

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/cosrev

Book review

Phillip Kaye, Raymond Laflamme, Michele Mosca. *An Introduction to Quantum Computing*. Oxford University Press (2007). xii+276 pp., Paperback \$54, ISBN: 019857049X

1. Background

Since the last decade of the 20th century, quantum computing and quantum information have been increasingly fascinating researchers in physics, computer science, and mathematics. It may be illustrative to look at the references of the book being reviewed to discover that out of more than 110 references to articles on quantum computing or quantum information theory, less than 5 were published before 1990.

Quantum computing is a cross-disciplinary area in the intersection of theoretical computer science and quantum physics. The education on quantum computing therefore involves studies in both areas at least to some extent, and this is probably one of the reasons why a regular quantum computing course seems to be an exception rather than a rule in universities everywhere.

As a new research area, quantum computing, as well as quantum information theory, includes a wide range of open research problems varying from practical matters to highly theoretical questions, from engineering to Hilbert spaces. As a new research area, quantum computing also suffers from an unestablished and very unstable education tradition. To educate a new generation of scientists familiar with quantum computing, accessible books on the topic are needed. Maybe they are not needed inevitably, but anyway they make education substantially easier.

2. About the authors

All the authors are affiliated with the Institute for Quantum Computing (IQC), Waterloo, Ontario, and they have been giving a course in quantum computing at the University of Waterloo since 1999. In a relatively short time, IQC has gained stature in the quantum computing community, which also raises high expectations for the book: the director and the deputy director of IQC are among the authors.

As the authors mention, the intention of the book is to serve as an introduction for a reader having “an undergraduate education in some scientific field, and [they]

should particularly have a solid background in linear algebra, including vector spaces and inner products. Prior familiarity with topics such as tensor products and spectral decomposition is not required, but may be helpful”.

3. The first impression

The book consists of ten chapters and an appendix. The topics of the chapters are chosen to represent current (mainstream) quantum computing: In addition to introductory chapters on computation and linear algebra, the book includes basics on the quantum mechanics framework, superdense coding and teleportation, quantum algorithms in three chapters, quantum lower bounds, and quantum error correction.

The chapters are 17–25 pages long excluding the shortest one (superdense coding and teleportation, 8 pages) and the two longest ones (algorithms and error correction, 42 and 37 pages, respectively). Considering the topics of the exceptional chapters, this structure makes the book fairly well balanced, and the topics are presented in an expedient order.

The book contains numerous pictures, diagrams, and tables, which make understanding easier. However, sometimes the layout seems quite dense, giving the feeling that the amount of information presented on some pages is overwhelming.

With only some exceptions (most notably chapter 9 and the appendix), the results and theorems are presented without proofs or leaving the proofs as exercises; the book is clearly focused on explaining the logical structure behind the topics studied, not on the formal mathematical derivations.

4. The content briefly

The book begins with an introduction containing the background information on the Church–Turing Thesis, models of computing, and a motivation for quantum computing presented in a concise manner. The introduction is followed by a chapter designed to supplement the reader’s knowledge on linear algebra, with the Dirac notation and the foundational results in operator theory such as The Spectral Theorem.

In the third chapter, the authors formulate the connection between quantum physics and representing information. The

notion of a quantum bit is introduced, the time evolution and the measurement are described, and finally the formalism is extended to cover also the mixed quantum states.

The fourth chapter is to present the quantum circuit model. The authors introduce the circuit diagrams and quantum gates. The notion of the universal set of gates is presented together with an example of such a set. In this chapter, the authors also introduce the important Solovay–Kitaev theorem guaranteeing an efficient approximation to quantum circuits.

In the fifth chapter, the protocols for superdense coding and for its complementary task, quantum teleportation, are presented briefly with an application to correct the failures of the controlled not-gate.

Chapter 6 begins with a discussion on interference, an important resource in quantum computing, which is helpful in explaining the efficiency of many quantum algorithms. As the first example of quantum algorithms, the Deutsch algorithm, as well as its extensions the Deutsch–Josza and Simon algorithms, are presented.

The theme of quantum Fourier transform-based algorithms is continued in an extensive chapter 7, where the most famous quantum computing result, Shor’s factoring algorithm, is presented. The core of Shor’s algorithm is the Order-Finding Procedure, which is presented after the Phase Estimation and the Eigenvalue Estimation Procedures. The Eigenvalue Estimation Procedure is close to one of the authors’ research areas, and supposedly this is the reason for selecting this perspective for the first presentation of the Order-Finding Procedure. To give another perspective, Shor’s original approach is also presented. The chapter is concluded with the quantum algorithm for the discrete logarithm and a more general Hidden Subgroup Problem.

