Lecture 3. The Quantum Supremacy Experiment

Pedram Roushan, Adam Zalcman, Charles Neill
Google Quantum AI, California, USA
Started: June 2019, last modified: September 18, 2019

Fundamental questions in chemistry and physics are very difficult to answer due to the exponential complexity of the underlying quantum phenomena. In the early 1980's, Richard Feynman proposed that a quantum computer would be an effective tool to solve such difficult problems [1]. A key step toward building a quantum computer will be performing a computation beyond the capabilities of any classical computer, thus demonstrating the so-called quantum supremacy [2]. We define the quantum supremacy experiment as performing a well-defined computational task in a programmable computer, where the running cost of doing the same computation on a classical computer is exponentially more than to do it on a quantum processor. This poses significant experimental and theoretical challenges:

> Complexity criterion: can we formulate a problem that is hard for a classical computer but easy for a quantum computer?
>
> Fidelity criterion: can a quantum system be engineered to perform a computation in a large enough state-space with low enough errors to provide a speedup?

In this lecture, I discuss the requirements for demonstrating quantum supremacy by presenting experimental data illustrative of these two ingredients of complexity and fidelity.

1. THE COMPUTATIONAL TASK: SAMPLING FROM A RANDOM QUANTUM CIRCUIT

Computational task The computational task that we choose to demonstrate quantum supremacy (QS) is to sample from the output distribution of random quantum circuits (RQCs) [3–5]. In our experiment, the circuits are implemented by repeated application of single- and two-qubit logical operations chosen from a universal set of quantum gates to a group of qubits (a set of gates is said to be universal if any unitary evolution can be approximated using the gates from the set). Sampling the circuit's output produces a collection of bitstrings, e.g. {0000101, 1011100, 0110001, . . . }.

Quantum speckle pattern The evolution of wavefunction under the Schrödinger equation can be understood as the interference of probability amplitudes. Due to this interference, the probability distribution of the measured bitstrings resembles a speckled intensity pattern produced by light interference in laser scatter, such that some bitstrings are much more likely to occur than others.

Computational hardness A computational problem is classically easy if there is a probabilistic classical algorithm which for every input (i.e. an instance of the problem) runs in time *polynomial* in the size of the input and solves the problem with high probability. A computational problem is hard if no such algorithm exists. In other words, a claim that a problem is computationally hard is essentially a negative one - that an algorithm does not exist. This makes proving computational hardness generally difficult.

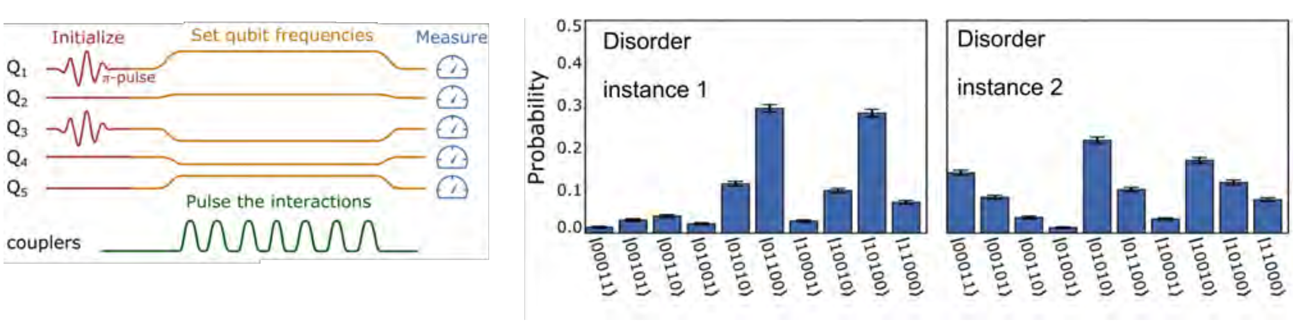

Figure 1. A five qubit example of the RQC protocol. First, the qubits are initialized using microwave pulses (red, to $|10100\rangle$ in this case). Next, the qubit frequencies are randomly chosen and set using rectangular pulses (orange). During this time, all the couplings are pulsed (green). Lastly, we measure the state of every qubit. The measurement is repeated many times in order to estimate the probability of each output state. Here, we plot the measured probabilities for two instances after 10 coupler pulses (cycles).

