

Quantum Computational Supremacy

Aram W. Harrow* Ashley Montanaro†

September 21, 2018

The original motivation for quantum computing was the apparent exponential overhead in simulating quantum mechanics on a classical computer. Following this idea led to the field of quantum algorithms, which aims to find quantum speedups for useful problems, some without obvious relation to quantum mechanics. A key milestone in this field will be when a universal quantum computer performs a computational task that is beyond the capability of any classical computer, an event known as *quantum supremacy*. This would be easier to achieve experimentally than full-scale quantum computing but involves new theoretical challenges.

As a goal, quantum supremacy⁴⁷ is unlike most algorithmic tasks since it is defined not in terms of a particular problem to be solved but in terms of what classical computers *cannot* do. This is like the situation in cryptography, where the goal is not only for the authorized parties to perform some task, but to do so in a way that restricts the capabilities of unauthorized parties. Understanding the fundamental limitations of computation is the remit of the theory of computational complexity.⁴⁶ A basic goal of this theory is to classify problems (such as integer factorisation) into complexity classes (such as the famous classes P and NP), and then to rigorously prove that these classes are unequal. In both the case of cryptography and quantum supremacy, computational complexity theory is extremely far from being able to unconditionally prove the desired no-go theorems. Just as we cannot yet prove that $P \neq NP$, we currently cannot unconditionally prove that quantum mechanics cannot be simulated classically. Instead claims of quantum supremacy will need to rely on complexity-theoretic assumptions, which in turn can be justified heuristically.

Requirements for quantum supremacy

Any proposal for a quantum supremacy experiment must have the following ingredients: (1) a well-defined computational task, (2) a plausible quantum algorithm for the problem, (3) an amount of time/space allowed to any classical competitor, (4) a complexity-theoretic assumption (as we will discuss below), and optionally (5) a verification method that can efficiently distinguish the quantum algorithm from any classical competitor using the allowed resources. Here “plausible” means ideally on near-term hardware and will likely include the need to handle noise and experimental imperfections. We will briefly describe some of the leading approaches to quantum supremacy in these terms.

Note that we do not require that the computational task is of practical interest. When discussing quantum supremacy, it is natural to ask what this gives us that other quantum algorithms, such as factoring⁵⁰ or quantum simulation,^{20,29} do not. Indeed, both could be said to be routes to quantum supremacy in their own right.

For factoring the computational assumption is simply that classical computers cannot factor quickly (say, faster than the current best known algorithm) and the successful operation of a quantum factoring device could be easily verified. However, the best current estimates suggest that the quantum algorithm requires ≈ 4000 qubits and $\approx 10^9$ gates³² to factor a 2048-bit number, and if overheads from fault-tolerance or architectural restrictions are added they could further raise the qubit and gate costs significantly.

*Center for Theoretical Physics, MIT. aram@mit.edu

†School of Mathematics, University of Bristol, UK. ashley.montanaro@bristol.ac.uk

Algorithm	Difficulty for quantum computers	Assumption implying no classical simulation	Easy to verify?	Useful?
factoring ⁵⁰	hard	RSA secure	yes	yes
boson sampling ²	easy	PH infinite or approx. counting ≠ exact counting	no	no
low-depth circuits ⁵⁴	moderate		no*	no
IQP ⁴⁹	moderate		sometimes	no
QAOA ²⁵	moderate		no*	maybe
random circuits ⁸	moderate	QUATH (see ⁴)	no	no
adiabatic optimization ²⁴	easy	unknown	no*	maybe
analog simulation ^{20,29}	easy	idiosyncratic	no	often

Table 1: **Approaches to quantum supremacy.** Boson sampling, IQP and random circuits are discussed in Boxes 2 and 3. Low-depth circuits are quantum circuits on many qubits, but with only a few layers of quantum gates. QAOA (“Quantum Approximate Optimization Algorithm”) and adiabatic optimization are quantum algorithms for finding reasonably good solutions to optimization problems. Analog simulation is the engineering of one quantum Hamiltonian to directly reproduce the behaviour of another. The “difficulty” column can be viewed as a very crude estimate of how far we would need to proceed towards building a universal quantum computer in order to carry out each algorithm. For verification we write “no*” to mean that we cannot fully verify the validity of the algorithm but we can check some properties such as few-qubit statistics, or the value of the objective function. We note that outputs of IQP circuits cannot be verified in general, but there is an IQP-based protocol for hiding a string in a code which does have an efficient verification procedure.⁴⁹

Analog quantum simulators,^{19,20,29} on the other hand, are already being used to estimate properties of quantum systems which we do not know how to efficiently calculate classically. If we believe that these properties *cannot* be calculated classically, then these experiments could already be considered demonstrations of quantum supremacy. However, our confidence in conjectures such as “The correlation functions of a strongly coupled Fermi gas are hard to calculate” is much lower than our confidence in the hardness of factoring or (as we will discuss below) the non-collapse of the Polynomial Hierarchy. We will discuss this point in more detail in the subsequent section on complexity theory.

Modern supremacy proposals include constant-depth circuits,⁵⁴ single photons passing through a linear-optical network (aka “boson sampling”)² and random quantum circuits containing gates which either all commute^{12,49} (a model known as “IQP”) or do not commute.⁸ These latter two are described in Boxes 1 and 2. In each of these cases, we will describe below arguments why classical simulation is hard even though these models are believed not to be capable of universal quantum computing. These occupy a sweet spot between factoring and analog simulation: they can be implemented with much less effort than factoring, including using a non-universal architecture, while the complexity-theoretic evidence for their superiority over classical computing is stronger than the evidence in favor of specific simulations. In the sections below we will describe the arguments from complexity theory, and discuss the complications that arise from experimental imperfections and the problem of verification. We summarize some of the main proposals for quantum supremacy in Table 1.

