Review of N. David Mermin's Quantum Computer Science: An Introduction

Eleanor Rieffel
FX Palo Alto Laboratory
rieffel@fxpal.com

1 Introduction

Over the years I have enjoyed Mermin's colorful, idiosyncratic, and insightful papers. His interest in the foundations of quantum mechanics has led him to discover alternative explanations for various quantum mechanical puzzles and protocols. These explanations are often superior to previous explanations in both simplicity and insight, and even when they are not outright better, they provide a valuable alternative point of view. His book is filled with such explanations, and with strong, sometimes controversial, opinions on the *right way* of seeing something, which make his book both valuable and entertaining.

Quantum computation is currently theory. Ardent efforts [2] are underway to build quantum computers, but it is too early to say which efforts will be successful. Mermin does not discuss implementation efforts. He is primarily interested in the artistry of the field: "there is a beauty to the theory of quantum computation that gives it a powerful appeal as a lovely branch of mathematics, and as a strange generalization of the paradigm of classical computer science." Quantum computation is more than that, however; it is the result of thoughtful inquiry into the computational implications of the physical theory that best describes our world. For this reason, I wish that he had included more physics in the book, particularly in motivating the basic concepts. For example, I would have liked to see the definition of a qubit, and the behavior of a qubit under measurement, grounded in the physics of photon polarization.

Quantum computing explores the implications of replacing the fundamental notions of information and computation with quantum mechanical ones. The supremely successful abstraction of computer science means that we can design algorithms without considering how the operations are carried out physically, obscuring the fact that our notion of computation is grounded in classical physics. The book starts with "It is tempting to say that a quantum computer is one whose operation is governed by the laws of quantum mechanics. But since the laws of quantum mechanics govern the behavior of all physical systems, this temptation must be resisted." For decades quantum mechanics has improved modern computers, but computers continue to encode information as bits and

perform the same logical operations. Quantum computing changes that.

Shor's discovery of fast quantum algorithms for the cryptographically important problems of factoring and the discrete logarithm propelled the field from a curious side water into mainstream research. The two algorithms together mean that all standard public key encryption systems are insecure. This stunning, practical result was followed by Grover's discovery of a search algorithm that, while less spectacular in its speed up, was the first quantum algorithm of practical significance that was provably better than any classical algorithm, known or unknown. After Grover's algorithm, there was a hiatus of five years in which no significantly new quantum algorithms were discovered, only variations on existing algorithms. From 2002 on, a variety of novel algorithms have been discovered [7]. Mermin's book, unfortunately, does not cover any results, as opposed to novel explanations, discovered after 1999. Readers will have to look elsewhere for more recent developments in quantum computation.

Mermin, who is famous for fanciful titles such as *Is the moon there when nobody looks? Reality and the quantum theory*, chose the pedestrian title *Quantum Computer Science* to emphasize the tight connection between classical (traditional, non-quantum) computer science and quantum computing. One of the joys of the books is watching him explain well known quantum protocols as elaborations on the simplest of classical manipulations. These original arguments make a strong case that classical computer science informs quantum computing. I am sorry he missed discussing the other side of the argument, the ever increasing evidence that quantum computing informs classical computer science. Drucker and de Wolf survey [1] a wealth of purely classical computational results, in such diverse fields as polynomial approximations, matrix theory, and computational complexity, that resulted from taking a quantum computational view. I know of two additional examples, not in their survey: Kuperberg's proof of Johansson's theorem, and Gentry's fully homomorphic encryption scheme.

2 A summary of the contents

Chapters 3, 4 and 5 give clear accounts of Shor's algorithm, Grover's algorithm, and quantum error correction respectively, adding to the many excellent expositions of these subjects. This review concentrates on Chapters 2 and 6 which contain the best expositions I've seen of some topics. After discussing these two chapters, I return to Chapter 1 to discuss its strengths and weaknesses. While Mermin uses completely nonstandard terminology, *Qbit* and *Cbit*, instead of *qubit* and *bit*, I will use the standard terminology except when quoting Mermin.

2.1 The contents of Chapter 2

Chapter 2 begins with a discussion of quantum parallelism, a commonly misunderstood concept. It is the favored explanation of journalists for the speed up enabled by quantum computation. This explanation, that quantum computers compute all values of a function at once, has lost favor among quantum

computer scientists. One reason is lower bound results that prove that, for many problems, quantum computation cannot provide any significant benefit over classical computation. Mermin's exposition of the "apparent miracle" of quantum parallelism is enjoyable and accurate: "A major part of the miracle is only apparent. One cannot say that the result of the calculation is 2^n evaluations of f, though some practitioners of quantum computation are rather careless about making such a claim. ... Before drawing extravagant practical, or even only metaphysical, conclusions from quantum parallelism, it is essential to remember that when you have a collection of Qbits in a definite but unknown state, *there is no way to find out what that state is*."