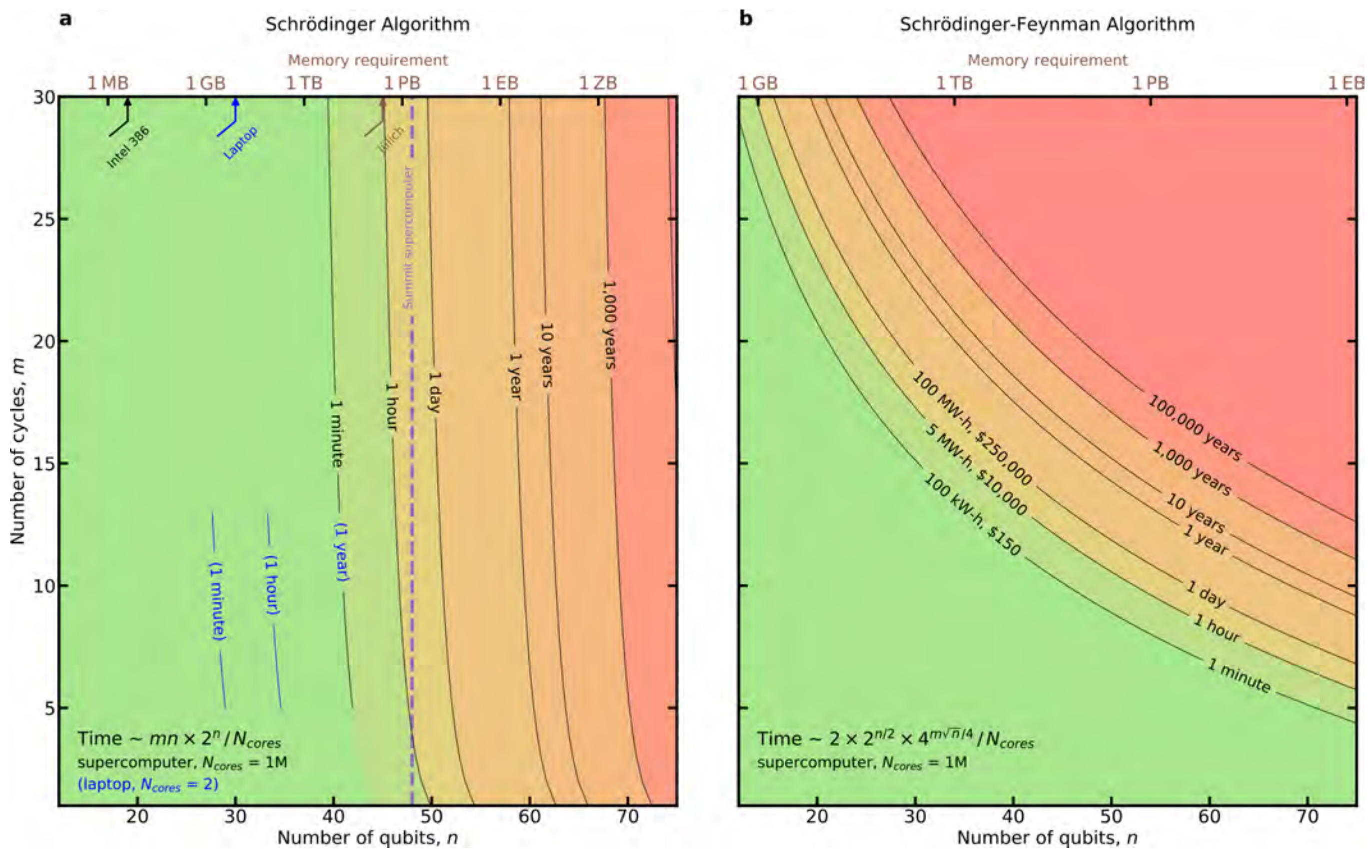

Figure 2. **The computational cost of simulating Schrödinger equation. a,** For a Schrödinger algorithm, the limitation is RAM size, shown as vertical dashed line for the Summit supercomputer, the largest supercomputer in the world. **b,** The Schrödinger-Feynman algorithm [6] breaks the circuit up into two patches of qubits and efficiently simulates each patch using a Schrödinger method, before connecting them using an approach reminiscent of the Feynman path-integral. While it is more memory-efficient, it becomes exponentially more computationally expensive with increasing circuit depth due to the exponential growth of paths with the number of gates connecting the patches.

Classical simulation Since random circuits do not possess structure that could be exploited to classically compute the probabilities of the output bitstrings, it is believed that obtaining the probabilities classically requires a high fidelity simulation.

