Experimental comparison of two quantum computing architectures

Norbert M. Linke[a,b,1], Dmitri Maslov[c], Martin Roetteler[d], Shantanu Debnath[a,b], Caroline Figgatt[a,b], Kevin A. Landsman[a,b], Kenneth Wright[a,b], and Christopher Monroe[a,b,e,1]

[a]Joint Quantum Institute, Department of Physics, University of Maryland, College Park, MD 20742; [b]Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, MD 20742; [c]National Science Foundation, Arlington, VA 22230; [d]Microsoft Research, Redmond, WA 98052; and [e]IonQ Inc., College Park, MD 20742

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected in 2016.

Contributed by Christopher Monroe, February 1, 2017 (sent for review November 1, 2016; reviewed by Eric Hudson and Igor L. Markov)

We run a selection of algorithms on two state-of-the-art 5-qubit quantum computers that are based on different technology platforms. One is a publicly accessible superconducting transmon device (www.research.ibm.com/ibm-q) with limited connectivity, and the other is a fully connected trapped-ion system. Even though the two systems have different native quantum interactions, both can be programed in a way that is blind to the underlying hardware, thus allowing a comparison of identical quantum algorithms between different physical systems. We show that quantum algorithms and circuits that use more connectivity clearly benefit from a better-connected system of qubits. Although the quantum systems here are not yet large enough to eclipse classical computers, this experiment exposes critical factors of scaling quantum computers, such as qubit connectivity and gate expressivity. In addition, the results suggest that codesigning particular quantum applications with the hardware itself will be paramount in successfully using quantum computers in the future.

quantum computing | quantum information | quantum information science | quantum physics | quantum computing architecture

Inspired by the vast computing power a universal quantum computer could offer, several candidate systems are being explored. They have allowed experimental demonstrations of quantum gates, operations, and algorithms of ever-increasing sophistication. Recently, two architectures, superconducting transmon qubits (1–5) and trapped ions (6, 7), have reached a new level of maturity. They have become fully programmable multiqubit machines that provide the user with the flexibility to implement arbitrary quantum circuits from a high-level interface. This makes it possible for the first time to test quantum computers irrespective of their particular physical implementation.

Whereas the quantum computers considered here are still small scale and their capabilities do not currently reach beyond small demonstration algorithms, this line of inquiry can still provide useful insights into the performance of existing systems and the role of architecture in quantum computer design. These will be crucial for the realization of more advanced future incarnations of the present technologies.

The standard abstract model of quantum computation assumes that interactions between arbitrary pairs of qubits are available. However, physical architectures will in general have certain constraints on qubit connectivity, such as nearest-neighbor couplings only. These restrictions do not in principle limit the ability to perform arbitrary computations, because swap operations may be used to effect gates between arbitrary qubits using the connections available. For a general circuit, reducing a fully connected system to the more sparse star-shaped or linear nearest-neighbor connectivity requires an increase in the number of gates of $O(n)$, where n is the number of qubits (8). How much overhead is incurred in practice depends on the connections used in a particular circuit and how efficiently they can be matched to the physical qubit-to-qubit interaction graph.

In this article, we make use of the public access recently granted by IBM to a 5-qubit superconducting device (illustrated in Fig. 1*A*) via their "Quantum Experience" cloud service (www.research.ibm.com/ibm-q). This allows us to repeat algorithms that we perform in our own ion-trap experiment on an independent quantum computer of identical size and comparable capability but with a different physical implementation at its core.

Physical Systems

The ion-trap system consists of five $^{171}Yb^{+}$ ions that are confined in a linear Paul trap and laser cooled close to their motional ground state (Fig. 1*B*) (6). The qubits are magnetic-field–insensitive pairs of states in the hyperfine-split $^2S_{1/2}$ ground level of each atom, which gives a qubit frequency of 12.642821 GHz. All control and measurement are performed optically. State preparation and readout are accomplished by optical pumping and state-dependent fluorescence detection (9). Qubit operations are realized via pairs of Raman beams, derived from a single 355-nm mode-locked laser. These optical controllers consist of an array of individual addressing beams and a counterpropagating global beam that illuminates the entire chain (6). Single-qubit rotations are driven by a Raman beat note

Significance

Quantum computers are an emerging technology promising to be vastly more powerful at solving certain problems than any conventional computer. These devices are now moving out of the laboratory and becoming generally programmable. This allows identical quantum tasks or algorithms to be implemented on radically different technologies to inform further development and scaling. We run a series of algorithms on the two leading platforms: trapped atomic ions and superconducting circuits. Whereas the superconducting system offers faster gate clock speeds and a solid-state platform, the ion-trap system features superior qubits and reconfigurable connections. The performance of these systems is seen to reflect the topology of connections in the base hardware, supporting the idea that quantum computer applications and hardware should be codesigned.