Specific proposals for quantum supremacy

Boson sampling

Boson sampling² is a formalisation of the problem of simulating noninteracting photons in linear optics; see Figure 1. n coincident photons are input into a linear-optical network on $m \gg n$ modes (usually generated at random), with detectors positioned at the output of the network. The challenge is to sample from the




Figure 1: Schematic of a boson sampling experiment. Photons are injected (on left-hand side) into a network of beamsplitters that are set up to generate a random unitary transformation. They are detected on the right-hand side according to a probability distribution conjectured to be hard to sample from classically.

distribution on detection outcomes. Following the initial theoretical proposal of Aaronson and Arkhipov,² several experimental groups quickly demonstrated small-scale examples of boson sampling experiments, with up to 4 coincident photons in up to 6 modes.^{15,21,53,56} Subsequent work has experimentally validated boson sampling, in the sense of implementing statistical tests that distinguish the boson sampling distribution from other particular distributions.^{17,51} The current records for implementation of arbitrary linear-optical transformations are 6 modes with up to 6 photons¹⁶ or 9 modes with up to 5 photons.^{17,51,58}

Initial boson-sampling experiments used single-photon sources based on spontaneous parametric down-conversion. This is a randomised process which has inherently poor scaling with the number of photons, requiring exponential time in the number of photons for each valid experimental run. A variant of boson sampling known as “scattershot” boson sampling has therefore been proposed. This uses many sources, each of which produces a photon with some small probability, and it is known in which modes a photon has been produced. Scattershot boson sampling has been implemented with 6 sources and 13 modes.⁷ An alternative approach is to use a high-performance quantum dot source.⁵⁸

Challenges faced by experimental implementations of boson sampling include handling realistic levels of loss in the network, and the possibility of the development of more efficient classical sampling techniques.

Random quantum circuits

Unlike boson sampling, some quantum supremacy proposals remain within the standard quantum circuit model. In the model of commuting quantum circuits,⁴⁹ known as IQP (for “Instantaneous Quantum Polynomial-time”), one considers circuits made up of gates which all commute, and in particular are all diagonal in the X basis; see Figure 2. Although these diagonal gates may act on the same qubit many times, as they all commute, in principle they could be applied simultaneously. The computational task is to sample from the distribution on measurement outcomes for a random circuit of this form, given a fixed input state. Such circuits are both potentially easier to implement than general quantum circuits, and have appealing theoretical properties which make them simpler to analyse.^{12,13} However, this very simplicity may make them easier to simulate classically too.

Of course, one need not restrict to commuting circuits to demonstrate supremacy. The quantum-AI group at Google has recently suggested an experiment based on superconducting qubits and noncommuting gates.⁸ The proposal is to sample from the output distributions of random quantum circuits, of depth around 25, on a system of around 49 qubits arranged in a 2d square lattice structure (see Fig. 3). It is suggested in⁸ that this should be hard to simulate, based on (a) the absence of any known simulation requiring less than a petabyte of storage, (b) IQP-style theoretical arguments¹³ suggesting that larger versions of this system should be asymptotically hard to simulate, and (c) numerical evidence⁸ that such circuits have properties that we would expect in hard-to-simulate distributions. If this experiment were successful it would come




Figure 2: Example of an IQP circuit. Between two layers of Hadamard gates is a collection of diagonal gates. Although these diagonal gates may act on the same qubit many times, they all commute so in principle could be applied simultaneously.

very close to being out of reach of current classical simulation (or validation, for that matter) using current hardware and algorithms.




Figure 3: The architecture proposed by the quantum-AI group at Google to demonstrate quantum supremacy consists of a 2d lattice of superconducting qubits. This figure depicts two illustrative timesteps in this proposal. At each timestep, 2-qubit gates (blue) are applied across some pairs of neighbouring qubits, and random single-qubit gates (red) are applied on other qubits.

Why supremacy?

Before proceeding, we should discuss why a demonstration of quantum supremacy is worthwhile. The field of quantum computing is based on the premise that quantum mechanics has changed the definitions of information and computation, with implications that are both philosophical and practical. For example, entanglement is a useful form of correlation that would not exist in a classical theory of information and its existence can be demonstrated with experiments designed to test Bell-inequality violations. Supremacy experiments can be thought of as the computational analogue of Bell experiments. Just as Bell experiments refute Local Hidden Variable models, supremacy experiments refute the old “Extended Church-Turing (ECT) thesis”, which asserts that classical computers can simulate any physical process with polynomial overhead. Such a demonstration would be convincing evidence confirming the consensus model of quantum mechanics, showing that the world contains not only entanglement but also is capable of computational feats beyond the reach of classical computers. Validating the standard picture of quantum mechanics in this way would be valuable for foundational reasons because quantum mechanics is so far the only physical theory to change our model of computing; and for practical reasons because it would greatly increase our confidence in the eventual feasibility of large-scale quantum computing.

The ECT thesis also motivates our focus on quantum mechanics, as opposed to hard-to-simulate classical systems such as fluid dynamics or protein folding. With these examples the difficulties are “merely” from issues such as separations of scales in time or space, and these in principle could be simulated with effort linear

in the energy and space-time volume of the system. This means that a protein-folding problem which would require 10^{50} steps for a naive simulation is not an instance of a family that includes problems requiring 10^{100} or 10^{1000} steps. By contrast, a quantum supremacy experiment that barely surpasses our existing classical computers would be significant in part because it would imply that vastly greater separations in computational power are likely to soon follow, as we will explore further in the next section.

Complexity-theoretic basis for quantum supremacy

Since quantum supremacy is ultimately about comparison between quantum and classical computers, demonstrating it will require some computational assumption about the limits to the power of classical computers. At a minimum, we need to assume that quantum mechanical systems cannot be simulated efficiently (i.e. with polynomial overhead) by classical computers. But just as cryptography always needs assumptions stronger than $P \neq NP$, each quantum supremacy proposal needs its own assumption. Although such assumptions must ultimately be at least as strong as the lack of efficient classical simulation of quantum computers, we may hope for them to be based on different principles and to be believable in their own right.

As discussed above, if we use the quantum computer for factoring or simulation, then our assumption should simply be that those problems are hard for classical computers. Our belief that factoring is hard is based on many mathematician-hours put into solving it; on the other hand, the best known algorithms are only from ca. 1990 and are significantly faster than brute-force search, so further improvements may well exist.

The complexity of quantum simulation is much murkier. One difference is the great diversity of quantum systems and of methods for treating them, which are often adapted to specific features of the system. The complexity of a simulation can also vary with parameters such as temperature and coupling strengths in nonobvious ways. Finally, when analog quantum simulators cannot address individual qubits, this limits their ability to encode a wide range of problem instances, and makes the complexity of the problem they do solve even less clear. The problems solved by quantum simulators can certainly teach us about physics and often in ways that we do not know how to do classically; however, our confidence that they cannot be classically simulated is rather weak.

