QUANTUM SIMULATION

Quantum optimization of maximum independent set using Rydberg atom arrays

S. Ebadi[1]†, A. Keesling[1,2]†, M. Cain[1]†, T. T. Wang[1], H. Levine[1]‡, D. Bluvstein[1], G. Semeghini[1], A. Omran[1,2], J.-G. Liu[1,2], R. Samajdar[1], X.-Z. Luo[2,3,4], B. Nash[5], X. Gao[1], B. Barak[5], E. Farhi[6,7], S. Sachdev[1,8], N. Gemelke[2], L. Zhou[1,9], S. Choi[7], H. Pichler[10,11], S.-T. Wang[2], M. Greiner[1]*, V. Vuletić[12]*, M. D. Lukin[1]*

Realizing quantum speedup for practically relevant, computationally hard problems is a central challenge in quantum information science. Using Rydberg atom arrays with up to 289 qubits in two spatial dimensions, we experimentally investigate quantum algorithms for solving the maximum independent set problem. We use a hardware-efficient encoding associated with Rydberg blockade, realize closed-loop optimization to test several variational algorithms, and subsequently apply them to systematically explore a class of graphs with programmable connectivity. We find that the problem hardness is controlled by the solution degeneracy and number of local minima, and we experimentally benchmark the quantum algorithm's performance against classical simulated annealing. On the hardest graphs, we observe a superlinear quantum speedup in finding exact solutions in the deep circuit regime and analyze its origins.

Combinatorial optimization is ubiquitous in many areas of science and technology. Many such problems have been shown to be computationally hard and form the basis for understanding complexity classes in modern computer science (*1*). The use of quantum machines to accelerate solving such problems has been theoretically explored for over two decades with a variety of quantum algorithms (*2*–*4*). Typically, a relevant cost function is encoded in a quantum Hamiltonian (*5*), and its low-energy state is sought starting from a generic initial state, either through an adiabatic evolution (*2*) or a variational approach (*3*), via closed optimization loops (*6*, *7*). The computational performance of such algorithms has been investigated theoretically (*4*, *8*–*13*) and experimentally (*14*–*16*) in small quantum systems with shallow quantum circuits, or in systems lacking the many-body coherence believed to be central for quantum advantage (*17*, *18*). However, these studies offer only limited insights into algorithms' performances in the most interesting regime involving large system sizes and high circuit depths (*19*, *20*).

Here we use a quantum device based on coherent, programmable arrays of neutral atoms trapped in optical tweezers to investigate quantum optimization algorithms for systems ranging from 39 to 289 qubits, and effective depths sufficient for the quantum correlations to spread across the entire graph. Specifically, we focus on maximum independent set, a paradigmatic NP-hard optimization problem (*21*). It involves finding the largest independent set of a graph—a subset of vertices such that no edges connect any pair in the set. An important class of such maximum independent set problems involves unit disk graphs, which are defined by vertices on a two-dimensional plane with edges connecting all pairs of vertices within a unit distance of one another (Fig. 1, A and B). Such instances arise naturally in problems associated with geometric constraints that are important for many practical applications, such as modeling wireless communication networks (*22*, *23*). Although there exist polynomial-time classical algorithms to find approximate solutions to the maximum independent set problem on such graphs (*24*), solving the problem exactly is known to be NP-hard in the worst case (*23*, *25*).

Maximum independent set on Rydberg atom arrays

Our approach uses a two-dimensional atom array described previously (*26*). Excitation from a ground state $|0\rangle$ into a Rydberg state $|1\rangle$ is utilized for hardware-efficient encoding of the unit disk maximum independent set problem (*27*). For a particular graph, we create a geometric configuration of atoms using optical tweezers such that each atom represents a vertex. The edges are drawn according to the unit disk criterion for a unit distance given by the Rydberg blockade radius R_b (Fig. 1C), the distance within which excitation of more than one atom to the Rydberg state is prohibited because of strong interactions (*28*). The Rydberg blockade mechanism thus restricts the evolution primarily to the subspace spanned by the states that obey the independent set constraint of the problem graph. Quantum algorithms for optimization are implemented via global atomic excitation using homogeneous laser pulses with a time-varying Rabi frequency (and a time-varying phase) $\Omega(t)e^{i\phi(t)}$ and detuning $\Delta(t)$ (Fig. 1D). The resulting quantum dynamics is governed by the Hamiltonian $H = H_q + H_{cost}$, with the quantum driver H_q and the cost function H_{cost} given by

$$H_q = \frac{h}{2}\sum_i \left[\Omega(t)e^{i\phi(t)}|0\rangle_i\langle 1| + \text{h.c.}\right],$$
$$H_{cost} = -h\Delta(t)\sum_i n_i + \sum_{i<j} V_{ij} n_i n_j \qquad (1)$$

where $n_i = |1\rangle_i\langle 1|$, and $V_{ij} = V_0/(|r_i - r_j|)^6$ is the interaction potential that sets the blockade radius R_b and determines the connectivity of the graph. For a positive laser detuning Δ, the many-body ground state of the cost function Hamiltonian maximizes the total number of qubits in the Rydberg state under the blockade constraint, corresponding to the largest independent set MIS(G) (hereafter MIS) of the underlying unit disk graph G (*27*) (Fig. 1E). Even with the finite blockade energy and long-range interaction tails, we empirically find that the ground states of H_{cost} still encode an MIS for the ensemble of graphs studied here [see (*25*, *27*)].

Variational optimization via a closed quantum-classical loop

In the experiment, we deterministically prepare graphs with vertices occupying 80% of an underlying square lattice, with the blockade extending across nearest and next-nearest (diagonal) neighbors (Fig. 1C). This allows us to explore a class of nonplanar graphs for which finding the exact solution of MIS is NP-hard for worst-case instances (*25*). To prepare quantum states with a large overlap with the MIS solution space, we use a family of variational quantum optimization algorithms using a quantum-classical optimization loop. We place atoms at positions defined by the vertices of the chosen graph, initialize them in state $|0\rangle$, and implement a coherent quantum evolution corresponding to the specific choice of variational parameters (Fig. 1D). Subsequently, we sample the wave function with a projective measurement and determine the

[1]Department of Physics, Harvard University, Cambridge, MA 02138, USA. [2]QuEra Computing Inc., Boston, MA 02135, USA. [3]Department of Physics and Astronomy, University of Waterloo, Waterloo N2L 3G1, Canada. [4]Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada. [5]School of Engineering and Applied Science, Harvard University, Cambridge, MA 02138, USA. [6]Google Quantum AI, Venice, CA 90291, USA. [7]Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. [8]School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA. [9]Walter Burke Institute for Theoretical Physics, California Institute of Technology, Pasadena, CA 91125, USA. [10]Institute for Theoretical Physics, University of Innsbruck, A-6020 Innsbruck, Austria. [11]Institute for Quantum Optics and Quantum Information, Austrian Academy of Sciences, A-6020 Innsbruck, Austria. [12]Department of Physics and Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
*Corresponding author. Email: greiner@physics.harvard.edu (M.G.); vuletic@mit.edu (V.V.); lukin@physics.harvard.edu (M.D.L.)
†These authors contributed equally to this work. ‡Present address: AWS Center for Quantum Computing, Pasadena, CA 91125, USA.

