


CAMBRIDGE

国际著名物理图书——影印版系列

Michael A. Nielsen, Isaac L. Chuang

量子计算 与量子信息

(10周年版)



清华大学出版社

CAMBRIDGE

国际著名物理图书——影印版系列

量子计算与量子信息

(10周年版)

**Quantum Computation
and Quantum Information**

(MICHAEL A. NIELSEN and ISAAC L. CHUANG)

Michael A. Nielsen
Isaac L. Chuang

清华大学出版社
北京

北京市版权局著作权合同登记号 图字：01-2014-5097

This is a reprint edition of the following title(s) published by Cambridge University Press:

Quantum Computation and Quantum Information, 10th Anniversary Edition

ISBN: 978-1-107-00217-3

This publication is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

This reprint edition for the People's Republic of China (excluding Hong Kong SAR, Macao SAR and Taiwan) is published by arrangement with the Press Syndicate of the University of Cambridge, Cambridge, United Kingdom.

© Cambridge University Press and Tsinghua University Press 2015.

This reprint edition is authorized for sale in the People's Republic of China (excluding Hong Kong SAR, Macao SAR and Taiwan) only. Unauthorized export of this reprint edition is a violation of the Copyright Act. No part of this publication may be reproduced or distributed by any means, or stored in a database or retrieval system, without the prior written permission of Cambridge University Press and Tsinghua University Press.

此版本仅限在中华人民共和国境内（不包括中国香港、澳门特别行政区及中国台湾地区）销售。

版权所有，侵权必究。侵权举报电话：010-62782989 13701121933

图书在版编目(CIP)数据

量子计算与量子信息：10周年版=Quantum computation and quantum information:10th anniversary edition: 英文/(美)尼尔森(Nielsen, M.A.), (美)庄(Chuang, I.L.)著.--北京:清华大学出版社, 2015

(国际著名物理图书·影印版系列)

ISBN 978-7-302-39485-3

I. ①量… II. ①尼… ②庄… III. ①量子力学—计算方法—教材—英文②第五代计算机—教材—英文 IV. ①O413.1②TP387

中国版本图书馆CIP数据核字(2015)第035990号

责任编辑: 石磊

封面设计: 常雪影

责任印制: 沈露

出版者: 清华大学出版社

网 址: <http://www.tup.com.cn>, <http://www.wqbook.com>

地 址: 北京清华大学学研大厦A座 邮 编: 100084

社总机: 010-62770175 邮 购: 010-62786544

投稿与读者服务: 010-62776969, c-service@tup.tsinghua.edu.cn

质量反馈: 010-62772015, zhiliang@tup.tsinghua.edu.cn

印刷者: 北京鑫丰华彩印有限公司

装订者: 北京市密云县京文制本装订厂

发行者: 全国新华书店

开 本: 175mm×245mm

印 张: 44.5

版 次: 2015年10月第1版

印 次: 2015年10月第1次印刷

定 价: 99.00元

产品编号: 049383-01

影印版序

量子信息处理可以带来许多意想不到的具有特殊优势的结果。量子算法可以有效地进行大数质因数分解。在量子算法背景下，很多经典保密通信协议的安全性受到威胁，然而量子保密通信可以抵抗来自包括量子计算机在内的任何针对通道的攻击。由于存在诸多特殊应用，量子计算与量子信息科学近年来得到蓬勃发展。

对于这一相对年轻但又具有广阔发展前景的学科，优秀的系统化的教材比较缺乏。本书是本领域公认的最权威的系统化教材之一，也几乎是本学科研究人员的必备基本材料，有学者曾将之称为量子计算与量子信息学科教材中的“圣经”。

本书分三大部分。

第 1 部分为总论和量子物理学基础，还包括少量计算机科学基础知识。量子物理学基础部分主要介绍学习量子计算与量子信息所必需的量子力学基础知识，这一部分采用了侧重于数学框架和公理化体系的讲述方法，从而更便于为非物理专业的读者所理解。

第 2 部分讲述量子计算，包括量子算法和实现量子算法的一些物理系统的基础内容。这部分首先介绍了实现普适量子计算所需要的基本逻辑门——单量子比特门与条件非门（第 4 章）。之后较详细地介绍了现有的两个主要量子算法问题，即质因数分解问题（第 5 章）和搜索问题（第 6 章）。第 7 章讲物理实现，介绍了几个主要的实现条件非门的物理系统；这些物理系统虽然目前尚未实现大规模量子计算，但是大多数已经实现或基本实现了普适量子计算所需要的基本逻辑门。