We now turn to the modern supremacy proposals. These are often based around sampling problems⁴⁰ rather than decision problems, where the task is to output samples from a desired distribution, rather than to output a deterministic answer. The strength of these is that, despite working with a restricted model of quantum computing (boson sampling, low-depth circuits, etc.), they do not need to assume that this specific model is hard to simulate. Indeed the complexity assumption can be expressed in terms of concepts that have been studied since the 1970s and are thought to be hard for reasons that do not rely on any beliefs about quantum mechanics. One assumption that will work is known as the “non-collapse of the polynomial hierarchy,” which we explain in a sidebar. Another possible assumption is that exact counting of exponentially large sets is harder than approximate counting. Stronger assumptions are also possible, and these can be used to rule out larger classes of classical simulations or in some cases to enable more efficient verification of the quantum device.

Why are these complexity assumptions relevant to simulating quantum computers? The main idea is to use a technique called “post-selection,” which refers to the following scenario. A computation, which could be either classical or quantum, takes input string x and outputs strings y and z . The string y is used to represent the output, while we condition (“post-select”) on the string z taking some fixed value, say $00\dots 0$. Many experiments post-select in practice (e.g. on coincident detection events) but usually on events whose probability is not too small. We will allow post-selection even on exponentially unlikely outcomes, which will make the ability to post-select extremely powerful. The purpose of post-selection is two-fold. First, an efficient classical simulation of a quantum computation implies that a classical computer with post-selection can efficiently simulate a quantum computer with post-selection. However, this latter claim would contradict our assumption that the polynomial hierarchy doesn’t collapse, as we will explain in the sidebar. Second, many non-universal models of quantum computation become universal once post-selection is allowed. Thus even an efficient classical simulation of one of these restricted models of quantum computing would lead to

the same contradictions.

In a sidebar we describe a somewhat indirect argument which implies that an efficient *exact* classical simulation (we discuss approximate simulations below) of any of these restricted models of quantum computing would lead to several surprises, including the collapse of the polynomial hierarchy and exact counting being roughly as hard as approximate counting. Neither of these is believed to be true; conversely, this consensus of complexity theorists implies that there is no efficient exact classical simulation of boson sampling, IQP or the other models. How strong are these beliefs? The non-collapse of the polynomial hierarchy is a conjecture that is stronger than $P \neq NP$ but that is plausible for similar reasons. Unlike factoring, there are essentially no non-trivial algorithms which would suggest that the ability to perform approximate counting would yield an exact counting algorithm in anything less than exponential time. On the other hand, these questions have been studied by a smaller and less diverse group of researchers. Nevertheless these are conjectures that we should have high confidence in. Going further, we will describe several stronger (i.e. less plausible) conjectures that will let us rule out more classical simulations.

Sidebar: The Polynomial Hierarchy and Post-Selection

The polynomial hierarchy. To explain the polynomial hierarchy (denoted PH), we start with NP. A boolean function $f(x)$ is in NP if it can be expressed as $f(x) = \bigvee_y g(x, y)$ for some poly-time computable function g . For example, in the 3-coloring problem, x specifies a graph, and $f(x)$ is true if and only if the graph is 3-colorable, i.e. the vertices can be each colored red, green or blue in a way such that no edge connects two vertices with the same color. If y is a list of colors for each vertex then we can easily compute $g(x, y)$, which is 1 if each edge connects vertices of different colors, and 0 if not; then $f(x) = \bigvee_y g(x, y)$ so the 3-coloring problem is in NP.

The k^{th} level of PH is defined to be the set of functions that can be written in terms of k alternating quantifiers, i.e. $f(x) = \bigvee_{y_1} \bigwedge_{y_2} \bigvee_{y_3} \cdots \bigwedge_{y_k} g(x, y_1, \dots, y_k)$, where $g(\cdot)$ is poly-time computable, and where the length of each y_1, \dots, y_k grows at most polynomially with the length of x . Just as it is conjectured that $P \neq NP$, it is conjectured that the k^{th} level of the PH is not equal to the $k + 1^{\text{st}}$ level for any k . If this were to be false and there existed some k for which the k^{th} level equalled the $k + 1^{\text{st}}$ level then the k^{th} level would also equal the $k + 2^{\text{nd}}$ and all later levels as well – a scenario we refer to as the “collapse” of the PH to the k^{th} level.

Post-selection. Allowing post-selection (described in the main text) dramatically increases the power of both classical and quantum computation. The corresponding complexity classes are called PostBPP and PostBQP respectively. It turns out that PostBPP is somewhere between the first and third levels of the PH³¹ and PostBQP corresponds to the class PP,¹ which appears to be much stronger. Indeed, any problem in PH can be reduced to solving a polynomial number of problems in PP, or formally $\text{PH} \subseteq \text{P}^{\text{PP}}$.⁵⁷ This means that if PostBPP were to equal PostBQP then it would imply the collapse the PH. Conversely, the non-collapse of the PH implies that post-selected classical and quantum computing are far apart in their computational power.

Sidebar: Counting

Another way to think about the difference between PostBPP and PostBQP is in terms of counting problems. Consider a problem in NP of the form $f(x) = \bigvee_y g(x, y)$, where again g is poly-time computable, and we represent True and False with 1 and 0 respectively. Rather than asking whether or not there is a y such that $g(x, y) = 1$, we can instead ask how many such y there are. This corresponds to the function $f_{\text{count}}(x) := \sum_y g(x, y)$. The class NP corresponds to determining whether $f_{\text{count}}(x)$ is $= 0$ or ≥ 1 . We can express PostBPP and PostBQP as well in terms of f_{count} . PostBQP corresponds to being given some threshold T and determining whether $f_{\text{count}}(x)$ is $> T$ or $\leq T$, a task known as *exact counting*. Keep in mind here that y is a string of $\text{poly}(n)$ bits, so that T can be exponentially large in n . By contrast, PostBPP is equivalent to *approximate counting*:⁴⁵ formally, given a threshold T and an accuracy parameter $\epsilon = 1/\text{poly}(n)$, determine whether $f_{\text{count}}(x)$ is $\geq T(1 + \epsilon)$ or $< T$ given that one of these is the case.

