Quantum advantage for computations with limited space

Dmitri Maslov,[1] Jin-Sung Kim,[2] Sergey Bravyi,[1] Theodore J. Yoder,[1] and Sarah Sheldon[2]

[1]*IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA*
[2]*IBM Quantum, Almaden Research Center, San Jose, CA 95120, USA*

Quantum computations are studied for their potential to offer an advantage over regular classical computations. The extent and provability of such advantage depend on the computational model selected. A simple example of a computational model can be a game, such as CHSH game [1] (Bell's inequality [2]). The best classical probability of winning this game, $\frac{3}{4}$, can be improved to $\frac{2+\sqrt{2}}{4}$ with the use of a quantum computer. While this gap allows to experimentally demonstrate quantumness, there is very little quantum computation involved, and Bell's inequality can be attributed to the property of quantum states rather than computations. A second model studies computations with black boxes, including algorithms such as Deutsch-Jozsa [3], Bernstein–Vazirani [4], and Grover's [5]). A practical experimental use of Grover's search, agruably the most practical of these, is likely far in the future given a mere quadratic speedup. A third computational model studies white box computations, and allows superpolynomial advantage for solving problems such as Hamiltonian dynamics simulation [6–8]. In this case, separations are not established formally. A quantum computer capable of outperforming a classical computer will likely need to be large—about 70 qubits and 650,000 gates in some of the shortest known quantum circuits solving a computational problem that is believed to be intractable for classical hardware [9]. Finally, a provable quantum advantage was established for the parallel model of computation [10–15]. It remains to be seen whether this type of advantage can be demonstrated experimentally with near-term devices due to the large number of qubits required.

Here we study a simple computational model that allows to both establish a provable separation between classical and quantum computational models and validate it experimentally. Our model is designed to highlight the superiority of quantum computational space, resulting in a different type of advantage compared to those examples highlighted in the previous paragraph. A related space advantage should be possible to exploit to improve computations beyond those explicitly discussed here.

Formally, we consider classical and quantum circuits where the input (also called primary input to distinguish from the constant qubit called the computational space) is a read-only memory (input cannot be written on), and the computational space is restricted to s bits. In the classical case, computations proceed by arbitrary s+1-input s-output Boolean functions/gates g, where exactly one bit of the input to g is from the primary input, and all outputs are computational bits. For s=1 this means 2-input 1-output Boolean gates, being the staple gate library for classical computations. The closest analog to such transformations in the quantum world is the controlled-U gates, where the unitary operation U is applied to the computational register and controlled by a primary input. We call this model limited-space computation.

The set of functions uncomputable by 1-bit limited-space classical computations includes symmetric functions with nontrivial Fourier spectra (equivalently, those that cannot be written as fixed polarity Reed-Muller expression with degree 0, 1, and n terms only). This implies that most symmetric functions may not be computed classically in this model. However, they can be computed by a quantum circuit with $O(n^2)$ entangling gates and 1 qubit of computational space, as discussed later. Other than symmetric Boolean functions, polynomial-size 1-qubit limited-space quantum computations include at least those functions in the NC^1 class, such as Boolean components of the integer addition, integer multiplication, and matrix determinant [16], as well as all linear combinations $f(x) \oplus g(y)$ where f and g are polynomial-size computable; most of these functions are uncomputable by 1-bit limited-space classical computations.

When the computational space is increased to 2 bits, the classical model can compute any Boolean function (e.g., by Disjunctive Normal Form), although circuit complexity may be high. Here, we focus on the case s=1 to maximize the gap between classical and quantum computations.

With perfect quantum computers, we would be able to demonstrate that the quantum computer always succeeds at computing those functions uncomputable by the classical 1-bit limited-space circuits. Unfortunately, current quantum computers are noisy and sometimes fail. This failure is often modeled probabilistically. To demonstrate quantum advantage using noisy quantum computers over (perfect) classical computers in an experiment, it would be fair to arm classical computations with free access to randomness. Specifically, we allow the classical computer to randomly select a limited-space circuit to run or, equivalently, replace Boolean gates $g(x_i, s)$ in it with Boolean gates $g(x_i, s, r)$, where r is a random number. We furthermore allow the classical limited-space computer to evaluate functions with probability p, which is equal to the normalized Hamming distance between truth vectors of the computable and desired functions. The value p for classical computations is thus analogous to ASP (Algorithmic Success Probability) in quantum computations. Computational machinery that achieves ASP above the maximal classical value p performs a computation unreachable by classical means and is thereby super-classical. Here we demonstrate a selection of experiments that achieve this.

