

Sparse-graph optimization using weighted quantum wires in Rydberg-atom arrays

A. G. de Oliveira^{1,*}, J. Kombe¹, G. Pelegrí¹, P. Schroff¹, M. T. Wells-Pestell¹, D. M. Walker¹,
A. J. Daley^{1,2} and J. D. Pritchard^{1,†}

¹*Department of Physics and SUPA, University of Strathclyde, Glasgow, G4 0NG, United Kingdom*

²*Department of Physics, Clarendon Laboratory, University of Oxford, OX1 3PU Oxford, United Kingdom*

 (Received 20 November 2025; revised 27 April 2026; accepted 6 July 2026; published 5 August 2026)

Neutral atom arrays provide a versatile platform to implement coherent quantum annealing as an approach to solving hard combinatorial optimization problems. Here we present and experimentally demonstrate an efficient encoding scheme based on chains of Rydberg-blockaded atoms, which we call quantum wires, to natively embed maximum weighted independent set (MWIS) and quadratic unconstrained binary optimization (QUBO) problems on a neutral atom architecture. For graphs with quasi-unit-disk connectivity, in which only a few long-range edges are required, our approach requires a significantly lower overhead in the number of ancilla qubits than previous proposals, facilitating the implementation on currently available hardware. To demonstrate the approach, we perform annealing of weighted graphs on a programmable atom array using local light shifts to encode problem-specific weights across graphs of varying sizes. This approach successfully identifies the solutions to the original MWIS and QUBO graph instances. Our work expands the operational toolkit of near-term neutral atom arrays, enhancing their potential for scalable quantum optimization.

DOI: [10.1103/bxsl-flzc](https://doi.org/10.1103/bxsl-flzc)

I. INTRODUCTION

Combinatorial optimization problems are of great relevance in a wide range of scientific fields and industry sectors [1]. In many cases, these problems have been shown to be NP-hard, and finding new algorithms to tackle them more efficiently remains an outstanding challenge. Over the last two decades, quantum optimization has attracted attention as a potential route toward achieving computational speedups for certain problems [2]. Quantum optimization approaches typically rely on preparing a many-body quantum system in a ground state that encodes the solution of the combinatorial optimization problem at hand. Typically, this state preparation is achieved through quantum annealing algorithms (QAA) [3,4] that rely on the quantum adiabatic principle, or using hybrid quantum-classical approaches [5,6] such as the variational quantum eigensolver [7–10]. These approaches are particularly well suited for combinatorial optimization problems that admit formulations as spin models, with examples including maximum (weighted) independent set (MIS/MWIS) and

quadratic unconstrained binary optimization (QUBO) problems [11].

Because of their ability to realize programmable spin models, neutral atom quantum computers have recently emerged as a promising platform for quantum optimization [12,13]. In these systems, the Rydberg blockade mechanism naturally enforces the independence constraint for atoms spaced below a given radius, resulting in a native embedding of the MIS problem on a particular family of geometrical graphs called unit-disk graphs (UDGs) [14,15]. The use of QAA to solve the MIS problem on UDGs with King's connectivity using hundreds of atoms has been experimentally demonstrated recently [16,17], with an indication toward a superlinear quantum speedup for hard graph instances. However, subsequent work [18] showed that MIS instances on UDGs with King's connectivity can be optimally solved for up to thousands of nodes within minutes by classical approaches (custom or generic commercial solvers) on commodity hardware, without instance-specific fine-tuning. This result highlights the significance of extending the range of problems that can be embedded in neutral atom arrays beyond UDG-MIS, not only to widen the scope of possible real-world applications of this technology but also to identify instances of quantum advantage. In this spirit, several proposals have been put forward to map non-UDG MIS [19–21], QUBO [22], higher-order unconstrained binary optimization (HUBO) [23], satisfiability [24], and general combinatorial optimization [25] problems onto the

*Contact author: andre.oliveira@strath.ac.uk

†Contact author: jonathan.pritchard@strath.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

UDG-MIS problem using ancillary atoms and requiring only global control.

Another possibility to encode problems beyond UDG-MIS in neutral atom hardware is to use local control fields [26]. This enables the realization of differential weights in the vertices represented by the atoms, permitting the embedding of MWIS problems defined on UDGs [27]. Of note, UDG-MWIS problems can be further mapped into more general combinatorial optimization problems such as non-UDG MWIS, QUBO, or integer factorization using an approach based on modular gadgets to copy and delocalize the information of the logical variables, and create arbitrary nonlocal connections [28,29]. Alternatively, an implementation based on the parity architecture [30] allows conversion of QUBO and HUBO problems into UDG-MWIS [31]. Both of these approaches require at worst an $\mathcal{O}(N^2)$ overhead for a problem with N logical variables [28,32], significantly reducing the ancillary qubit count from $\mathcal{O}(N^6)$ of previous proposals using only global addressing [16]. However, while there are recent proposals to optimize the layout of the embedded graphs [33], the overhead required for the known mappings into UDG-MWIS still poses a challenge for the embedding of large-scale combinatorial optimization problems onto near-term neutral atom hardware.