Just as we could start with the assumption that the PH does not collapse, we could also conjecture that exact counting is much more difficult than approximate counting. This assumption would equally well imply that $\text{PostBPP} \neq \text{PostBQP}$ and in turn that there does not exist an efficient classical simulation of even restricted models of quantum computing.

Fine-grained complexity assumptions

If we equate “efficient” with “polynomial-time”, then conjecturing that $\text{PostBPP} \neq \text{PostBQP}$ is enough to show that classical computers cannot exactly simulate quantum computers “efficiently.” However, these asymptotic statements tell us little about the actual quantum computers being built in the coming years, or about the ability of existing classical competitors to simulate them. For this, we would like statements of the form “a quantum circuit with 500 gates on 100 qubits cannot be perfectly simulated using a classical computer with 10^9 bits of memory using fewer than 10^{12} operations.” Only in this way could we point to the outcomes of a specific experiment and conclude that quantum supremacy had been achieved. A crucial ingredient here is a concrete or “fine-grained” complexity assumption. An example of such an assumption is the Exponential Time Hypothesis (ETH)³⁵ which asserts that 3-SAT instances on n bits require time $\geq 2^{cn}$ (for some constant $c > 0$) to solve on a classical computer. Concrete bounds require an even more explicit hypothesis: for example, we might assert that the ETH holds with $c = 0.386$ for random 3-SAT instances with a planted solution, corresponding to the best known algorithm. These bounds might be based on analysis of the best known algorithm or on provable lower bounds if we replace the 3-SAT instance with a black-box “oracle” function. Work in progress²² uses oracle arguments to devise a plausible conjectured variant of ETH which will in turn imply that ≈ 1700 qubits can carry out a computation which cannot be simulated classically in less than a day by a supercomputer performing 10^{17} operations per second. Improving this is an important open problem, and possible routes to progress include stronger assumptions, better algorithms or a sharper analysis.

Average-case assumptions

The best-studied conjectures and results in complexity theory consider *worst-case* hardness, meaning that to compute a function one must be able to do so for all possible inputs. However, in terms of quantum supremacy worst-case hardness assumptions only translate into the statement that there exists *some* quantum circuit of a given size which cannot be efficiently simulated, or equivalently that no classical algorithm works for all quantum circuits. What can an experimentalist do with such statements? There is no clear guidance on which quantum circuits to implement, and once a circuit has been chosen, no way to certify it as hard.

The same issues arise in cryptography, where we would like to argue that a given cryptosystem is hard to break not merely in the worst case but for *most* choices of random key. Such statements are the domain

of *average-case complexity*, in which we say that a function f is hard to compute on some distribution D if there is no efficient algorithm whose output equals $f(x)$ for most values of x sampled from D . Average-case hardness conjectures are stronger than (i.e. less plausible than) worst-case hardness conjectures and cannot be reduced to each other as readily as can the worst-case hardness of NP-complete problems. Nevertheless there are well-studied distributions of instances of NP-complete problems, such as random 3-SAT,^{18,42} where no algorithms are known to run in less than exponential time, and there are distributional versions of NP-completeness.³⁹

The benefits of using average-case assumptions are two-fold. First, they give us a concrete family of experiments for which we can say that most such experiments are hard to simulate. For example, we might consider random quantum circuits of a specified size. Second, and less obviously, they allow us to rule out a larger class of classical simulations.^{2,13,28} Suppose the quantum circuit outputs string z with probability $q(z)$ and a classical simulator does so with probability $p(z)$. Worst-case assumptions allow us to argue that for any specific z , say 0^n , it is not possible to find a classical simulator with $p(0^n) = q(0^n)$ for all circuits, or even with low *multiplicative* error:

$$0.9q(0^n) \leq p(0^n) \leq 1.1q(0^n). \quad (1)$$

While this does indeed rule out highly accurate simulations, we would like to rule out even simulations with what is called low *additive* error

$$\sum_z |p(z) - q(z)| \leq 0.001. \quad (2)$$

This notion of approximation is natural because, if p and q are close under this distance measure, they cannot be distinguished without taking many samples. If we make an average-case hardness assumption and prove one more technical assumption known as anticoncentration, we can rule out additive-error simulations (i.e. satisfying (2)). (Anticoncentration means that the distribution $q(z)$ is reasonably close to uniform. It is known to hold for random circuits⁴ and for IQP¹³ and is conjectured to hold for boson sampling,² although it does not hold for constant-depth random circuits.³⁴)

One disadvantage of average-case assumptions is that they cannot easily be reduced to each other. By contrast, if a problem is NP-hard in the worst case then we know that an algorithm that works for all inputs would yield algorithms for thousands of other problems in NP, which collectively have been studied for decades by researchers across all of science and engineering. But for average-case hardness, we may have different hardness for each distribution of instances. For example, for 3-SAT a natural distribution is to choose n variables and αn random clauses. It is believed that random instances are likely to be satisfiable for $\alpha < \alpha_c$ and unsatisfiable for $\alpha > \alpha_c$, for some critical value $\alpha_c \approx 4.2667$.⁴² Based on this, it is reasonable to conjecture that choosing $\alpha = \alpha_c$ yields a hard distribution, but this conjecture is far flimsier than the worst-case conjectures even for this relatively well-studied problem.

In rare cases, a problem will have the same average-case and worst-case complexity, and it is a major open question to establish quantum supremacy based on such a problem. Boson sampling takes steps in that direction,² by using a conjecture about the average-case complexity of estimating the permanent, while an average-to-worst-case reduction is known only for the exact case. Indeed the known reduction is based on polynomial interpolation and its numerical instability means that new ideas will be needed to argue that estimating the permanent is hard on average. More generally, a major open problem is to base the hardness of approximate classical simulations of the form of (2) merely on well-believed classical complexity assumptions, such as non-collapse of the polynomial hierarchy.

