Efficient classical simulation of random shallow 2D quantum circuits

John Napp, Rolando La Placa, Alexander Dalzell, Fernando Brandão, Aram Harrow

Overview. One of the most basic questions in quantum computing is whether a given class of quantum computations admits an efficient classical simulation. Due to the exponential blowup in Hilbert space dimension associated with an extensive quantum system, brute-force simulation quickly becomes intractable. However, this does not preclude the existence of clever classical simulation algorithms which exploit extra structure of the problem at hand to obtain efficiency: classic examples include 1D quantum circuits with low entanglement (simulatable via MPS) and Clifford circuits (simulatable via the Gottesman-Knill theorem).

In this work, we add another class of quantum circuits to this list: namely, random, sufficiently shallow, 2D quantum circuits. Besides the fact that delineating the boundary between classically simulatable and classically hard-to-simulate processes is a basic question in computational complexity theory, random circuits are of use in physics due to their being models of chaotic quantum dynamics, and in quantum computation due to (among other applications) the fact that Random Circuit Sampling (RCS) — that is, sampling from the output distribution produced by a random circuit — is a leading candidate for demonstrating so-called "quantum computational supremacy", and in fact was the basis for Google's recent claim of achieving this [1]. In this context, it becomes crucial to understand the asymptotic hardness of the RCS problem. Unfortunately, while there are some hardness results for RCS under additional conjectures which have not received widespread scrutiny, strong theoretical evidence remains sparse. A major advance was made by Bouland, Fefferman, Nirkhe, and Vazirani (QIP 2019) [2], who proved that the alternative problem of computing output probabilities of random circuits admits a certain worst-to-average case reduction; since computing output probabilities of generic instances is intractable under widely believed assumptions, this work essentially showed that computing output probabilities of random circuits is intractable. Subsequently, Movassagh made a technical improvement to this result (QIP 2020) [3].

Despite these strong hardness results for computing output probabilities, we describe an efficient classical algorithm for the RCS problem in the setting of sufficiently shallow 2D circuits. This algorithm is only efficient for *shallow* circuits, and therefore does not refute Google's quantum computational supremacy claims which are based on *deep* circuits. Nonetheless, the existence of such an algorithm has some bearing on the line of work attempting to establish the hardness of RCS. The conventional wisdom, partially informed by [2, 3], has been that for a given architecture, random gates should be almost the hardest possible. Our work challenges this intuition and cautions against making bold average-case hardness conjectures by showing that there is a natural setting where simulation is exponentially easier for typical instances than for the worst case.

As a more concrete implication, our work demonstrates limitations of the hardness results of [2, 3] in providing evidence for the hardness of RCS. In particular, these hardness results do apply to the shallow random circuit families we study, showing that near-exactly computing output probabilities is average-case hard. Yet, we find that it is classically tractable to solve the RCS problem, and even to compute typical output probabilities when some small additive error is allowed. While the hardness results of [2, 3] apply to computing output probabilities, they have been widely cited as evidence for the hardness of RCS; however, our work implies that there are natural settings in which classically solving the RCS problem is far easier than precisely computing output probabilities in the average case. Therefore, hardness results for the latter task should not in isolation be viewed as evidence for hardness of the former.

Interestingly, we also find evidence that our algorithms experience computational phase transitions from polynomial-time to exponential-time when the circuit depth or local dimension exceeds some critical, constant value. As elaborated upon below, we relate these computational phase transitions to (1) measurement-driven entanglement phase transitions in 1D chaotic quantum dynamics, and (2) phase transitions in classical statistical mechanical models.

Simulation algorithms. We propose two classical simulation algorithms for (noiseless) random shallow 2D quantum circuits. To the best of our knowledge, these algorithms represent the first simulation algorithms for 2D random circuits which go beyond (exponential-time) methods based on tensor network contraction. The first algorithm, which is also the algorithm we study in greatest depth, is based on a reduction of the 2D simulation problem to the problem of simulating a certain 1D dynamics evolving in time. In turn, this effective 1D dynamics is simulated via Matrix Product State (MPS) methods. More precisely, the effective 1D dynamics is simulated using the Time Evolving Block Decimation (TEBD) algorithm of Vidal [4], which involves periodically truncating (i.e. compressing) the MPS; we therefore refer to this algorithm as Space Evolving Block Decimation (`SEBD`). While MPS are used to simulate the effective 1D dynamics, we note that `SEBD` is not merely a tensor network contraction scheme, but crucially exploits the unitarity of the circuit for its effectiveness. This differentiates it from simulation proposals for random circuits based on truncated tensor network contractions (e.g. [5]). Also unlike such truncated tensor network contraction approaches, the algorithm is *self-certifying* in the sense that it can bound the sampling error it's making as a function of the Schmidt coefficients discarded in the MPS compression steps, even though exact simulation is hard and therefore the exact solution is unknown. This self-certification feature allows us to numerically study the performance of the algorithm. Our second proposed algorithm, which we call `Patching`, is based on first exactly sampling from the marginal distributions of small causally disconnected regions, before "stitching" these patches together via recovery maps to obtain a global sample. The efficiency of the first algorithm hinges on the effective 1D process having low entanglement, while the second hinges on the classical output distribution being approximately Markovian in the sense that the conditional mutual information (CMI) decays exponentially quickly with respect to shielded regions.