Two basic approaches to such simulation are provided by the Schrödinger and Feynman algorithms. In the former, the classical computer stores the full wavefunction in memory and propagates it under the evolution induced by the quantum circuit. The memory required to store the full wavefunction of n qubits scales exponentially in n and for a 47-qubit system reaches about a petabyte which is at the limit of the most powerful classical supercomputers.

In practice, one does not need to compute all 2^n amplitudes since for large n only a small fraction of the output bitstrings can be observed in any given experiment. Feynman algorithm computes a single output amplitude by summing the contributions from all Feynman paths and uses only polynomial amount of memory. However, its runtime scales exponentially in the number of gate cycles m (roughly proportional to the evolution time).

There are intermediate approaches that combine ideas from both algorithms and enable classical simulation for moderate values of n and m which in practice are too large for both Schrödinger and Feynman algorithms. However, all known high fidelity classical algorithms for the simulation of RQCs with gates from a universal set require resources *exponential* in n, m or both. The non-existence of polynomial algorithms suggests that sampling from the output probabilities of universal RQCs constitutes a viable approach to demonstrating quantum supremacy.

2. MEASURING FIDELITY: CROSS-ENTROPY BENCHMARKING

Evaluation of small systems Demonstrating QS requires performing a computational task with high fidelity. For the RQC algorithm, to verify that the quantum processor is working with high fidelity one can compare measured probabilities P_{measured} of observed bitstrings with their corresponding theoretical probabilities P_{expected} computed by the simulation on a classical computer

$$Fidelity = \frac{S(P_{\text{incoherent}}, P_{\text{expected}}) - S(P_{\text{measured}}, P_{\text{expected}})}{S(P_{\text{incoherent}}, P_{\text{expected}}) - S(P_{\text{expected}})}, \tag{1}$$

where $P_{\text{incoherent}}$ stands for an incoherent mixture with each output state given equal likelihood and S is the cross-entropy of two sets of probabilities $S(P,Q) = -\sum_i P_i \log(Q_i)$ or the entropy $S(P) = -\sum_i P_i \log(P_i)$ when $P = Q$. $P_{\text{incoherent}}$ describes the behavior observed after a large number of cycles [3, 5]. When the distances between the measured and expected probabilities is small, the fidelity approaches 1. When the measured probabilities approach an incoherent mixture, the fidelity approaches 0.

Evaluation of large systems When the number of qubits n is large, there are many more bitstrings (2^n) than the number of measurements one can afford to perform in the lab ($\sim 10^6$). Therefore, we are unlikely to observe any given bitstring twice, let alone enough times to measure its probability. In this limit and under chaotic dynamics, the above relation takes a simpler form. The linear cross-entropy benchmarking (XEB) estimates the fidelity as the mean of the theoretical probabilities of the measured bitstrings

$$\mathcal{F}_{\text{XEB}} = 2^n \langle P(x_i) \rangle_i - 1, \tag{2}$$

where n is the number of qubits, $P(x_i)$ is the probability of bitstring x_i computed for the ideal quantum circuit, and the average is over the observed bitstrings. Intuitively, $\mathcal{F}_{\text{XEB}}$ is correlated with how often we sample high probability bitstrings. When there are no errors in the quantum circuit, sampling the probability distribution will produce $\mathcal{F}_{\text{XEB}} \simeq 1$. On the other hand, sampling from the uniform distribution corresponding to the maximally mixed state (i.e. complete decoherence) will give $\langle P(x_i) \rangle_i = 1/2^n$ and produce $\mathcal{F}_{\text{XEB}} \simeq 0$. The probabilities $P(x_i)$ must be obtained from classically simulating the quantum circuit, and thus computing $\mathcal{F}_{\text{XEB}}$ is intractable in the QS regime.

3. CHAOTIC DYNAMICS AND THE PORTER-THOMAS DISTRIBUTION

Necessary and sufficient conditions for computational hardness? Quantum processor with n qubits has a Hilbert space with 2^n dimensions. Therefore, it is tempting to imagine that system evolution which involves all 2^n basis states performs a classically intractable computation. It is easy to see that this is not the case. Hadamard transform $H^{\otimes n}$ on n qubits creates a superposition of all 2^n computational basis states yet its output distribution is uniform and therefore easy to sample from classically. Thus, superposition over exponentially many states is by itself insufficient for computational hardness. Hadamard transform generates no entanglement and its output state has a trivial description in the $|+\rangle, |-\rangle$ basis. What about entanglement then? It turns out that entanglement is also insufficient. This is demonstrated by stabilizer circuits, i.e. quantum circuits built from the Hadamard, Phase and CNOT gates. Stabilizer circuits can generate maximally entangled states between any two subsystems of the quantum processor and yet they admit a polynomial-time classical simulation [7]. This is related to the fact that stabilizer circuits explore a large, but discrete set of states in the Hilbert space. In fact, even certain continuously-parametrized quantum gates have the property that circuits composed of them admit polynomial time classical simulation [8].