第 3 部分为量子信息论，主要介绍量子纠错码和量子信息论的数学框架。这里包括了非常重要的量子密码的基础内容，即理想条件下的安全性证明（第 12.6 节）。第 3 部分未包含物理实现内容。

本书的特点是全面包含量子计算和量子信息的核心内容，且系统性强，结构清晰，深入浅出。这使其很适合作为相关专业的本科生和研究生教材，也适合于用作本领域研究人员的基础参考资料；对于那些准备从其他研究领域转行投入本领域研究的具有物理学、信息科学或数学等专业背景的研究人员，本书也是一本非常合适的入门书籍。

如同任何其他书籍一样，本书的内容也不可能面面俱到。本书几乎未涉及连续变量量子信息处理的有关内容。这方面，目前有多篇综述论文和一些专门教材可供参考。另外，尽管同多数同类教材比较起来，本书已经较深入地介绍了量子密码内容，但是相比于其重要性，量子密码方面的内容量还是有些偏少，对这方面感兴趣的读者可参阅相关专著。

王向斌
清华大学物理系
2015 年 9 月

Quantum Computation and Quantum Information

10th Anniversary Edition

One of the most cited books in physics of all time, *Quantum Computation and Quantum Information* remains the best textbook in this exciting field of science. This 10th Anniversary Edition includes a new Introduction and Afterword from the authors setting the work in context.

This comprehensive textbook describes such remarkable effects as fast quantum algorithms, quantum teleportation, quantum cryptography, and quantum error-correction. Quantum mechanics and computer science are introduced, before moving on to describe what a quantum computer is, how it can be used to solve problems faster than “classical” computers, and its real-world implementation. It concludes with an in-depth treatment of quantum information.

Containing a wealth of figures and exercises, this well-known textbook is ideal for courses on the subject, and will interest beginning graduate students and researchers in physics, computer science, mathematics, and electrical engineering.

MICHAEL NIELSEN was educated at the University of Queensland, and as a Fulbright Scholar at the University of New Mexico. He worked at Los Alamos National Laboratory, as the Richard Chace Tolman Fellow at Caltech, was Foundation Professor of Quantum Information Science and a Federation Fellow at the University of Queensland, and a Senior Faculty Member at the Perimeter Institute for Theoretical Physics. He left Perimeter Institute to write a book about open science and now lives in Toronto.

ISAAC CHUANG is a Professor at the Massachusetts Institute of Technology, jointly appointed in Electrical Engineering & Computer Science, and in Physics. He leads the quanta research group at the Center for Ultracold Atoms, in the MIT Research Laboratory of Electronics, which seeks to understand and create information technology and intelligence from the fundamental building blocks of physical systems, atoms, and molecules.

In praise of the book 10 years after publication

Ten years after its initial publication, “Mike and Ike” (as it’s affectionately called) remains the quantum computing textbook to which all others are compared. No other book in the field matches its scope: from experimental implementation to complexity classes, from the philosophical justifications for the Church–Turing Thesis to the nitty-gritty of bra/ket manipulation. A dog-eared copy sits on my desk; the section on trace distance and fidelity alone has been worth many times the price of the book to me.

Scott Aaronson, Massachusetts Institute of Technology

Quantum information processing has become a huge interdisciplinary field at the intersection of both, theoretical and experimental quantum physics, computer science, mathematics, quantum engineering and, more recently, even quantum metrology. The book by Michael Nielsen and Isaac Chuang was seminal in many ways: it paved the way for a broader, yet deep understanding of the underlying science, it introduced a common language now widely used by a growing community and it became the standard book in the field for a whole decade. In spite of the fast progress in the field, even after 10 years the book provides the basic introduction into the field for students and scholars alike and the 10th anniversary edition will remain a bestseller for a long time to come. The foundations of quantum computation and quantum information processing are excellently laid out in this book and it also provides an overview over some experimental techniques that have become the testing ground for quantum information processing during the last decade. In view of the rapid progress of the field the book will continue to be extremely valuable for all entering this highly interdisciplinary research area and it will always provide the reference for those who grew up with it. This is an excellent book, well written, highly commendable, and in fact imperative for everybody in the field.