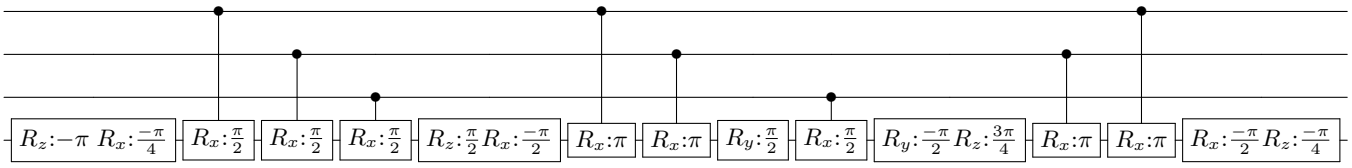

FIG. 1. True SLSB3 with 8 entangling gates, obtained using signal processing technique and local optimization. The gates used are axial rotations $R_{\{x,y,z\}}{:}\theta$ by the angle θ and their controlled versions.

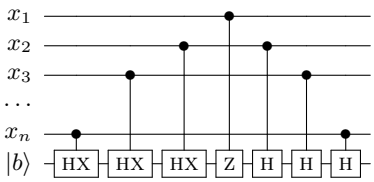

FIG. 2. Relative-phase SLSBn using $2n-1$ entangling gates. The gates used are the controlled versions of the Clifford gate HX, Pauli-Z gate, and the Hadamard gate H.

The simplest function not computable in the 1-bit limited-space classical model is SLSB3$(x_1, x_2, x_3) = x_1x_2 \oplus x_2x_3 \oplus x_1x_3$. In general, SLSB$n$ is defined as the value of the Second Least Significant Bit of the input weight $|x|$. The maximal classical probability p of computing SLSB3 using limited-space computations is $\frac{7}{8} = 0.875$, meaning the truth vector distance to a computable function is 1. We developed two quantum circuits to compute SLSB3, one with 5 entangling gates (Fig. 2) and one with 8 entangling gates (Fig. 1), achieved with the use of Quantum Signal Processing (QSP). The quantum computer ASPs are 0.9429 ± 0.0011 and 0.9280 ± 0.0013, respectively. For 4 bits, the function SLSB4 achieves the minimal among maximal classical values $p = \frac{13}{16} = 0.8125$ across all symmetric Boolean functions. We developed a quantum circuit with 7 entangling gates, Fig. 2, that maps into a quantum circuit with 13 entangling gates over the experiment, due to the requirement to use two SWAP gates. The measured ASP is 0.8743 ± 0.0035. For 5 bits, the function SLSB5 is most difficult to approximate classically, with the threshold of $\frac{23}{32} = 0.71875$; we achieved quantum ASP of 0.8460 ± 0.0053 by a quantum circuit with 9 entangling gates (21 in the experiment), Fig. 2. For 6 bits, the most difficult function is SLSB6, featuring the threshold value $\frac{43}{64} = 0.671875$; we implemented it with fidelity 0.7984 ± 0.0047 over quantum circuit with 11 gates (29 in the experiment), Fig. 2. In each of these experiments, we beat the classical threshold, thus demonstrating a quantum advantage.

For arbitrary n, SLSBn as well as any symmetric Boolean function, can be computed using $O(n^2)$ entangling gates by a quantum limited-space circuit, constructed using QSP. SLSBn may furthermore be computed by a specialized circuit using $2n-1$ gates (see Fig. 2), showing that QSP gives a loose upper bound. The classical probability p of evaluating SLSBn correctly within the limited-space computational model approaches the theoretical minimum of $\frac{1}{2}$ exponentially fast, namely, $p \leq 1/2 + O(n/\sqrt{2}^n)$. This presents an opportunity to demonstrate larger quantum advantage with a higher number of qubits. Formal proofs of the above statements can be found in the full paper.

Our goal is the construction of a quantum circuit implementation of the n-bit Boolean function $f(x)$, expressed by an $n+1$-qubit unitary $U : |x\rangle |b\rangle \rightarrow e^{i\theta(x,b)} |x\rangle |b \oplus f(x)\rangle$ for some real-valued function $\theta(x, b)$. In the 1-qubit limited-space model we may write $U = \sum_x |x\rangle \langle x| \otimes U(x)$, where $U(x)$ is the product of single-qubit gates, each controlled by a single qubit of the input register $|x\rangle$. We show in the full paper that the simplest implementation of U in which $\theta(x, b)$ is constant and $U(x) = e^{i\theta(x,b)} X^{f(x)}$ is impossible. The closest we can get to such a phaseless implementation is $U(x) = (iX)^{f(x)}$, which we call a *true* implementation. Any other case we regard as a *relative phase* implementation. Note that both true and relative phase implementations faithfully compute $f(x)$ upon measurement in the computational basis. An advantage of true implementation comes from the ability to remove the phase entirely through introducing a new ancilla qubit.