Author contributions: N.M.L., D.M., M.R., S.D., C.F., and C.M. designed research; N.M.L., D.M., M.R., S.D., C.F., K.A.L., K.W., and C.M. performed research; N.M.L., D.M., M.R., S.D., C.F., K.A.L., K.W., and C.M. analyzed data; and N.M.L., D.M., M.R., S.D., C.F., K.A.L., K.W., and C.M. wrote the paper.

Reviewers: E.H., University of California, Los Angeles; and I.L.M., The University of Michigan.

Conflict of interest statement: C.M. is the cofounder and Chief Scientist of IonQ, Inc.

Freely available online through the PNAS open access option.

[1]To whom correspondence may be addressed. Email: monroe@umd.edu or linke@umd.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1618020114/-/DCSupplemental.

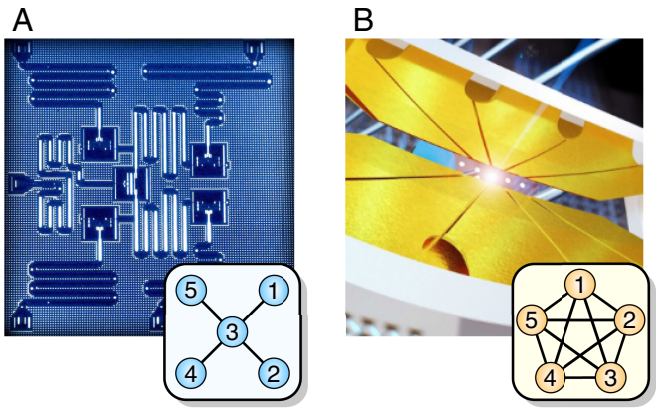

Fig. 1. Graphic representations of the two systems. (*A*) the superconducting qubits connected by microwave resonators (credit: IBM Research). (*B*) The linear chain of trapped ions connected by laser-mediated interactions. (*A* and *B*, *Insets*) Qubit connectivity graphs: (*A*) star shaped and (*B*) fully connected.

of defined amplitude, phase, and duration resonant with the qubit frequency. Two-qubit operations are produced by applying Raman beams to a pair of ions, with beat-note frequencies near the motional sidebands. This creates an effective XX–Ising interaction between the spins mediated by all modes of motion (10–12). We use a pulse-shaping scheme to ensure spin and motion are disentangled at the end of the operation (13, 14). Because all ions partake in the collective motion of the chain, gates between any pair can be invoked in this way (Fig. 1*B*, *Inset*). The addressing during operations and the distinction between qubits during readout are both achieved by spatially resolving the ions. The fidelities for single- and 2-qubit gates are typically $99.1(5)\%$ and $97(1)\%$, respectively. The single-qubit readout fidelity is $99.7(1)\%$ for state $|0\rangle$ and $99.1(1)\%$ for state $|1\rangle$. The latter is lower because off-resonant excitation during readout predominantly causes $|1\rangle \rightarrow |0\rangle$ pumping. The average readout fidelity for an entire 5-qubit state is $95.7(1)\%$. This is lower than one would expect from the average single-qubit readout fidelity, because there is crosstalk that leads to $|0\rangle \rightarrow |1\rangle$ errors on adjacent channels. Typical gate times are $20\,\mu$s for single-cubit and $250\,\mu$s for 2-qubit gates. Spin depolarization is negligible for hyperfine ground-level qubits ($T_1 \sim \infty$). The spin-dephasing time (T_2^*) is ~ 0.5 s in the current setup and can be easily extended by suppressing magnetic-field noise.

In analogy to atoms given by nature, the man-made superconducting circuits in the IBM quantum computer can be thought of as "artificial atoms" (16). They are transmon qubits (17) or superconducting islands connected by Josephson junctions and shunt capacitors that provide superpositions of charge states that are insensitive to charge fluctuations. The device used here has a range of qubit frequencies between 5 GHz and 5.4 GHz (www.research.ibm.com/ibm-q). The qubits are connected to each other and the classical control system by microwave resonators. State preparation (18) and readout, as well as single- (19) and 2-qubit gates (20), are achieved by applying tailored microwave signals to this network and measuring the response. Qubits are resolved in the frequency domain during addressing and readout. In the Quantum Experience hardware, the qubits are connected in a star-shaped pattern that provides four 2-qubit interactions (Fig. 1*A*, *Inset*), which are controlled-NOT (CNOT) gates targeting the central qubit. Single-qubit readout fidelities are typically $\sim 96\%$ (www.research.ibm.com/ibm-q), and the average readout fidelity for an arbitrary 5-qubit state is $\sim 80\%$ (www.research.ibm.com/ibm-q). Typical gate fidelities are 99.7% and 96.5% for single- and 2-qubit gates, respectively. Typical gate times are 130 ns for single-qubit and 250–450 ns for 2-qubit gates, whereas coherence times are $\sim 60\,\mu$s for both depolarization (T_1) and spin dephasing (T_2). The publicly accessible system runs autonomously, not requiring any human intervention over many weeks (www.research.ibm.com/ibm-q). This level of reliability may come at a cost due to drifts between periodic calibrations. Higher connectivity can in general be achieved by coupling three to four transmons to one resonator, limited by spectral resolution. The present layout could be modified to provide connections from qubit 1 to 5 and from qubit 2 to 4 (www.research.ibm.com/ibm-q). Furthermore, other superconducting architectures involving multimode resonators (3) can offer higher connectivity.