Results and techniques. We give two classes of results: first, a rigorous proof that `SEBD` is efficient in a specific situation that is hard in the worst case (Theorem 1), and second, numerical and analytical evidence that the runtime of our algorithms is efficient more generally when the circuits are sufficiently shallow, transitioning to inefficient when the qudit local dimension or circuit depth becomes too large.

`SEBD` involves a reduction of the 2D simulation problem to a 1D simulation problem evolving over time. We find that, after performing this reduction, the effective 1D dynamics are highly similar to those of "unitary-and-measurement" processes, which have seen an explosion in interest within the condensed matter physics community in the past few years [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. This line of work has found strong evidence that, when random unitary dynamics are interspersed with weak measurements, the 1D dynamics experiences an entanglement phase transition from an area-law to volume-law phase as measurement strength is tuned. We find that, roughly, increasing the depth or the local dimension corresponds to decreasing the measurement strength in this picture, hence the phase transition from an efficient to inefficient regime.

Unfortunately, the aforementioned measurement-driven entanglement phase transition has so far eluded formal proof, and it is similarly hard to formally prove that `SEBD` is generally efficient for random shallow 2D circuits, as the efficiency of `SEBD` hinges on its associated effective 1D dynamics lying in an entanglement area-law phase. Nonetheless, we are able to formally prove that `SEBD` runs in polynomial time for some special-case 2D architectures for which exact simulation is known (from prior works) to be hard.

Theorem 1 (Informal)**.** *There exists a 2D circuit architecture A defined on n qubits such that, if C_A is the Haar-random circuit family associated with A, under standard hardness conjectures there does not exist a polynomial-time classical algorithm for sampling from the output distribution of arbitrary instances of C_A, or for near-exactly (i.e. up to $e^{-\Omega(n^2)}$ precision) computing output probabilities of typical instances of C_A. However, `SEBD` runs in time $O(n)$ and, with probability $1-2^{-n^{0.99}}$ over choice of circuit instance, samples from the output distribution of C_A up to error at most $2^{-n^{0.99}}$ in total variation distance, and estimates a fixed output probability of C_A with additive error $2^{-n}/2^{n^{0.99}}$.*

While the architecture used to prove this result is contrived, it suffices to demonstrate a formal separation between the hardness of approximately simulating average-case circuits and worst-case simulation. The proof of the above theorem lies partly on a technical lemma on the typical behavior of entanglement after measurement; namely, we show that if a contiguous block of m qubits are measured after applying a random shallow 1D circuit, the expected entanglement entropy of the induced bipartite pure state is exponentially small in m. While this theorem applies for a contrived architecture for which a formal proof is feasible, it is desirable to understand the performance of these simulation algorithms for more general and more natural 2D circuit architectures. To this end, we make conjectures on the general performance of the algorithms, which are summarized informally below.

Conjecture 1 (Informal)**.** *`SEBD` and `Patching` are asymptotically efficient for any family of 2D quantum circuits with Haar-random gates of sufficiently low depth, but experience a computational phase transition to an inefficient (exponential runtime) regime when the depth or local Hilbert space dimension exceeds some (constant) critical value. The critical values are architecture-dependent.*

We collect numerical and heuristic, analytical evidence for these conjectures. Numerically, we implement `SEBD` for two natural 2D circuit architectures for which worst-case simulation is known to be hard. For the so-called "brickwork" architecture of depth three, we find that `SEBD` running on a laptop with non-optimized code can sample from the output distribution of typical instances on a 400×400 grid of qubits with total variation distance error less than 0.01, requiring approximately one minute per sample. In contrast, we estimate that for previously best-known algorithms based on tensor network contraction, simulating such instances would be roughly as hard as simulating a depth-40 circuit on a 20×20 grid, which is highly intractable. Numerical investigation of the typical entanglement spectrum in the effective 1D dynamics indicates that the algorithm remains efficient asymptotically. Moreover, to further understand the entanglement spectrum we analytically study a toy model for a unitary-and-measurement dynamic in an area-law phase, and find that its entanglement spectrum is consistent with that observed numerically. The entanglement spectrum derived for this toy model leads us to conjecture that, when `SEBD` is in the efficient regime, its runtime scales asymptotically like $n^{1+o(1)} \exp\left(\Theta\left(\sqrt{\log(1/\varepsilon\delta)}\right)\right)$ to simulate a $1-\delta$ fraction of n-qubit random circuit instances up to total variation distance error ε. This toy model could be useful in other scenarios where "unitary-and-measurement" processes have been studied. Indeed, in that context, our work might be regarded as a new and distinct reason to study these processes, which previously have been studied primarily from the condensed matter and quantum chaos angles.