Our structured approach to computing symmetric Boolean functions $f(x)$ makes use of QSP [8]. Suppose that we only access the input bits with a unitary $S = \sum_x |x\rangle \langle x| \otimes R_x(\phi_x)$, where $R_P(\chi) = \cos(\chi/2)I - i \sin(\chi/2)P$ is a single-qubit rotation for any Pauli operator $P \in \{X, Y, Z\}$. Letting $\phi_x = \Delta|x| - \delta$ for real parameters δ and Δ, it is clear that we can implement S with n controlled-$R_x(\Delta)$ gates and an $R_x(\delta)$ gate. QSP is a method to create U using the S operation several (say, L) times, interspersed with single-qubit gates on the computational qubit. In the simplest case, sufficient for our purposes, these additional single-qubit gates are Z-rotations $R_z(\xi)$. To be more concrete, suppose we write $U(x) = A(x)I + iB(x)X + iC(x)Y + iD(x)Z$ for real-valued functions A, B, C, and D. Provided that these functions satisfy $A^2 + B^2 + C^2 + D^2 = 1$ and have certain symmetries, QSP guarantees the existence of angles ξ_j such that $U = R_z(\xi_0) \prod_{j=1}^{L} R_z(\xi_j) S R_Z^\dagger(\xi_j)$ and gives an efficient method to find these angles [17, 18]. See the full paper for more details.

To compute a symmetric function of n bits, we choose $A(x) = 1$ when $f(x) = 0$ and $B(x) = 1$ when $f(x) = 1$. These