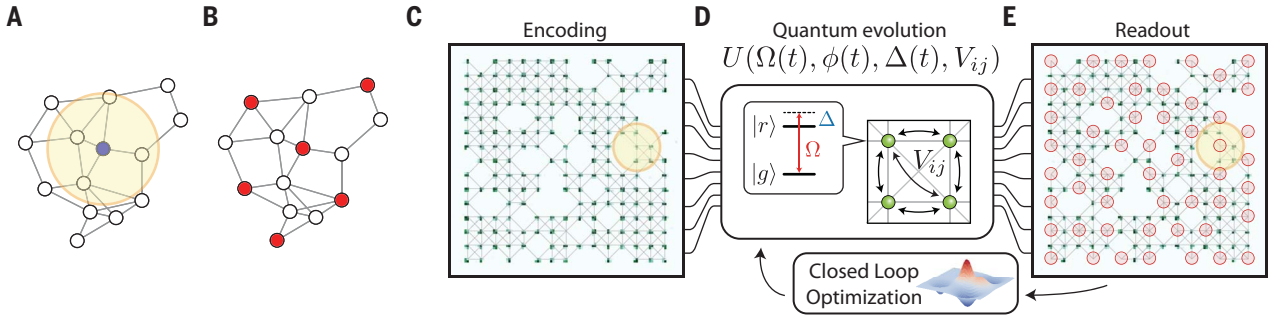

Fig. 1. Hardware-efficient encoding of the maximum independent set using Rydberg atom arrays. (**A**) An example of a unit disk graph, with any single vertex (e.g., the blue vertex) being connected to all other vertices within a disk of unit radius. (**B**) A corresponding MIS solution (denoted by the red nodes). (**C**) The maximum independent set problem is encoded with atoms placed at the vertices of the target graph and with interatomic spacing chosen such that the unit disk radius of the graph corresponds to the Rydberg blockade radius. Shown is an example fluorescence image of atoms, with gray lines added to indicate edges between connected vertices. (**D**) The system undergoes coherent quantum many-body evolution under a programmable laser drive [$\Omega(t)$, $\phi(t)$, $\Delta(t)$] and long-range Rydberg interactions V_{ij}. (**E**) A site-resolved projective measurement reads out the final quantum many-body state, with atoms excited to the Rydberg state (red circles) corresponding to vertices forming an independent set. A classical optimizer uses the results to update the parameters of the quantum evolution [$\Omega(t)$, $\phi(t)$, $\Delta(t)$] to maximize a figure of merit for finding an MIS.

size of the output independent set by counting the number of qubits in $|1\rangle$, using classical post-processing to remove blockade violations and reduce detection errors (*25*) (Fig. 1E). This procedure is repeated multiple times to estimate the mean independent set size $\langle\sum_i n_i\rangle$ of the sampled wave function, the approximation ratio $R \equiv \langle\sum_i n_i\rangle/|\mathrm{MIS}|$, and the probability P_{MIS} of observing an MIS (where |MIS| denotes the size of an MIS of the graph). The classical optimizer tries to maximize $\langle\sum_i n_i\rangle$ by updating the variational parameters in a closed-loop hybrid quantum-classical optimization protocol (*25*) (Fig. 1D).

We test two algorithm classes, defined by different parametrizations of the quantum driver and the cost function in Eq. 1. The first approach consists of resonant ($\Delta = 0$) laser pulses of varying durations τ_i and phases ϕ_i (Fig. 2A). This algorithm closely resembles the canonical quantum approximate optimization algorithm (QAOA) (*3*), but instead of exact single-qubit rotations, resonant driving generates an effective many-body evolution within the subspace of independent sets associated with the blockade constraint (*25*). Phase jumps between consecutive pulses implement a global phase gate (*29*), with a phase shift proportional to the cost function of the maximum independent set problem in the subspace of independent sets (see eq. S2). Taken together, these implement the QAOA, where each pulse duration τ_i and phase ϕ_i are used as a variational parameters.

The performance of QAOA as a function of depth p (the number of pulses) is shown in Fig. 2B for an instance of a 179-vertex graph embedded in a 15 × 15 lattice. We find that the approximation ratio grows as a function of the number of pulses up to $p = 4$, and increasing the depth further does not appear to lead to better performance (Fig. 2B). As discussed in (*25*), we attribute these performance limitations to the difficulty of finding the optimal QAOA parameters for large depths within a limited number of queries to the experiment, leakage out of the independent set subspace during resonant excitation due to imperfect blockade associated with the finite interaction energy between next-nearest neighbors, and laser pulse imperfections.

The second approach is a variational quantum adiabatic algorithm (VQAA) (*2*, *30*), related to methods previously used to prepare quantum many-body ground states (*26*, *31*, *32*). In this approach, we sweep the detuning Δ from an initial negative detuning Δ_0 to a final large positive value Δ_f at constant Rabi frequency Ω, along a piecewise-linear schedule characterized by a total number of segments f, the duration τ_i of each, and the end detuning Δ_i of each segment. Moreover, we turn on the coupling Ω in duration τ_Ω and smoothen the detuning sweep using a low-pass filter with a characteristic filter time τ_Δ (Fig. 2C), both of which minimize nonadiabatic excitations and serve as additional variational parameters. For this evolution, we define an effective circuit depth $\tilde{p}$ as the duration of the sweep ($T = \tau_1 + \ldots + \tau_f$) in units of the π-pulse time τ_π, which is the time required to perform a spin flip operation.

We find that with only three segments optimized for an effective depth of $\tilde{p} = 10$ (Fig. 2D inset), the optimizer converges to a pulse that substantially outperforms the QAOA approach described above. Furthermore, the optimized pulse shows a better performance compared to a linear (one-segment) detuning sweep of the same $\tilde{p}$ (Fig. 2D). We find that similar pulse shapes produce high approximation ratios for a variety of graphs (see, e.g., fig. S8C), consistent with theoretical predictions of pulse shape concentration (*20*, *25*, *33*, *34*). At large sweep times ($\tilde{p} > 15$), we observe a turn-around in the performance likely associated with decoherence (*25*). For the remainder of this work, we focus on the quantum adiabatic algorithm for solving maximum independent set.