In this paper, we introduce and experimentally realize a hardware-efficient scheme to map non-UDG MWIS and QUBO problems into a natively embeddable MWIS problem defined on a UDG. Our approach, schematically depicted in Fig. 1, is based on chains of atoms in the blockade regime, which we will call *quantum wires*, to mediate interactions between physically separated vertices. For problems that can be formulated on graphs with quasi-UDG connectivity, and which require only a handful of long-range interactions, this allows one to significantly lower the overhead in ancillary atoms with respect to the general embedding schemes discussed above. By engineering the weights of the atoms that form the wires, we further ensure the preservation of the spectrum of the original problem.

This goes beyond the wire architecture presented in Refs. [19,20], which is designed to encode unweighted problems and only preserves the ground state, requiring postselection of the antiferromagnetic state of the wires to identify valid solutions. Our wire constructions rely only on the blockade condition being fulfilled between neighboring atoms, easing the experimental requirements with respect to proposals that require precise control of the relative strength of the interactions [26,34].

The paper is organized as follows. In Sec. II, we review the implementation of UDG-MWIS problems on neutral atom hardware. Section III introduces wire constructions that allow for the encoding of MWIS problems on non-UDG geometries, while Sec. IV explains how these wires can be adapted to encode QUBO problems as




FIG. 1. Quantum wire approach. (a) An example MWIS problem defined on a graph with non-unit-disk connectivity imposed by a sparse number of long-range edges (left) can be efficiently embedded on neutral atom hardware by introducing weighted quantum wires (right) created from chains of ancilla atoms as indicated by red circles. (b) An MWIS problem is mapped to a UDG-MWIS and solved on the QPU, with the solution indicated by dashed gray circles. (c) This strategy can be extended to QUBO problems by transforming them into equivalent UDG-MWIS instances where weighted edges are embedded using quantum wires.

UDG-MWIS. Section V discusses the hardware implementation and chosen annealing protocol [27], while Sec. VI presents examples of non-UDG MWIS and QUBO problems embedded into UDG-MWIS using these wires, followed by experimental demonstrations of quantum annealing to validate the approach. We conclude in Sec. VII with a summary of our results and an outlook for future work.

II. BACKGROUND

We consider a weighted graph $G(V, E)$, which consists of a set of vertices V , each with weight $w_i > 0$, and a collection of edges E between pairs of vertices. The maximum weighted independent set is defined as the subset of vertices not connected by an edge that maximizes the total sum of the weights. Introducing binary variables $n_i = \{0, 1\}$ to indicate whether vertex i belongs to the set, we can identify the MWIS with the bitstring $\{n_i\}$ that minimizes the following classical cost function

$$H_{\text{MWIS}} = -\sum_i w_i n_i + \sum_{(i,j) \in E} U_{ij} n_i n_j, \quad (1)$$

where $U_{ij} > \max\{w_i\}$ is a penalty term that ensures that no two vertices connected by an edge simultaneously belong to the MWIS. Let us also consider a unit-disk graph, which is a graph that can be embedded in a

two-dimensional plane and such that there is an edge between a vertex pair (i, j) if and only if their separation fulfills the condition $r_{ij} < R$, where R is the so-called unit-disk radius. The cost function of the UDG-MWIS problem, i.e., the MWIS problem defined on UDG graphs, can be naturally mapped to the many-body Hamiltonian of a neutral atom array with a laser driving between a ground state $|g\rangle \equiv |0\rangle$ and a Rydberg state $|r\rangle \equiv |1\rangle$

$$H_{\text{Ryd}}/\hbar = \frac{\Omega}{2} \sum_i \sigma_i^x - \sum_i \Delta_i n_i + \sum_{i<j} V(r_{ij}) n_i n_j, \quad (2)$$

where Ω is the Rabi frequency, Δ_i is a site-dependent detuning from the $|0\rangle \leftrightarrow |1\rangle$ transition frequency, $n_i \equiv |1\rangle_i \langle 1|$ is an operator that counts the number of Rydberg excitations on site i , and $V(r) = C_6/r^6$ is the dipole-dipole van der Waals (vdW) interaction strength between two atoms separated by a distance r . Because of this interaction, each atom has a blockade radius $r_B^i = \sqrt[6]{C_6/\sqrt{\Omega^2 + \Delta_i^2}}$ in which only one excitation can occur. In the limit $\Omega = 0$, defining the smallest blockade radius as $r_B^{\min} = \sqrt[6]{C_6/\max\{\Delta_i\}}$ and assuming $V(r > r_B^{\min}) \approx 0$, the ground state of H_{Ryd} encodes the solution of a UDG-MWIS problem with weights $0 \leq \tilde{w}_i = \Delta_i/\Delta_0 \leq 1$, and unit-disk radius $R \approx r_B^{\min}$, realizing the constraint $V(r_{ij}) > \max\{\Delta_k\}$ for all (i, j) sharing an edge. Although the assumption $V(r > r_B) \approx 0$ is well justified due to the rapid decay of the dipole interactions with the atom separation, in practical implementations the interaction tails can have a non-negligible effect on the ground state of the system. Therefore, it is necessary to further impose the condition that the maximal unwanted interaction is less than the smallest weight in the problem, $\max_{(i,j) \notin E} V(r_{ij}) < \min_k \Delta_k$. This guarantees that all pairs of vertices not connected by an edge are allowed to belong to the MWIS.