We also give analytical evidence for both algorithms' efficiency using methods from statistical mechanics. We map 2D shallow circuits with Haar-random gates to classical statistical mechanical models, utilizing techniques developed in [17, 18, 19, 20, 13, 14], such that the free energy cost incurred by twisting boundary conditions of the stat mech model corresponds to entanglement-entropy-like quantities in the quantum circuit we call the "quasi-entropy". Increasing the qudit dimension q of the random circuit model corresponds to strengthening the interactions in the associated stat mech model. For depth-3 circuits with brickwork architecture, we show that this drives a phase transition from a disordered to an ordered phase at some constant critical value of q. We also argue that there should be a similar phase transition driven by the circuit depth d that occurs at a constant critical value of d. These phase transitions are accompanied by phase transitions in the quasi-entropy from an area-law to a volume-law phase. While an area law for the quasi-entropy does not prove that our algorithms are efficient (this is implied by an area law for different entropic quantities), we take this as evidence in favor of that conjecture, and the phase transition in quasi-entropy supports our conjecture of a computational phase transition in our algorithms.

References

[1] Frank Arute et al. "Quantum supremacy using a programmable superconducting processor". In: *Nature* 574.7779 (2019), pp. 505–510.

[2] Adam Bouland et al. "On the complexity and verification of quantum random circuit sampling". In: *Nature Physics* 15.2 (2019), p. 159. arXiv: 1803.04402.

[3] Ramis Movassagh. *Cayley path and quantum computational supremacy: A proof of average-case #P–hardness of Random Circuit Sampling with quantified robustness*. 2019. arXiv: 1909.06210.

[4] Guifré Vidal. "Efficient simulation of one-dimensional quantum many-body systems". In: *Physical Review Letters* 93.4 (2004), p. 040502. arXiv: quant-ph/0310089.

[5] Feng Pan et al. *Contracting Arbitrary Tensor Networks: general approximate algorithm and applications in graphical models and quantum circuit simulations*. 2019. arXiv: 1912.03014 [physics.comp-ph].

[6] Yaodong Li, Xiao Chen, and Matthew PA Fisher. "Quantum Zeno effect and the many-body entanglement transition". In: *Physical Review B* 98.20 (2018), p. 205136. arXiv: 1808.06134.

[7] Amos Chan et al. *Weak Measurements Limit Entanglement to Area Law*. 2018. arXiv: 1808.05949.

[8] Brian Skinner, Jonathan Ruhman, and Adam Nahum. "Measurement-Induced Phase Transitions in the Dynamics of Entanglement". In: *Physical Review X* 9.3 (2019), p. 031009. arXiv: 1808.05953.

[9] Yaodong Li, Xiao Chen, and Matthew Fisher. *Measurement-driven entanglement transition in hybrid quantum circuits*. 2019. arXiv: 1901.08092.

[10] Marcin Szyniszewski, Alessandro Romito, and Henning Schomerus. *Entanglement transition from variable-strength weak measurements*. 2019. arXiv: 1903.05452.

[11] Soonwon Choi et al. *Quantum error correction and entanglement phase transition in random unitary circuits with projective measurements*. 2019. arXiv: 1903.05124.

[12] Michael J Gullans and David A Huse. *Dynamical purification phase transition induced by quantum measurements*. 2019. arXiv: 1905.05195.

[13] Yimu Bao, Soonwon Choi, and Ehud Altman. *Theory of the Phase Transition in Random Unitary Circuits with Measurements*. 2019. arXiv: 1908.04305.

[14] Chao-Ming Jian et al. *Measurement-induced criticality in random quantum circuits*. 2019. arXiv: 1908.08051.

[15] Michael J Gullans and David A Huse. *Scalable probes of measurement-induced criticality*. 2019.

[16] Aidan Zabalo et al. *Critical properties of the measurement-induced transition in random quantum circuits*. 2019. arXiv: 1911.00008.

[17] Adam Nahum, Sagar Vijay, and Jeongwan Haah. "Operator spreading in random unitary circuits". In: *Physical Review X* 8.2 (2018), p. 021014. arXiv: 1705.08975.

[18] CW Von Keyserlingk et al. "Operator hydrodynamics, OTOCs, and entanglement growth in systems without conservation laws". In: *Physical Review X* 8.2 (2018), p. 021013. arXiv: 1705.08910.

[19] Tianci Zhou and Adam Nahum. "Emergent statistical mechanics of entanglement in random unitary circuits". In: *Physical Review B* 99.17 (2019), p. 174205. arXiv: 1804.09737.

[20] Nicholas Hunter-Jones. *Unitary designs from statistical mechanics in random quantum circuits*. 2019. arXiv: 1905.12053.