On these two machines, we compare a selection of composite gates and algorithms that represent a variety of circuit connectivities. In each case, we map the algorithms to the device by breaking them down into circuits made up of gates native to the specific hardware. We rely on an optimization protocol (21) to accomplish this task for the trapped ions and CNOT + T/Z^a algebra (22) with further manual optimization to compose the experiments for the IBM machine (23). The available gate set for the ion-trap system consists of the 2-qubit XX gate, as well as arbitrary single-qubit R_α^θ gate rotations by an angle θ about any axis (given by α) on the equator of the Bloch sphere. We call this the R/XX library. The IBM system makes available the family of gates [X, Y, Z, H, S, CNOT, and T (15)], known as the Clifford + T library. Because each gate is subject to errors, the circuits are optimized to minimize the number of operations used. The resulting gate numbers are optimal for 2-qubit gates and either optimal or close to optimal for single-qubit gates. The total number of single- and 2-qubit gates for each algorithm is shown in Table 1. The R/XX library offers a better overall expressive power. However, we note that the Clifford +T library was likely chosen for didactic reasons and is not native to superconducting systems, which do in principle offer continuous parameters for single- and 2-qubit gates.

In addition to the two systems considered here, Table 1 also gives the numbers for an LNN connectivity architecture as used, e.g., in superconducting qubits (1) as well as semiconductor gated quantum dots (24). The numbers in Table 1 show that the 2-qubit gate count strongly depends on the matching between the circuit and the qubit connectivity graph. The LNN architecture is as efficient as the fully connected system for the hidden shift

Table 1. Single- and 2-qubit gate counts for the circuits on the superconducting (star-shaped) and the ion-trap (fully connected) system after mapping to the respective hardware using the respective gate libraries

Connectivity	Star		LNN		Full	
Hardware	Superconductor		Superconductor		Ion Trap	
Gate type	1-qubit	2-qubit	1-qubit	2-qubit	1-qubit	2-qubit
Margolus	20	3	20	3	11	3
Toffoli	17	10	9	10	9	5
Bernstein–Vazirani	10	0–4	10	0–10	14–26	0–4
Hidden shift	28–34	10	20–26	4	42–50	4
QFT-3	42	19	11	7	8	3
QFT-5	*	*	35	28	22	10

For comparison, the gate counts for a linear nearest-neighbor (LNN) architecture as implemented in ref. 1 are included. We also note the gate count for the quantum Fourier transform (QFT) for 3 and 5 qubits. The latter was implemented in ref. 6, using a sequence of modular gates that was not optimized for gate count. The QFT-5 cannot be implemented exactly using the current IBM gate library.
*If we assume Z^a operations are possible, the counts shown are 47 for single-cubit and 29 for 2-qubit gates.