The correspondence between H_{MWIS} and H_{Ryd} opens up the possibility to solve general optimization tasks with

neutral atom quantum computers by first mapping them to UDG-MWIS instances [22,23,25,28,31,32], preparing the ground state via quantum annealing, and finally undoing the mapping to recover the solution to the original problem.

III. QUANTUM WIRES FOR MWIS PROBLEMS

Figure 2(a) illustrates a basic wire construction to realize an MWIS constraint between two physically separated vertices with weights α and β (with $\alpha > \beta$ in this example). The allowed configurations in the logical MWIS problem are $|\alpha, \beta\rangle = \{|10\rangle, |01\rangle, |00\rangle\}$, with corresponding energies $\{-\alpha, -\beta, 0\}$. The wire (highlighted in red) consists of an even number of atoms L with a uniform weight c , arranged in such a way that every atom is within the blockade radius of its two neighbors, forming an effective one-dimensional chain. As shown in the energy diagram in Fig. 2(a), choosing ancilla weights $c > \alpha + \beta$ yields the correct ordering of the states of the logical nodes in the three lowest-energy configurations of the system formed by the logical nodes and the wire, with the logical state $|11\rangle$ being energetically forbidden. Even though we only show one possible configuration for the $|00\rangle$ and $|11\rangle$ states, we note that, without considering the effect of vdW interactions, they have respective degeneracies of $L/2 + 1$ and $L/2$ due to the degree of freedom corresponding to the localization of the domain wall [29]. All the states of the wire are shifted by a constant offset $-Lc/2$ with respect to their corresponding logical configurations, ensuring the preservation of the MWIS solution when the wire is embedded in a larger graph with more atoms. We emphasize that this construction *preserves* the spectrum of the logical problem and truly facilitates an interaction between physically spatially separated nodes.

The construction that enables us to realize an edge between two spatially separated weighted atoms can also be generalized to delocalize other connectivity structures involving a larger number of atoms. In Fig. 2(b), we




FIG. 2. Weighted quantum wires. (a) Basic construction of a wire to connect two nodes with weights α and β in MWIS and QUBO problems. The energy diagram shows the ordering of the eigenstates for an MWIS implementation. (b),(c) Wire constructions to delocalize triangular and all-to-all square interactions between vertices in MWIS problems. (d) Generalization of the crossing gadget introduced in Ref. [28] to allow for arbitrary weights of the nodes.

illustrate a wire construction that realizes an all-to-all connectivity between three atoms with weights α , β , and γ , where atoms α and β are within the blockade radius of each other but outside the one of atom γ . The wire is constructed in the same way as in Fig. 2(a), but now with vertices α and β connected to one end of the wire, and node γ to the opposite end. The analysis of the possible configurations of the wire shows that in order to preserve the correct ordering of the energy spectrum, one can choose a weight for the ancillary atom of $c > \max\{\alpha, \beta\} + \gamma$. Similarly, in Fig. 2(c), we show how to use a wire of even length L to stretch an all-to-all connected square formed by four atoms with weights α, β, γ , and δ into two separated regions, in one of which there is a direct edge between α and β and in the other an edge between γ and δ . In this situation, a choice of the wire weight $c > \max\{\alpha, \beta\} + \max\{\gamma, \delta\}$ ensures the preservation of the logical spectrum. Analogously, this quantum wire method can be extended to delocalize clique structures of larger size than those depicted in Figs. 1(b) and 1(c).

In order to implement more complicated structures with long-range connectivity, it might be necessary for quantum wires to cross. This can be achieved with a generalization of the crossing gadget introduced in Ref. [28], as shown in Fig. 2(d). In the most general case, in which the four vertices involved in the crossing have different weights $\alpha, \beta, \gamma, \delta$, choosing the weight of the crossing gadget ancillary nodes as $c > \alpha + \beta + \gamma + \delta$ guarantees the preservation of the spectrum of the logical problem. Further details and a discussion of the robustness of the proposed quantum wires can be found in Appendix A.