Quantum optimization on different graphs

The experimentally optimized quasi-adiabatic sweep (depicted in Fig. 2D) was applied to 115 randomly generated graphs of various sizes ($N = 80$ to 289 vertices). For graphs of the same size ($N = 180$), the approximation error $1 - R$ decreases and the probability of finding an MIS solution P_{MIS} increases with the effective circuit depth at early times, with the former showing a scaling consistent with a power-law relation for short effective depths (Fig. 3A and fig. S15) (*25*). We find a strong correlation between the performance of the quantum algorithm on a given graph and its total number of MIS solutions, which we refer to as the MIS degeneracy $D_{|\mathrm{MIS}|}(G)$ (hereafter $D_{|\mathrm{MIS}|}$). This quantity is calculated classically using a tensor network algorithm (*25*, *35*) and varies by nine orders of magnitude across different 180-vertex graphs. We observe a clear logarithmic relation between $D_{|\mathrm{MIS}|}$ and the approximation error $1 - R$, accompanied by a nearly three-orders-of-magnitude variation of P_{MIS} at a fixed depth $\tilde{p} = 20$ (Fig. 2B). P_{MIS} does not scale linearly with the MIS degeneracy, as would be the case for a naive algorithm that samples solutions at random. Figure 2C shows the sharp collapse of $1 - R$ as a function of the logarithm of the MIS degeneracy normalized by the graph size, $\rho \equiv \log(D_{|\mathrm{MIS}|})/N$. This quantity, a measure of MIS degeneracy density, determines the hardness in approximating solutions for the quantum algorithm at shallow depths.

These observations can be modeled as resulting from a Kibble-Zurek–type mechanism

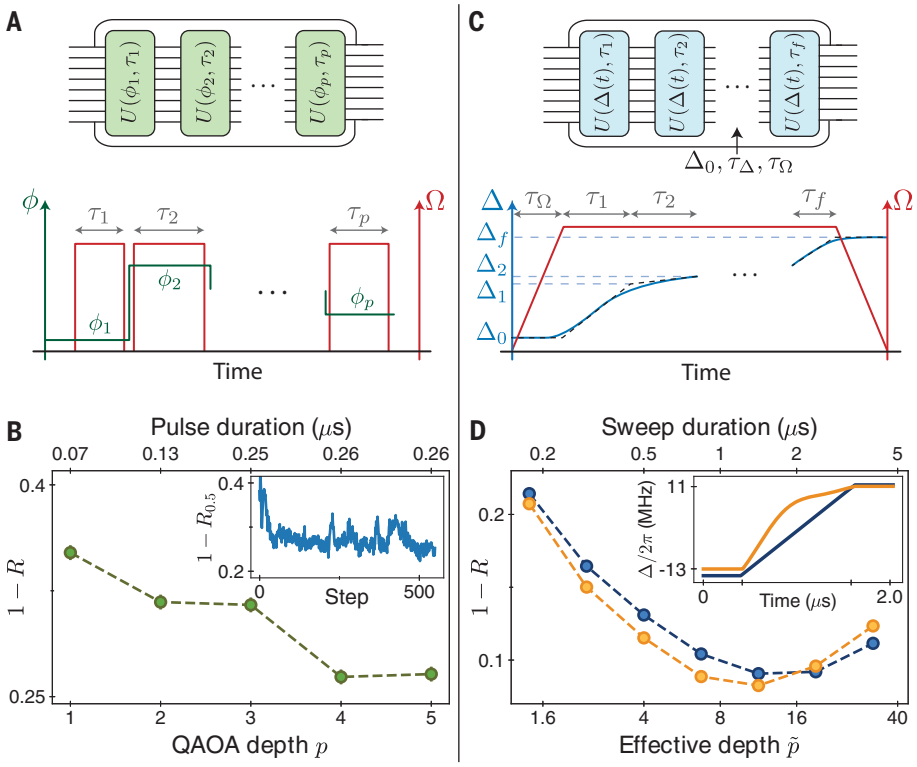

Fig. 2. Testing variational quantum algorithms. (**A**) Implementation of the quantum approximate optimization algorithm (QAOA), consisting of sequential layers of resonant pulses with variable duration τ_i and laser phase ϕ_i. (**B**) Variational optimization of QAOA parameters results in a decrease in approximation error $1 - R$, up to depth $p = 4$ (inset: example performance of quantum-classical closed-loop optimization at $p = 5$). Approximation error calculated using the top 50 percentiles of independent set sizes ($1 - R_{0.5}$) is used as the figure of merit to reduce effects of experimental imperfections on the optimization procedure (*25*). (**C**) Quantum evolution can also be parametrized as a variational quantum adiabatic algorithm (VQAA) using a quasi-adiabatic pulse with a piecewise-linear sweep of detuning $\Delta(t)$ at constant Rabi coupling $\Omega(t)$. $\Omega(t)$ is turned on and off within τ_Ω, and a low-pass filter with time scale τ_Δ is used to smoothen the $\Delta(t)$ sweep. (**D**) Performance of a rescaled piecewise-linear sweep as a function of its effective depth $\tilde{p} = (\tau_1 + \ldots + \tau_f)/\tau_\pi$. Variational optimization of a three-segment (orange) piecewise-linear pulse (optimized for $\tilde{p} = 10$) improves on the performance of a simple one-segment linear (blue) pulse, as well as the best results from QAOA (inset: detuning sweep profiles for one-segment (blue) and three-segment (orange) optimized pulses, for a total pulse duration of 2.0 μs). Error bars for approximation ratio R are the SEM here and throughout the text and are smaller than the points.

where the quantum algorithm locally solves the graph in domains whose sizes are determined by the evolution time and speed at which quantum information propagates (*36*, *37*). We show that the scaling of the approximation error with depth can originate from the conflicts between local solutions at the boundaries of these independent domains (*25*). In graphs with a large degeneracy density ρ, there may exist many MIS configurations that are compatible with the local ordering in these domains. This provides a possible mechanism to reduce domain walls at their boundaries (fig. S14) and decrease the approximation error. Such a scenario would predict a linear relation between $1 - R$ and ρ at a fixed depth, which is consistent with our observations (Fig. 2C and fig. S15).

Benchmarking against simulated annealing

To benchmark the results of the quantum optimization against a classical algorithm, we use simulated annealing (SA) (*38*). It seeks to minimize the energy of a cost Hamiltonian by thermally cooling a system of classical spins while maintaining thermal equilibrium. Although some specifically tailored state-of-the-art algorithms (*24*, *39*) may have better performance than SA in solving the maximum independent set problem, we have chosen SA for extensive benchmarking because similar to the quantum algorithms used, it is a general-purpose algorithm that only relies on information from the cost Hamiltonian for solving the problem. Our highly optimized variant of SA stochastically updates local clusters of spins using the Metropolis-Hastings (*40*) update rule, rejecting energetically unfavorable updates with a probability dependent on the energy cost and the instantaneous temperature (*25*). We use collective updates under the MIS Hamiltonian cost function (eq. S15), which applies an optimized uniform interaction energy to each edge, penalizing states that violate the independent set criterion (*25*). The annealing depth p_{SA} is defined as the average number of attempted updates per spin.