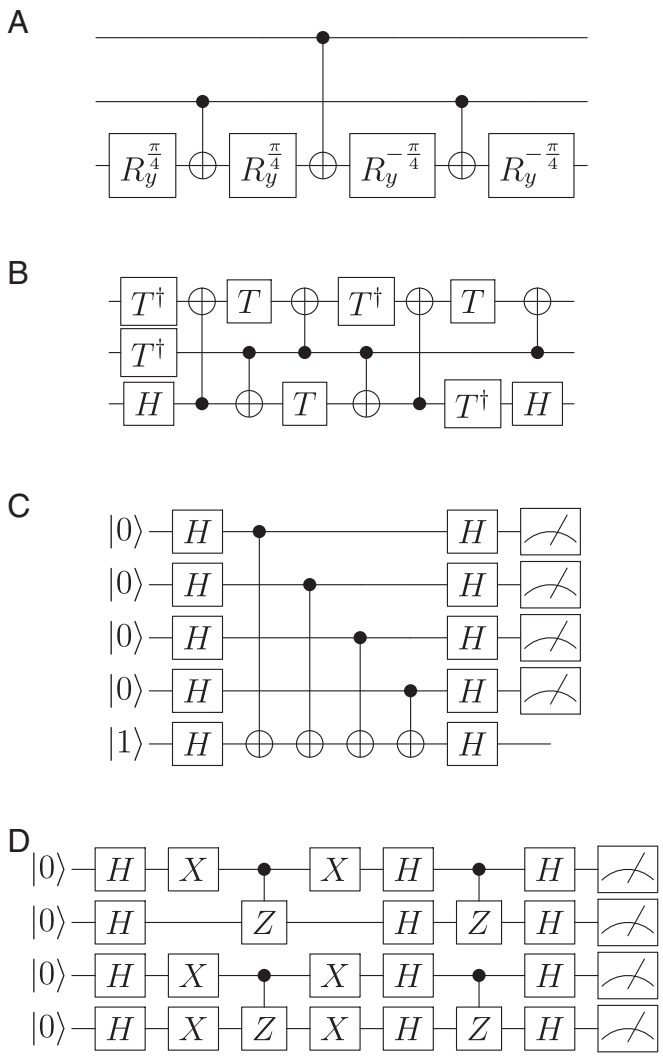

Fig. 2. High-level circuits of the implemented example computations (gates defined in ref. 15). (*A–D*) Margolus gate (*A*), Toffoli gate (*B*), Bernstein–Vazirani (*C*), and hidden shift (*D*). The Bernstein–Vazirani algorithm is shown for the oracle $c = (1111)$, where all CNOTs are present. The hidden shift circuit is shown for the shift pattern $s = (1011)$, where X operations are present on qubits 1, 3, and 4.

algorithm, whereas the star-shaped system incurs overheads; the reverse is true for the Bernstein–Vazirani algorithm (Fig. 2).

Algorithms

Margolus and Toffoli Gates. The Toffoli gate is a 3-qubit controlled-controlled–NOT gate that requires six CNOT gates (25, 26). It is possible to implement the Toffoli gate with five entangling gates if the square root of the CNOT operation is available (15), which is the case with the trapped-ion XX gate. The Margolus gate is a simplified version of the Toffoli operation, which introduces an additional phase on the state corresponding to $|100\rangle$. It can be realized with just 3 CNOT gates (27, 28). The circuits are shown in Fig. 2 *A* and *B*. Note that for the Margolus gate, all entangling operations connect to the same qubit, which means that this circuit can be realized efficiently with star-shaped qubit connectivity. The systems perform this circuit at success probability $74.1(7)\%$ for superconductors and $90.1(2)\%$ for ions (Fig. 3 *A1* and *B1*).

The full Toffoli circuit uses the same 3 qubits as the Margolus implementation so that preparation and measurement errors remain the same. The optimized circuit for the fully connected ion-trap system contains five 2-qubit gates and the additional operations lower the fidelity to $85.0(2)\%$ (Fig. 3*B2*). For the star-shaped system, an additional seven 2-qubit gates are needed to effect the swap operations necessary to go from the Margolus to the full Toffoli gate. This leads to a reduced success rate of $52.6(8)\%$ for the superconducting system (Fig. 3*A2*). Note that the transformation $|a, b, c\rangle \rightarrow |c \oplus ab, b, a\rangle$ may be obtained with the Clifford + T library on a star-shaped graph with the provably minimal number of seven CNOT gates. We do not consider such input-to-output mappings of the composite gates in this work. However, we always choose the most favorable input-to-output mapping for the IBM star and LNN architectures when executing entire quantum algorithms, which is merely a classical swap between physically measured signals.

Bernstein–Vazirani and Hidden Shift Algorithms. In the Bernstein–Vazirani algorithm, an oracle implements the function $f_c(x) = x \cdot c$. The algorithm finds the unknown bit string $\mathbf{c}$ in a single shot. In the oracle, $\mathbf{c}$ is encoded in a pattern of CNOT gates, all of which target the ancilla qubit (29). As can be seen from the circuit in Fig. 2*C*, the entire algorithm maps well onto a star-shaped architecture. This algorithm is very similar to a parity check circuit used in error correction applications, and indeed the IBM system was laid out with this application in mind (5). The single-shot success probabilities are $72.8(5)\%$ for the star-shaped superconducting system and $85.1(1)\%$ for the fully connected ion-trap system (Fig. 4 *A1* and *B1*).