While a systematic scheme to automate problem embeddings using quantum wires is outside of the scope of this paper, we can make a heuristic resource estimation to identify in which cases the wire approach could provide a more favorable ancilla scaling than the $\mathcal{O}(N^2)$ that the gadget and parity architecture incur into [28,31]. Assuming a graph with N atoms arranged in a square lattice, a single wire would utilize in the worst case $\mathcal{O}(\sqrt{N})$ ancillary qubits to realize an interaction between distant atoms. Accounting for the fact that when N_{wires} are present, it might be necessary to introduce crossing gadgets that increase the ancillary atom count by up to $\mathcal{O}(N_{\text{wires}}(N_{\text{wires}} - 1))$. A conservative estimate suggests that the wire embedding would be advantageous for graphs in which most of the edges admit a unit-disk representation and $N_{\text{wires}} \sim \mathcal{O}(\sqrt{N})$ edges require long-range connections. This will serve as our operational definition of quasi-UDGs. Quantum wires would also be particularly well suited for problems composed of several UDG units that are connected between each other only through a very small number of nodes. Additionally, quantum wire constructions as presented in our work could be added to the toolkit for overhead reduction in the gadget embedding presented in Ref. [33].

IV. QUANTUM WIRES FOR QUBO PROBLEMS

Quadratic unconstrained binary optimization is a well-known NP-hard problem with a myriad of applications [11,35]. The solution of a QUBO problem can be formulated in terms of a bitstring $\{n_i\}$, with $n_i \in \{0, 1\}$, that minimizes the classical cost function

$$H_{\text{QUBO}} = \sum_i h_i n_i + \sum_{i,j} J_{ij} n_i n_j, \quad (3)$$

where h_i and J_{ij} are coefficients that define the specific optimization problem to be solved. In the case when $h_i = -|h_i| < 0 \forall i$ and $J_{ij} \geq 0 \forall (i, j)$, H_{QUBO} can be implemented within the quantum wire framework presented here. To understand this, let us consider the simplest case of a QUBO problem defined over two variables described by the cost function

$$f(n_1, n_2) = -|h_1|n_1 - |h_2|n_2 + J_{12}n_1n_2. \quad (4)$$

Comparing this expression with the wire spectrum shown in Fig. 2(a), we can see that by choosing $\alpha = |h_1|$, $\beta = |h_2|$ and $c = J_{12}$, the states of the wire exactly mimic the QUBO cost function Eq. (4) up to a constant energy offset $-Lc/2$ (for $L \geq 2$). Note that the constraint $c > \alpha + \beta$, necessary for the MWIS implementation, is now relaxed so the state $|11\rangle$ can appear in any order in the spectrum depending on the value of J_{12} .

Following this basic construction for a 2-bit QUBO, larger problems can be implemented by assigning an atom with relative weight $w_i = h_i$ to each logical variable n_i , and connecting the logical qubits with wires with homogeneous weights $c_{ij} = J_{ij}$. The building blocks, shown in Figs. 2(b) and 2(c), could also be used for triangles or squares with uniform J couplings, as depicted in Fig. 1(c). The crossing gadget of Fig. 2(d) could be used in the same way as for MWIS problems. While this approach for embedding QUBO problems is only scalable for instances with mostly local connectivity, in these scenarios it can provide an overhead reduction with respect to the proposals in Refs. [22,28,31], which are designed for general graphs.

V. HARDWARE IMPLEMENTATION

To implement MWIS annealing on neutral atom hardware, we employ the protocol introduced in Ref. [27], using local light shifts to implement node weights as illustrated in Fig. 3 and detailed in Appendix B. This method integrates a global annealing schedule with site-specific, time-dependent light shifts, dynamically generated via an acousto-optic modulator (AOM) and a spatial light modulator (SLM). The protocol is defined by a global Rabi frequency Ω and local detuning pulses $\Delta_i(t) = \Delta(t) + \delta_{\text{ac}}^i(t)$, where $\Delta_i(t)$ is the global detuning of the




FIG. 3. Weighted annealing on an atomic QPU. (a) Local light shifts encode vertex weights via $\delta_{ac}^i = w_i \delta_{ac}$. (b) The system undergoes annealing by globally ramping the detuning to resonance, followed by ramping of the local detunings to implement differential vertex weights. (c) Atoms are rearranged geometrically to encode MWIS vertices, with site-dependent light shift for weights. Rydberg excitation is performed using a pair of counterpropagating global excitation laser beams far detuned from the intermediate excited state ($\Delta/2\pi = 1$ GHz) with an effective Rabi frequency $\Omega = \Omega_1 \Omega_2 / 2\Delta$. Together, this realizes the Hamiltonian described in Eq. (2). (d) Postannealing, Rydberg excitations are detected via atom loss.

Rydberg lasers and $\delta_{ac}^i(t) = w_i \delta_{ac}(t)$ encodes vertex weights through local light shifts.