Mermin's exposition of Deutsch's algorithm emphasizes that the algorithm yields "less information than we get in answering the question with a classical computer," and that "by renouncing the possibility of acquiring that part of the information which is irrelevant to the question we wish to answer, we can get the answer with only a single application of the black box." It is a lovely section, containing multiple views on the algorithm, some original to Mermin. I wish he had defined *black box* for readers new to the concept, but in all other respects it is the best exposition of Deutsch's problem I have seen.

Mermin's insight shines brightly when he discusses the Bernstein-Vazirani algorithm. The standard argument is phrased in terms of quantum parallelism and quantum interference: compute a function on all inputs in superposition and use a trick to make the bad answers cancel, leaving only the good answer. Explaining the power of quantum computation in terms of quantum parallelism has gone out of favor as other, more insightful, explanations have been found. Mermin's explanation of the Bernstein-Vazirani algorithm, originally published in his paper *Copenhagen Computation: How I Learned to Stop Worrying and Love Bohr*, contributed to this enlightenment. He was the first to see that, without changing the algorithm at all, just viewing it in a different light, the algorithm becomes clear, almost obvious, and definitely not a trick. The key insight is to view the algorithm in a different basis. Part of the magic of quantum computation, as Mermin likes to say, is that the role of control and target qubits in a controlled operation can reverse in a different basis. By changing viewpoint, the algorithm goes from one in which a calculation is needed to see that it gives the desired result, to one in which the outcome is evident.

The Bernstein-Vazirani algorithm, and Mermin's argument in particular, deserves to be better known because of the insight it has given into quantum computation. The algorithm is not discussed, for example, in Nielsen-Chuang [8]. Mermin could have written an even stronger section on this problem had he chose to include Meyer's recognition [6] that there is no entanglement in the Bernstein-Vazirani algorithm. This observation leads directly into the fascinating, and still evolving, question of the role of entanglement in quantum computation's superior computing capabilities. Entanglement remains the most popular explanation among quantum computer scientist for the power of quantum computing, in spite of increasingly serious questions as to how satisfactory an answer it can provide. Jozsa and Linden [3] show that entanglement is required in order for a quantum algorithm to achieve an exponential speed up

over classical algorithms. In that same paper, however, they end their abstract with "we argue that it is nevertheless misleading to view entanglement as a key resource for quantum-computational power." I had hoped that, with Mermin's strong interest in foundations and why things work, he would discuss the role of entanglement, but he does not.

2.2 The contents of Chapter 6

Mermin's colorful and clear account of quantum key distribution could have benefited from a mention that it must be combined with an authentication protocol to defeat man-in-the-middle attacks. In other respects it is a pleasure to read, with scattered shrewd observations such as the following comment on the usefulness of key distribution mechanisms, whether quantum or classical: "What is bizarre is that human ingenuity combined with human perversity has succeeded in inventing a context in which the need to hide information from a third party actually provides a purpose for such an otherwise useless exchange of random strings of bits."

It is unfortunate and surprising that Mermin continues to perpetuate the impression that quantum cryptography is synonymous with quantum key distribution. Quantum cryptography is a broad field with protocols for tasks such as secret sharing, fingerprinting, and authentication. Mermin makes a rare mistake in claiming that "Nobody has figured out how to exploit quantum mechanics to provide a secure means for directly exchanging meaningful messages," when Gottesman's unclonable encryption does just that. He does discuss the status of quantum bit commitment, concluding that "the structure of quantum mechanics might be uniquely determined by requiring it to enable the secure exchange of random strings of bits ..., but not to enable bit commitment."

Mermin is at his best when he develops both quantum dense coding and quantum teleportation as elaborations of a simple classical operation. This way of looking at these protocols appeared in two earlier papers of Mermin. It is good to see them collected here. His discussions of the GHZ puzzle and the Hardy paradox, both of which deserve to be better known, are also strong.

2.3 The contents of Chapter 1

I now return to discuss the first chapter which, while containing a number of gems, also has some failings. Sections 1.2 through 1.4 walk the reader through the most basic notions of classical computer science, bits and operations on bits, represented in an unusual way that reflects how the basic notions of quantum computing, qubits and operations on qubits, will be represented. These sections are the most problematic of the book. They require the reader to take on faith Mermin's claim that "Playing unfamiliar and somewhat silly games with Cbits will enable you to become acquainted with much of the quantum mechanical formalism in a familiar setting." I worry that some readers will find his short introduction insufficiently motivating to make it through these fifteen pages of classical computer science in an odd notation. In addition, to benefit from these

sections, readers must take the initiative to play with this new notation on their own. These early sections call out for exercises, but none are provided.