To compare this to a similar algorithm with different connectivity requirements, we implement the hidden shift algorithm (30) for a black box bent function (31, 32). An oracle implements the shifted version $f(\mathbf{x}+\mathbf{s})$ of the known Boolean function f. We want to determine the n-bit string $\mathbf{s}$ that constitutes the "hidden shift." For a subset of Boolean functions, there exists a quantum algorithm that can solve this problem in a single oracle query, whereas classical algorithms require $\Omega(\sqrt{2^n})$ queries. This subset

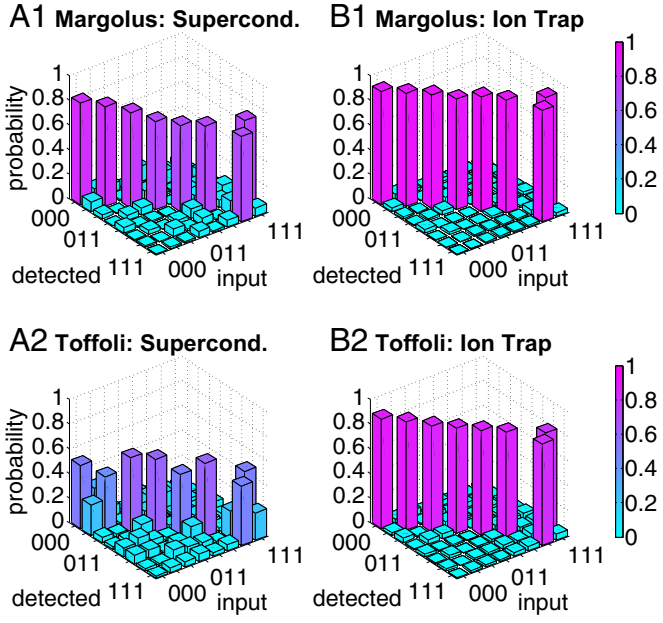

Fig. 3. Margolus gate results from the star-shaped superconductor (*A1*) and the fully connected ion trap system (*B1*). The fidelities are 74.1(7)% and 90.1(2)%, respectively. The full Toffoli gate results give success probabilities of 52.6(8)% for the superconducting (*A2*) and 85.0(2)% for the ion-trap (*B2*) system. The axes represent states as 3-bit binary numbers. For each input state, the probabilities of detecting each state are shown.