We compare the quantum algorithm and SA on two metrics: the approximation error $1 - R$, and the probability of sampling an exact solution P_{MIS}, which determines the inverse of time-to-solution. As shown in Fig. 4A, for relatively shallow depths and moderately hard graphs, optimized SA results in approximation errors similar to those observed on the quantum device. In particular, we find that the hardness in approximating the solution for short SA depths is also controlled by degeneracy density ρ (fig. S18, A and B). However, some graph instances appear to be considerably harder for SA compared to the quantum algorithm at higher depths (see, e.g., gold and purple curves in Fig. 4A).

Detailed analysis of the SA dynamics for graphs with low degeneracy densities ρ reveals that for some instances, the approximation ratio displays a plateau at $R = (|\text{MIS}| - 1)/|\text{MIS}|$, corresponding to independent sets with one less vertex than an MIS (Fig. 4A, gold and purple solid lines). Graphs displaying this behavior have a large number of local minima with independent set size $|\text{MIS}| - 1$, in which SA can be trapped up to large depths. By analyzing the dynamics of SA at low temperatures as a random walk among $|\text{MIS}| - 1$ and $|\text{MIS}|$ configurations (Fig. 4D), we show in (*25*) that the ability of SA to find a global optimum is limited by the ratio of the number of suboptimal independent sets of size $|\text{MIS}| - 1$ to the number of ways to reach global minima, resulting in a "hardness parameter" $\mathcal{HP} = D_{|\text{MIS}|-1}/(|\text{MIS}|D_{|\text{MIS}|})$ (Fig. 4E). This parameter lower bounds the mixing time for the Markov chain describing the SA dynamics at low temperatures (eq. S19), and it appears to increase exponentially with the square root of the system size for the hardest graphs (fig. S11). This suggests that a large number of local minima cause SA to take an exponentially long time to find an MIS for the hardest cases as N grows. If SA performance saturates this lower bound, consistent with numerics (fig. S19), its runtime to find an MIS is polynomially related to the best known exact classical algorithms (*41*).

Quantum speedup on the hardest graphs

We now turn to study the algorithms' ability to find exact solutions on the hardest graphs (with up to $N = 80$), chosen from graphs in the top two percentile of the hardness parameter $\mathcal{HP}$ (fig. S11). We find that for some

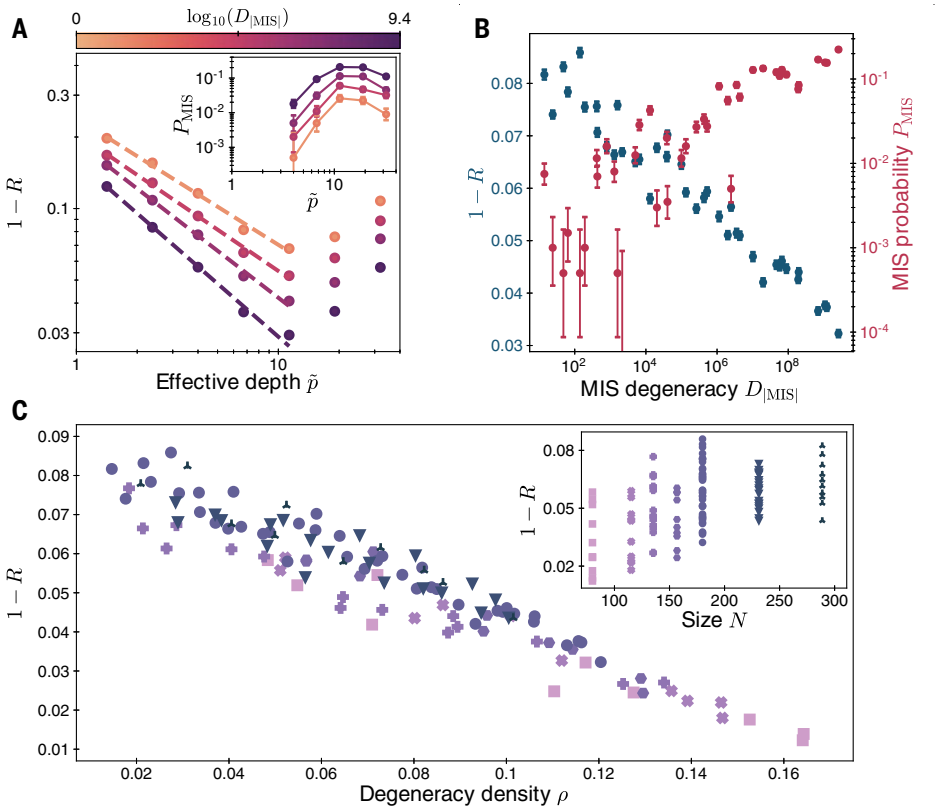

Fig. 3. Quantum algorithm performance across different graphs. (**A**) The approximation error $1 - R$ for an optimized quasi-adiabatic sweep plotted as a function of effective depth $\tilde{p}$ on four graphs of the same size (N = 180 vertices), showing strong dependence on the number of MIS solutions (MIS degeneracy) $D_{|\mathrm{MIS}|}$ (inset: corresponding MIS probability P_{MIS} versus $\tilde{p}$). (**B**) At a fixed depth $\tilde{p}$ = 20, $1 - R$ and P_{MIS} for various 180-vertex graphs are strongly correlated with $D_{|\mathrm{MIS}|}$. (**C**) At the same effective depth $\tilde{p}$ = 20, $1 - R$ for 115 graphs of different sizes (N = 80 to 289) and MIS degeneracies $D_{|\mathrm{MIS}|}$ exhibit universal scaling with the degeneracy density $\rho \equiv \log(D_{|\mathrm{MIS}|})/N$ (inset: data plotted as a function of N). Error bars for P_{MIS}, here and throughout the text, denote the 68% confidence interval.

of these graphs (e.g., gold curves in Fig. 4, A to C), the quantum algorithm quickly approaches the correct solutions, reducing the average Hamming distance (number of spin flips normalized by N) to the closest MIS and increasing P_{MIS}, while SA remains trapped in local minima at a large Hamming distance from any MIS. For other instances (e.g., purple curves in Fig. 4, A to C), both the quantum algorithm and SA have difficulty finding the correct solution. Moreover, in contrast to our earlier observations suggesting variational parameter concentration for generic graphs, we find that for these hard instances, the quantum algorithm needs to be optimized for each graph individually by scanning the slowdown point of the detuning sweep $\Delta(t)$ to maximize P_{MIS} (Fig. 5, A and B, and fig. S9) (*25*).

Figure 4E shows the resulting highest P_{MIS} reached within a depth of 32 for each hard graph instance as a function of the classical hardness parameter $\mathcal{HP}$. For simulated annealing, we find the scaling $P_{\mathrm{MIS}} = 1 - \exp(-C\mathcal{HP}^{-1.03(4)})$, where C is a positive fitted constant, which is in good agreement with theoretical expectations (*25*). Although for many instances the quantum algorithm outperforms SA, there are significant instance-by-instance variations, and on average, we observe a similar scaling $P_{\mathrm{MIS}} = 1 - \exp(-C\mathcal{HP}^{-0.95(15)})$ (dashed red line).

