
Preface

The first half of the twentieth century witnessed the foundation of three pillars of modern science: quantum mechanics, information theory, and computer science. In the latter half of the century, scientists began to connect these fields, first by exploring the implications of information itself being quantum—leading to the birth of quantum information theory. The nonclassical features of quantum information, such as no-cloning and entanglement, were identified as resources for novel applications, for instance, information theoretically secure communication. Alongside these developments were the observations of Benioff [120], Feynman [391], and Manin [737] that models of computation and simulation could be formulated within quantum mechanics—and that in some cases these models appeared exponentially challenging to simulate using a classical computer.

In 1985, Deutsch [346] further developed these early models of quantum computation and presented what was essentially the first quantum algorithm—a simple procedure that, with just one black-box query, could accomplish a task that classically requires two queries. Over the next decade, larger black-box separations were discovered, such as the Deutsch–Jozsa [347], Bernstein–Vazirani [129], and Simon’s [940] algorithms, and finally, in 1994, the first truly *end-to-end* quantum algorithm was developed: Shor’s algorithm [937] for factoring integers and computing discrete logarithms, bringing extensive ramifications for cryptography. This breakthrough demonstrated that quantum computers could not only speed up the solution of contrived black-box problems but, at least in theory, could provide faster solutions to important real-world problems. The discovery of Shor’s algorithm transformed the field of quantum algorithms from a relatively niche topic into a major research area.

During the three decades since Shor’s seminal discovery, the field of quantum algorithms matured significantly. For example, our knowledge of upper and lower bounds on the quantum query complexity of black-box

problems—often deduced through sophisticated, nonconstructive mathematical arguments—has been greatly expanded. Moreover, many additional quantum algorithms and subroutines—for example, primitives for quantum simulation and linear algebra—have been discovered, optimized, and subsequently generalized multiple times. Meanwhile, advances in hardware and the theory of fault-tolerant quantum computation have reached the point where it is conceivable that (some of) these algorithms might soon become implementable at scales large enough to surpass what can be done classically.

Nevertheless, the magnitude of available quantum speedups for real-world applications is often hard to assess and can be obscured by technical caveats, assumptions, and limitations in the underlying quantum algorithmic primitives. Despite being one of the oldest, Shor's algorithm for factoring arguably remains the cleanest example of a substantial quantum speedup with minimal caveats that targets a problem of significant real-world relevance. This survey aims to elucidate the true resource requirements of end-to-end quantum computing applications, and thereby aid in identifying the most likely applications for fault-tolerant quantum computers. Through this distinct perspective, the survey is intended to complement the wealth of existing quantum algorithms resources, including a number of review articles, lecture notes, textbooks, and the quantum algorithm zoo [586].

We highlight both the opportunities and challenges of currently known quantum algorithms. To truly understand the potential advantage of a quantum algorithm, it is necessary to consider its resource requirements in an end-to-end fashion. By this, we mean the cost of solving the full problem of interest to the user, not only the cost of running a given quantum circuit that is a subroutine of the full solution. One must consider all quantum and classical overheads: keeping track of classical precomputation and postprocessing, explicitly instantiating quantum oracles and data access structures, and ideally computing the constant prefactors of all quantum subroutines (including those overheads associated with fault-tolerant protocols and quantum error correction). We note, however, that this task is a major undertaking for complex quantum algorithms, and so has only been achieved for a minority of quantum algorithms in the literature. In addition to studying end-to-end quantum complexities, it is also necessary to compare any quantum results to the state-of-the-art classical solutions of the same problem, as well as known complexity-theoretic limitations.

We summarize the end-to-end complexities of several leading quantum application proposals (by which we mean quantum algorithms applied to a well-defined, real-world problem). The complexities of these applications are deduced from the complexities of their underlying primitives, which we review in detail. The modular structure of the book aids the high-level understanding

of the costs and tradeoffs coming from the various choices one makes when designing and compiling a quantum algorithm, as well as identifying the bottlenecks for a given application. On the technical front, this book does not attempt to advance the state of the art; rather, it aims to collect, synthesize, and contextualize key results in the literature. We consider algorithms in the quantum circuit model, which is arguably the best-studied model for quantum computation and renders the presented complexities hardware agnostic (although the overhead associated with executing these circuits in a fault-tolerant fashion can, of course, depend on details of the hardware). In order to obtain concrete bounds, we require oracles to be explicitly instantiated. We generally assume that quantum error correction of some form will be necessary, due to unavoidable imperfections inherent to all known quantum hardware modalities. As such, we typically consider the non-Clifford cost of quantum algorithms as the dominant cost, in keeping with leading quantum fault-tolerance schemes. Due to the general lack of application-scale experimental data, we focus on elucidating provable speedups, and we only mention noisy, intermediate-scale quantum (NISQ) algorithms in passing, where appropriate, since they are typically heuristic.

