- 10.2 Choosing a Functional, 215
- 10.3 Examples of Physical Accuracy, 220
 - 10.3.1 Benchmark Calculations for Molecular Systems—Energy and Geometry, 220
 - 10.3.2 Benchmark Calculations for Molecular Systems—Vibrational Frequencies, 221
 - 10.3.3 Crystal Structures and Cohesive Energies, 222
 - 10.3.4 Adsorption Energies and Bond Strengths, 223
- 10.4 DFT+X Methods for Improved Treatment of Electron Correlation, 224
 - 10.4.1 Dispersion Interactions and DFT-D, 225
 - 10.4.2 Self-Interaction Error, Strongly Correlated Electron Systems, and DFT+U, 227
- 10.5 Larger System Sizes with Linear Scaling Methods and Classical Force Fields, 229
- 10.6 Conclusion, 230

References, 231

Further Reading, 232

Index

235

PREFACE

The application of density functional theory (DFT) calculations is rapidly becoming a "standard tool" for diverse materials modeling problems in physics, chemistry, materials science, and multiple branches of engineering. Although a number of highly detailed books and articles on the theoretical foundations of DFT are available, it remains difficult for a newcomer to these methods to rapidly learn the tools that allow him or her to actually perform calculations that are now routine in the fields listed above. This book aims to fill this gap by guiding the reader through the applications of DFT that might be considered the core of continually growing scientific literature based on these methods. Each chapter includes a series of exercises to give readers experience with calculations of their own.

We have aimed to find a balance between brevity and detail that makes it possible for readers to realistically plan to read the entire text. This balance inevitably means certain technical details are explored in a limited way. Our choices have been strongly influenced by our interactions over multiple years with graduate students and postdocs in chemical engineering, physics, chemistry, materials science, and mechanical engineering at Carnegie Mellon University and the Georgia Institute of Technology. A list of Further Reading is provided in each chapter to define appropriate entry points to more detailed treatments of the area. These reading lists should be viewed as identifying highlights in the literature, not as an effort to rigorously cite all relevant work from the thousands of studies that exist on these topics. One important choice we made to limit the scope of the book was to focus solely on one DFT method suitable for solids and spatially extended materials, namely plane-wave DFT. Although many of the foundations of plane-wave DFT are also relevant to complementary approaches used in the chemistry community for isolated molecules, there are enough differences in the applications of these two groups of methods that including both approaches would only have been possible by significantly expanding the scope of the book. Moreover, several resources already exist that give a practical "handson" introduction to computational chemistry calculations for molecules.

Our use of DFT calculations in our own research and our writing of this book has benefited greatly from interactions with numerous colleagues over an extended period. We especially want to acknowledge J. Karl Johnson (University of Pittsburgh), Aravind Asthagiri (University of Florida), Dan Sorescu (National Energy Technology Laboratory), Cathy Stampfl (University of Sydney), John Kitchin (Carnegie Mellon University), and Duane Johnson (University of Illinois). We thank Jeong-Woo Han for his help with a number of the figures. Bill Schneider (University of Notre Dame), Ken Jordan (University of Pittsburgh), and Taku Watanabe (Georgia Institute of Technology) gave detailed and helpful feedback on draft versions. Any errors or inaccuracies in the text are, of course, our responsibility alone.

DSS dedicates this book to his father and father-in-law, whose love of science and curiosity about the world are an inspiration. JAS dedicates this book to her husband, son, and daughter.

DAVID SHOLL

Georgia Institute of Technology, Atlanta, GA, USA

JAN STECKEL

National Energy Technology Laboratory, Pittsburgh, PA, USA

1

WHAT IS DENSITY FUNCTIONAL THEORY?