These examples illustrate that dynamics that may initially appear complex can in fact possess hidden structure that enables efficient classical algorithms. Therefore, an experimental demonstration of QS requires a careful choice of system behavior that is known to resist classical simulation.

Role of chaos and ergodicity The application of RQCs with universal quantum gates leads to chaotic dynamics. Quantum butterfly effect ensures that the output is highly sensitive to minor perturbations in the circuit. This sensitivity defeats approaches to the classical simulation that seek to reduce computational cost by neglecting certain aspects of the RQC (e.g. approximating some gates with similar ones known to be easier to simulate such as Matchgates or Cliffords). Instead, expensive high fidelity classical simulation is believed to be required to approximate the output distribution.

In addition, chaotic dynamics under RQCs leads to ergodicity of the wavefunction. In other words, the state of the quantum processor is not confined to any portion of the Hilbert space, but instead explores the full state space

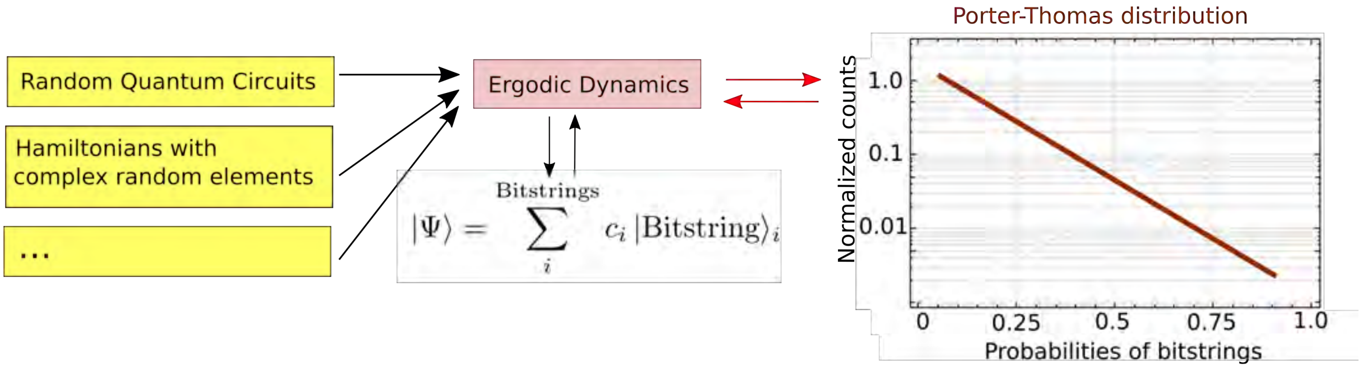

Figure 3. Ergodic dynamics could result from different approaches; as a result of it, the wavefunction uniformly explores entire Hilbert-space and takes the form given in Eqn. (3). The probability of bitstrings associated with this wavefunction takes an exponential form, known as the Porter-Thomas distribution.

uniformly. This helps to defeat approaches to the classical simulation that seek to approximate the output distribution by running a computation on a truncated state space.

Emergence of Porter-Thomas output statistics This brings us to the question of what an ergodic wavefunction looks like? At any instant of time, we can write it as

$$|\Psi\rangle = \sum_i c_i\,|i\rangle \tag{3}$$

where i ranges over n-bit bitstrings and c_i are the complex amplitudes. Uniform exploration of the Hilbert space means that the real vector $[\mathrm{Re}(c_0), \mathrm{Im}(c_0), \mathrm{Re}(c_1), \ldots, \mathrm{Re}(c_{2^n-1}), \mathrm{Im}(c_{2^n-1})]$ is distributed uniformly over the unit sphere in $\mathbb{R}^{2^{n+1}}$. By multiplying the probability density functions of independent standard (i.e. zero mean and unit variance) Gaussian random variables, one observes that the angular part of a random vector with components drawn independently from the standard Gaussian distribution depends only on the vector's length. In other words, it is uniform over all spatial directions, i.e. uniform over the unit sphere. Therefore, for large n the real and imaginary parts of the amplitudes c_i are approximately Gaussian and independent. Also, by symmetry, the Gaussian distributions have the same variance. Now, it is well-known that the sum of squares of two independent and identically distributed Gaussian random variables is an exponential random variable (a special case of the χ^2-distributed random variable). We conclude that the distribution of probabilities $p_i = |c_i|^2$ under chaotic dynamics, known as the Porter-Thomas distribution (PT), is approximately exponential.