Atoms are spatially arranged to represent the unit-disk graph instance of the MWIS problem, with edge-connected qubits placed within a blockade radius ($r_B \sim 12 \mu\text{m}$) to enforce interaction constraints. Initially, the system is prepared in the ground state $|0\rangle^{\otimes N}$ with $\Omega = 0$ and $\Delta_i = \Delta < 0$. The Rabi frequency is increased to its maximum and kept constant while the global detuning $\Delta(t)$ is swept to resonance ($\Delta_i = \Delta = 0$). After crossing resonance, the local detunings $\Delta_i(t) = \delta_{ac}^i(t) = w_i \delta_{ac}(t)$ are introduced for every atom while keeping the global Rydberg lasers at resonance ($\Delta = 0$). During annealing, the maximum local detuning is limited to $\Delta_i^{\text{max}}(t) < 0.9C_6/R^6$ to maintain blockade constraints between edge-connected atoms. Annealing typically occurs over a timescale of about $8 \mu\text{s}$. This procedure is illustrated in Figs. 3(a) and 3(b). The local detunings are calibrated via closed-loop feedback of the SLM hologram using spectroscopic measurements, typically achieving $\lesssim 7\%$ rms error after five iterations. The annealing schedule ensures that blockade constraints are satisfied as detunings increase, and the final Hamiltonian corresponds directly to the classical MWIS cost function of Eq. (1).

After annealing, Rydberg excitations are detected via atom loss. Repetition of the annealing sequence multiple times, combined with excluding trials that showed imperfect loading, generates statistics, and configurations are

ranked by occurrences. When the annealing protocol is executed properly, the most frequent configuration typically corresponds to the many-body ground state of the system. Errors arising from blockade violations, losses, and imperfect state preparation and readout can be identified and corrected through classical postprocessing, following the protocol introduced in Ref. [16]. We first apply vertex reduction to eliminate vertices that violate the edge constraints of the graph to recover a valid independent set solution string. In many cases, this corrected configuration permits additional excitations without violating blockade constraints. The insertion of those excitations is performed by a fixed-depth vertex addition algorithm. See Appendix D for further details. Closed-loop optimization can be used to tune the annealing parameters by seeking to minimize the weighted average cost $\langle H_{\text{MWIS}} \rangle$ calculated over the ensemble of observed configurations.

VI. RESULTS AND DISCUSSION

A. Quasi-UDG MWIS

The first sparse graph example for which we want to find the MWIS solution is shown in Fig. 4(a) featuring ten vertices. The set of vertex weights used in this problem instance is given by $\{w_0, w_1, \dots, w_9\} = \{0.22, 0.51, 0.36, 0.21, 0.46, 0.22, 0.1, 0.57, 0.48, 0.1\}$. By rearranging the positions of the nodes, all the edges can be made to satisfy the UDG criterion except for the triangle between nodes $\{0, 4, 5\}$, and the all-to-all connected cluster formed by nodes $\{0, 1, 8, 9\}$. As shown in Fig. 4(b), these links can be realized with quantum wire constructions, resulting in an embedded graph formed by a total of 20 atoms with exclusively UDG connectivity. For comparison, we note that the implementation of the same logical problem using the gadget encoding described in Ref. [28] with additional optimization takes a total of 211 atoms using the open software package `UnitDiskMapping.jl` [36]. Naturally, this could be further optimized, for example using graph reduction techniques [33,37], but for quasi-UDGs our technique provides a robust and efficient alternative that can be readily solved with weighted-graph annealing in our neutral atom hardware.

The ground state of the encoded problem is degenerate, with three distinct configurations corresponding to the minimum MWIS cost. Among these, the configuration highlighted with dashed gray circles in Fig. 4(b) represents the ground state obtained using the experimentally calibrated interaction weights. To optimize the annealing schedule, we employed a closed-loop procedure that iteratively tuned the parameters of the annealing schedule. Classical postprocessing of the measurement outcomes reveals a clear anticorrelation between the ground state relative frequency and the average MWIS cost, validating the effectiveness of the annealing parameter optimization,




FIG. 4. Solving non-UDG MWIS problems using quantum wire encoding. (a) A non-unit-disk graph instance of the maximum weighted independent set problem is first rearranged, and quantum wires are inserted to transform it into a UDG. (b) The resulting UDG can be natively implemented on a neutral atom array. In this example, two weighted quantum wires are added to realize the connections highlighted in red. The MWIS solution is indicated by dashed lines around Rydberg excited nodes in the logical and embedded graphs. (c) Relative frequency of ground state observations versus average MWIS cost observed during annealing parameter optimization. Classical postprocessing of the measured configurations enhances the anticorrelation between these quantities. (d) Histograms of the 20 most frequent states and their cost from 230 (out of 308) defect-free realizations of the annealing run yielding the lowest (vertex addition) average cost. Ground-state configurations, which correspond to MWIS solutions, are highlighted with red bars, while their cost is highlighted by a red horizontal line. The top histograms show raw experimental data, while the bottom displays frequencies and costs after applying both vertex reduction and addition.

as shown in Fig. 4(c). To further illustrate the performance of our approach, we highlight a specific measurement corresponding to the lowest average cost $\langle H_{\text{MWIS}} \rangle$ observed during the optimization cycle. This measurement, shown in Fig. 4(d), confirms that the most frequent

configuration is a valid ground state solution of the full MWIS graph. The histogram data is sorted first by frequency of occurrence, then by state index. We display only the first 20 bins for clarity, along with the cost of each measured state. In the data, we observe that while there should be three degenerate solutions, the system preferentially returns one of the configurations. This is likely due to small calibration errors in the relative weights ($< 7\%$ rms) and trap positions that break the degeneracy associated with different configurations of excitations along the quantum wire. However, since these all contain the same configuration of logical qubits, the broken degeneracy does not prevent observation of a valid solution to the target problem instance.