Maximal assumptions

Another reasonable possibility is to make our complexity assumptions as strong as possible without contradicting known algorithms. Here the high-level strategy is to try to improve our (often exponential-time) classical simulations as far as possible and then to conjecture that they are essentially optimal. Aaronson and Chen⁴ have recently carried out this program. Among other contributions, they developed classical simulations for n -qubit, depth- d circuits that calculate matrix elements in time $O((2d)^n)$ and nearly linear

space (note that with 2^n space, $O(d2^n)$ time is possible). An easier task than classical simulation is to distinguish likely from unlikely outcomes of a quantum circuit with some exponentially small advantage over random guessing. The “QUATH” conjecture⁴ asserts that poly-time classical algorithms cannot perform this task for quantum circuits whose depth $d = \Omega(n)$. The advantage of this approach is that it enables a “semi-efficient” verification procedure which uses the quantum device only a polynomial number of times but still requires exponential time on a classical computer.

Making these conjectures as strong as possible makes our confidence in them as low as possible; essentially any non-trivial improvement in simulating quantum mechanics would refute them. But so what? Unlike the case of cryptographic assumptions, a too-strong conjecture would not create any vulnerabilities to hackers. In this view, hardness conjectures are just ways of guessing the complexity of simulating quantum systems, and these estimates are always subject to revision as new evidence (in the form of algorithms) appears. Further, these conjectures highlight the limits of our current simulation algorithms, so that refuting them would be both plausible and a significant advance in our current knowledge.

Physical noise and simulation errors

Any realistic quantum experiment will be affected by noise, i.e. undesired interactions with the environment. Dealing with this noise is a major challenge for both theorists and experimentalists. The general theory of quantum fault-tolerance^{37,38} allows quantum computations to be protected against a sufficiently small amount of physically reasonable noise. However, although the asymptotic overhead of fault-tolerance is relatively minor, the constant factors involved are daunting: to produce a fault-tolerant logical qubit may require $10^3 - 10^4$ physical qubits,²⁶ an overhead far too great for short-term quantum supremacy experiments. As excessive noise can render a hard probability distribution easy to simulate, it is an important question to determine to what extent these experiments remain hard to simulate classically, even in the presence of uncorrected noise.

A related issue is that classical simulation algorithms of quantum circuits will have errors of their own. This could be seen as analogous to the fact that realistic quantum computers only implement ideal quantum circuits imperfectly. Classical noise could be multiplicative as in (1) or additive as in (2). Methods based on representing the exact state⁴¹ can achieve low enough error rates that we can think of them as low multiplicative error, while methods based on sampling (e.g.¹⁰) naturally achieve low additive error. For multiplicative noise it is relatively easy to show hardness results. IQP circuits remain hard to simulate under this notion of noise,¹² and similar results have since been shown for the one clean qubit model⁴³ and other restricted classes of circuits. However, additive noise is arguably a more natural model, and ruling out such simulations would be a stronger result.

Addressing this question was one of the major steps forward taken by Aaronson and Arkhipov² in their work on boson sampling. Based on two reasonable (yet currently unproven) conjectures, they argued that sampling from the boson sampling distribution should still remain classically hard if the classical sampler is allowed to only approximately sample from the distribution. That is, the classical sampler is asked to output a sample from any distribution whose total variation distance from the true boson sampling distribution is at most a small constant. Assuming their conjectures, as long as the quantum experiment experiences a total amount of noise below this threshold, its output is still hard to sample from classically.

One of the conjectures is a technical claim about anticoncentration of the permanent of a random matrix, with strong numerical evidence for its truth. The other (known as the “permanent-of-Gaussians” conjecture) is an average-case hardness assumption asserting that the permanent of a matrix of Gaussian random variables should be hard to approximate up to small relative error. This hardness property can be shown to hold for exact computation of the permanent of such random matrices,² but extending it to small relative error seems to be beyond the reach of current techniques.

Another step forward was the proof of a similar result for the IQP model.¹³ In this case, two conjectures occur which are analogous to those for boson sampling; however, in the setting of IQP the analogue of the anticoncentration conjecture can be proven. The permanent-of-Gaussians conjecture is replaced with equivalent conjectures about either the average-case hardness of approximately computing the partition function of

the Ising model, or the average-case hardness of approximately computing a natural property of low-degree polynomials over finite fields.¹³ Anticoncentration and average-case hardness conjectures naturally occur in the setting of noisy quantum supremacy experiments because approximating a probability distribution up to small total variation distance is similar to approximating most of the probabilities up to a very small additive error. If most of the probabilities are hard to approximate up to small relative error, and most of them are rather large (i.e. the distribution is not too concentrated) then a good classical approximation in total variation distance leads to a contradiction.

The question of how to model simulability in the presence of noise is subtle and still under debate. For example, a counterpoint to these hardness results is provided by recent work showing that, if an *arbitrarily small* constant amount of noise occurs on each qubit at the end of an IQP circuit, the class of random IQP circuits which is conjectured hard to simulate¹³ can be simulated classically in polynomial time up to small variational-distance (as in (2)) error¹¹. This contrasts with another recent result showing that classical simulation of the noisy distribution up to small *relative* error (as in (1)) can be hard.²⁷ As noise at the end of an IQP circuit can be dealt with using simple classical error-correction techniques with low overhead,¹¹ this suggests that quantum supremacy experiments may need to make use of some form of error-correction, but this might be substantially simpler than the machinery required for full quantum fault-tolerance.

Verification

A key issue for any proposed quantum supremacy experiment is verification of the results of the experiment. In order to claim quantum supremacy, we must have confidence that the experiment has indeed done something which is hard for a classical computer. By definition, quantum supremacy experiments cannot be simulated efficiently classically, so we must seek another means of checking that such an experiment has succeeded. If we had a large-scale quantum computer that could run Shor’s algorithm, verification would be easy: we could challenge the experimenter to factor a 2048-bit RSA key, then check that the claimed factors multiplied to the correct number. However, integer factorisation is a rare example of a problem which is both tractable on a quantum computer (in the complexity class BQP⁵⁹), checkable on a classical computer (in the complexity class NP⁴⁶), yet not known to be efficiently solvable on a classical computer. Very few such problems are known, and none are currently known which would be classically intractable for instance sizes small enough to be solved by a quantum computer with, say, 100 logical qubits.

In the short term, then, verification of quantum supremacy needs to use different methods, none of which is yet as simple and powerful as checking integer factorisation. Which approach is preferred may depend on the assumptions one wishes to make about the experiment being performed. This is analogous to the setting of experimental tests of Bell-inequality violations: different techniques can be used to rule out different loopholes, but it is very challenging to rule out all loopholes simultaneously.