In these sections, Mermin introduces the quantum computationally important phase change and Hadamard operators. I question whether introducing them in a classical context, in which they cannot be given meaning, will make the reader more comfortable with them. On the other hand, it is here that Mermin introduces the Pauli operators through the classical Swap operation. As Mermin says, "It is pleasing to find them here, buried in the interior of the operator that simply swaps two classical bits." It certainly is!

Sections 1.8 through 1.10 discuss measurement. Mermin emphasizes that the state of an n qubit system is "not associated with any ascertainable property of those qubits," and that "To the extent that it suggests that some preexisting property is being revealed, "measurement" is a dangerously misleading term, but it is hallowed by three quarters of a century of use by quantum physicists." Mermin does not define general quantum measurements, only the measurement of a single qubit, and only one type of single qubit measurement from among the infinite number of possibilities. Mermin has emphasized, in papers such as *On the Absence of a Measurement Problem in Quantum Computer Science* that because quantum computation requires only one type, the simplest type, of measurement, the theory of quantum computation avoids one of the trickiest conceptual issues of quantum mechanics.

Quantum computing is intricately connected, however, to the tricky concepts of tensor product decompositions and entanglement, concepts that are key to the difference between classical and quantum mechanics. Readers would have benefited from a more extended introduction to tensor products and their central role in quantum computation. Similarly, Mermin only briefly defines what it means for a state to be entangled, without giving examples, or explaining its dependence on which tensor product decomposition is under consideration.

Mermin devotes an entire section to the use of measurement for state preparation, a use he correctly says plays a crucial role that is not often emphasized. In his discussion of the common use of the word "collapse" to describe the effect of measurement on a quantum state, he cautions that the state is "nothing more than an abstract symbol, used ... to calculate probabilities of measurement outcomes." I would have loved to see him go further, making explicit the ties with classical probability theory, and elucidating the differences between classical probability theory and quantum mechanics. This viewpoint has appeared in a number of places including [4], but is waiting for an popular, elementary explanation of the sort Mermin does so well.

2.4 The index, and the lack of references

Not all topics mentioned in the index are actually covered. For example, *Schmidt decomposition*, *mixed state*, and *density operator* appear, but none of these concepts are described; they are only mentioned in passing. Most surprising is that Bell's Theorem, while appearing in the index, is never described even though Mermin has written eloquently on the subject [5].

The omission I found most dismaying is that the book has no reference section. A few references are given in footnotes, but in most cases the reader is left with little indication of how to get more information on a given topic. As one example, Mermin comments that Grover's algorithm requires knowledge of the number of solutions in order to know how many iterations to apply. He then mentions that "a clever application" provides a means to estimate this number, but no reference is given. Authors of introductory books intend to intrigue readers enough that they will want to pursue the subject further. Mermin succeeds. It is unfortunate that he does not make it easier for readers to do so.

3 Concluding thoughts

This book provides an enjoyable and insightful read that will enhance both the novice and expert reader's knowledge of quantum computation. I would not recommend it as a sole source of information on quantum computation; it leaves out too many important topics such as fault tolerance, all algorithmic results more recent than 1999, the known limits on quantum computation, and any study of quantum subsystems and the insight they give into entanglement. Its lack of exercises and a reference section also limit its use as a single source. But anyone with an interest in quantum computation will enjoy reading Mermin's highly personal account of the subject.

References

[1] Andrew Drucker and Ronald de Wolf. Quantum proofs for classical theorems. arXiv:0910.3376, 2009.

[2] Richard Hughes and et al. Quantum cryptography roadmap, version 1.1. http://qist.lanl.gov, July 2004.

[3] Richard Jozsa and Noah Linden. On the role of entanglement in quantum computational speed-up. *Proceedings of the Royal Society of London Ser. A*, 459:2011–2032, 2003.

[4] Alexei Kitaev, Alexander Shen, and Mikhail N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, Providence, RI, 2002.

[5] N. David Mermin. Hidden variables and the two theorems of John Bell. *Reviews of Modern Physics*, 65:803–815, 1993.

[6] David A. Meyer. Sophisticated quantum search without entanglement. *Physical Review Letters*, 85:2014–2017, 2000.

[7] Michele Mosca. Quantum algorithms. arXiv:0808.0369, 2008.

[8] Michael Nielsen and Isaac L. Chuang. *Quantum Computing and Quantum Information*. Cambridge University Press, Cambridge, 2001.