1.1 HOW TO APPROACH THIS BOOK

There are many fields within the physical sciences and engineering where the key to scientific and technological progress is understanding and controlling the properties of matter at the level of individual atoms and molecules. Density functional theory is a phenomenally successful approach to finding solutions to the fundamental equation that describes the quantum behavior of atoms and molecules, the Schrödinger equation, in settings of practical value. This approach has rapidly grown from being a specialized art practiced by a small number of physicists and chemists at the cutting edge of quantum mechanical theory to a tool that is used regularly by large numbers of researchers in chemistry, physics, materials science, chemical engineering, geology, and other disciplines. A search of the *Science Citation Index* for articles published in 1986 with the words "density functional theory" in the title or abstract yields less than 50 entries. Repeating this search for 1996 and 2006 gives more than 1100 and 5600 entries, respectively.

Our aim with this book is to provide just what the title says: an *introduction* to using density functional theory (DFT) calculations in a *practical* context. We do not assume that you have done these calculations before or that you even understand what they are. We do assume that you want to find out what is possible with these methods, either so you can perform calculations

Density Functional Theory: A Practical Introduction. By David S. Sholl and Janice A. Steckel Copyright © 2009 John Wiley & Sons, Inc.

yourself in a research setting or so you can interact knowledgeably with collaborators who use these methods.

An analogy related to cars may be useful here. Before you learned how to drive, it was presumably clear to you that you can accomplish many useful things with the aid of a car. For you to use a car, it is important to understand the basic concepts that control cars (you need to put fuel in the car regularly, you need to follow basic traffic laws, etc.) and spend time actually driving a car in a variety of road conditions. You do not, however, need to know every detail of how fuel injectors work, how to construct a radiator system that efficiently cools an engine, or any of the other myriad of details that are required if you were going to actually build a car. Many of these details may be important if you plan on undertaking some especially difficult car-related project such as, say, driving yourself across Antarctica, but you can make it across town to a friend's house and back without understanding them.

With this book, we hope you can learn to "drive across town" when doing your own calculations with a DFT package or when interpreting other people's calculations as they relate to physical questions of interest to you. If you are interested in "building a better car" by advancing the cutting edge of method development in this area, then we applaud your enthusiasm. You should continue reading this chapter to find at least one surefire project that could win you a Nobel prize, then delve into the books listed in the Further Reading at the end of the chapter.

At the end of most chapters we have given a series of exercises, most of which involve actually doing calculations using the ideas described in the chapter. Your knowledge and ability will grow most rapidly by doing rather than by simply reading, so we strongly recommend doing as many of the exercises as you can in the time available to you.

1.2 EXAMPLES OF DFT IN ACTION

Before we even define what density functional theory is, it is useful to relate a few vignettes of how it has been used in several scientific fields. We have chosen three examples from three quite different areas of science from the thousands of articles that have been published using these methods. These specific examples have been selected because they show how DFT calculations have been used to make important contributions to a diverse range of compelling scientific questions, generating information that would be essentially impossible to determine through experiments.

1.2.1 Ammonia Synthesis by Heterogeneous Catalysis

Our first example involves an industrial process of immense importance: the catalytic synthesis of ammonia (NH₃). Ammonia is a central component of

fertilizers for agriculture, and more than 100 million tons of ammonia are produced commercially each year. By some estimates, more than 1% of all energy used in the world is consumed in the production of ammonia. The core reaction in ammonia production is very simple:

$$N_2 + 3H_2 \rightarrow 2NH_3$$

To get this reaction to proceed, the reaction is performed at high temperatures (>400°C) and high pressures (>100 atm) in the presence of metals such as iron (Fe) or ruthenium (Ru) that act as catalysts. Although these metal catalysts were identified by Haber and others almost 100 years ago, much is still not known about the mechanisms of the reactions that occur on the surfaces of these catalysts. This incomplete understanding is partly because of the structural complexity of practical catalysts. To make metal catalysts with high surface areas, tiny particles of the active metal are dispersed throughout highly porous materials. This was a widespread application of nanotechnology long before that name was applied to materials to make them sound scientifically exciting! To understand the reactivity of a metal nanoparticle, it is useful to characterize the surface atoms in terms of their local coordination since differences in this coordination can create differences in chemical reactivity; surface atoms can be classified into "types" based on their local coordination. The surfaces of nanoparticles typically include atoms of various types (based on coordination), so the overall surface reactivity is a complicated function of the shape of the nanoparticle and the reactivity of each type of atom.