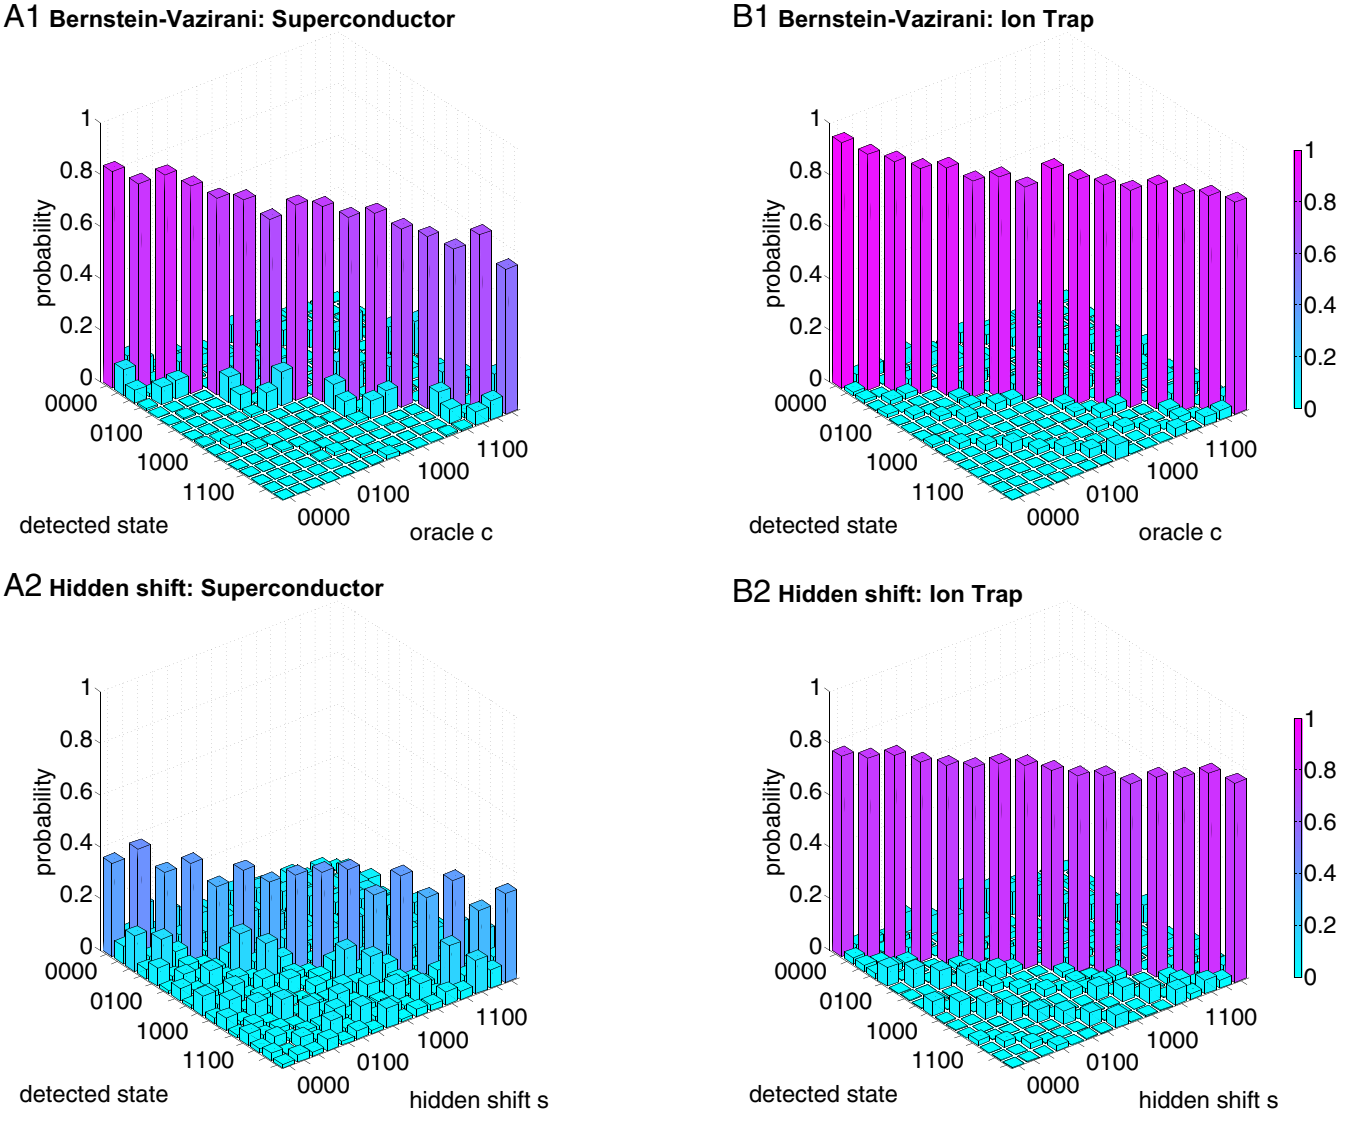

Fig. 4. Results from the Bernstein–Vazirani algorithm implementing the oracle function $f_c(x) = x_0c_0 \oplus c_1x_1 \oplus c_2x_2 \oplus c_3x_3$ for all possible 4-bit oracles **c** performed on the star-shaped (*A1*) and the fully connected (*B1*) systems. The average success probabilities are 72.8(5)% for the superconductor and 85.1(1)% for the ion-trap system. The hidden shift algorithm for $f(x) = x_0x_1 \oplus x_2x_3$. All possible 4-bit shifted oracle functions are implemented on the superconducting system (*A2*) as well as the ion trap (*B2*). The average success probabilities are 35.1(6)% and 77.1(2)%, respectively. The axes represent states and oracle parameters as 4-bit binary numbers.

contains functions that have a flat Fourier spectrum and whose dual $f^{\sim}$ can be calculated efficiently, i.e., so-called bent functions of the Maiorana–McFarland class (32). Here we choose the 4-bit function $f(\mathbf{x}) = x_1x_2 \oplus x_3x_4$ for which $f = f^{\sim}$. We implement all possible 4-bit shift patterns **s** using the circuit shown in Fig. 2*D*. The algorithm output state directly corresponds to the hidden shift **s**. The circuit involves gates between two disconnected pairs of qubits, which creates an overhead of six two-qubit gates for a star-shaped architecture. The results are shown in Fig. 4 *A2* and *B2*. The fidelity of the fully connected ion-trap implementation is 77.1(2)%, compared with 35.1(6)% for the superconducting device. The numerical values of the data plotted in Figs. 3 and 4 are displayed in Figs. S1–S3.

The errors in both devices appear concentrated in certain sets of states, leading to patterns in the off-diagonal elements of the result plots (Fig. 4). These highly structured signatures suggest that systematic errors dominate, especially readout errors. The grouped patterns such as in Fig. 4*A1* indicate flips of the least-significant bits, whereas parallel lines correspond to the most significant bits changing their state. In the trapped-ion results, these lines can be modulated in height due to readout crosstalk and are more pronounced on the lower-numbered state side due to $1 \to 0$ being the dominant detection error channel. Finally, we stress that comparing quantum computations across systems depends on the specifics of error propagation, which will vary between different hardware implementations, through their particular connectivity and physical errors. We summarize the success probabilities for the implemented circuits on both machines in Table 2. We also show the expected values for two simple error propagation models based on the errors of the individual gates ϵ_g and of M-qubit single-shot readout ϵ_M for both systems. The first model assumes random error propagation per operation with overall error $(1-\epsilon_M)^M(1-\epsilon_g)^{\sqrt{N}}$, where N is the number of gates. Because the errors for each step are independent and comparable to a random walk, the overall error involves