Both the input graph and the postprocessed data are shown, providing direct evidence of the system's ability to identify MWIS solutions through quantum annealing on the neutral atom hardware. In the input data, we observe a target ground state solution with a relative frequency of 0.05 ± 0.02 . Without the annealing sequence (see Appendix B for details), the expected probability of observing a target state solely due to atom loss is $p_{\text{target}} = p^{(N-n)}(1-p)^n$, with target state length N , number of lost atoms n , and measured baseline survival rate p . Using $N = 20$, $n = 8$, and $p = 0.985$, we find $p_{\text{target}} = 2 \times 10^{-15}$ for the MWIS solution. These results show that the coherent annealing process dominates over atom losses and validate the approach of using weighted quantum wires for finding MWIS solutions on near term hardware.

B. QUBO problem graphs

To demonstrate that quantum wires can be used to natively embed QUBO problems on neutral atom arrays, we investigate small QUBO problems structured as square-shaped graphs, which act as modular components for constructing more complex instances. For each QUBO, the problems are generated using randomized node weightings and edge couplings. As illustrated in Fig. 5(a), these blocks can be combined up to three at a time to generate larger problem configurations. Using a top-left to bottom-right indexing convention, the QUBO problem depicted in Fig. 5(a) is characterized by the linear coefficients $\{h_0, h_1, \dots, h_7\} = \{-0.3, -0.8, -0.6, -0.7, -1.0, -0.7, -0.8, -0.9\}$ and couplings $\{J_{01}, J_{04}, J_{12}, J_{15}, J_{23}, J_{26}, J_{37}, J_{45}, J_{56}, J_{67}\} = \{0.8, 0.5, 0.1, 0.4, 0.4, 0.3, 0.4, 0.5, 0.6, 1.0\}$. Each QUBO instance is then translated into a UDG embedding using quantum wires as shown in Fig. 5(b), ensuring operational compatibility with the neutral atom array. For each configuration, we perform weighted-graph annealing and gather measurement statistics over multiple experimental runs, with the total number of defect-free realizations (those with all atoms present) recorded in the caption.

As shown in Figs. 5(c)–5(g), the experimental outcomes consistently favor the ground-state configurations, which




FIG. 5. Block-wise decomposition of a QUBO problem. (a) A QUBO problem consisting of eight nodes and ten edges is shown. (b) This problem is mapped to a UDG formulation of the MWIS, enabling native implementation on a Rydberg quantum processing unit. The QUBO graph can be partitioned into blocks of varying sizes by selectively removing edges. (c)–(e) Annealing results for each individual block derived from (a), shown from left to right. Block (e) exhibits two degenerate ground state solutions. Data shown for 870, 763, and 749 defect-free realizations, respectively, out of 1000 trials. (f) MWIS results corresponding to the QUBO problem composed of the two leftmost blocks (664/1000 defect-free realizations). (g) MWIS results for the full three-block QUBO problem for 669 defect-free realizations. In (c)–(g), the ground-state configurations are highlighted in red.

emerge as the most frequently observed results and correspond to the logical QUBO solutions. The histograms show the ten most common configurations without any postprocessing of the data for each graph and confirm that the MWIS solutions dominate the distributions. The postprocessed data are presented in Fig. 13 for reference.

To further demonstrate the capabilities of our approach, we now consider a fully connected QUBO problem consisting of four weighted logical nodes $\{h_0, h_1, h_2, h_3\} = \{-0.4, -0.3, -0.375, -0.25\}$ linked by six weighted edges $\{J_{01}, J_{02}, J_{03}, J_{12}, J_{13}, J_{23}\} = \{0.275, 0.2, 0.175, 0.325, 0.2, 0.225\}$, as illustrated in Fig. 6(a). This densely




FIG. 6. Solving a fully connected QUBO problem via UDG-MWIS mapping. (a) A QUBO instance with four nodes and all-to-all connectivity is shown. (b) This problem is mapped to a UDG-MWIS using quantum wires and a modified crossing gadget. The resulting system exhibits six degenerate ground states with distinct wire configurations but identical QUBO solutions. (c) Ground state relative frequency versus average MWIS cost observed during annealing parameter optimization. (d) Histogram showing the raw frequency of all the observed states out of 728 (out of 2000) valid repetitions of a representative annealing run. (e) Histogram for the relative frequencies and costs of the 20 most observed states after applying both vertex reduction and addition to the data. The two most likely configurations are ground states (red) and can be mapped back to the solution of the all-to-all QUBO problem. (f) Histograms showing the raw frequency and cost of all 16 possible states for the four logical sites enumerated in (b).

connected graph is mapped onto a UDG-MWIS embedding using a combination of constructions from Figs. 2(a) and 2(d). The resulting layout requires 28 atoms: 4 for the logical vertices and 24 auxiliary atoms forming both standard quantum wires and a crossing gadget for overlapping connections. Figure 6(b) shows the UDG-MWIS embedding used to enable graph optimization on the QPU.