The discussion above raises a fundamental question: Can a direct connection be made between the shape and size of a metal nanoparticle and its activity as a catalyst for ammonia synthesis? If detailed answers to this question can be found, then they can potentially lead to the synthesis of improved catalysts. One of the most detailed answers to this question to date has come from the DFT calculations of Honkala and co-workers,¹ who studied nanoparticles of Ru. Using DFT calculations, they showed that the net chemical reaction above proceeds via at least 12 distinct steps on a metal catalyst and that the rates of these steps depend strongly on the local coordination of the metal atoms that are involved. One of the most important reactions is the breaking of the N₂ bond on the catalyst surface. On regions of the catalyst surface that were similar to the surfaces of bulk Ru (more specifically, atomically flat regions), a great deal of energy is required for this bond-breaking reaction, implying that the reaction rate is extremely slow. Near Ru atoms that form a common kind of surface step edge on the catalyst, however, a much smaller amount of energy is needed for this reaction. Honkala and co-workers used additional DFT calculations to predict the relative stability of many different local coordinations of surface atoms in Ru nanoparticles in a way that allowed them to predict the detailed shape of the nanoparticles as a function of particle size. This prediction makes a precise connection between the diameter of a Ru nanoparticle and the number of highly desirable reactive sites for breaking N_2 bonds on the nanoparticle. Finally, all of these calculations were used to develop an overall model that describes how the individual reaction rates for the many different kinds of metal atoms on the nanoparticle's surfaces couple together to define the overall reaction rate under realistic reaction conditions. At no stage in this process was any experimental data used to fit or adjust the model, so the final result was a truly predictive description of the reaction rate of a complex catalyst. After all this work was done, Honkala et al. compared their predictions to experimental measurements made with Ru nanoparticle catalysts under reaction conditions similar to industrial conditions. Their predictions were in stunning quantitative agreement with the experimental outcome.

1.2.2 Embrittlement of Metals by Trace Impurities

It is highly likely that as you read these words you are within 1 m of a large number of copper wires since copper is the dominant metal used for carrying electricity between components of electronic devices of all kinds. Aside from its low cost, one of the attractions of copper in practical applications is that it is a soft, ductile metal. Common pieces of copper (and other metals) are almost invariably polycrystalline, meaning that they are made up of many tiny domains called grains that are each well-oriented single crystals. Two neighboring grains have the same crystal structure and symmetry, but their orientation in space is not identical. As a result, the places where grains touch have a considerably more complicated structure than the crystal structure of the pure metal. These regions, which are present in all polycrystalline materials, are called grain boundaries.

It has been known for over 100 years that adding tiny amounts of certain impurities to copper can change the metal from being ductile to a material that will fracture in a brittle way (i.e., without plastic deformation before the fracture). This occurs, for example, when bismuth (Bi) is present in copper (Cu) at levels below 100 ppm. Similar effects have been observed with lead (Pb) or mercury (Hg) impurities. But how does this happen? Qualitatively, when the impurities cause brittle fracture, the fracture tends to occur at grain boundaries, so something about the impurities changes the properties of grain boundaries in a dramatic way. That this can happen at very low concentrations of Bi is not completely implausible because Bi is almost completely insoluble in bulk Cu. This means that it is very favorable for Bi atoms to segregate to grain boundaries rather than to exist inside grains, meaning that the local concentration of Bi at grain boundaries can be much higher than the net concentration in the material as a whole.

Can the changes in copper caused by Bi be explained in a detailed way? As you might expect for an interesting phenomena that has been observed over many years, several alternative explanations have been suggested. One class of explanations assigns the behavior to electronic effects. For example, a Bi atom might cause bonds between nearby Cu atoms to be stiffer than they are in pure Cu, reducing the ability of the Cu lattice to deform smoothly. A second type of electronic effect is that having an impurity atom next to a grain boundary could weaken the bonds that exist across a boundary by changing the electronic structure of the atoms, which would make fracture at the boundary more likely. A third explanation assigns the blame to size effects, noting that Bi atoms are much larger than Cu atoms. If a Bi atom is present at a grain boundary, then it might physically separate Cu atoms on the other side of the boundary from their natural spacing. This stretching of bond distances would weaken the bonds between atoms and make fracture of the grain boundary more likely. Both the second and third explanations involve weakening of bonds near grain boundaries, but they propose different root causes for this behavior. Distinguishing between these proposed mechanisms would be very difficult using direct experiments.