To understand these observations, we carried out detailed analyses of both classical and quantum algorithms' performance for hard graph instances. Specifically, in (*25*) we show that for a broad class of SA algorithms with both single-vertex and correlated updates, the scaling is at best $P_{\mathrm{MIS}} = 1 - \exp(-C\mathcal{HP}^{-1})$ (where C generally could have polynomial dependence on the system size), indicating that the observed scaling of our version of SA is close to optimal. To gain insight into the origin of the quantum scaling, we numerically compute the minimum energy gap $\delta_{\min}$ during the adiabatic evolution using density-matrix renormalization group (Fig. 5A) (*25*). Figure 5C shows that the performance of the quantum algorithm is mostly well described by quasi-adiabatic evolution with transition probability out of the ground state governed by the minimum energy gap, according to the Landau-Zener formula $P_{\mathrm{MIS}} = 1 - \exp\left(-A\delta_{\min}^{\eta}\right)$ for a constant A, and η = 1.2(2) (*42*). This observation suggests that our quantum algorithm achieves near-maximum efficiency, consistent with the smallest possible value of $\eta = 1$ obtained for optimized adiabatic following (*43*).

By focusing only on instances with large enough spectral gaps such that the evolution time T obeys the "speed limit" determined by the uncertainty principle ($\delta_{\min} > 1/T$) associated with Landau-Zener scaling (*42*), we find an improved quantum algorithm scaling $P_{\mathrm{MIS}} = 1 - \exp(-C\mathcal{HP}^{-0.63(13)})$ (Fig. 4E, solid red line). Because $1/[-\log(1 - P_{\mathrm{MIS}})] \approx 1/P_{\mathrm{MIS}}$ is proportional to the runtime sufficient to find a solution by repeating the experiment, the smaller exponent observed in the scaling for quantum algorithm ($\sim\mathcal{HP}^{1.03(4)}$ for SA and $\sim\mathcal{HP}^{0.63(13)}$ for the quantum algorithm) suggests a superlinear [with a ratio in scaling of 1.6(3)] speedup in the runtime to find an MIS, for graphs where the deep-circuit-regime ($T > 1/\delta_{\min}$) is reached. Moreover, the observed scaling is not altered by the postprocessing used on the experimental data (*25*). We emphasize that achieving this speedup requires an effective depth large enough to probe the lowest-energy many-body states of the system; by contrast, no speedup is observed for graph instances where this depth condition is not fulfilled.

Discussion and outlook

Several mechanisms for quantum speedup in combinatorial optimization problems have been previously proposed. Grover-type algorithms are known to have a quadratic speedup in comparison to brute-force classical search over all possible solutions (*44*, *45*). A quadratic quantum speedup has also been suggested for quantized SA based on discrete quantum walks (*46*, *47*). However, these methods use