One straightforward approach is to build confidence that the experiment (which is hard to test in its entirety) is working correctly by testing smaller parts of it. This could involve testing individual components within a quantum circuit – a task likely to be required for any experiment anyway – or running quantum computations which are small or simple enough to be classically simulable. A non-trivial example of this is executing computations which are mostly or entirely comprised of Clifford gates, which are known to be efficiently classically simulable^{5,10} despite displaying such quantum characteristics as allowing the creation of large-scale entanglement. Another example is replacing the random linear-optical transformation used in boson sampling with a highly structured one, such as a quantum Fourier transform.⁵⁵ The risk here is the so-called “Volkswagen problem” in which the diagnostic runs of the experiment are systematically different from when we run the algorithm of interest.

Another natural thought is to apply statistical tests to samples from the output distribution of a quantum supremacy experiment, to attempt to determine whether it is consistent with the desired distribution. A challenge for this approach is that many such tests require calculation of individual probabilities, which is assumed to be classically hard in the post-selection-based strategies. Indeed a classical simulator with post-selection could simply guess a random output string and then post-select on it passing the verification. This is an obstacle to simultaneously (a) basing our hardness on assuming $\text{PostBPP} \neq \text{PostBQP}$, (b) efficiently

verifying the output, and (c) having our verifier depend solely on the input and output to the quantum device. However, examples exist of verification procedures that drop each of these. Using factoring as a hard problem avoids using (a). Using exponential classical time but polynomial quantum time^{4,8} means dropping (b). And we can drop (c) by generating a secret string x and only providing some derived string $y = f(x)$ to the quantum computer.

In other words, we can attempt to find a task for which we know the answer in advance, but where we are able to hide it in such a way that it is likely hard to determine classically, while still being accessible to a quantum computer. In the case of IQP, Shepherd and Bremner⁴⁹ proposed a task of this form based on, roughly speaking, hiding a bit string within a seemingly random linear code. The verifier knows how the code has been scrambled but the quantum device (or classical simulator) sees only a version of the code from which the hidden string is not obvious. A hidden correlation within the output probability distribution of the IQP circuit then encodes the identity of the hidden string. If the string is known in advance, one can test from a few samples whether the output distribution from the experiment corresponds to the correct string. However, if the hidden string is unknown, there is conjectured to be no structure in the code which would allow it to be determined by a classical algorithm. It is an interesting open problem to try to develop further challenges of this form which are based on more familiar cryptographic hardness assumptions.

One can also define efficient tests which rule out particular distributions as alternative explanations for the output of the supremacy experiment.^{3,17,52} This can be reasonable if one has prior knowledge for suspecting that some distributions are likely, say if they correspond to the effects of decoherence.

The gold standard of verification would be direct certification of the quantum computation, where we check directly that the computation has worked, perhaps using information beyond merely the classical output of the experiment. In each case known so far, this approach requires more resources than performing the original computation. One example is a proposition by Hangleiter et al.³³ that IQP computations could be verified by encoding them into the ground state of a local Hamiltonian based on a universality construction for adiabatic quantum computation,³⁰ and then checking that the state was prepared correctly using local measurements. This idea in fact works for any quantum computation. However, it loses the simplicity and noise-tolerance of the original IQP circuit, and requires one to believe that the experimenter has correctly implemented the local measurement operations. Another approach is the use of a distributed protocol to certify that a remote computer (or multiple computers) has performed an arbitrary quantum computation. This is exemplified by the model of blind quantum computation,¹⁴ which requires the verifier to be able to create single-qubit states and send them to the prover. If we could make the verifier fully classical then such a verified quantum computation could be a way of experimentally confirming the computational predictions of quantum mechanics, analogous to the way that Bell experiments test the theory of entanglement.⁶

In summary, all known verification techniques for quantum supremacy experiments have drawbacks: they are either inefficient in terms of the classical or quantum resources required, or assume that the behaviour of the experiment under test conditions corresponds to the real experiment, or make computational hardness assumptions which are not yet well understood. Developing verification techniques which avoid these issues is a pressing open question in quantum supremacy research.

Outlook

In just a few years, quantum supremacy-type experiments have progressed from demonstrating boson sampling with 3 photons to a proposed implementation of random quantum circuits on 49 qubits. Each of the diverse quantum supremacy proposals meets the five requirements we described at the start of this article, except for verification. In parallel with experimental progress towards demonstrating quantum supremacy, improved classical simulation results have been proven for noisy and imperfect quantum supremacy experiments.^{11,36,48} Thus the early stages of quantum supremacy experiments are likely to be characterised by an iterative process where proposed supremacy experiments are challenged by efficient classical simulations. Nevertheless, given the speed of recent experimental developments, it seems plausible that quantum supremacy could be convincingly demonstrated in a matter of years.

There are many important open theoretical questions in the area of quantum supremacy. The most

pressing in our view is to develop a scheme that can be efficiently verified, by analogy with the way that the statistics for Bell tests can be easily checked. Developing good classical simulations (or even attempting to and failing) would also help clarify the quantum/classical boundary. The hardness assumptions could also be simplified and strengthened. One ambitious goal in this direction would be to show that simulation with even low variational distance (cf. (2)) would imply the collapse of the PH. Theorists can and should also do more work to come to terms with two other models that appear to be non-universal for quantum computing but where we lack good classical simulations: finite-temperature adiabatic evolution with stoquastic Hamiltonians^{9,23,44} (as used in the quantum annealers available from the company D-Wave) and analog quantum simulation,^{20,29} for example of lattice models.

We close by noting that supremacy is not a long-term goal but rather a necessary step in the development of quantum computers. Eventually we expect that quantum computers will justify themselves by solving important problems which we do not know how to otherwise solve. But in these early days of the field, the focus on quantum supremacy is a way to ensure that quantum computers solve clearly defined problems for which the classical competition can be well understood.

Acknowledgements

AWH was funded by NSF grants CCF-1629809 and CCF-1452616. AM was supported by EPSRC Early Career Fellowship EP/L021005/1. AM would like to thank Mick Bremner and Dan Shepherd for introducing him to this topic, and for many debates concerning it over the years.

