Quantum Computing without Magic

Devices

Zdzisław Meglicki

The MIT Press
Cambridge, Massachusetts
London, England

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in LaTeX by the author and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Meglicki, Zdzisław, 1953–
Quantum computing without magic : devices / Zdzisław Meglicki.
p. cm. — (Scientific and engineering computation)
Includes bibliographical references and index.
ISBN 978-0-262-13506-1 (pbk. : alk. paper)
1. Quantum computers. I. Title.
QA76.889.M44 2008
004.1—dc22
2008017033

10 9 8 7 6 5 4 3 2 1

Preface

This book arose from my occasional discussions on matters related to quantum computing with my students and, even more so, with some of my quite distinguished colleagues, who, having arrived at this juncture from various directions, would at times reveal an almost disarming lack of understanding of quantum physics fundamentals, while certainly possessing a formidable aptitude for skillfully juggling mathematics of quantum mechanics and therefore also of quantum computing—a dangerous combination. This lack of understanding sometimes led to perhaps unrealistic expectations and, on other occasions, even to research suggestions that were, well, unphysical.

Yet, as I discovered in due course, their occasionally pointed questions were very good questions indeed, and their occasional disbelief was well enough founded, and so, in looking for the right answers to them I was myself forced to revise my often canonical views on these matters.

From the perspective of a natural scientist, the most rewarding aspect of quantum computing is that it has reopened many of the issues that had been swept under the carpet and relegated to the dustbin of history back in the days when quantum mechanics had finally solidified and troublemakers such as Einstein, Schrödinger, and Bohm were told in no uncertain terms to "put up or shut up." And so, the currently celebrated Einstein Podolsky Rosen paradox [38] lingered in the dustbin, not even mentioned in the Feynman's Lectures on Physics [42] (which used to be my personal bible for years and years), until John Stewart Bell showed that it could be examined experimentally [7].

Look what the cat dragged in!

When Aspect measurements [4] eventually confirmed that quantum physics was indeed "nonlocal," and not just on a microscopic scale, but over large macroscopic distances even, some called it the "greatest crisis in the history of modern physics." Why should it be so? Newtonian theory of gravity is nonlocal, too, and we have been living with it happily since its conception in 1687. On the other hand, others spotted an opportunity in the crisis. "This looks like fun," they said. "What can we do with it?" And this is how quantum computing was born. An avalanche of ideas and money that has since tumbled into physics laboratories has paid for many wonderful experiments and much insightful theoretical work.

But some of the money flowed into departments of mathematics, computer science, chemistry, and electronic engineering, and it is for these somewhat bewildered colleagues of mine that I have written this book. Its basic purpose is to explain how quantum differs from "classical," how quantum devices are supposed to work, and even why and how the apparatus of quantum mechanics comes into being. If this text were compared to texts covering classical computing, it would be a text about the most basic classical computing devices: diodes, transistors, gates. Such books are well known to people who are called *electronic device engineers. Quantum Computing without Magic* is a primer for future *quantum device engineers.*

In classical digital computing everything, however complex or sophisticated, can be ultimately reduced to Boolean logic and NAND or NOR gates. So, once we know how to build a NAND or a NOR gate, all else is a matter of just connecting enough of these gates to form such circuitry as is required. This, of course, is a simplification that omits power conditioning and managing issues, and mechanical issues—after all, disk drives rotate, have bearings and motors, as do cooling fans; and then we have keyboards, mice, displays, cameras, and so on. But the very heart of it all is Boolean logic and simple gates. Yet, the gates are no longer this simple when their functioning is scrutinized in more detail. One could write volumes about gates alone.

Quantum computing, on the other hand, has barely progressed beyond a single gate concept in practical terms. Although numerous learned papers exist that contemplate large and nontrivial algorithms, the most advanced quantum computers of 2003 comprised mere two "qubits" and performed a single gate computation. And, as of early 2007, there hasn't been much progress. Quantum computing is extremely hard to do. Why? This is one of the questions this book seeks to answer.

Although *Quantum Computing without Magic* is a simple and basic text about qubits and quantum gates, it is not a "kindergarten" text. The readers are assumed to have mathematical skills befitting electronic engineers, chemists, and, certainly, mathematicians. The readers are also assumed to know enough basic quantum physics to not be surprised by concepts such as energy levels, Josephson junctions, and tunneling. After all, even entry-level students nowadays possess considerable reservoirs of common knowledge—if not always very detailed—about a great many things, including the world of quantum physics and enough mathematics to get by.