We applied the closed-loop optimization protocol described earlier to refine the annealing ramp parameters, as shown in Fig. 6(c). Analysis of the measurement data through classical postprocessing revealed a pronounced anticorrelation between the ground state relative frequency and the average cost throughout the optimization cycle. In contrast, the raw data show consistently flat histograms with each experimental shot returning a single unique bitstring. Surprisingly, even for the dataset giving the highest postprocessed target state relative frequency [shown in Figs. 6(d) and 6(e)], we never observe any of the six possible target solutions for the full 28 qubit MWIS in the input data. Despite not returning a full ground state solution, tracing over the four logical qubits we measure the correct logical state with a relative frequency of 0.183 ± 0.01 from the raw input data [see Fig. 6(f)].

While this example is far from a demonstration of quantum advantage, as a classical solver could readily solve the underlying 4-bit QUBO, it is instructive to quantify the performance of the neutral atom QPU in providing MWIS solution samples to the postprocessing algorithm. We compare the QPU performance to an equivalent number of samples (728) chosen uniformly from the 2^{28} possible output strings as inputs to the classical postprocessing algorithm. Averaging over multiple repeats, we find that the QPU provides ground state solutions with a relative frequency of 0.11 ± 0.01 , while the random sampling provides only 0.02 ± 0.005 , showing the QPU is providing higher-quality samples despite experimental imperfections. As the time-to-solution scales as $-1/\log(1 - p_{\text{solution}})$ [18], the ability to postprocess the raw data to obtain solutions with higher p_{solution} enables us to recover useful sample data even where the full state is never perfectly prepared. See Appendix F for more details.

These results reinforce the relevance of classical postprocessing in extracting meaningful solutions from noisy quantum annealing outputs from the QPU at moderate system sizes, as have been previously used in other demonstrations of quantum annealing for atomic [16] and superconducting hardware platforms [38].

VII. CONCLUSION

In this paper, we proposed and experimentally demonstrated a novel quantum wire architecture to map certain hard combinatorial optimization problems on general graphs to a UDG-MWIS problem. This quantum wire encoding serves as a building block for solving optimization

problems on neutral atom quantum processors, allowing for an efficient encoding and greater flexibility in problem graphs that can be solved using this technology.

We demonstrated how our quantum wire scheme is able to efficiently encode MWIS problems beyond the UDG connectivity by mediating effective interactions between distant vertices. Furthermore, by carefully tuning the weights of the ancillary wire atoms, we are able to faithfully preserve, and tune, the entire logical spectrum of the target problem, which allows for successful native encodings of QUBO problems within the same architecture. Because we do not encode a single qubit in the entire wire, but rather we use it to facilitate long-range connections between distant nodes, we find that, relative to previous general schemes [28,31], our protocol requires substantially fewer resources.

Beyond the theoretical framework, we experimentally implement our method by solving non-UDG maximum weight independent set and QUBO problems via quantum annealing on neutral atom arrays. These results confirm the feasibility of our embedding strategy and its compatibility with current neutral atom hardware.

Finally, we note that while we discuss this embedding scheme from a physical implementation perspective, it can also be used as a subroutine for classical solvers to efficiently map non-unit-disk graphs into weighted UDG problem instances. An intriguing direction for future work is the optimization and automation of the embedding layout, as well as the potential integration of the wire architecture with existing toolkits for the reduction of MWIS problems and the identification of hard instances that would benefit from the implementation of quantum solutions [33,37].

ACKNOWLEDGMENTS

We thank Gillian Marshall and Kristina Bell from QinetiQ for useful discussions and provision of some of the example problems that underpin the results presented above. We also thank C. J. Picken for his technical support. This work was supported by EPSRC through the Prosperity Partnership *SQuAre* (EP/T005386/1) and Grants No. EP/Y005058/2 and No. EP/Z53318X/1, as well as Innovate UK through the Innovation Accelerator for Neutral Atom Quantum Optimization (Grant No. 10059444).

DATA AVAILABILITY

The data that support the findings of this article are openly available [39].

APPENDIX A: ROBUSTNESS OF THE WIRE ARCHITECTURE

While the proposed wire architecture is guaranteed to give the correct ground state solution, it is *a priori* not

clear how robust the annealing ramps and the embedding scheme are to experimental imperfections and uncertainties. In this Appendix, we separately consider two aspects of this—(1) how robust and scalable the annealing performance of the quantum wires is as the system sizes are increased, and (2) how resilient the embedding is to variations in the vertex weights.