Rainer Blatt, Universität Innsbruck

My well-perused copy of Nielsen and Chuang is, as always, close at hand as I write this. It appears that the material that Mike and Ike chose to cover, which was a lot, has turned out to be a large portion of what will become the eternal verities of this still-young field. When another researcher asks me to give her a clear explanation of some important point of quantum information science, I breathe a sigh of relief when I recall that it is in this book – my job is easy, I just send her there.

David DiVincenzo, IBM T. J. Watson Research Center

If there is anything you want to know, or remind yourself, about quantum information science, then look no further than this comprehensive compendium by Ike and Mike. Whether you are an expert, a student or a casual reader, tap into this treasure chest of useful and well presented information.

Artur Ekert, Mathematical Institute, University of Oxford

Nearly every child who has read Harry Potter believes that if you just say the right thing or do the right thing, you can coerce matter to do something fantastic. But what adult would believe it? Until quantum computation and quantum information came along in the early 1990s, nearly none. The quantum computer is the Philosopher’s Stone of our century, and Nielsen and Chuang is our basic book of incantations. Ten years have passed since its publication, and it is as basic to the field as it ever was. Matter will do wonderful things if asked to, but we must first understand its language. No book written since (there was no before) does the job of teaching the language of quantum theory’s possibilities like Nielsen and Chuang’s.

Chris Fuchs, Perimeter Institute for Theoretical Physics

Nielsen and Chuang is the bible of the quantum information field. It appeared 10 years ago, yet even though the field has changed enormously in these 10 years – the book still covers most of the important concepts of the field.

Lov Grover, Bell Labs

Quantum Computation and Quantum Information, commonly referred to as “Mike and Ike,” continues to be a most valuable resource for background information on quantum information processing. As a mathematically-impaired experimentalist, I particularly appreciate the fact that armed with a modest background in quantum mechanics, it is possible to pick up at any point in the book and readily grasp the basic ideas being discussed. To me, it is still “the” book on the subject.

David Wineland, National Institute of Standards and Technology, Boulder, Colorado

Endorsements for the original publication

Chuang and Nielsen have produced the first comprehensive study of quantum computation. To develop a robust understanding of this subject one must integrate many ideas whose origins are variously within physics, computer science, or mathematics. Until this text, putting together the essential material, much less mastering it, has been a challenge. Our Universe has intrinsic capabilities and limitations on the processing of information. What these are will ultimately determine the course of technology and shape our efforts to find a fundamental physical theory. This book is an excellent way for any scientist or graduate student – in any of the related fields – to enter the discussion.

Michael Freedman, Fields Medalist, Microsoft

Nielsen and Chuang's new text is remarkably thorough and up-to-date, covering many aspects of this rapidly evolving field from a physics perspective, complementing the computer science perspective of Gruska's 1999 text. The authors have succeeded in producing a self-contained book accessible to anyone with a good undergraduate grounding in math, computer science or physical sciences. An independent student could spend an enjoyable year reading this book and emerge ready to tackle the current literature and do serious research. To streamline the exposition, footnotes have been gathered into short but lively History and Further Reading sections at the end of each chapter.

Charles H Bennett, IBM

This is an excellent book. The field is already too big to cover completely in one book, but Nielsen and Chuang have made a good selection of topics, and explain the topics they have chosen very well.

Peter Shor, Massachusetts Institute of Technology

Introduction to the Tenth Anniversary Edition

Quantum mechanics has the curious distinction of being simultaneously the most successful and the most mysterious of our scientific theories. It was developed in fits and starts over a remarkable period from 1900 to the 1920s, maturing into its current form in the late 1920s. In the decades following the 1920s, physicists had great success applying quantum mechanics to understand the fundamental particles and forces of nature, culminating in the development of the standard model of particle physics. Over the same period, physicists had equally great success in applying quantum mechanics to understand an astonishing range of phenomena in our world, from polymers to semiconductors, from superfluids to superconductors. But, while these developments profoundly advanced our understanding of the natural world, they did only a little to improve our understanding of quantum mechanics.

This began to change in the 1970s and 1980s, when a few pioneers were inspired to ask whether some of the fundamental questions of computer science and information theory could be applied to the study of quantum systems. Instead of looking at quantum systems purely as phenomena to be explained as they are found in nature, they looked at them as systems that can be *designed*. This seems a small change in perspective, but the implications are profound. No longer is the quantum world taken merely as presented, but instead it can be created. The result was a new perspective that inspired both a resurgence of interest in the fundamentals of quantum mechanics, and also many new questions combining physics, computer science, and information theory. These include questions such as: what are the fundamental physical limitations on the space and time required to construct a quantum state? How much time and space are required for a given dynamical operation? What makes quantum systems difficult to understand and simulate by conventional classical means?