Recently, Schweinfest, Paxton, and Finnis used DFT calculations to offer a definitive description of how Bi embrittles copper; the title of their study gives away the conclusion.² They first used DFT to predict stress-strain relationships for pure Cu and Cu containing Bi atoms as impurities. If the bond stiffness argument outlined above was correct, the elastic moduli of the metal should be increased by adding Bi. In fact, the calculations give the opposite result, immediately showing the bond-stiffening explanation to be incorrect. In a separate and much more challenging series of calculations, they explicitly calculated the cohesion energy of a particular grain boundary that is known experimentally to be embrittled by Bi. In qualitative consistency with experimental observations, the calculations predicted that the cohesive energy of the grain boundary is greatly reduced by the presence of Bi. Crucially, the DFT results allow the electronic structure of the grain boundary atoms to be examined directly. The result is that the grain boundary electronic effect outlined above was found to not be the cause of embrittlement. Instead, the large change in the properties of the grain boundary could be understood almost entirely in terms of the excess volume introduced by the Bi atoms, that is, by a size effect. This reasoning suggests that Cu should be embrittled by any impurity that has a much larger atomic size than Cu and that strongly segregates to grain boundaries. This description in fact correctly describes the properties of both Pb and Hg as impurities in Cu, and, as mentioned above, these impurities are known to embrittle Cu.

1.2.3 Materials Properties for Modeling Planetary Formation

To develop detailed models of how planets of various sizes have formed, it is necessary to know (among many other things) what minerals exist inside planets and how effective these minerals are at conducting heat. The extreme conditions that exist inside planets pose some obvious challenges to probing these topics in laboratory experiments. For example, the center of Jupiter has pressures exceeding 40 Mbar and temperatures well above 15,000 K. DFT calculations can play a useful role in probing material properties at these extreme conditions, as shown in the work of Umemoto, Wentzcovitch, and Allen.³ This work centered on the properties of bulk MgSiO₃, a silicate mineral that is important in planet formation. At ambient conditions, MgSiO₃ forms a relatively common crystal structure known as a perovskite. Prior to Umemoto et al.'s calculations, it was known that if MgSiO₃ was placed under conditions similar to those in the core-mantle boundary of Earth, it transforms into a different crystal structure known as the CaIrO₃ structure. (It is conventional to name crystal structures after the first compound discovered with that particular structure, and the naming of this structure is an example of this convention.)

Umemoto et al. wanted to understand what happens to the structure of MgSiO₃ at conditions much more extreme than those found in Earth's core-mantle boundary. They used DFT calculations to construct a phase diagram that compared the stability of multiple possible crystal structures of solid MgSiO₃. All of these calculations dealt with bulk materials. They also considered the possibility that MgSiO₃ might dissociate into other compounds. These calculations predicted that at pressures of ~11 Mbar, MgSiO₃ dissociates in the following way:

$$MgSiO_3$$
 [CaIrO₃ structure] $\rightarrow MgO$ [CsCl structure]

+ SiO₂ [cotunnite structure].

In this reaction, the crystal structure of each compound has been noted in the square brackets. An interesting feature of the compounds on the right-hand side is that neither of them is in the crystal structure that is the stable structure at ambient conditions. MgO, for example, prefers the NaCl structure at ambient conditions (i.e., the same crystal structure as everyday table salt). The behavior of SiO_2 is similar but more complicated; this compound goes through several intermediate structures between ambient conditions and the conditions relevant for MgSiO₃ dissociation. These transformations in the structures of MgO and SiO₂ allow an important connection to be made between DFT calculations and experiments since these transformations occur at conditions that can be directly probed in laboratory experiments. The transition pressures

predicted using DFT and observed experimentally are in good agreement, giving a strong indication of the accuracy of these calculations.