Explicitly, for large n the fraction of bitstrings with probability near p is approximately $2^n e^{-2^n p}$. There is a small number of bitstrings with high probabilities while the majority of the bitstrings have small probabilities. This is the origin of the quantum speckle pattern at the output of RQCs described above.

An important caveat concerns the relationship between the PT distribution and the output distribution of the circuit. Whereas the latter maps the output bitstring i to its probability p_i, the former maps the output probability p to the number of bitstrings with probability near p. In other words, the PT distribution *describes the shape of the histogram* of the output probabilities. It does not *specify* the output distribution.

Hallmark of classical hardness? Universal RQCs induce chaotic quantum evolution which is believed hard to simulate classically. Meanwhile, their output distributions produce the characteristic speckles of the PT statistics. This concurrence raises the question whether the presence of the speckles signifies classically intractable system dynamics generally. It turns out the answer to this question is negative. There are systems which are easy to simulate classically and whose output distribution approximates PT [9].

Therefore, the presence of quantum behavior cannot be inferred solely based on the shape of the histogram of the output probabilities. Instead, one must look at the *relationship* between the observed and theoretical output of the system. This suggests that generally verification of quantum supremacy experiments requires classical simulation of large quantum systems and is therefore computationally expensive. In particular, cross-entropy benchmarking accounts for this relationship since it computes fidelity estimate using theoretical probabilities of observed bitstrings.

Computational complexity According to formal results in the computational complexity theory, the problem of computing the probabilities in the output distribution of a RQC is #**P**-hard [10]. This is a class of counting problems which is considered even harder than the more well-known class of **NP**-complete problems.

A decision (i.e. "yes"-"no") problem is in the class **NP** if for every "yes" answer there is a piece of information called a witness that one may exhibit to efficiently prove that the answer to the question posed by the problem is indeed "yes". For example, the question "is a given integer composite" is an **NP** problem, because one may exhibit a divisor which then provides an efficient means of verifying a number is indeed composite. Similarly, the question "does a given Boolean formula have a satisfying assignment" is an **NP** problem, because one may exhibit a satisfying assignment.

A subclass of **NP** called **NP**-complete consists of the most general and difficult problems in **NP** in the sense that we know how to modify any efficient algorithm for any **NP**-complete problem to obtain efficient algorithms for all problems in **NP**. After decades of unsuccessful attempts at solving any of them, the **NP**-complete problems are considered very difficult. It is thought that exhaustively trying all possible witnesses is the best algorithm available to solve any **NP**-complete problem. Therefore, they are sometimes described as the problems of "looking for a needle in a haystack".

Now, #**P**-hard problems are yet more difficult. The solution to a #**P** problem is the *number of witnesses* for the corresponding **NP** problem. Therefore, if one describes an **NP** problem as "looking for a needle in a haystack" then the corresponding #**P** problem is "exactly counting all the needles in a haystack". Problems in #**P**-hard are at least as difficult as the problems in #**P**.

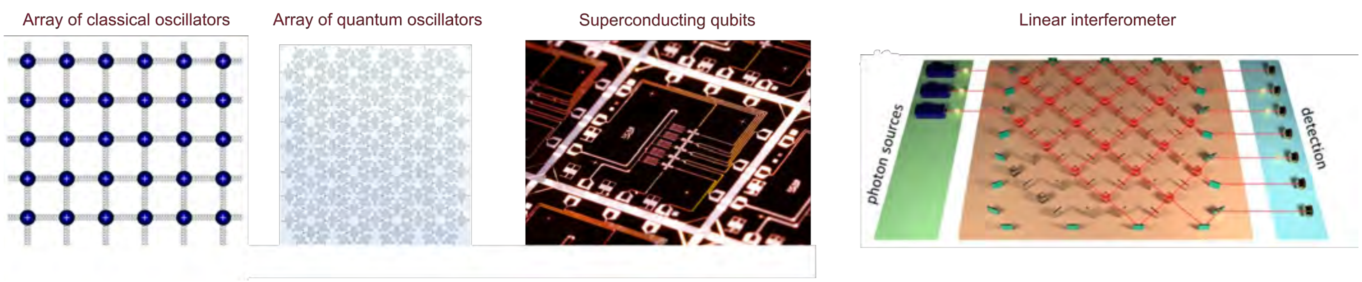

Figure 4. The Boson sampling problem asserts that simulating photon dynamics in a linear interferometer is hard. What is the essence of hardness? If interference is what makes it hard? Can we do this experiment with classical coherent waves? Are quantum systems essential for complexity? how about non-linearity?