On the other hand, the text attempts to explain everything in sufficient detail to avoid unnecessary magic—including detailed derivations of various formulas, that may appear tedious to a professional physicist but that should help a less experienced reader understand how they come about. In the spirit of stripping quantum computing of magic, we do not leave such results to exercises.

For these reasons, an adventurous teacher might even risk complementing an introductory course in quantum mechanics with selected ideas and materials derived from this text. The fashionable subject of quantum computing could serve here as an added incentive for students to become acquainted with many important and interesting concepts of quantum physics that traditionally have been either put on the back burner or restricted to more advanced classes. The familiarity gained with the density operator theory at this early stage, as well as a good understanding of what is actually being measured in quantum physics and how, can serve students well in their future careers.

Quantum mechanics is a probability theory. Although this fact is well known to physicists, it is often swept under the carpet or treated as somehow incidental. I have even heard it asked, “Where do quantum probabilities come from?”—as if this question could be answered by unitary manipulations similar to those invoked to explain decoherence. In the days of my youth a common opinion prevailed that quantum phenomena could not be described in terms of probabilities alone and that quantum mechanics itself could not be formulated in a way that would not require use of complex numbers. Like other lore surrounding quantum mechanics this opinion also proved untrue, although it did not become clear until 2000, when Stefan Weigert showed that every quantum system and its dynamics could be characterized fully in terms of nonredundant probabilities [145]. Even this important theoretical discovery was not paid much attention until Lucien Hardy showed a year later that quantum mechanics of discrete systems could be derived from “five reasonable axioms” all expressed in terms of pure theory of probability [60].

Why should it matter? Isn’t it just a question of semantics? I think it matters if one is to understand where the power of quantum mechanics as a theory derives from. It also matters in terms of expectations. Clearly, one cannot reasonably expect that a theory of probability can explain the source of probability, if such exists at all—which is by no means certain in quantum physics, where probabilities may be fundamental.

This book takes probability as a starting point. In Chapter 1 we discuss classical bits and classical registers. We look at how they are implemented in present-day computers. Then we look at randomly fluctuating classical registers and use this example to develop the basic formalism of probability theory. It is here that we introduce concepts of fiducial states, mixed and pure states, linear forms representing measurements, combined systems, dimensionality, and degrees of freedom. Hardy’s theorem that combines the last two concepts is discussed as well, as it expresses most succinctly the difference between classical and quantum physics. This chapter also serves as a place where we introduce basic linear algebra, taking care to distinguish between vectors and forms, and introducing the concept of tensor product.

We then use this apparatus in Chapter 2, where we introduce a qubit. We describe it in terms of its fiducial vector and show how the respective probabilities can be measured by using the classical Stern-Gerlach example. We show a difference between fully polarized and mixed states and demonstrate how an act of measurement breaks an initial pure state of the beam, converting it to a mixture. Eventually we arrive at the Bloch ball representation of qubit states. Then we introduce new concepts of Pauli vectors and forms. These will eventually map onto Pauli matrices two chapters later. But at this stage they will help us formulate laws of qubit dynamics in terms of pure probabilities—following Hardy, we call this simple calculus the *fiducial formalism.* It is valuable because it expresses qubit dynamics entirely in terms of directly measurable quantities. Here we discuss in detail Larmor precession, Rabi oscillations, and Ramsey fringes—these being fundamental to the manipulation of qubits and quantum computing in general. We close this chapter with a detailed discussion of quantronium, a superconducting circuit presented in 2002 by Vion, Aassime, Cottet, Joyez, Pothier, Urbina, Esteve, and Devoret, that implemented and demonstrated the qubit [142].

Chapter 3 is short but pivotal to our exposition. Here we introduce quaternions and demonstrate a simple and natural mapping between the qubit's fiducial representation and quaternions. In this chapter we encounter the von Neumann equation, as well as the legendary *trace formula,* which turns out to be the same as taking the arithmetic mean over the statistical ensemble of the qubit. We learn to manipulate quaternions by the means of commutation relations and discover the sole source of their power: they capture simultaneously in a single formula the cross and the dot products of two vectors. The quaternion formalism is, in a nutshell, the density operator theory. It appears here well before the *wave function* and follows naturally from the qubit's probabilistic description.