The dissociation reaction predicted by Umemoto et al.'s calculations has important implications for creating good models of planetary formation. At the simplest level, it gives new information about what materials exist inside large planets. The calculations predict, for example, that the center of Uranus or Neptune can contain MgSiO₃, but that the cores of Jupiter or Saturn will not. At a more detailed level, the thermodynamic properties of the materials can be used to model phenomena such as convection inside planets. Umemoto et al. speculated that the dissociation reaction above might severely limit convection inside "dense-Saturn," a Saturn-like planet that has been discovered outside the solar system with a mass of ~67 Earth masses.

A legitimate concern about theoretical predictions like the reaction above is that it is difficult to envision how they can be validated against experimental data. Fortunately, DFT calculations can also be used to search for similar types of reactions that occur at pressures that are accessible experimentally. By using this approach, it has been predicted that NaMgF₃ goes through a series of transformations similar to MgSiO₃; namely, a perovskite to postperovskite transition at some pressure above ambient and then dissociation in NaF and MgF₂ at higher pressures.⁴ This dissociation is predicted to occur for pressures around 0.4 Mbar, far lower than the equivalent pressure for MgSiO₃. These predictions suggest an avenue for direct experimental tests of the transformation mechanism that DFT calculations have suggested plays a role in planetary formation.

We could fill many more pages with research vignettes showing how DFT calculations have had an impact in many areas of science. Hopefully, these three examples give some flavor of the ways in which DFT calculations can have an impact on scientific understanding. It is useful to think about the common features between these three examples. All of them involve materials in their solid state, although the first example was principally concerned with the interface between a solid and a gas. Each example generated information about a physical problem that is controlled by the properties of materials on atomic length scales that would be (at best) extraordinarily challenging to probe experimentally. In each case, the calculations were used to give information not just about some theoretically ideal state, but instead to understand phenomena at temperatures, pressures, and chemical compositions of direct relevance to physical applications.

1.3 THE SCHRÖDINGER EQUATION

By now we have hopefully convinced you that density functional theory is a useful and interesting topic. But what is it exactly? We begin with the observation that one of the most profound scientific advances of the twentieth century was the development of quantum mechanics and the repeated experimental observations that confirmed that this theory of matter describes, with astonishing accuracy, the universe in which we live.

In this section, we begin a review of some key ideas from quantum mechanics that underlie DFT (and other forms of computational chemistry). Our goal here is not to present a complete derivation of the techniques used in DFT. Instead, our goal is to give a clear, brief, introductory presentation of the most basic equations important for DFT. For the full story, there are a number of excellent texts devoted to quantum mechanics listed in the Further Reading section at the end of the chapter.

Let us imagine a situation where we would like to describe the properties of some well-defined collection of atoms-you could think of an isolated molecule or the atoms defining the crystal of an interesting mineral. One of the fundamental things we would like to know about these atoms is their energy and, more importantly, how their energy changes if we move the atoms around. To define where an atom is, we need to define both where its nucleus is and where the atom's electrons are. A key observation in applying quantum mechanics to atoms is that atomic nuclei are much heavier than individual electrons; each proton or neutron in a nucleus has more than 1800 times the mass of an electron. This means, roughly speaking, that electrons respond much more rapidly to changes in their surroundings than nuclei can. As a result, we can split our physical question into two pieces. First, we solve, for fixed positions of the atomic nuclei, the equations that describe the electron motion. For a given set of electrons moving in the field of a set of nuclei, we find the lowest energy configuration, or state, of the electrons. The lowest energy state is known as the ground state of the electrons, and the separation of the nuclei and electrons into separate mathematical problems is the Born-Oppenheimer approximation. If we have M nuclei at positions $\mathbf{R}_1, \ldots, \mathbf{R}_M$, then we can express the ground-state energy, E, as a function of the positions of these nuclei, $E(\mathbf{R}_1, \ldots, \mathbf{R}_M)$. This function is known as the adiabatic potential energy surface of the atoms. Once we are able to calculate this potential energy surface we can tackle the original problem posed above-how does the energy of the material change as we move its atoms around?