4. OTHER PROPOSALS FOR THE DEMONSTRATION OF QUANTUM SUPREMACY

How do we know if dynamics in a platform is complex ? What is the role of non-linearity ? Is it essential ? What about interaction between degrees of freedom ? We explore these question by showing similarities and differences between our choice of platform and algorithms with others.

Linear optical quantum computing Photons have many characteristics, such as long coherence times, that make them promising basis for quantum computing. However, it is difficult to make photons interact. Traditionally, approaches to building two-qubit quantum gates would seek to exploit the Kerr effect, i.e. change in refractive index with light intensity. The effect is observed in non-linear optical media, i.e. media in which electrical polarization varies non-linearly with the electric field. This remains technically challenging as the existing non-linear media suffer from weak non-linearity and high absorption.

Linear optical quantum computing (LOQC) seeks to perform quantum computation with single-photon sources and detectors and linear optical elements such as beamsplitters and phaseshifters. Various linear optical setups were studied with the goal of running a quantum algorithm to solve a classically intractable problem. A series of results proved that some of those interferometers would have to be exponentially large and that others could be simulated on a classical computer in polynomial time. This situation contributed to the belief that linear interferometry alone could not provide a path to universal quantum computation.

The belief that non-linear optical elements may be required to exceed the computational power of classical computers was overturned by Knill, Laflamme and Milburn who proposed an ingenious scheme, now called the KLM protocol, to perform universal quantum computation with linear interferometer. The scheme employs ancilla photons and post-selection to perform non-deterministic two-qubit quantum gates whose success probability is subsequently boosted with gate teleportation.

Non-universal, yet computationally hard While KLM protocol obviates the need for the use of non-linear optical media, it does require active elements such as photo-detectors and fast feed-forward mechanism which re-configures the downstream portion of a quantum circuit based on the upstream measurement results. This presents a formidable technical challenge in practice.

A further boost for LOQC occurred when Aaronson and Arkhipov showed that the operation of a much simpler linear optics setup without active elements cannot likely be simulated by a classical computer [11]. They provided evidence that boson sampling is an intractable problem as it is related to the evaluation of permanents of random complex matrices, a problem known to be #**P**-hard (similarly to the problem of computing output probabilities of a RQC we have seen earlier).

Interestingly, boson sampling shows that simpler quantum computing setups which do not achieve the full computational power of a universal quantum computer can nonetheless perform computational tasks thought to be classically intractable. Since such setups are often easier to realize experimentally they constitute natural candidates for the demonstration of quantum supremacy.

In summary, linear optical interferometers can achieve universal quantum computation by employing the (very complex) KLM protocol. Moreover, even simpler interferometers that fall short of the full power of a universal quantum computer can still perform computational tasks believed to be classically intractable.

Hilbert space of the optical field A solution of the wave equation can be expanded in the sum of orthogonal plane waves known as modes. In the quantization of the electromagnetic field each mode is quantized independently and admits a description equivalent to the quantum harmonic oscillator. Consequently, the Hilbert space of the optical field has a basis, called the occupation number basis, consisting of states with well-defined number of photons in each mode, such as for example $|3, 2, 1\rangle$ with three photons in the first mode, two in the second mode and one in the third. These states are also know as Fock states.

Boson sampling problem The boson sampling problem consists of sampling from the output distribution of n indistinguishable bosons (in practice, photons) that interfere during the evolution through a linear interferometer. The initial state of the system is a Fock state with n single photons in m modes $|\psi_0\rangle = |1_1, ..., 1_n, 0_{n+1}, ...0_m\rangle$. The input state is evolved via a passive linear optics network: an n-mode interferometer described by a random n-dimensional unitary operator. The output state is a superposition of the different configurations of how the m photons could have arrived in the output modes. The measurement in the lab returns n-long (?) bitstrings, by placing photon detectors on every output, similar to what we measure with qubits.