Writing this book in the late 1990s, we were fortunate to be writing at a time when these and other fundamental questions had just crystallized out. Ten years later it is clear such questions offer a sustained force encouraging a broad research program at the foundations of physics and computer science. Quantum information science is here to stay. Although the theoretical foundations of the field remain similar to what we discussed 10 years ago, detailed knowledge in many areas has greatly progressed. Originally, this book served as a comprehensive overview of the field, bringing readers near to the forefront of research. Today, the book provides a basic foundation for understanding the field, appropriate either for someone who desires a broad perspective on quantum information science, or an entryway for further investigation of the latest research literature. Of course,

many fundamental challenges remain, and meeting those challenges promises to stimulate exciting and unexpected links among many disparate parts of physics, computer science, and information theory. We look forward to the decades ahead!

– Michael A. Nielsen and Isaac L. Chuang, March, 2010.

Afterword to the Tenth Anniversary Edition

An enormous amount has happened in quantum information science in the 10 years since the first edition of this book, and in this afterword we cannot summarize even a tiny fraction of that work. But a few especially striking developments merit comment, and may perhaps whet your appetite for more.

Perhaps the most impressive progress has been in the area of experimental implementation. While we are still many years from building large-scale quantum computers, much progress has been made. Superconducting circuits have been used to implement simple two-qubit quantum algorithms, and three-qubit systems are nearly within reach. Qubits based on nuclear spins and single photons have been used, respectively, to demonstrate proof-of-principle for simple forms of quantum error correction and quantum simulation. But the most impressive progress of all has been made with trapped ion systems, which have been used to implement many two- and three-qubit algorithms and algorithmic building blocks, including the quantum search algorithm and the quantum Fourier transform. Trapped ions have also been used to demonstrate basic quantum communication primitives, including quantum error correction and quantum teleportation.

A second area of progress has been in understanding what physical resources are required to quantum compute. Perhaps the most intriguing breakthrough here has been the discovery that quantum computation can be done via measurement alone. For many years, the conventional wisdom was that coherent superposition-preserving unitary dynamics was an essential part of the power of quantum computers. This conventional wisdom was blown away by the realization that quantum computation can be done without any unitary dynamics at all. Instead, in some new models of quantum computation, quantum measurements alone can be used to do arbitrary quantum computations. The only coherent resource in these models is quantum memory, i.e., the ability to store quantum information. An especially interesting example of these models is the one-way quantum computer, or cluster-state computer. To quantum compute in the cluster-state model requires only that the experimenter have possession of a fixed universal state known as the cluster state. With a cluster state in hand, quantum computation can be implemented simply by doing a sequence of single-qubit measurements, with the particular computation done being determined by which qubits are measured, when they are measured, and how they are measured. This is remarkable: you're given a fixed quantum state, and then quantum compute by "looking" at the individual qubits in appropriate ways.

A third area of progress has been in *classically* simulating quantum systems. Feynman's pioneering 1982 paper on quantum computing was motivated in part by the observation that quantum systems often seem hard to simulate on conventional classical computers. Of course, at the time there was only a limited understanding of how difficult it is to simulate different quantum systems on ordinary classical computers. But in the 1990s and, especially, in the 2000s, we have learned much about which quantum systems are easy

to simulate, and which are hard. Ingenious algorithms have been developed to classically simulate many quantum systems that were formerly thought to be hard to simulate, in particular, many quantum systems in one spatial dimension, and certain two-dimensional quantum systems. These classical algorithms have been made possible by the development of insightful classical descriptions that capture in a compact way much or all of the essential physics of the system in question. At the same time, we have learned that some systems that formerly seemed simple are surprisingly complex. For example, it has long been known that quantum systems based on a certain type of optical component – what are called linear optical systems – are easily simulated classically. So it was surprising when it was discovered that adding two seemingly innocuous components – single-photon sources and photodetectors – gave linear optics the full power of quantum computation. These and similar investigations have deepened our understanding of which quantum systems are easy to simulate, which quantum systems are hard to simulate, and why.