Chapter 4 continues the story, beginning with a search for a simplest matrix representation of quaternions, which yields Pauli matrices. We then build the Hilbert space, which the quaternions, represented by Pauli matrices, act on and discover within it the images of the basis states of the qubit we saw in Chapter 2. We discover the notion of state superposition and *derive* the probabilistic interpretation of transition amplitudes. We also look at the transformation properties of spinors, something that will come handy when we get to contemplate Bell inequalities in Chapter 5. We rephrase the properties of the density operator in the unitary language and then seek the unitary equivalent of the quaternion von Neumann equation, which is how we arrive at the Schrödinger equation. We study its general solution and revisit and reinterpret the phenomenon of Larmor precession. We investigate single qubit *gates*, a topic that leads to the discussion of Berry phase [12], which is further illustrated by the beautiful 1988 experiment of Richardson, Kilvington, Green, and Lamoreaux [119].

In Chapter 5 we encounter the simplest bipartite quantum system, the biqubit.

We introduce the reader to the notion of entanglement and then illustrate it with experimental examples. We strike while the iron is hot; otherwise who would believe such weirdness to be possible? We begin by showing a Josephson junction biqubit made by Berkley, Ramos, Gubrud, Strauch, Johnson, Anderson, Dragt, Lobb, and Wellstood in 2003 [11]. Then we show an even more sophisticated Josephson junction biqubit made in 2006 by Steffen, Ansmann, Bialczak, Katz, Lucero, cDermott, Neeley, Weig, Cleland, and Martinis [134]. In case the reader is still not convinced by the functioning of these quantum microelectronic devices, we discuss a very clean example of entanglement between an ion and a photon that was demonstrated by Blinov, Moehring, Duan, and Monroe in 2004 [13]. Having (we hope) convinced the reader that an entangled biqubit is not the stuff of fairy tales, we discuss its representation in a rotated frame and arrive at Bell inequalities. We discuss their philosophical implications and possible ontological solutions to the puzzle at some length before investigating yet another feature of a biqubit—its single qubit expectation values, which are produced by partial traces. This topic is followed by a quite detailed classification of biqubit states, based on Englert and Metwally [39], and discussion of biqubit separability that is based on the Peres-Horodeckis criterion [113, 66].

Mathematics of biqubits is a natural place to discuss nonunitary evolution and to present simple models of important nonunitary phenomena such as depolarization, dephasing, and spontaneous emission. To a future quantum device engineer, these are of fundamental importance, inasmuch as every classical device engineer must have a firm grasp of thermodynamics. One cannot possibly design a working engine, or a working computer, while ignoring the fundamental issue of heat generation and dissipation. Similarly, one cannot possibly contemplate designing working quantum devices while ignoring the inevitable loss of unitarity in every realistic quantum process.

We close this chapter with the discussion of the Schrödinger cat paradox and a beautiful 1996 experiment of Brune, Hagley, Dreyer, Maitre, Maali, Wunderlich, Raimond, and Haroche [18]. This experiment clarifies the muddled notion of what constitutes a quantum measurement and, at the same time, is strikingly “quantum computational” in its concepts and methodology.

The last major chapter of the book, Chapter 6. puts together all the physics and mathematics developed in the previous chapters to strike at the heart of quantum computing: the controlled-NOT gate. We discuss here the notion of quantum gate universality and demonstrate, following Deutsch [29], Khaneja and Glaser [78], and Vidal and Dawson [140], that the controlled-NOT gate is universal for quantum computation. Then we look closely at the Cirac-Zoller idea of 1995 [22] and its elegant 2003 implementation by Schmidt-Kaler, Häffner, Riebe, Gulde, Lancaster, Deuschle, Bechner, Roos, Eschner, and Blatt [125]. On this occasion we also discuss the functioning of the linear Paul trap, electron shelving technique, laser cooling, and side-band transitions, which are all crucial in this experiment. We also look at the 2007 superconducting controlled-NOT gate developed by Plantenberg, de Groot, Harmans, and Mooij [114] and at the 2003 all-optical controlled-NOT gate demonstrated by O'Brien, Pryde, White, Ralph, and Branning [101].