Selected references

Some of the key readings on the topic are:

- Terhal and DiVincenzo:⁵⁴ This paper gave the first complexity-theoretic argument that a simple class of quantum circuits should be hard to simulate classically.
- Aaronson and Arkhipov:² This seminal paper introduced the boson sampling problem.
- Bremner, Jozsa and Shepherd:¹² This paper gave evidence that Instantaneous Quantum Polynomial-time (IQP) circuits are hard to simulate classically.
- Boixo et al.:⁸ This paper describes a proposal for a near-term quantum supremacy experiment.

References

- [1] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. R. Soc. A*, 461:3473, 2005, [arXiv:quant-ph/0412187](#).
- [2] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013, [arXiv:1011.3245](#).
- [3] S. Aaronson and A. Arkhipov. Bosonsampling is far from uniform. *Quantum Inf. Comput.*, 14(15&16):1383–1423, 2014, [arXiv:1309.7460](#).
- [4] S. Aaronson and L. Chen. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the 32nd Computational Complexity Conference, CCC '17*, pages 22:1–22:67, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, [arXiv:1612.05903](#).
- [5] S. Aaronson and D. Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, 2004, [arXiv:quant-ph/0406196](#).

- [6] D. Aharonov and U. Vazirani. Is quantum mechanics falsifiable? a computational perspective on the foundations of quantum mechanics. In *Computability: Turing, Gödel, Church, and Beyond*. MIT Press, 2013, [arXiv:1206.3686](#).
- [7] M. Bentivegna, N. Spagnolo, C. Vitelli, F. Flamini, N. Viggianiello, L. Latmiral, P. Mataloni, D. J. Brod, E. F. Galvão, A. Crespi, R. Ramponi, R. Osellame, and F. Sciarrino. Experimental scattershot boson sampling. *Science Advances*, 1(3), 2015, [arXiv:1505.03708](#).
- [8] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices. *Nature Physics*, 14(6):595–600, 2018, [arXiv:1608.00263](#).
- [9] S. Bravyi, D. DiVincenzo, R. Oliveira, and B. Terhal. The complexity of stoquastic local Hamiltonian problems. *Quant. Inf. Comput.*, 8(5):0361–0385, 2008, [arXiv:quant-ph/0606140](#).
- [10] S. Bravyi and D. Gosset. Improved classical simulation of quantum circuits dominated by Clifford gates. *Phys. Rev. Lett.*, 116:250501, 2016, [arXiv:1601.07601](#).
- [11] M. Bremner, A. Montanaro, and D. Shepherd. Achieving quantum supremacy with sparse and noisy commuting quantum circuits, 2016, [arXiv:1610.01808](#).
- [12] M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2010, [arXiv:1005.1407](#).
- [13] M. J. Bremner, A. Montanaro, and D. J. Shepherd. Average-case complexity versus approximate simulation of commuting quantum computations. *Phys. Rev. Lett.*, 117:080501, Aug 2016, [arXiv:1504.07999](#).
- [14] A. Broadbent, J. Fitzsimons, and E. Kashefi. Universal blind quantum computation. In *Proc. 50th Annual Symp. Foundations of Computer Science*, pages 517–526, 2009, [arXiv:0807.4154](#).
- [15] M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White. Photonic boson sampling in a tunable circuit. *Science*, 339(6121):794–798, 2013, [arXiv:1212.2234](#).
- [16] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh, G. D. Marshall, M. G. Thompson, J. C. F. Matthews, T. Hashimoto, J. L. O’Brien, and A. Laing. Universal linear optics. *Science*, 349(6249):711–716, 2015, [arXiv:1505.01182](#).
- [17] J. Carolan, J. Meinecke, P. Shadbolt, N. Russell, N. Ismail, K. Wörhoff, T. Rudolph, M. Thompson, J. O’Brien, J. Matthews, and A. Laing. On the experimental verification of quantum complexity in linear optics. *Nature Photonics*, 8:621–626, 2014, [arXiv:1311.2913](#).
- [18] P. Cheeseman, B. Kanefsky, and W. Taylor. Where the really hard problems are. In *Proc. IJCAI ’91*, pages 331–337, 1991.
- [19] L. W. Cheuk, M. A. Nichols, K. R. Lawrence, M. Okan, H. Zhang, E. Khatami, N. Trivedi, T. Paiva, M. Rigol, and M. W. Zwierlein. Observation of spatial charge and spin correlations in the 2d fermi-hubbard model. *Science*, 353(6305):1260–1264, 2016, [arXiv:1606.04089](#).
- [20] J. I. Cirac and P. Zoller. Goals and opportunities in quantum simulation. *Nature Physics*, 8:264–266, 2012.
- [21] A. Crespi, R. Osellame, R. Ramponi, D. J. Brod, E. F. Galvao, N. Spagnolo, C. Vitelli, E. Maiorino, P. Mataloni, and F. Sciarrino. Integrated multimode interferometers with arbitrary designs for photonic boson sampling. *Nature Photonics*, 7(7):545–549, 2013, [arXiv:1212.2783](#).