Throughout this book, we attempt to be thorough, but not exhaustive in presentation; we only aim to give a representative collection of references, rather than providing a complete list. Generally, we try to explain how asymptotic complexity statements arise from their underlying primitives, but technical results are typically presented without explicit derivation or proofs, for which we refer the reader to the cited references. Additionally, we often quote resource estimates from the literature without covering all of the application-specific optimizations to the underlying primitives that are required to arrive at the reported constant prefactors. We survey a number of quantum applications, primitives, and fault-tolerance schemes; however, the omission of other approaches does not indicate that they are unimportant. Also, the primary scope of this work excludes substantial topics, such as quantum sensing or communications, measurement-based quantum computing, adiabatic quantum computing and quantum annealing, analog quantum simulators, quantum-inspired (“dequantization”) methods, and tensor network algorithms—comments on these topics are provided in instances where they are relevant to our primary discussion.

An overarching takeaway of this survey is that the current literature generally lacks fully end-to-end analyses for concrete quantum applications. Consequently, in several parts of this survey, a fully satisfactory end-to-end accounting is not achieved. In part, this is due to certain technical aspects of the relevant quantum algorithms being underexplored, and in some cases also due to a lack of specific details on how the output of the quantum algorithm will

integrate into concrete computational workflows for future quantum computing users. Quantum algorithms research often works upward from algorithmic primitives to identify computational tasks with maximal quantum speedups, but these may not align with the tasks most relevant to the user. On the other hand, potential users themselves may not yet know exactly how they would use a new capability to advance their high-level goals. Yet, we find ourselves at a point in the history of quantum computing at which it behooves us to fill in these details and adopt this end-to-end lens. As more end-to-end applications are found, and with small fault-tolerant quantum computers now on the horizon, we expect the story to continue to evolve—this survey provides a snapshot of the state of play in roughly mid-2024. While improved quantum algorithms and approaches to quantum error correction and fault tolerance are likely to be discovered, classical computers continue to grow in scale and speed, and classical algorithms are also constantly refined and developed, thereby moving the goalposts for end-to-end quantum speedups. We hope the reader will find this book a valuable guide for navigating this complex and dynamic landscape.

How to use this book

This book does not need to be read from cover to cover. Instead, it has a modular structure, which enables readers to directly explore the applications and primitives relevant for their use case. To the extent possible, each numbered chapter and section has been written in a self-contained fashion and can be read independently from the rest of the book. Readers looking for a quick introduction to (or refresher on) the common notation, conventions, and background concepts that underlie the technical exposition in the main text are advised to begin with the Appendix, where we provide information on quantum mechanics, bra-ket notation, quantum circuits (and quantum computing more generally), big- O notation, and complexity theory.

Readers of the link-enabled online version are encouraged to navigate the book using the hyperlinks embedded throughout, which connect interdependent material. Both print and online readers may also find utility in the index, which organizes mentions of the important topics discussed in the book, including computational tasks and problems, quantum algorithmic tools and primitives used to solve those tasks, and finally competing classical methods for those tasks (in addition to other miscellaneous topics).

Acknowledgments

We thank Joao Basso, J. Kyle Brubaker, Christopher Chamberland, Andrew Childs, Isabel Franco Garrido, Helmut G. Katzgraber, Eric M. Kessler, Robin Kothari, Péter Kutas, Yi-Kai Liu, Pavel Lougovski, Carl Miller, Oskar Painter,

Nicola Pancotti, Simone Severini, Sophia Simon, James D. Whitfield, and Xi-aodi Wu for helpful comments and conversations on various aspects of this survey.

After the first version of this material was released online, we conducted a self-managed, nonanonymized peer-review process. We are grateful to colleagues who agreed to help us by reviewing a subset of the survey. At the beginning of each chapter, we acknowledge those who reviewed that chapter as part of this process, and we also list them alphabetically here: Dong An, Eric Anschuetz, Ryan Babbush, Matthew Campagna, Earl Campbell, Marco Cerezo, Andrew Cross, Zohreh Davoudi, Vedran Dunjko, Glen Evenbly, Di Fang, Joshua Goings, Johnnie Gray, Sander Gribling, Thomas Häner, Matthew Hastings, Samuel Jaques, Robin Kothari, Richard Kueng, Lin Lin, Daniel Malz, Ashley Milsted, Ashley Montanaro, John Preskill, Patrick Rall, Patrick Rebentrost, Rolando Somma, Nikitas Stamatopoulos, Damian Steiger, Yuan Su, Ewin Tang, Ronald de Wolf, Nobuyuki Yoshioka, and Xiao Yuan.

Finally, we acknowledge the staff from across the AWS Center for Quantum Computing that enabled this project, with special thanks to Mike Sadowitz and Wendy Yu for their legal support. We are also grateful to the Institute for Quantum Information and Matter at Caltech, which is an NSF Physics Frontier Center. We thank Harry Atwater, Fiona Harrison, Tom Rosenbaum, and David Tirrell at Caltech, and Nafea Bshara, Peter DeSantis, James Hamilton, Andy Jassy, Simone Severini, and Bill Vass at AWS, for their involvement and support of the research activities at the AWS Center for Quantum Computing.