A fourth area of progress has been a greatly deepened understanding of quantum communication channels. A beautiful and complete theory has been developed of how entangled quantum states can assist classical communication over quantum channels. A plethora of different quantum protocols for communication have been organized into a comprehensive family (headed by “mother” and “father” protocols), unifying much of our understanding of the different types of communication possible with quantum information. A sign of the progress is the disproof of one of the key unsolved conjectures reported in this book (p. 554), namely, that the communication capacity of a quantum channel with product states is equal to the unconstrained capacity (i.e., the capacity with any entangled state allowed as input). But, despite the progress, much remains beyond our understanding. Only very recently, for example, it was discovered, to considerable surprise, that two quantum channels, each with zero quantum capacity, can have a positive quantum capacity when used together; the analogous result, with classical capacities over classical channels, is known to be impossible.

One of the main motivations for work in quantum information science is the prospect of fast quantum algorithms to solve important computational problems. Here, the progress over the past decade has been mixed. Despite great ingenuity and effort, the chief algorithmic insights stand as they were 10 years ago. There has been considerable technical progress, but we do not yet understand what exactly it is that makes quantum computers powerful, or on what class of problems they can be expected to outperform classical computers.

What is exciting, though, is that ideas from quantum computation have been used to prove a variety of theorems about classical computation. These have included, for example, results about the difficulty of finding certain hidden vectors in a discrete lattice of points. The striking feature is that these proofs, utilizing ideas of quantum computation, are sometimes considerably simpler and more elegant than prior, classical proofs. Thus, an awareness has grown that quantum computation may be a more natural model of computation than the classical model, and perhaps fundamental results may be more easily revealed through the ideas of quantum computation.

Preface

This book provides an introduction to the main ideas and techniques of the field of quantum computation and quantum information. The rapid rate of progress in this field and its cross-disciplinary nature have made it difficult for newcomers to obtain a broad overview of the most important techniques and results of the field.

Our purpose in this book is therefore twofold. First, we introduce the background material in computer science, mathematics and physics necessary to understand quantum computation and quantum information. This is done at a level comprehensible to readers with a background at least the equal of a beginning graduate student in one or more of these three disciplines; the most important requirements are a certain level of mathematical maturity, and the desire to learn about quantum computation and quantum information. The second purpose of the book is to develop in detail the central results of quantum computation and quantum information. With thorough study the reader should develop a working understanding of the fundamental tools and results of this exciting field, either as part of their general education, or as a prelude to independent research in quantum computation and quantum information.

Structure of the book

The basic structure of the book is depicted in Figure 1. The book is divided into three parts. The general strategy is to proceed from the concrete to the more abstract whenever possible. Thus we study quantum computation before quantum information; specific quantum error-correcting codes before the more general results of quantum information theory; and throughout the book try to introduce examples before developing general theory.

Part I provides a broad overview of the main ideas and results of the field of quantum computation and quantum information, and develops the background material in computer science, mathematics and physics necessary to understand quantum computation and quantum information in depth. Chapter 1 is an introductory chapter which outlines the historical development and fundamental concepts of the field, highlighting some important open problems along the way. The material has been structured so as to be accessible even without a background in computer science or physics. The background material needed for a more detailed understanding is developed in Chapters 2 and 3, which treat in depth the fundamental notions of quantum mechanics and computer science, respectively. You may elect to concentrate more or less heavily on different chapters of Part I, depending upon your background, returning later as necessary to fill any gaps in your knowledge of the fundamentals of quantum mechanics and computer science.

Part II describes quantum computation in detail. Chapter 4 describes the fundamen-