In the closing chapter of the book we outline a roadmap for readers who wish to learn more about quantum computing and, more generally, about quantum information theory. Various quantum computing algorithms as well as error correction procedures are discussed in numerous texts that have been published as far back as 2000, many of them "classic." The device physics background provided by this book should be sufficient to let its readers follow the subject and even read professional publications in technical journals.

But there is another aspect of the story we draw the reader's attention to in this chapter. How "quantum" is quantum computing? Is "quantum" really so unique and different that it cannot be faked at all by classical systems? When comparisons are made between quantum and classical algorithms and statements are made along the lines that "no classical algorithm can possibly do this," the authors, rather narrow-mindedly, restrict themselves to comparisons with classical *digital* algorithms. But the principle of superposition, which makes it possible for quantum algorithms to attain exponential speedup, is not limited to quantum physics only. The famous Grover search algorithm can be implemented on a classical analog computer, as Grover himself demonstrated together with Sengupta in 2002 [57]. It turns out that a great many features of quantum computers can be implemented by using classical analog systems, even entanglement [24, 133, 103, 104, 105]. For a device engineer this is a profound revelation. Classical analog systems are far easier to construct and operate than are quantum systems. If similar computational efficiencies can be attained this way, may not this be an equally profitable endeavor? We don't know the full answer to this question, perhaps because it has not been pursued with as much vigor as has quantum computing itself. But it is an interesting fundamental question in its own right, even from a natural scientist's point of view.

Throughout the whole text and in all quoted examples, I have continuously made the point that everything in quantum physics is about probabilities. A single detection is meaningless and useless, even in those rare situations when theoretical reasoning lets us reduce a problem's solution to such. Experimental realities ensure that we must repeat our detection many times to provide us with classical, not necessarily quantum, error estimates. When a full characterization of a quantum state is needed, the whole statistical ensemble that represents the state must be explored. After all, what is a "quantum state" if not an abstraction that refers to the vector of probabilities that characterize it [60]? And there is but one way to arrive at this characterization. One has to measure and record sometimes hundreds of thousands of detections in order to estimate the probabilities with such error as the context requires.

And don't you ever forget it!

This, of course, does have some bearing on the cost and efficiency of quantum computation, even if we were to overlook quantum computation's energetic inefficiency [49], need for extraordinary cooling and isolation techniques, great complexity and slowness of multiqubit gates, and numerous other problems that all derive from . . . physics. This is where quantum computing gets stripped of its magic and dressed in the cloak of reality. But this is not a drab cloak. It has all the coarseness and rich texture of wholemeal bread, and wholemeal bread is good for you.

Acknowledgments

This book owes its existence to many people who contributed to it in various ways, often unknowingly, over the years. But in the first place I would like to thank the editors for their forbearance and encouragement, and to Gail Pieper, who patiently read the whole text herself and through unquestionable magic of her craft made it readable for others.

To Mike McRobbie of Indiana University I owe the very idea that a book could be made of my early lecture notes, and the means and opportunity to do so. Eventually, little of my early notes made it into the book, which is perhaps for the better.

I owe much inspiration, insights and help to my professional colleagues, Zhenghan Wang, Lucien Hardy, Steven Girvin, and Mohammad Amin, whose comments, suggestions, ideas, and questions helped me steer this text into what it has eventually become.

My interest in the foundations of quantum mechanics was awoken many years ago by Asher Peres, who visited my alma mater briefly and talked about the field, and by two of my professors Bogdan Mielnik and Iwo Birula-Białynicki. Although I understood little of it at the time, I learned that the matter was profound and by no means fully resolved. It was also immensely interesting.

To my colleagues and friends at the University of Western Australia, Armenag Nassibian, Laurie Faraone, Paul McCormick, Armando Scolaro, Zig Budrikis, and Yianni Attikiouzel, I owe my electronic engineering background and common sense that, ultimately, helped me navigate through the murky waters of quantum computing and quantum device engineering.

To my many colleagues and friends at the Australian National University, among them Bob Maier, John Slaney, Bob Dewar, Dayal Wickramasighe, and Bob Gingold, I owe my return to physics and deeper interest in computing, as well as countless hours of discussions on completely unrelated topics, though enjoyable nevertheless.

And to my cats, who did all they could to stop me from writing this book, I owe it that they gave up in the end.