- [22] A. M. Dalzell, A. W. Harrow, D. E. Koh, and R. L. La Placa. How many qubits are needed for quantum computational supremacy?, 2018, [arXiv:1805.05224](#).
- [23] N. G. Dickson, M. Johnson, M. Amin, R. Harris, F. Altomare, A. Berkley, P. Bunyk, J. Cai, E. Chapple, P. Chavez, et al. Thermally assisted quantum annealing of a 16-qubit problem. *Nature communications*, 4:1903, 2013.
- [24] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. Quantum computation by adiabatic evolution. Technical Report MIT-CTP-2936, MIT, 2000, [arXiv:quant-ph/0001106](#).
- [25] E. Farhi and A. W. Harrow. Quantum supremacy through the quantum approximate optimization algorithm, 2016, [arXiv:1602.07674](#).
- [26] A. Fowler, M. Mariantoni, J. Martinis, and A. Cleland. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86:032324, 2012, [arXiv:1208.0928](#).
- [27] K. Fujii and S. Tamate. Computational quantum-classical boundary of noisy commuting quantum circuits. *Scientific Reports*, 6:25598, 2016, [arXiv:1406.6932](#).
- [28] X. Gao, S.-T. Wang, and L.-M. Duan. Quantum supremacy for simulating a translation-invariant Ising spin model. *Phys. Rev. Lett.*, 118, 2017, [arXiv:1607.04947](#).
- [29] I. Georgescu, S. Ashhab, and F. Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153, 2014, [arXiv:1308.6253](#).
- [30] D. Gosset, B. Terhal, and A. Vershynina. Universal adiabatic quantum computation via the space-time circuit-to-Hamiltonian construction. *Phys. Rev. Lett.*, 114:140501, 2015, [arXiv:1409.7745](#).
- [31] Y. Han, L. Hemaspaandra, and T. Thierauf. Threshold computation and cryptographic security. *SIAM Journal on Computing*, 26(1):59–78, 1997.
- [32] T. Häner, M. Roetteler, and K. Svore. Factoring using $2n + 2$ qubits with Toffoli based modular multiplication, 2016, [arXiv:1611.07995](#).
- [33] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert. Direct certification of a class of quantum simulations, 2016, [arXiv:1602.00703](#).
- [34] A. W. Harrow and S. Mehraban. Approximate unitary t -designs by short random quantum circuits using nearest-neighbor and long-range gates, 2018, [arXiv:1809.06957](#).
- [35] R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [36] G. Kalai and G. Kindler. Gaussian noise sensitivity and bosonsampling, 2014, [arXiv:1409.3093](#).
- [37] E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434:39–44, 2005, [arXiv:quant-ph/0410199](#).
- [38] E. Knill, R. Laflamme, and W. Zurek. Resilient quantum computation. *Science*, 279(5349):342–345, 1998.
- [39] L. A. Levin. Average case complete problems. *SIAM Journal on Computing*, 15(1):285–286, 1986.
- [40] A. P. Lund, M. J. Bremner, and T. C. Ralph. Quantum sampling problems, bosonsampling and quantum supremacy. *npj Quantum Information*, 3(1):15, 2017, [arXiv:1702.03061](#).
- [41] I. L. Markov and Y. Shi. Simulating quantum computation by contracting tensor networks. *SIAM Journal on Computing*, 38(3):963–981, 2008, [arXiv:quant-ph/0511069](#).

- [42] S. Mertens, M. Mézard, and R. Zecchina. Threshold values of random k-SAT from the cavity method. *Random Struct. Algorithms*, 28(3):340–373, May 2006, [arXiv:cs/0309020](#).
- [43] T. Morimae, K. Fujii, and J. Fitzsimons. On the hardness of classically simulating the one-clean-qubit model. *Phys. Rev. Lett.*, 112:130502, 2014, [arXiv:1312.2496](#).
- [44] K. Nishimura, H. Nishimori, A. J. Ochoa, and H. G. Katzgraber. Retrieving the ground state of spin glasses using thermal noise: Performance of quantum annealing at finite temperatures. *Phys. Rev. E*, 94:032105, Sep 2016, [arXiv:1605.03303](#).
- [45] R. O’Donnell and A. C. C. Say. The weakness of CTC qubits and the power of approximate counting. *ACM Trans. Comput. Theory*, 10(2):5:1–5:22, May 2018. <https://eccc.weizmann.ac.il/report/2016/147/>.
- [46] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [47] J. Preskill. Quantum computing and the entanglement frontier, 2012, [arXiv:1203.5813](#).
- [48] S. Rahimi-Keshari, T. C. Ralph, and C. M. Caves. Sufficient conditions for efficient classical simulation of quantum optics. *Physical Review X*, 6(2):021039, 2016, [arXiv:1511.06526](#).
- [49] D. Shepherd and M. J. Bremner. Temporally unstructured quantum computation. *Proceedings of the Royal Society A*, 465:1413–1439, 2009, [arXiv:0809.0847](#).
- [50] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In S. Goldwasser, editor, *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, pages 124–134, Los Alamitos, CA, 1994. IEEE Computer Society.
- [51] N. Spagnolo, C. Vitelli, M. Bentivegna, D. J. Brod, A. Crespi, F. Flamini, S. Giacomini, G. Milani, R. Ramponi, P. Mataloni, R. Osellame, E. F. Galvao, and F. Sciarrino. Experimental validation of photonic boson sampling. *Nature Photonics*, 8(8):615–620, 2014, [arXiv:1311.1622](#).
- [52] N. Spagnolo, C. Vitelli, L. Sansoni, E. Maiorino, P. Mataloni, F. Sciarrino, D. Brod, E. Galvaõ, A. Crespi, R. Ramponi, and R. Osellame. General rules for bosonic bunching in multimode interferometers. *Phys. Rev. Lett.*, 111:130503, 2013, [arXiv:1305.3188](#).
- [53] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, X.-M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys, J. C. Gates, B. J. Smith, P. G. R. Smith, and I. A. Walmsley. Boson sampling on a photonic chip. *Science*, 339(6121):798–801, 2013, [arXiv:1212.2622](#).
- [54] B. M. Terhal and D. P. DiVincenzo. Adaptive quantum computation, constant-depth quantum circuits and Arthur-Merlin games. *Quantum Info. Comput.*, 4(2):134–145, Mar. 2004, [arXiv:quant-ph/0205133](#).
- [55] M. Tichy, K. Mayer, A. Buchleitner, and K. Mølmer. Stringent and efficient assessment of boson-sampling devices. *Phys. Rev. Lett.*, 113:020502, 2014, [arXiv:1312.3080](#).
- [56] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther. Experimental boson sampling. *Nature Photonics*, 7(7):540–544, 2013, [arXiv:1212.2240](#).
- [57] S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- [58] H. Wang, Y. He, Y.-H. Li, Z.-E. Su, B. Li, H.-L. Huang, X. Ding, M.-C. Chen, C. Liu, J. Qin, J.-P. Li, Y.-M. He, C. Schneider, M. Kamp, C.-Z. Peng, S. Hoeffling, C.-Y. Lu, and J.-W. Pan. Multi-photon boson-sampling machines beating early classical computers, 2016, [arXiv:1612.06956](#).
- [59] J. Watrous. Quantum computational complexity. In *Encyclopedia of Complexity and Systems Science*, pages 7174–7201. Springer New York, 2009, [arXiv:0804.3401](#).