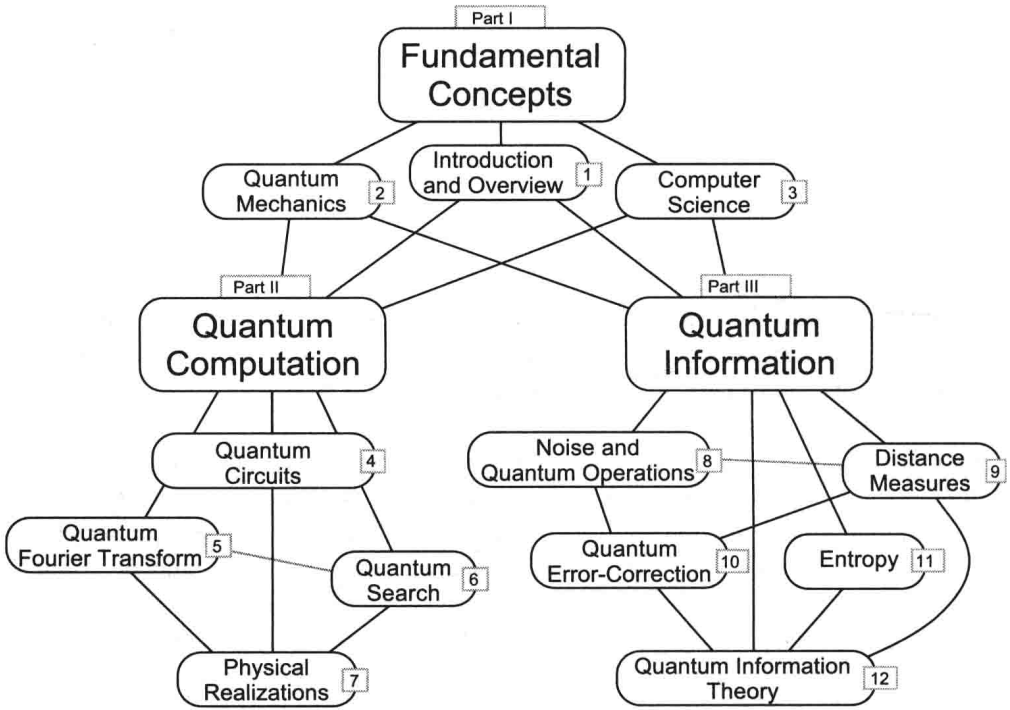


Figure 1. Structure of the book.

tal elements needed to perform quantum computation, and presents many elementary operations which may be used to develop more sophisticated applications of quantum computation. Chapters 5 and 6 describe the quantum Fourier transform and the quantum search algorithm, the two fundamental quantum algorithms presently known. Chapter 5 also explains how the quantum Fourier transform may be used to solve the factoring and discrete logarithm problems, and the importance of these results to cryptography. Chapter 7 describes general design principles and criteria for good physical implementations of quantum computers, using as examples several realizations which have been successfully demonstrated in the laboratory.

Part III is about quantum information: what it is, how information is represented and communicated using quantum states, and how to describe and deal with the corruption of quantum and classical information. Chapter 8 describes the properties of *quantum noise* which are needed to understand real-world quantum information processing, and the *quantum operations formalism*, a powerful mathematical tool for understanding quantum noise. Chapter 9 describes *distance measures* for quantum information which allow us to make quantitatively precise what it means to say that two items of quantum information are similar. Chapter 10 explains quantum error-correcting codes, which may be used to protect quantum computations against the effect of noise. An important result in this chapter is the *threshold theorem*, which shows that for realistic noise models, noise is *in principle* not a serious impediment to quantum computation. Chapter 11 introduces the fundamental information-theoretic concept of *entropy*, explaining many properties of entropy in both classical and quantum information theory. Finally, Chapter 12 discusses the information carrying properties of quantum states and quantum communication chan-

nels, detailing many of the strange and interesting properties such systems can have for the transmission of information both classical and quantum, and for the transmission of secret information.

A large number of exercises and problems appear throughout the book. Exercises are intended to solidify understanding of basic material and appear within the main body of the text. With few exceptions these should be easily solved with a few minutes work. Problems appear at the end of each chapter, and are intended to introduce you to new and interesting material for which there was not enough space in the main text. Often the problems are in multiple parts, intended to develop a particular line of thought in some depth. A few of the problems were unsolved as the book went to press. When this is the case it is noted in the statement of the problem. Each chapter concludes with a summary of the main results of the chapter, and with a ‘History and further reading’ section that charts the development of the main ideas in the chapter, giving citations and references for the whole chapter, as well as providing recommendations for further reading.

The front matter of the book contains a detailed Table of Contents, which we encourage you to browse. There is also a guide to nomenclature and notation to assist you as you read.

The end matter of the book contains six appendices, a bibliography, and an index.