Linear and non-interacting, yet computationally hard. The interferometer commonly described as a passive linear device with non-interacting photons. However, there is an effective interaction at each beam-splitter; the presence or absence of a photon in one input mode radically changes the output state of a second photon in another input mode. This "interaction" between indistinguishable particles, known as the exchange interaction, arises simply from the demand that the multi-particle wavefunction be properly symmetrized. It is therefore a misnomer to describe these interferometers as linear devices with non-interacting bosons. Interestingly, in the ideal case we ignore this effect and yet the problem is hard to solve classically. Note: just to be clear by interaction we mean the presence or absence of photon at a beam-splitter or phase-shifter changes the state of other photons. By linear we mean that phase of two photon after they entered a splitter or shifter device is a linear combination of their initial phases before they enter this device.

5. COMMENTS

The goal of the QS experiment is to achieve a high enough fidelity for a circuit with sufficient width and depth such that the classical computing cost is prohibitively large. In this section we discuss this matter in more details and bring up practical considerations.

5.1. Requirements for demonstrating quantum supremacy

The notion of quantum supremacy was originally introduced by John Preskill [2]. He conceived of it as *"the day when well-controlled quantum systems can perform tasks surpassing what can be done in the classical world"*. For the purpose of an experimental demonstration we would like to refine this concept:

1. A well-defined computational task, i.e. a mathematical specification of a computational problem with a well defined solution. This requirement, standard in computer science, excludes tasks such as simulate a glass of water. However, it would include finding the ground state energy of a water molecule to a given precision governed by a specific Hamiltonian. Note that this requirement poses a difficult experimental challenge, since logic gates are imperfect and the quantum states we intend to create are sensitive to errors. for instance in RQC, a single bit or phase flip over the course of the algorithm will completely shuffle the speckle pattern. Therefore, in order to claim quantum supremacy we need a quantum processor that executes the program with sufficiently low error rates.

2. Programmable computational device. Many physics experiments estimate the values of observables to a precision which can not be obtained numerically. But those do not involve a freely programmable computational device. Ideally, we would even restrict ourselves to devices that are computationally universal. However, this would exclude proposals to demonstrate quantum supremacy with Boson Sampling [11] or IQP circuits [12].

3. A scaling run time difference between the quantum and classical computational processes that can be made large enough as a function of problem size so that it becomes impractical for a supercomputer to solve the task using any known classical algorithm. One may wonder to what extent algorithmic innovation can enhance classical simulations. What is impractical for classical computers today may become tractable in ten years. Indeed, simulation methods have improved steadily over the past few years. We expect that simulation costs will steadily reduce in the future, but we also expect they will be consistently outpaced by hardware improvements on larger quantum processors. So the quantum supremacy frontier will be moving towards larger and larger problems. But if a task is chosen such that the scaling for the quantum processors is polynomial while for the classical computer is exponential then this shift will be small. Sampling the output of random quantum circuits is likely to exhibit this scaling separation as a function of the number of qubits for large enough depth.

5.2. The Church-Turing thesis

The Church-Turing (CT) thesis expresses the belief that any real-world physical process can be simulated on a probabilistic Turing machine (a formal model of the classical computer). In other words, if there is an algorithm to make some problem computable, then that problem can be done on a Turing machine. CT concerns the question of the theoretical feasibility of the simulation and neglects the computational resources (time and memory) that affect its practical realizability.

By contrast, the Extended Church-Turing (ECT) thesis, formulated by Bernstein and Vazirani [13], expresses the stronger belief that any behavior of a real-world physical system can be simulated on a probabilistic Turing machine using *computational resources polynomial* in the size of the system.

The promise of quantum computation does not undermine our confidence in CT. However, a successful experimental demonstration of QS would strongly suggest that a model of computation exists that violates ECT.

5.3. Complexity class of RQC. where Adam will shine

In the specific case of QS demonstration using RQC algorithm, this is because a quantum processor can sample bitstrings in polynomial time and yet no efficient method is known to exist for classical computing machinery as all the classical algorithms are exponential in time.

What is the complexity class of Simulating RQC? is NP hard?

Our assumption, based on insights from complexity theory, is that the cost of this algorithmic task is exponential in the number of qubits and as well as the depth of the circuit or evolution time.

We note that formal complexity claims are asymptotic, and therefore assume an arbitrarily large number of qubits. This is only possible with a fault tolerant quantum computer and therefore near term practical demonstrations of quantum supremacy must rely on a careful comparison with highly optimized classical algorithms on state-of-the-art supercomputers.