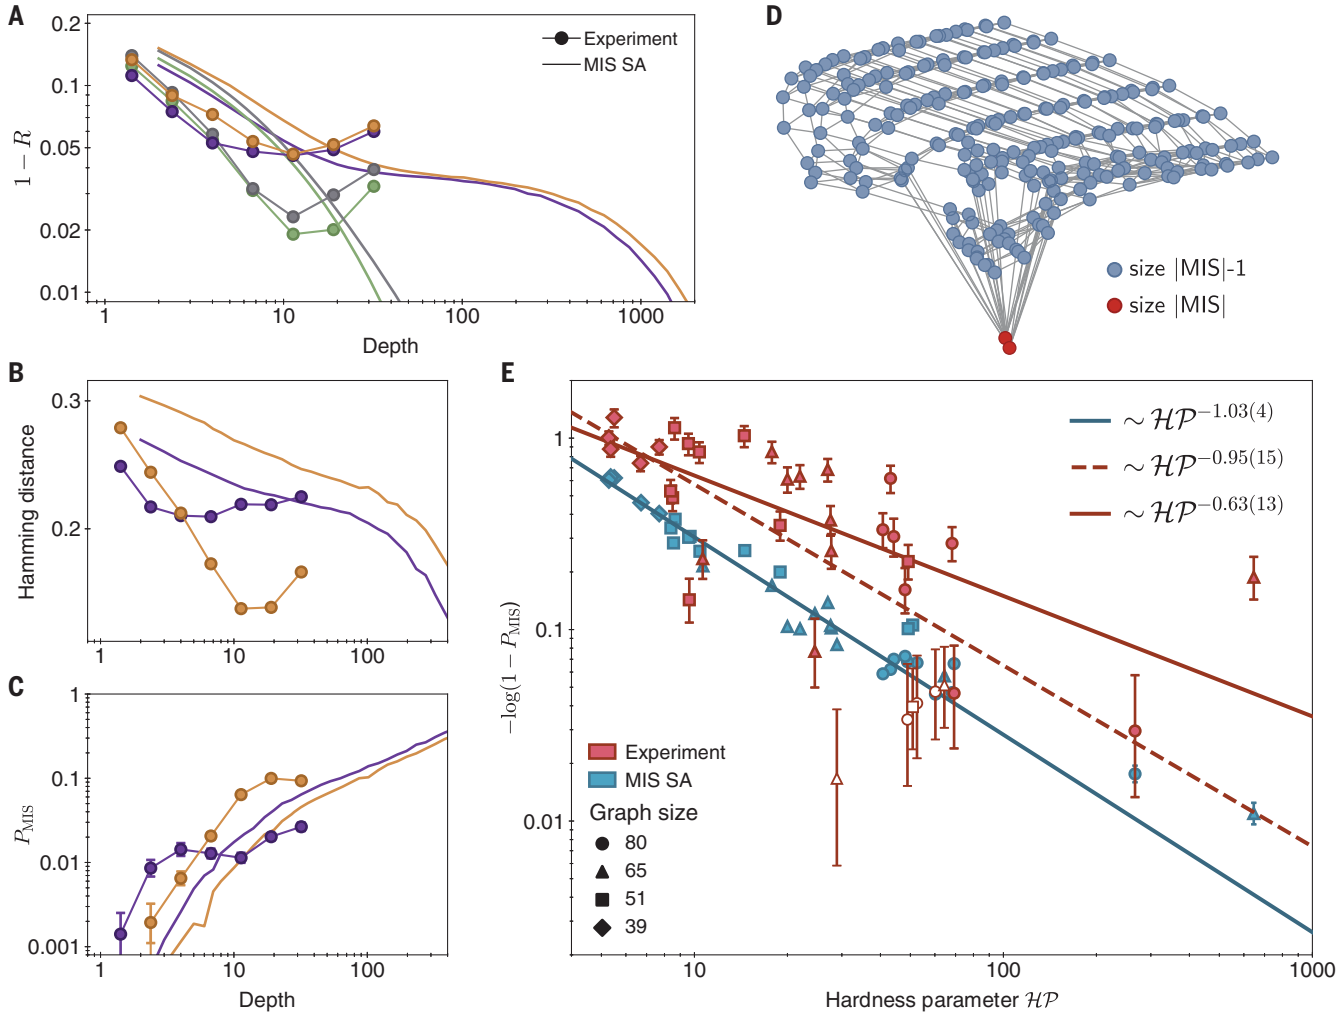

Fig. 4. Benchmarking the quantum algorithm against classical simulated annealing. (**A**) Performance of the quantum algorithm, and the optimized simulated annealing with the MIS Hamiltonian, shown as a function of depth ($\tilde{p}$ for quantum algorithm and p_{SA} for simulated annealing) for four 80-vertex graphs. Green ($\mathcal{HP}$ = 1.8, ρ = 0.13) and gray ($\mathcal{HP}$ = 2.1, ρ = 0.11) graphs are easy for the quantum and classical algorithm; however, purple ($\mathcal{HP}$ = 69, ρ = 0.08) and gold ($\mathcal{HP}$ = 68, ρ = 0.06 are significantly harder and show a plateau at R = (|MIS| − 1)/|MIS|, i.e., independent sets with one less vertex than an MIS. (**B** and **C**) One of the hard graphs (gold) shows much better quantum scaling of average normalized Hamming distance to the closest MIS, and MIS probability (P_{MIS}) compared to the other graph (purple). By contrast, the performance of SA (lines) remains similar between the two graphs. (**D**) Configuration graph of independent sets of size |MIS| and |MIS| − 1 for an example 39-vertex graph ($\mathcal{HP}$ = 5), where the edges connect two configurations if they are separated by one step of simulated annealing. At low temperatures, simulated annealing finds an MIS solution by a random walk on this configuration graph. (**E**) −log(1 − P_{MIS}) for instance-by-instance optimized quantum algorithm (crimson) and simulated annealing (teal) reached within a depth of 32, for 36 graphs selected from the top two percentile of hardness parameter $\mathcal{HP}$ for each size. Power-law fits to the SA (teal, $\sim\mathcal{HP}^{-1.03(4)}$) and the quantum data (dashed crimson line, $\sim\mathcal{HP}^{-0.95(15)}$) are used to compare scaling performance with graph hardness $\mathcal{HP}$. The error in the power-law exponents from the fit is the combination of statistical errors and the error in the least-squares fit. If only graphs with minimum energy gaps large enough to be resolved in the duration of the quantum evolution are considered ($\delta_{min} > 1/T$, excluding hollow data points), the fit (solid crimson line) shows a superlinear speedup $\sim\mathcal{HP}^{-0.63(13)}$ over optimized simulated annealing.

specifically constructed circuits and are not directly applicable to the algorithms implemented here. In addition, the following mechanisms can contribute to the speedup observed in our system. The quantum algorithm's performance in the observed regime appears to be mostly governed by the minimum energy gap δ_{min} (Fig. 5C). We show that under certain conditions, one can achieve coherent quantum enhancement for the minimum gap resulting in a quadratic speedup via $\delta_{min} \sim \mathcal{HP}^{-1/2}$ (*25*). In practice, however, we find that the minimum energy gap does not always correlate with the classical hardness parameter $\mathcal{HP}$, as is evident in the spread of the quantum data in Fig. 4E (see also fig. S21). Some insights into these effects can be gained by a more direct comparison of the quantum algorithm with SA using the same cost function corresponding to the Rydberg Hamiltonian (*25*) (Fig. 5D). Although the observed power-law scaling supports the possibility of a nearly quadratic speedup for instances in the deep circuit regime ($\delta_{min} > 1/T$), it is an open question whether such a speedup can be extended, with a guarantee, in all instances. Finally, it is possible that δ_{min} alone does not fully determine the quantum performance, as suggested by the data points that deviate from the Landau-Zener prediction in Fig. 5C, where

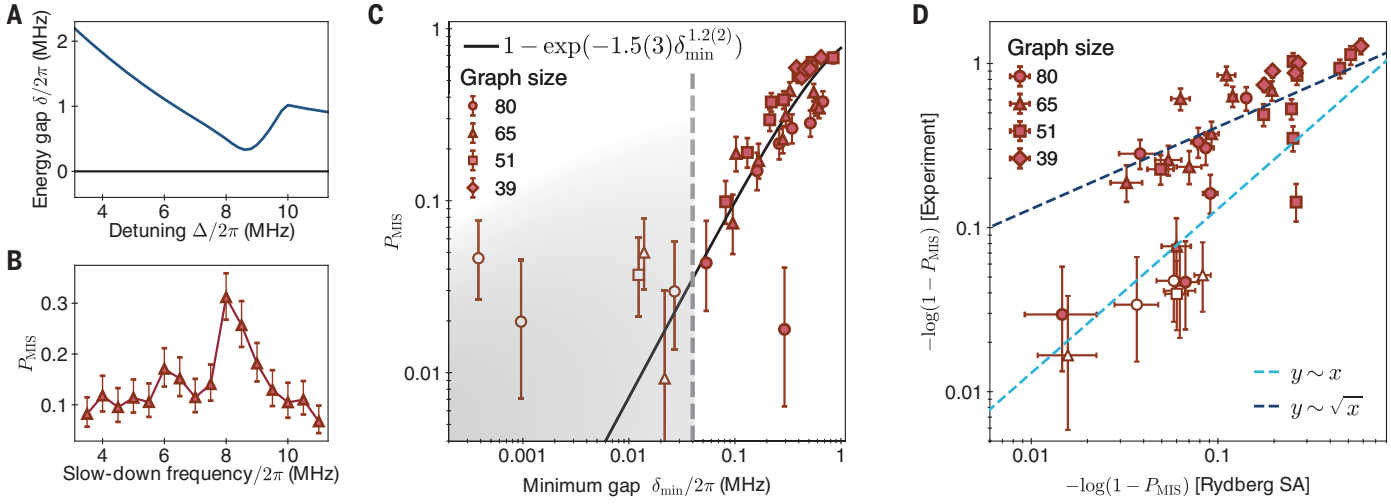

Fig. 5. Understanding hardness for the quantum algorithm. (**A**) Energy gap between the ground (black) and first-excited (blue) states, calculated using the density matrix renormalization group (DMRG) for a graph of 65 atoms. (**B**) To maximize P_{MIS} for hard graphs, the frequency at which the detuning sweep is slowed down is varied (see fig. S9). The largest P_{MIS} corresponds to a slow-down frequency close to the location of the minimum gap. (**C**) Measured P_{MIS} for a fixed effective depth $\tilde{p}$ = 32 as a function of the calculated minimum gap δ_{min}. For many instances, the relation is well described by the Landau-Zener prediction for quasi-adiabatic ground state preparation. The shaded region corresponds to when the gap is too small ($\delta_{\mathrm{min}} < 1/T$) to be properly resolved relative to the quantum evolution time, and points in this region are excluded from the fit both here and in the solid crimson line in Fig. 4E. (**D**) Scaling of $-\log(1 - P_{\mathrm{MIS}})$ observed in the experiment versus in simulated annealing under the classical Rydberg cost function, eq. S14, for best P_{MIS} reached within a depth of 32. These results are consistent with a nearly quadratic speedup for a subset of graphs where $\delta_{\mathrm{min}} > 1/T$.

enhancement through diabatic effects could be possible (*34*, *48*).