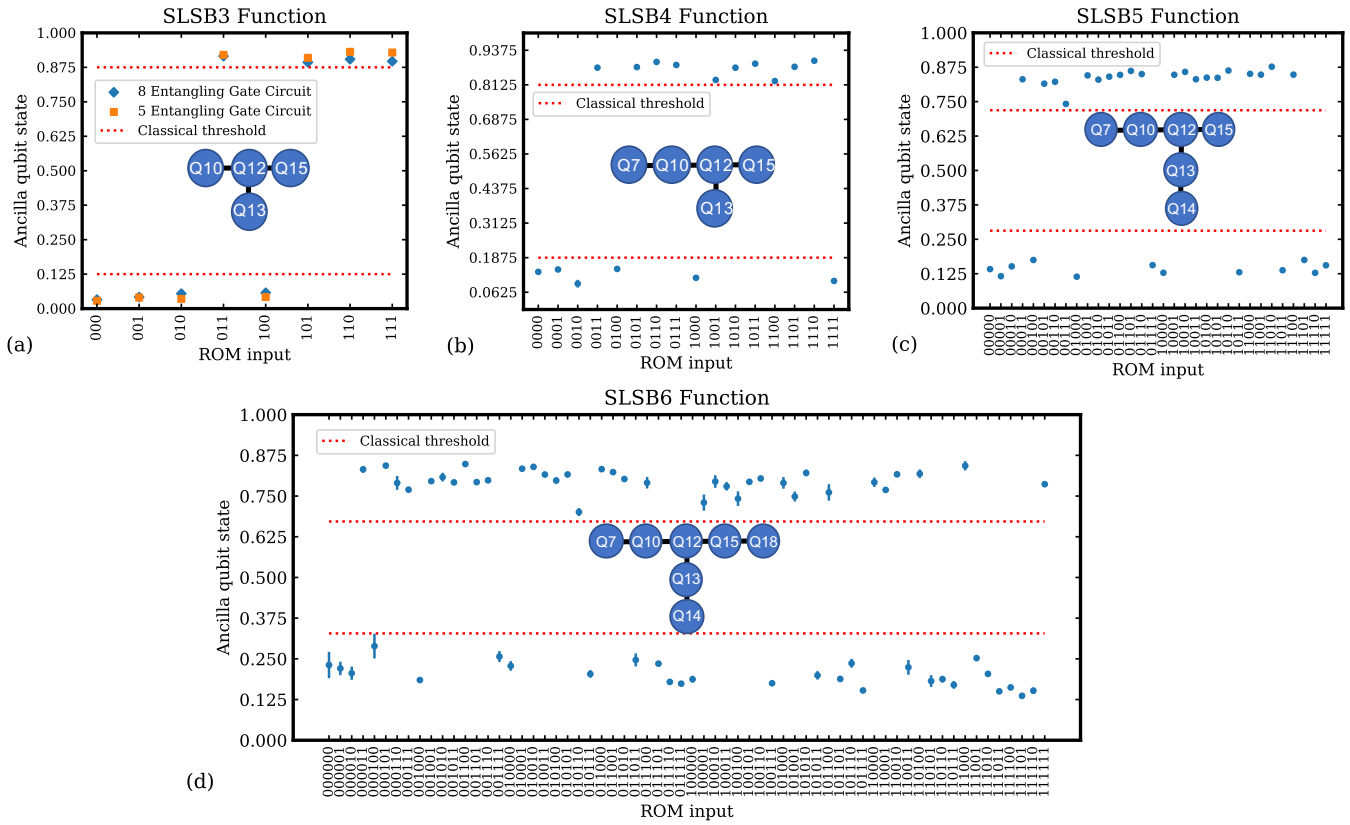

FIG. 3. Ancilla qubit (Q12) excited state population for each input ROM state combination for the (a) SLSB3, (b) SLSB4 (c) SLSB5 function and (d) SLSB6. The 5 and 8 entangling gate SLSB3 circuits achieve the ASP of 0.9429±0.0011 and 0.9280±0.0013, while the SLSB4 circuit achieves the ASP of 0.8743±0.0035, the SLSB5 circuit achieves the ASP of 0.8460±0.0053, and the SLSB6 circuit achieves the ASP of 0.7984±0.0047. Experimental error bars for each function are smaller than the plot marker in most cases. The maximal classical ASP is illustrated with the red dotted lines. In addition, the qubit layout for each function is displayed in the inset of each plot.

constraints are satisfiable for $L = 4n+1$ uses of S (see the full paper). Since each instance of S uses $n+1$ gates, the total gate-complexity of this approach is $O(n^2)$. For certain functions f, symmetries and circuit simplifications can reduce this gate count. For instance, the QSP approach calculates true SLSB3 = MAJ3 using 9 entangling gates, and a simple gate merging simplification reduces their number to 8, Fig. 1(a). MAJn, the majority function, evaluates to one iff more than half the inputs equal one. The true 5-bit majority MAJ5 implementation by QSP takes 25 entangling gates. For more general Boolean functions that lack the symmetry present in MAJn, gate counts are larger. For instance, an unoptimized QSP circuit for the true implementation of SLSB4 function, which operates over fewer bits than MAJ5, takes 52 entangling gates. In contrast, a relative-phase implementation constructed directly has only 7 entangling gates, Fig. 2.

Our experimental results are summarized in Fig. 3. We highlight that all experimental data points stay outside the classical band indicated by the red dotted lines; this means that quantum advantage was achieved not just on average, but for any input. More data is available in the full manuscript.

[1] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880, 1969.

[2] John S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.

[3] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[4] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997.

[5] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996.

[6] Hale F. Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.

[7] Masuo Suzuki. Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51(2):183–190, 1976.

[8] Guang Hao Low and Isaac L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118(1):010501, 2017.

[9] Yunseong Nam and Dmitri Maslov. Low-cost quantum circuits for classically intractable instances of the Hamiltonian dynamics simulation problem. *npj Quantum Information*, 5(1):1–8, 2019.

[10] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018.

[11] Sergey Bravyi, David Gosset, Robert König, and Marco Tomamichel. Quantum advantage with noisy shallow circuits. *Nature Physics*, pages 1–6, 2020.

[12] François Le Gall. Average-case quantum advantage with shallow circuits. In *Proceedings of the 34th Computational Complexity Conference*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.

[13] Matthew Coudron, Jalex Stark, and Thomas Vidick. Trading locality for time: certifiable randomness from low-depth circuits. *arXiv preprint arXiv:1810.04233*, 2018.

[14] Adam Bene Watts, Robin Kothari, Luke Schaeffer, and Avishay Tal. Exponential separation between shallow quantum circuits and unbounded fan-in shallow classical circuits. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 515–526, 2019.

[15] Daniel Grier and Luke Schaeffer. Interactive shallow Clifford circuits: quantum advantage against NC^1 and beyond. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 875–888, 2020.

[16] Farid Ablayev, Aida Gainutdinova, Marek Karpinski, Cristopher Moore, and Christopher Pollett. On the computational power of probabilistic and quantum branching program. *Information and Computation*, 203(2):145–162, 2005.

[17] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang. Methodology of resonant equiangular composite quantum gates. *Physical Review X*, 6(4):041067, 2016.

[18] Jeongwan Haah. Product decomposition of periodic functions in quantum signal processing. *Quantum*, 3:190, 2019.