One simple form of the Schrödinger equation—more precisely, the timeindependent, nonrelativistic Schrödinger equation—you may be familiar with is $H\psi = E\psi$. This equation is in a nice form for putting on a T-shirt or a coffee mug, but to understand it better we need to define the quantities that appear in it. In this equation, *H* is the Hamiltonian operator and ψ is a set of solutions, or eigenstates, of the Hamiltonian. Each of these solutions, ψ_n , has an associated eigenvalue, E_n , a real number^{*} that satisfies the eigenvalue equation. The detailed definition of the Hamiltonian depends on the physical system being described by the Schrödinger equation. There are several well-known examples like the particle in a box or a harmonic oscillator where the Hamiltonian has a simple form and the Schrödinger equation can be solved exactly. The situation we are interested in where multiple electrons are interacting with multiple nuclei is more complicated. In this case, a more complete description of the Schrödinger is

$$\left[-\frac{\hbar^2}{2m}\sum_{i=1}^N \nabla_i^2 + \sum_{i=1}^N V(\mathbf{r}_i) + \sum_{i=1}^N \sum_{j < i} U(\mathbf{r}_i, \mathbf{r}_j)\right] \psi = E \psi.$$
(1.1)

Here, *m* is the electron mass. The three terms in brackets in this equation define, in order, the kinetic energy of each electron, the interaction energy between each electron and the collection of atomic nuclei, and the interaction energy between different electrons. For the Hamiltonian we have chosen, ψ is the electronic wave function, which is a function of each of the spatial coordinates of each of the *N* electrons, so $\psi = \psi(\mathbf{r}_1, \dots, \mathbf{r}_N)$, and *E* is the ground-state energy of the electrons.^{**} The ground-state energy is independent of time, so this is the time-independent Schrödinger equation.[†]

Although the electron wave function is a function of each of the coordinates of all *N* electrons, it is possible to approximate ψ as a product of individual electron wave functions, $\psi = \psi_1(\mathbf{r})\psi_2(\mathbf{r}), \dots, \psi_N(\mathbf{r})$. This expression for the wave function is known as a Hartree product, and there are good motivations for approximating the full wave function into a product of individual oneelectron wave functions in this fashion. Notice that *N*, the number of electrons, is considerably larger than *M*, the number of nuclei, simply because each atom has one nucleus and lots of electrons. If we were interested in a single molecule of CO₂, the full wave function is a 66-dimensional function (3 dimensions for each of the 22 electrons). If we were interested in a nanocluster of 100 Pt atoms, the full wave function requires more the 23,000 dimensions! These numbers should begin to give you an idea about why solving the Schrödinger equation for practical materials has occupied many brilliant minds for a good fraction of a century.

^{*}The value of the functions ψ_n are complex numbers, but the eigenvalues of the Schrödinger equation are real numbers.

^{**}For clarity of presentation, we have neglected electron spin in our description. In a complete presentation, each electron is defined by three spatial variables and its spin.

[†]The dynamics of electrons are defined by the time-dependent Schrödinger equation, $i\hbar(\partial\psi/\partial t) = H\psi$. The appearance of $i = \sqrt{-1}$ in this equation makes it clear that the wave function is a complex-valued function, not a real-valued function.

The situation looks even worse when we look again at the Hamiltonian, H. The term in the Hamiltonian defining electron–electron interactions is the most critical one from the point of view of solving the equation. The form of this contribution means that the individual electron wave function we defined above, $\psi_i(\mathbf{r})$, cannot be found without simultaneously considering the individual electron wave functions associated with all the other electrons. In other words, the Schrödinger equation is a many-body problem.