Appendix 1 reviews some basic definitions, notations, and results in elementary probability theory. This material is assumed to be familiar to readers, and is included for ease of reference. Similarly, Appendix 2 reviews some elementary concepts from group theory, and is included mainly for convenience. Appendix 3 contains a proof of the Solovay–Kitaev theorem, an important result for quantum computation, which shows that a finite set of quantum gates can be used to quickly approximate an arbitrary quantum gate. Appendix 4 reviews the elementary material on number theory needed to understand the quantum algorithms for factoring and discrete logarithm, and the RSA cryptosystem, which is itself reviewed in Appendix 5. Appendix 6 contains a proof of Lieb’s theorem, one of the most important results in quantum computation and quantum information, and a precursor to important entropy inequalities such as the celebrated strong subadditivity inequality. The proofs of the Solovay–Kitaev theorem and Lieb’s theorem are lengthy enough that we felt they justified a treatment apart from the main text.

The bibliography contains a listing of all reference materials cited in the text of the book. Our apologies to any researcher whose work we have inadvertently omitted from citation.

The field of quantum computation and quantum information has grown so rapidly in recent years that we have not been able to cover all topics in as much depth as we would have liked. Three topics deserve special mention. The first is the subject of *entanglement measures*. As we explain in the book, entanglement is a key element in effects such as quantum teleportation, fast quantum algorithms, and quantum error-correction. It is, in short, a resource of great utility in quantum computation and quantum information. There is a thriving research community currently fleshing out the notion of entanglement as a new type of physical resource, finding principles which govern its manipulation and utilization. We felt that these investigations, while enormously promising, are not yet complete enough to warrant the more extensive coverage we have given to other subjects in this book, and we restrict ourselves to a brief taste in Chapter 12. Similarly, the subject of distributed quantum computation (sometimes known as quantum communication complexity) is an enormously promising subject under such active development that we

have not given it a treatment for fear of being obsolete before publication of the book. The implementation of quantum information processing machines has also developed into a fascinating and rich area, and we limit ourselves to but a single chapter on this subject. Clearly, much more can be said about physical implementations, but this would begin to involve many more areas of physics, chemistry, and engineering, which we do not have room for here.

How to use this book

This book may be used in a wide variety of ways. It can be used as the basis for a variety of courses, from short lecture courses on a specific topic in quantum computation and quantum information, through to full-year classes covering the entire field. It can be used for independent study by people who would like to learn just a little about quantum computation and quantum information, or by people who would like to be brought up to the research frontier. It is also intended to act as a reference work for current researchers in the field. We hope that it will be found especially valuable as an introduction for researchers new to the field.

Note to the independent reader

The book is designed to be accessible to the independent reader. A large number of exercises are peppered throughout the text, which can be used as self-tests for understanding of the material in the main text. The Table of Contents and end of chapter summaries should enable you to quickly determine which chapters you wish to study in most depth. The dependency diagram, Figure 1, will help you determine in what order material in the book may be covered.

Note to the teacher

This book covers a diverse range of topics, and can therefore be used as the basis for a wide variety of courses.

A one-semester course on quantum computation could be based upon a selection of material from Chapters 1 through 3, depending on the background of the class, followed by Chapter 4 on quantum circuits, Chapters 5 and 6 on quantum algorithms, and a selection from Chapter 7 on physical implementations, and Chapters 8 through 10 to understand quantum error-correction, with an especial focus on Chapter 10.

A one-semester course on quantum information could be based upon a selection of material from Chapters 1 through 3, depending on the background of the class. Following that, Chapters 8 through 10 on quantum error-correction, followed by Chapters 11 and 12 on quantum entropy and quantum information theory, respectively.

A full year class could cover all material in the book, with time for additional readings selected from the ‘History and further reading’ section of several chapters. Quantum computation and quantum information also lend themselves ideally to independent research projects for students.

Aside from classes on quantum computation and quantum information, there is another way we hope the book will be used, which is as the text for an introductory class in quantum mechanics for physics students. Conventional introductions to quantum mechanics rely heavily on the mathematical machinery of partial differential equations. We believe this often obscures the fundamental ideas. Quantum computation and quantum informa-

tion offers an excellent conceptual laboratory for understanding the basic concepts and unique aspects of quantum mechanics, without the use of heavy mathematical machinery. Such a class would focus on the introduction to quantum mechanics in Chapter 2, basic material on quantum circuits in Chapter 4, a selection of material on quantum algorithms from Chapters 5 and 6, Chapter 7 on physical implementations of quantum computation, and then almost any selection of material from Part III of the book, depending upon taste.

Note to the student

We have written the book to be as self-contained as possible. The main exception is that occasionally we have omitted arguments that one really needs to work through oneself to believe; these are usually given as exercises. Let us suggest that you should at least attempt all the exercises as you work through the book. With few exceptions the exercises can be worked out in a few minutes. If you are having a lot of difficulty with many of the exercises it may be a sign that you need to go back and pick up one or more key concepts.