Chapter 8 turns the focus to quantum algorithms based on amplitude amplification. No quantum algorithm in this class offers an exponential advantage over the classical ones (the Hidden Subgroup algorithms *seem* to offer an exponential advantage, but this is unproven). The algorithms presented in chapter 8 include the famous Grover search in $O(\sqrt{N})$ time, as well as its various spin-offs.

The three previous chapters have contained positive results by showing that some particular tasks *can be achieved* with the help of a quantum computer. Chapter 9, in turn, is devoted to results that impose *lower bounds* on quantum computing. As in the general theory of computing, the strongest lower bound results concern oracle computing, which is usually presented as a *black box model* in quantum computing. This chapter introduces the most essential lower bound techniques: the polynomial method and the adversary methods.

Chapter 10 is devoted to a topic slightly different to those of the previous chapters: quantum error correction. In order to realize a quantum computer physically, it is of course essential to correct the errors introduced during the computation, and the elementary methods for that purpose are introduced in this chapter.

The appendix is 19 pages long, and it consists of various topics, basically including some background on probability theory and some extra information on the topics of the ten chapters. The appendix is actually the only structure in the

book where I would disagree with the authors: I believe that, possibly with the exception of A.1 (Tools for Analyzing Probabilistic Algorithms), the issues handled in the appendix could have been embedded into the chapters 1–10 in a very confluent manner.

5. Comparison

The purpose of this section is by no means to give a detailed analysis of the present book based on the perspective created by the previous ones. Rather, I would like to point out some differences and similarities, as well as some strengths and weaknesses of the books. For short, the book of Kaye, Laflamme, and Mosca being reviewed will be referred as to KLM.

For the comparison, I have selected the books by Gruska [1], Nielsen and Chuang [4], Kitaev, Shen, and Vyalys (KSV) [3], and by Hirvensalo [2] (the undersigned).

5.1. Gruska, and Nielsen and Chuang

The book being reviewed is very different from both Gruska’s monograph (more than 400 pages) and the volume of Nielsen and Chuang (more than 650 pages), measuring by magnitude, content, or purpose. Gruska’s book contains issues such as quantum automata theory and quantum complexity theory not presented in KLM, whereas the book by Nielsen and Chuang includes more details on the physical realization of quantum computers and quantum information theory. On the other hand, KLM contains lower bound techniques that are not presented by Gruska or by Nielsen and Chuang.

Both Gruska’s and Nielsen and Chuang’s books are more elaborate in the sense that they include more proofs and individual research results. On the other hand, Gruska’s and Nielsen and Chuang’s books serve a different purpose: The present book gives an introduction to quantum computing for a newcomer in a much more concise manner than [1] or [4], whereas those volumes can be recommended for a more advanced reader, or to serve as a handbook for researchers. For a teacher, the book by Kaye, Laflamme, and Mosca offers also a more comfortable way to prepare a lecture course than [1] or [4], but the latter volumes clearly explore quantum computing more deeply than KLM. Thus KLM does not compare very well to the books by Gruska and by Nielsen and Chuang.

On the contrary, books by the undersigned and by Kitaev, Shen, and Vyalys [3] are kindred to KLM, and hence it may be illustrative to compare them in more detail.

5.2. Kitaev, Shen, and Vyalys

The book by Kitaev, Shen, and Vyalys consists, as its name indicates, of two main parts, the first being devoted to classical computing, and the latter to quantum computing. It is not by any means inapt to present these two topics in a single exposition, but in KSV the joint presentation has left an open crack between the presentation styles. KLM does not contain an analogous part devoted to classical complexity theory, but knowing that Laflamme has a background in

theoretical physics, whereas Mosca and Kaye have their basic education in theoretical computer science, one might expect a similar gap between the presentation styles in KLM, too. However, the authors have been pretty successful in avoiding such a divergence; the presentation in KLM is considerably more uniform than that in KSV. Kitaev's background in theoretical physics comes through sometimes, for example in the notation chosen for KSV.

There are also some other differences between KLM and KSV, when focusing only on the quantum computing part of the latter book. KSV includes some quantum complexity theory (such as the quantum analogy of NP), which KLM does not have. On the other hand, KLM discusses the lower bound techniques not included in KSV. However, a more important difference between KLM and the quantum computing part of KSV can be found again in the general presentation style: KSV includes more technical details and assertions with a proof than KLM. As mentioned before, most of the results presented in KLM are without proofs, which is not the case in KSV. A noteworthy difference between KSV and KLM of highly pragmatic nature is that KSV includes the solutions to the exercises, whereas KLM does not.