Although solving the Schrödinger equation can be viewed as the fundamental problem of quantum mechanics, it is worth realizing that the wave function for any particular set of coordinates cannot be directly observed. The quantity that can (in principle) be measured is the probability that the *N* electrons are at a particular set of coordinates, $\mathbf{r}_1, \ldots, \mathbf{r}_N$. This probability is equal to $\psi^*(\mathbf{r}_1, \ldots, \mathbf{r}_N)\psi(\mathbf{r}_1, \ldots, \mathbf{r}_N)$, where the asterisk indicates a complex conjugate. A further point to notice is that in experiments we typically do not care which electron in the material is labeled electron 1, electron 2, and so on. Moreover, even if we did care, we cannot easily assign these labels. This means that the quantity of physical interest is really the probability that a set of *N* electrons in any order have coordinates $\mathbf{r}_1, \ldots, \mathbf{r}_N$. A closely related quantity is the density of electrons at a particular position in space, $n(\mathbf{r})$. This can be written in terms of the individual electron wave functions as

$$n(\mathbf{r}) = 2\sum_{i} \psi_{i}^{*}(\mathbf{r})\psi_{i}(\mathbf{r}).$$
(1.2)

Here, the summation goes over all the individual electron wave functions that are occupied by electrons, so the term inside the summation is the probability that an electron in individual wave function $\psi_i(\mathbf{r})$ is located at position \mathbf{r} . The factor of 2 appears because electrons have spin and the Pauli exclusion principle states that each individual electron wave function can be occupied by two separate electrons provided they have different spins. This is a purely quantum mechanical effect that has no counterpart in classical physics. The point of this discussion is that the electron density, $n(\mathbf{r})$, which is a function of only three coordinates, contains a great amount of the information that is actually physically observable from the full wave function solution to the Schrödinger equation, which is a function of 3N coordinates.

1.4 DENSITY FUNCTIONAL THEORY—FROM WAVE FUNCTIONS TO ELECTRON DENSITY

The entire field of density functional theory rests on two fundamental mathematical theorems proved by Kohn and Hohenberg and the derivation of a set of equations by Kohn and Sham in the mid-1960s. The first theorem, proved by Hohenberg and Kohn, is: *The ground-state energy from Schrödinger's equation is a unique functional of the electron density.*

This theorem states that there exists a one-to-one mapping between the ground-state wave function and the ground-state electron density. To appreciate the importance of this result, you first need to know what a "functional" is. As you might guess from the name, a functional is closely related to the more familiar concept of a function. A function takes a value of a variable or variables and defines a single number from those variables. A simple example of a function dependent on a single variable is $f(x) = x^2 + 1$. A *functional* is similar, but it takes a function and defines a single number from the function. For example,

$$F[f] = \int_{-1}^{1} f(x) \, dx,$$

is a functional of the function f(x). If we evaluate this functional using $f(x) = x^2 + 1$, we get $F[f] = \frac{8}{3}$. So we can restate Hohenberg and Kohn's result by saying that the ground-state energy *E* can be expressed as $E[n(\mathbf{r})]$, where $n(\mathbf{r})$ is the electron density. This is why this field is known as density functional theory.

Another way to restate Hohenberg and Kohn's result is that the ground-state electron density uniquely determines all properties, including the energy and wave function, of the ground state. Why is this result important? It means that we can think about solving the Schrödinger equation by finding a function of three spatial variables, the electron density, rather than a function of 3*N* variables, the wave function. Here, by "solving the Schrödinger equation" we mean, to say it more precisely, finding the ground-state energy. So for a nanocluster of 100 Pd atoms the theorem reduces the problem from something with more than 23,000 dimensions to a problem with just 3 dimensions.

Unfortunately, although the first Hohenberg–Kohn theorem rigorously proves that a functional of the electron density exists that can be used to solve the Schrödinger equation, the theorem says nothing about what the functional actually is. The second Hohenberg–Kohn theorem defines an important property of the functional: *The electron density that minimizes the energy of the overall functional is the true electron density corresponding to the full solution of the Schrödinger equation*. If the "true" functional form were known, then we could vary the electron density until the energy from the functional is minimized, giving us a prescription for finding the relevant electron density. This variational principle is used in practice with approximate forms of the functional.