Table 2. Summary of the achieved success probabilities for the implemented circuits, in percentages

Connectivity	Star shaped			Fully connected		
Hardware	Superconducting			Ion trap		
Success probability/%	Obs	Rand	Sys	Obs	Rand	Sys
Margolus	74.1(7)	82	75	90.1(2)	91	81
Toffoli	52.6(8)	78	59	85.0(2)	89	78
Bernstein–Vazirani	72.8(5)	80	74	85.1(1)	90	77
Hidden shift	35.1(6)	75	52	77.1(2)	86	57

The observed probabilities ("Obs") are tabulated alongside two simple error propagation models given the gate number N and the individual gate and readout errors of the two systems encapsulated in the parameters ϵ_g and ϵ_M, respectively (main text). The first estimate assumes random ("Rand") error propagation with overall error $(1-\epsilon_g)^{\sqrt{N}}$ whereas the second one is based on systematic ("Sys") coherent over- or underrotations with overall error $(1-\epsilon_g)^N$, where N is the number of gates. The readout error for M qubits is $(1-\epsilon_M)^M$ in both cases.

$\sqrt{N}$ factors. The second model is based on systematic (coherent) over- or underrotations with overall error $(1-\epsilon_M)^M(1-\epsilon_g)^N$, which accumulates with N factors. We see that the numbers are broadly consistent, with systematic errors better predicting the superconducting system, whereas the ion-trap performance falls in between the two. The superconducting hidden shift algorithm is the only example with a significantly lower experimental result, perhaps from inhomogeneous errors in the device.

Outlook

Comparing quantum computing architectures involves many interrelated factors. Quantum gate operation fidelities, qubit numbers, primitive gate speeds, and coherence times are obviously important low-level metrics in a large-scale quantum computer. The results presented here show higher absolute fidelities and coherence times in the trapped-ion system, with higher clock speeds for the superconducting system. However, these metrics are moving targets: Whereas these systems are the most advanced and versatile quantum computing platforms built to date, both technologies are currently advancing rapidly.

In any case, such metrics should not be considered in isolation. Our comparison points to important higher-level considerations in scaling a quantum computer. The overall performance of a quantum circuit and the "time to solution" will depend critically on architectural restrictions, qubit connectivity, gate reconfigurability, and gate expressivity, and these attributes will become ever more important as the system is scaled up. Even with 5-qubit systems, we find that the qubit connectivity graph is best codesigned to mirror the structure of the particular quantum circuit and that the choice of a more expressive gate library affects the efficiency of the computations.

The physical scaling of each of these leading technologies has many challenges, and how they will be connected and reconfigured at large scales is an open question. One of the biggest challenges is the management of the control complexity in larger systems and potential crosstalk from overlapping qubit interactions or control buses. In most superconducting designs, there are many current-carrying wires necessary for control and biasing the individual qubits, and this may be difficult to route through a large superconducting chip (1–5). It will likely become a great challenge to manage the dilution refrigerator heat budget with such circuitry. Alternative modular superconducting architectures improve connectivity by integrating qubits with microwave cavity modes, at the expense of significant added volume per qubit (33). Ion-trap designs will hinge upon the stable and accurate delivery of laser beams (or near-field microwave sources) to address each qubit individually in a vacuum chamber. The fully connected nature of the ion-trap architecture may not scale to arbitrarily large numbers of qubits, owing to the spectral overlap of collective normal modes of motion. However, full connectivity between 20 and 100 trapped-ion qubits appears possible (6) and a modular approach for scaling to much larger systems with high connectivity and distance-independent operations seems promising (34, 35). In any hardware, an automated calibration procedure and powerful user interface will likely provide a higher level of integration. Such system-level attributes will become even more important as quantum circuits grow in complexity, regardless of physical platform.

ACKNOWLEDGMENTS. We thank D. L. Moehring, J. Kim, and K. R. Brown for key discussions; J. Gambetta and J. Chow at IBM for their assistance in interfacing with the IBM Quantum Experience project; and E. Edwards for the ion trap image. This work was supported by the Army Research Office with funds from the Intelligence Advanced Research Projects Activity (IARPA) LogiQ program, the Air Force Office of Scientific Research (AFOSR) Multidisciplinary University Research Initiative (MURI) program on Optimal Quantum Circuits, and the National Science Foundation (NSF) Physics Frontier Center at JQI. D.M. acknowledges support by the NSF. Any opinion, finding, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF, IBM, or any of their employees.