	Classical oscillators	Quantum oscillators	Qubits	Qubits	Linear interferometer
	n masses (modes)	m photons, n resonators	m photons, n qubits	m photons, n qubits	m photons, n modes
Experiment	random initial position of n's	m photons randomly placed	RQC	only Clifford gates	m photons randomly placed
Dimension of computational space	$2n$ space and momentum	$\binom{n}{m}$	$\binom{n}{m}$	$\binom{n}{m}$	$\binom{n}{m}$
Ergodic	yes	yes	yes	no	yes
Non-linear	xx	xx	xx	xx	xx

6. HOMEWORKS

• HW1. Use bitstrings (Fock states, e.g. $\{0100101,\ 1010100,\ \dots\}$) as basis. Consider a wavefunction in a 2^n-dimensional Hilbert space where the coefficient of bitstrings are complex numbers, with real and imaginary part chosen from a Gaussian distribution (with mean and StD of ???). Show that output distribution, i.e. measurement of the probability of bitstrings, obeys the Porter-Thomas distribution. Plot this resulting distribution for several values of n.

• HW2. Make a random Hamiltonian H by starting with a random complex matrix. Be sure H is Hermitian. Create a unitary operator by exponentiation this matrix $U = e^{iH}$. Consider a system of qubits, initialized by preparing all possible bitstrings. Convince yourself that applying U to these initial states gives you the square of the elements of this unitary, $|u_{p,q}|^2$. Show that the distribution of $|u_{p,q}|^2$ obeys Porter-Thomas distribution.

• HW3. Consider $H_{XY} = \sum h_i \sigma_i^Z + J \sum_{<i,j>} \sigma_i^X \sigma_j^X + \sigma_i^Y \sigma_j^Y$, where each qubit is only coupled to its nearest neighbor. By exponentiating this Hamiltonian for a given time, examine if the the resulting evolution is ergodic or not computing the PT distribution. Do this procedure for 1D as well as 2D grid of qubits.

• HW4. Repeat the above exercise for $H = H_{XY} + \gamma \sum_{<i,j>} \sigma_i^Z \sigma_j^Z$. It would be impressive if your code can consider 50 instances of Hamiltonians in a 4×4 grid in a few hours.

Acknowledgments. I wish to acknowledge fruitful discussions with members of the Google Quantum AI team; in particular Sergio Boixo, Jimmy Chen, Kostyantyn Kechedzhi, Alexander Korotkov, Fedor Kostritsa, Daniel Sank, Kevin Satzinger, and Vadim Smelyanskiy.

[1] Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys.* **21**, 467–488 (1982).
[2] Preskill, J. Quantum computing and the entanglement frontier. *Rapporteur talk at the 25th Solvay Conference on Physics, Brussels* (2012).
[3] Boixo, S. *et al.* Characterizing quantum supremacy in near-term devices. *Nature Physics* **14**, 595 (2018).
[4] Aaronson, S. & Chen, L. Complexity-theoretic foundations of quantum supremacy experiments. In *32nd Computational Complexity Conference (CCC 2017)* (2017).
[5] Neill, C. *et al.* A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science* **360**, 195–199 (2018).
[6] Markov, I. L., Fatima, A., Isakov, S. V. & Boixo, S. Quantum Supremacy Is Both Closer and Farther than It Appears. *arXiv:1807.10749* (2018).
[7] Aaronson, S. & Gottesman, D. Improved simulation of stabilizer circuits. *Phys. Rev. A* **70**, 052328 (2004).

[8] Jozsa, R. & Miyake, A. Matchgates and classical simulation of quantum circuits. *Proc. R. Soc. A* **464**, 3089–3106 (2008).

[9] Bouland, A., Fefferman, B., Nirkhe, C. & Vazirani, U. Quantum supremacy and the complexity of random circuit sampling. *Preprint at https://arxiv.org/abs/1803.04402* (2018).

[10] Bouland, A., Fefferman, B., Nirkhe, C. & Vazirani, U. On the complexity and verification of quantum random circuit sampling. *Nature Physics* **15**, 159 (2019).

[11] Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 333–342 (ACM, 2011).

[12] Bremner, M. J., Montanaro, A. & Shepherd, D. J. Average-case complexity versus approximate simulation of commuting quantum computations. *Physical review letters* **117**, 080501 (2016).

[13] Bernstein, E. & Vazirani, U. Quantum complexity theory. *Proc. 25th Annual ACM Symposium on Theory of Computing* (1993).