5.3. *Hirvensalo*

I feel very much like I am walking on dangerous ground when beginning to compare KLM and my book [2]. Partially this is because of the fear of unfairly preferring my own text and losing neutrality. This fear increased, but in an opposite manner to what I had expected, when I began to feel a very great sense of familiarity with KLM: despite numerous differences, I believe that KLM and my book are very alike, and hence there is a potential danger of me commending KLM just because it seems familiar to me. More seriously, for the aforementioned reason, there is a danger that this whole review is slanted in favour of KLM, and I can only guarantee that I try to achieve neutrality.

The topics in KLM and my book [2] overlap very much (especially for the 2nd edition of [2]), but also some differences can be found. For instance KLM includes eigenvalue estimation procedure, adversary methods for lower bounds, and quantum error correction, which I did not present in my book. On the other hand, I included some background in physics such as Gleason's theorem, uncertainty relations, as well as their entropic versions, not discussed in KLM. Also, KLM has been published only very recently, and consequently it is more up to date than my book, but this is of course true when comparing KLM to any other book.

There are several differences, and the importance of each difference depends on the viewpoint. In my opinion, the most essential difference between KLM and my book [2] is that while the authors of KLM present the topics very clearly and consistently, the proofs and the mathematical derivations are mostly omitted or left as exercises. In contrast, I have chosen to include most of the mathematical proofs in the presentation, leaving only some as exercises. I believe that this difference is the most important one, and that giving a list of some minor ones is not very useful for anyone.

6. Conclusions

"An introduction to Quantum Computing" by Kaye, Laflamme, and Mosca, is a very recent book on quantum computing. In this review, I have briefly described it and compared it to four books on the same area published earlier.

In a book review, one may (legitimately) expect to find reviewer's opinions on questions such as "Is the book worth reading?", "What is the book good for?". It is not always easy to form an opinion on such questions, but in this case, there is no ambiguity; it is exactly as the authors state: the book is for a reader willing become acquainted with quantum computing, having sufficient knowledge in linear algebra. In my opinion, the authors have succeeded pretty well in giving a clear introduction to quantum computing for a beginner. The book may also be helpful for a teacher when preparing a lecture course in quantum computing, but the book by Nielsen and Chuang and that by Gruska clearly serve better when a handbook or a deep exploration in quantum computing is needed.

It may still be convenient to have some debate on which book, KLM, KSV, or that by the undersigned, is most useful for a beginner. That, however, is an issue which depends on many factors and is not straightforward at all. The intersection of the topics handled in these three books is large, as discussed above, and this unfortunately means that, as an introduction to quantum computing, the two other books will not bring much more added value after reading one of them. This however does not mean that after reading one book it would be absolutely useless to look at another.

From the triplet being discussed, KSV sticks out most clearly: It has an initial part designed to give an introduction to the classical theory of computing, including non-basic topics on complexity theory. I would therefore recommend KSV for a reader with little or no education in complexity theory, but a background in physics, since I also believe that a reader with some background in theoretical physics would find the quantum computing part in KSV most familiar.

For readers having no background in physics I would rather recommend KLM or my book. The main difference in the presentation is that in KLM most of the proofs are omitted or left as exercises. Hence the reader can either rapidly learn the basic principles behind quantum computing without getting embroiled in the technical details, or expend some more time and effort while reading, and enter deeper into the details by completing the exercises while studying the book. On the other hand, I have chosen a slightly more formal mathematical presentation with proofs in my book. Above, I have discussed the differences between the contents, but they overlap so much that it leads me to believe that the differences between the presentation styles mentioned above should be the most important criterion for a beginner when deciding which book to choose.

REFERENCES

- [1] Jozef Gruska, Quantum Computing, McGraw-Hill, 2000.

- [2] Mika Hirvensalo, Quantum Computing, 2nd edition, Springer, 2001-2004.
- [3] A.Yu. Kitaev, A.H. Shen, M.N. Vyalyi, Classical and Quantum Computation, American Mathematical Society, 2002.
- [4] Michael A. Nielsen, Isaac L. Chuang, Quantum Computation and Quantum Information, Cambridge University Press, 2000.

Mika Hirvensalo
Turku Centre for Computer Science,
Department of Mathematics,
University of Turku, FIN-20014 Turku,
Finland
E-mail address: mikhirve@utu.fi.