Although the scaling speedup observed here suggests a possibility of quantum advantage in runtime, to achieve practical runtime speedups over specialized state-of-the-art heuristic algorithms [e.g., (*39*)], qubit coherence, system size, and the classical optimizer loop need to be improved. The useful depth accessible via quantum evolution is limited by Rydberg-state lifetime and intermediate-state laser scattering, which can be suppressed by increasing the control laser intensity and intermediate-state detuning. Advanced error mitigation techniques such as STIRAP (*49*), as well as error correction methods, should also be explored to enable large-scale implementations. The classical optimization loop can be improved by speeding up the experimental cycle time and by using more advanced classical optimizers. Larger atom arrays can be realized by using improvements in vacuum-limited trap lifetimes and sorting fidelity.

Our results demonstrate the potential of quantum systems for the discovery of new algorithms and highlight a number of new scientific directions. It would be interesting to investigate whether instances with large Hamming distance between the local and global optima of independent set sizes |MIS| – 1 and |MIS| can be related to the overlap gap property of the solution space, which is associated with classical optimization hardness (*50*). In particular, our method can be applied to the optimization of "planted graphs," designed to maximize the Hamming distance between optimal and suboptimal solutions, which can provably limit the performance of local classical algorithms (*51*). Our approach can also be extended to beyond unit disk graphs by using ancillary atoms, hyperfine qubit encoding, and a reconfigurable architecture based on coherent transport of entangled atoms (*52*). Furthermore, local qubit addressing during the evolution can be used to both extend the range of optimization parameters and the types of optimization problems (*5*). Further analysis could elucidate the origins of classical and quantum hardness, for example, by using graph neural network approaches (*53*). Finally, similar approaches can be used to explore realizations of other classes of quantum algorithm [see, e.g., (*54*)], enabling a broader range of potential applications.

REFERENCES AND NOTES

1. M. Sipser, *Introduction to the Theory of Computation* (Course Technology, Boston, ed. 3, 2013).
2. E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, Quantum Computation by Adiabatic Evolution. arXiv:quant-ph/0001106 (2000).
3. E. Farhi, J. Goldstone, S. Gutmann, A Quantum Approximate Optimization Algorithm. arXiv:1411.4028 [quant-ph] (2014).
4. T. Albash, D. A. Lidar, *Rev. Mod. Phys.* **90**, 015002 (2018).
5. A. Lucas, *Front. Phys.* **2**, 5 (2014).
6. D. Wecker, M. B. Hastings, M. Troyer, *Phys. Rev. A* **94**, 022309 (2016).
7. C. Kokail *et al.*, *Nature* **569**, 355–360 (2019).
8. F. Barahona, *J. Phys. Math. Gen.* **15**, 3241–3253 (1982).
9. V. Bapst, L. Foini, F. Krzakala, G. Semerjian, F. Zamponi, *Phys. Rep.* **523**, 127–205 (2013).
10. E. Farhi *et al.*, *Science* **292**, 472–475 (2001).
11. E. Farhi *et al.*, *Phys. Rev. A* **86**, 052334 (2012).
12. S. Knysh, *Nat. Commun.* **7**, 12370 (2016).
13. A. P. Young, S. Knysh, V. N. Smelyanskiy, *Phys. Rev. Lett.* **104**, 020502 (2010).
14. M. P. Harrigan *et al.*, *Nat. Phys.* **17**, 332–336 (2021).
15. G. Pagano *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 25396–25401 (2020).
16. T. M. Graham *et al.*, Demonstration of multi-qubit entanglement and algorithms on a programmable neutral atom quantum computer. arXiv:2112.14589 [quant-ph] (2022).
17. T. F. Rønnow *et al.*, *Science* **345**, 420–424 (2014).
18. H. G. Katzgraber, F. Hamze, Z. Zhu, A. J. Ochoa, H. Munoz-Bauza, *Phys. Rev. X* **5**, 031026 (2015).
19. E. Farhi, D. Gamarnik, S. Gutmann, The Quantum Approximate Optimization Algorithm Needs to See the Whole Graph: A Typical Case. arXiv:2004.09002 [quant-ph] (2020).
20. C.-N. Chou, P. J. Love, J. S. Sandhu, J. Shi, Limitations of Local Quantum Algorithms on Random Max-k-XOR and Beyond. arXiv:2108.06049 [quant-ph] (2021).
21. M. R. Garey, D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, 1979).
22. J. Wurtz, P. Lopes, N. Gemelke, A. Keesling, S. Wang, arXiv:2205.08500 [quant-ph] (2022).
23. B. N. Clark, C. J. Colbourn, D. S. Johnson, *Discrete Math.* **86**, 165–177 (1990).
24. E. J. van Leeuwen, in *Graph-Theoretic Concepts in Computer Science*, D. Kratsch, Ed. (Springer, 2005), pp. 351–361.
25. See supplementary materials.
26. S. Ebadi *et al.*, *Nature* **595**, 227–232 (2021).
27. H. Pichler, S.-T. Wang, L. Zhou, S. Choi, M. D. Lukin, Quantum Optimization for Maximum Independent Set Using Rydberg Atom Arrays. arXiv:1808.10816 [quant-ph] (2018).
28. M. D. Lukin *et al.*, *Phys. Rev. Lett.* **87**, 037901 (2001).
29. D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, J. M. Gambetta, *Phys. Rev. A* **96**, 022330 (2017).
30. B. F. Schiffer, J. Tura, J. I. Cirac, Adiabatic Spectroscopy and a Variational Quantum Adiabatic Algorithm. arXiv:2103.01226 [quant-ph] (2021).
31. G. Semeghini *et al.*, *Science* **374**, 1242–1247 (2021).
32. P. Scholl *et al.*, *Nature* **595**, 233–238 (2021).
33. F. G. S. L. Brandao, M. Broughton, E. Farhi, S. Gutmann, H. Neven, For Fixed Control Parameters the Quantum