1. Barends R, et al. (2014) Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature* 508:500–503.
2. Córcoles AD, et al. (2015) Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nat Commun* 6:7979.
3. Ristè D, et al. (2015) Detecting bit-flip errors in a logical qubit using stabilizer measurements. *Nat Commun* 6:6983.
4. Ofek N, et al. (2016) Extending the lifetime of a quantum bit with error correction in superconducting circuits. *Nature* 536:441–445.
5. Takita M, et al. (2016) Demonstration of weight-four parity measurements in the surface code architecture. arXiv:1605.01351v2.
6. Debnath S, et al. (2016) Demonstration of a small programmable quantum computer module using atomic qubits. *Nature* 536:63–66.
7. Monz T, et al. (2016) Realization of a scalable Shor algorithm. *Science* 351(6277):1068–1070.
8. Cheung D, Maslov D, Severini S (2007) Translation techniques between quantum circuit architectures. *Workshop on Quantum Information Processing*. Available at www.umiacs.umd.edu/∼dmaslov/papers/qip07.pdf. Accessed July 6, 2016.
9. Olmschenk S, et al. (2007) Manipulation and detection of a trapped Yb^+ hyperfine qubit. *Phys Rev A* 76:052314.
10. Mølmer K, Sørensen A (1999) Multiparticle entanglement of hot trapped ions. *Phys Rev Lett* 82:1835–1838.
11. Solano E, de Matos Filho RL, Zagury N (1999) Deterministic bell states and measurement of the motional state of two trapped ions. *Phys Rev A* 59:R2539–R2543.
12. Milburn G, Schneider S, James D (2000) Ion trap quantum computing with warm ions. *Fortschr Phys* 48(9-11):801–810.
13. Zhu SL, Monroe C, Duan LM (2006) Trapped ion quantum computation with transverse phonon modes. *Phys Rev Lett* 97:050505.
14. Choi T, et al. (2014) Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement. *Phys Rev Lett* 112:190502.
15. Nielsen MA, Chuang IL (2011) *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge Univ Press, New York), 10th Ed.
16. Devoret MH, Schoelkopf RJ (2013) Superconducting circuits for quantum information: An outlook. *Science* 339(6124):1169–1174.
17. Koch J, et al. (2007) Charge-insensitive qubit design derived from the cooper pair box. *Phys Rev A* 76:042319.
18. Geerlings K, et al. (2013) Demonstrating a driven reset protocol for a superconducting qubit. *Phys Rev Lett* 110:120501.
19. Chow JM, et al. (2010) Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Phys Rev A* 82:040305.
20. Chow JM, et al. (2011) Simple all-microwave entangling gate for fixed-frequency superconducting qubits. *Phys Rev Lett* 107:080502.
21. Maslov D (2017) Basic circuit compilation techniques for an ion-trap quantum machine. *New J Phys* 19:023035.
22. Amy M, Maslov D, Mosca M, Roetteler M (2013) A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Trans Comput Aided Des Integrated Circ Syst* 32:818–830.
23. Wecker D, Svore, KM (2014) LIQUi|>: A software design architecture and domain-specific language for quantum computing. arXiv:1402.4467.
24. Zajac DM, Hazard TM, Mi X, Nielsen E, Petta JR (2016) Scalable gate architecture for a one-dimensional array of semiconductor spin qubits. *Phys Rev Applied* 6:054013.
25. Barenco A, et al. (1995) Elementary gates for quantum computation. *Phys Rev A* 52:3457–3467.

26. Shende VV, Markov IL (2009) On the CNOT-cost of Toffoli gates. *Quant Inform Comput* 9(5):461–486.
27. DiVincenzo DP (1998) Quantum gates and circuits. *Proc R Soc Lond Ser A* 454(1969):261–276.
28. Song G, Klappenecker A (2004) Optimal realizations of simplified Toffoli gates. *Quant Inform Comput* 4(5):361–372.
29. Bernstein E, Vazirani U (1997) Quantum complexity theory. *SIAM J Comput* 26:1411–1473.
30. van Dam W, Hallgreen S, Lawrence I (2006) Quantum algorithms for some hidden shift problems. *SIAM J Comput* 36(3):763–778.
31. Childs A, Kothari R, Ozols M, Roetteler M (2013) Easy and hard functions for the Boolean hidden shift problem. *Proceedings TQC 2013, Leibniz International Proceedings in Informatics (LIPIcs)*, eds Severini S, Brand FGSL (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Guelph, Canada), Vol 22, pp 50–79.
32. Roetteler M (2010) Quantum algorithms for highly non-linear Boolean functions. *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'10)*, ed Charikar M (Siam, Austin, TX), pp 448–457.
33. Brecht T, et al. (2016) Multilayer microwave integrated quantum circuits for scalable quantum computing. *Nature J Quant Inf* 2:16002.
34. Monroe C, et al. (2014) Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Phys Rev A* 89: 022317.
35. Kielpinski D, Monroe C, Wineland DJ (2002) Architecture for a large-scale ion-trap quantum computer. *Nature* 417:709–711.