Further reading

As already noted, each chapter concludes with a ‘History and further reading’ section. There are also a few broad-ranging references that might be of interest to readers. Preskill’s^[Pre98b] superb lecture notes approach quantum computation and quantum information from a somewhat different point of view than this book. Good overview articles on specific subjects include (in order of their appearance in this book): Aharonov’s review of quantum computation^[Aha99b], Kitaev’s review of algorithms and error-correction^[Kit97b], Mosca’s thesis on quantum algorithms^[Mos99], Fuchs’ thesis^[Fuc96] on distinguishability and distance measures in quantum information, Gottesman’s thesis on quantum error-correction^[Got97], Preskill’s review of quantum error-correction^[Pre97], Nielsen’s thesis on quantum information theory^[Nie98], and the reviews of quantum information theory by Bennett and Shor^[BS98] and by Bennett and DiVincenzo^[BD00]. Other useful references include Gruska’s book^[Gru99], and the collection of review articles edited by Lo, Spiller, and Popescu^[LSP98].

Errors

Any lengthy document contains errors and omissions, and this book is surely no exception to the rule. If you find any errors or have other comments to make about the book, please email them to: qci@squint.org. As errata are found, we will add them to a list maintained at the book web site: <http://www.squint.org/qci/>.

Acknowledgements

A few people have decisively influenced how we think about quantum computation and quantum information. For many enjoyable discussions which have helped us shape and refine our views, MAN thanks Carl Caves, Chris Fuchs, Gerard Milburn, John Preskill and Ben Schumacher, and ILC thanks Tom Cover, Umesh Vazirani, Yoshi Yamamoto, and Bernie Yurke.

An enormous number of people have helped in the construction of this book, both directly and indirectly. A partial list includes Dorit Aharonov, Andris Ambainis, Nabil Amer, Howard Barnum, Dave Beckman, Harry Buhrman, the Caltech Quantum Optics Foosballers, Andrew Childs, Fred Chong, Richard Cleve, John Conway, John Cortese, Michael DeShazo, Ronald de Wolf, David DiVincenzo, Steven van Enk, Henry Everitt, Ron Fagin, Mike Freedman, Michael Gagen, Neil Gershenfeld, Daniel Gottesman, Jim Harris, Alexander Holevo, Andrew Huijbers, Julia Kempe, Alesha Kitaev, Manny Knill, Shing Kong, Raymond Laflamme, Andrew Landahl, Ron Legere, Debbie Leung, Daniel Lidar, Elliott Lieb, Theresa Lynn, Hideo Mabuchi, Yu Manin, Mike Mosca, Alex Pines, Sridhar Rajagopalan, Bill Risk, Beth Ruskai, Sara Schneider, Robert Schrader, Peter Shor, Sheri Stoll, Volker Strassen, Armin Uhlmann, Lieven Vandersypen, Anne Verhulst, Debby Wallach, Mike Westmoreland, Dave Wineland, Howard Wiseman, John Yard, Xinlan Zhou, and Wojtek Zurek.

Thanks to the folks at Cambridge University Press for their help turning this book from an idea into reality. Our especial thanks go to our thoughtful and enthusiastic editor Simon Capelin, who shepherded this project along for more than three years, and to Margaret Patterson, for her timely and thorough copy-editing of the manuscript.

Parts of this book were completed while MAN was a Tolman Prize Fellow at the California Institute of Technology, a member of the T-6 Theoretical Astrophysics Group at the Los Alamos National Laboratory, and a member of the University of New Mexico Center for Advanced Studies, and while ILC was a Research Staff Member at the IBM Almaden Research Center, a consulting Assistant Professor of Electrical Engineering at Stanford University, a visiting researcher at the University of California Berkeley Department of Computer Science, a member of the Los Alamos National Laboratory T-6 Theoretical Astrophysics Group, and a visiting researcher at the University of California Santa Barbara Institute for Theoretical Physics. We also appreciate the warmth and hospitality of the Aspen Center for Physics, where the final page proofs of this book were finished.

MAN and ILC gratefully acknowledge support from DARPA under the NMRQC research initiative and the QUIC Institute administered by the Army Research Office. We also thank the National Science Foundation, the National Security Agency, the Office of Naval Research, and IBM for their generous support.