Approximate Optimization Algorithm's Objective Function Value Concentrates for Typical Instances. arXiv:1812.04170 [quant-ph] (2018).
34. L. Zhou, S.-T. Wang, S. Choi, H. Pichler, M. D. Lukin, *Phys. Rev. X* **10**, 021067 (2020).
35. J.-G. Liu, X. Gao, M. Cain, M. D. Lukin, S.-T. Wang, arxiv:2205.03718 [cond-mat-stat-mech] (2022).
36. W. H. Zurek, U. Dorner, P. Zoller, *Phys. Rev. Lett.* **95**, 105701 (2005).
37. E. H. Lieb, D. W. Robinson, *Commun. Math. Phys.* **28**, 251 (1972).
38. S. Kirkpatrick, C. D. Gelatt Jr., M. P. Vecchi, *Science* **220**, 671–680 (1983).
39. S. Lamm, P. Sanders, C. Schulz, D. Strash, R. F. Werneck, *J. Heuristics* **23**, 207–229 (2017).
40. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, *J. Chem. Phys.* **21**, 1087–1092 (1953).
41. F. V. Fomin, D. Kratsch, *Exact Exponential Algorithms* (Springer, ed. 1, 2010).
42. L. D. Landau, E. M. Lifshitz, *Quantum Mechanics Non-Relativistic Theory* (Pergamon, ed. 3, 1977).
43. J. Roland, N. J. Cerf, *Phys. Rev. A* **65**, 042308 (2002).
44. L. K. Grover, A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, Philadelphia, 1996.
45. C. Durr, P. Hoyer, A Quantum Algorithm for Finding the Minimum. arXiv quant-ph/9607014 (1999).
46. M. Szegedy, in *45th Annual IEEE Symposium on Foundations of Computer Science* (2004), pp. 32–41.
47. R. D. Somma, S. Boixo, H. Barnum, E. Knill, *Phys. Rev. Lett.* **101**, 130504 (2008).
48. E. Crosson, E. Farhi, C. Y.-Y. Lin, H.-H. Lin, P. Shor, Different Strategies for Optimization Using the Quantum Adiabatic Algorithm. arXiv:1401.7320 [quant-ph] (2014).
49. M. Fleischhauer, A. Imamoglu, J. P. Marangos, *Rev. Mod. Phys.* **77**, 633–673 (2005).
50. D. Gamarnik, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2108492118 (2021).
51. D. Gamarnik, I. Zadik, The Landscape of the Planted Clique Problem: Dense subgraphs and the Overlap Gap Property. arXiv:1904.07174 [math.ST] (2019).
52. D. Bluvstein *et al*., *Nature* **604**, 451–456 (2022).
53. A. Sohrabizadeh, Y. Bai, Y. Sun, J. Cong, Enabling Automated FPGA Accelerator Optimization Using Graph Neural Networks. arXiv:2111.08848 [cs.ARs] (2021).
54. D. S. Wild, D. Sels, H. Pichler, C. Zanoci, M. D. Lukin, *Phys. Rev. Lett.* **127**, 100504 (2021).
55. M. Fishman, S. R. White, E. M. Stoudenmire, The ITensor Software Library for Tensor Network Calculations. arXiv:2007.14822 [cs.MS] (2020).
56. S. Ebadi, Quantum Optimization of Maximum Independent Set using Rydberg Atom Arrays, Zenodo (2022); https://doi.org/10.5281/zenodo.6462687.

ACKNOWLEDGMENTS

We thank I. Cirac, J. Cong, S. Evered, M. Kalinowski, M. Lin, T. Manovitz, M. Murphy, B. Schiffer, J. Singh, A. Sohrabizadeh, J. Tura, and D. Wild for illuminating discussions and feedback on the manuscript. **Funding:** We acknowledge financial support from the DARPA ONISQ program (grant no. W911NF2010021), the Center for Ultracold Atoms, the National Science Foundation, the Vannevar Bush Faculty Fellowship, the US Department of Energy [DE-SC0021013 and DOE Quantum Systems Accelerator Center (contract no. 7568717)], the Army Research Office MURI, QuEra Computing, and Amazon Web Services. M.C. acknowledges support from DOE CSG award fellowship (DE-SC0020347). H.L. acknowledges support from the National Defense Science and Engineering Graduate (NDSEG) fellowship. D.B. acknowledges support from the NSF Graduate Research Fellowship Program (grant DGE1745303) and The Fannie and John Hertz Foundation. G.S. acknowledges support from a fellowship from the Max Planck/Harvard Research Center for Quantum Optics. R.S. and S.S. were supported by the U.S. Department of Energy under grant DE-SC0019030. X.G. acknowledges support from a fellowship from the Max Planck/Harvard Research Center for Quantum Optics. B.B. acknowledges support from DARPA grant W911NF2010021, a Simons investigator fellowships, NSF grants CCF 1565264 and DMS-2134157, and DOE grant DE-SC0022199. H.P. acknowledges support by the Army Research Office (grant no. W911NF-21-1-0367). The DMRG calculations in this paper were performed using the ITensor package (*54*) and were run on the FASRC Odyssey cluster supported by the FAS Division of Science Research Computing Group at Harvard University. **Author contributions:** S.E., A.K., M.C., T.T.W., H.L., D.B., G.S., A.O., and J.-G. L. contributed to building the experimental setup, performing the measurements, and data analysis. M.C., J.-G. L., R. S., X.-Z. L., B. N., X. G., L. Z., S. C., H. P., and S.-T. W. contributed to theoretical analysis and interpretation. B. B., E. F., S. S., and N. G. contributed to interpretation of the observations and benchmarking studies. All authors discussed the results and contributed to the manuscript. All work was supervised by M.G., V.V., and M.D.L. **Competing interests:** N.G., M.G., V.V., and M.D.L. are cofounders and shareholders of QuEra Computing. A.K. is a shareholder and an executive at QuEra Computing. A.O. and S.-T.W. are shareholders of QuEra Computing. Some of the techniques and methods used in this work are included in pending patent applications filed by Harvard University (Patent application nos. PCT/US2018/042080 and PCT/US2019/049115). **Data and materials availability:** Code for classical tensor network algorithms for graph characterization is available at (*35*). Data and code are available on Zenodo (*56*).

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.abo6587
Materials and Methods
Figs. S1 to S22
Table S1
References (*57*–*87*)

Submitted 17 February 2022; accepted 22 April 2022
Published online 5 May 2022
10.1126/science.abo6587

Science

Quantum optimization of maximum independent set using Rydberg atom arrays

S. EbadiA. KeeslingM. CainT. T. WangH. LevineD. BluvsteinG. SemeghiniA. OmranJ.-G. LiuR. SamajdarX.-Z. LuoB. NashX. GaoB. BarakE. FarhiS. SachdevN. GemelkeL. ZhouS. ChoiH. PichlerS.-T. WangM. GreinerV. Vuleti#M. D. Lukin

Science, 376 (6598), • DOI: 10.1126/science.abo6587

Solving hard graph problems

Realizing quantum speedup for solving practical, computationally hard problems is the central challenge in quantum information science. Ebadi *et al*. used Rydberg atom arrays composed of up to 289 coupled qubits in two spatial dimensions to investigate quantum optimization algorithms for solving the maximum independent set, a paradigmatic nondeterministic polynomial time–hard combinatorial optimization problem (see the Perspective by Schleier-Smith). A hardware-efficient encoding protocol associated with Rydberg blockade was used to realize a closed-loop optimization method to test several variational algorithms and subsequently apply them to systematically explore a class of nonplanar graphs with programmable connectivity. The results demonstrate the potential of quantum machines as a tool for the discovery of new promising algorithm classes. —ISO

Science (ISSN) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.