1. The annealing protocol

As indicated by the results discussed above, the quantum wire approach presented here allows for a flexible and robust encoding of combinatorial optimization problems that are close to UDG connectivity. In order to better understand the reason for this robustness, we numerically compute the spectral gap of the basic building blocks in this architecture as a function of system size.

The two leftmost panels of Fig. 7 show the spectral gap $\Delta E(t)$ of the Rydberg Hamiltonian Eq. (2) along the annealing path for logical wire weightings $(\alpha, \beta) = \{(0.7, 0.3), (0.51, 0.5)\}$ in (a) and (b), respectively. The different shadings indicate the system size (logical and ancilla atoms) of the wire $N = 2^k$ for $2 \leq k \leq 9$. The minimal spectral gap $\Delta E_{\min}$ along the path is marked by

black crosses for each system size. We used the same annealing schedule for all considered logical weights. It has been variationally optimized for the smallest $N = 4$ wire with logical weights $(0.7, 0.3)$ to isolate the effect of the wire length on the spectral gap. Using the matrix product state ansatz, we variationally find the spectral gap using the density matrix renormalization group (DMRG) algorithm [40–44]. Figure 7(c) shows the minimum spectral gap (as indicated by the black markers in the left plots) as a function of system size N for different wire weights. For the wires, we find that the spectral gap closes polynomially as N^{-z} with $z \lesssim 1$, indicating that these wires do not pose a fundamental problem for the adiabatic protocol. These findings are also in agreement with the recent analysis performed in Ref. [29].

The black markers in Figs. 7(a) and 7(b) indicate where the minimal spectral gap occurs along the annealing path. We have chosen two distinct scenarios of a strongly imbalanced (a) and a closely balanced (b) wire. We already saw above that the spectral gap closes more slowly (smaller z) for the former case (a). Interestingly, in this case the minimum gap appears to occur around the same detuning value, independent of the wire length. Meanwhile, for the closely balanced wire (b), the gap closes




FIG. 7. The left two plots show the spectral gap along the annealing path $\Delta E(t)$ for different system sizes $N = \{4, \dots, 512\}$ (light to dark blue) for wire weights $(\alpha, \beta) = (0.7, 0.3)$ (a) and $(\alpha, \beta) = (0.51, 0.5)$ (b). The black markers indicate the extracted minimum spectral gap depicted in the main plot. (c) Minimal spectral gap $\Delta E_{\min}$ as a function of system size for wire weightings $(\alpha, \beta) = \{(0.1, 0.9), (0.7, 0.3), (0.51, 0.5)\}$ (blue circle, orange inverted triangle, and green triangle, respectively). The spectral gap closes polynomially with system size approximately scaling as N^{-z} , with a linear fit giving an exponent of $z \lesssim 1$. These results were obtained using DMRG simulations using a bond dimension of $D = 200$, and a truncation error of $\epsilon = 10^{-12}$. The extracted energies were converged to a threshold of 10^{-4} . (d) Spectral gap along annealing path for the wire in (a) embedded in a 1D weighted logical graph of $N_{\log} = 9$ nodes for varying lengths L of the wire (note $L = 0$ shows the gap of just the logical graph). For clarity, only a subset of the data points are indicated by the markers. (e) Minimal gap along the annealing path for different system sizes extracted from (d). Here too, the minimal spectral gap of a logical problem embedded using a quantum wire closes polynomially as approximately N^{-z} as the wire length is increased. The full system (logical graph and quantum wire) is depicted in the inset for a wire length of $L = 4$. These results were obtained using exact diagonalization, and a linear fit to the minimal spectral gap yields an exponent of $z \sim 1.25$.

more quickly with system size, and there is a larger variation in the place where the minimal gap occurs along the annealing path. This is intuitively expected as the MWIS problem is computationally more difficult if the weights are all very similar due to smaller energy gaps between the different configurations. We obtain qualitatively very similar results for the triangle and square constructions of Figs. 2(b) and 2(c) with a scaling exponent of $z \lesssim 1.5$.

Finally, in order to study how these quantum wires perform in a real logical graph, we choose a simple one-dimensional MWIS problem of $N_{\text{log}} = 9$ nodes separated by a lattice spacing of $d = 8 \mu\text{m}$. The logical weights of the problem are given by the vector $\mathbf{w} = \{0.22, 0.87, 0.21, 0.7, 0.3, 0.92, 0.49, 0.61, 0.77\}$, and we used the quantum wire of Fig. 7(a) to connect the leftmost four to the rightmost five logical nodes. Figure 7(d) shows the spectral gap along the annealing path (same annealing schedule as used in the previous simulations) for varying lengths of the quantum wire $L = 0, \dots, 12$. Extracting the minimal spectral gap as above, we can see in (e) that as we make the wire longer, the spectral gap of the full system still closes polynomially with system size $N = N_{\text{log}} + L$, with a scaling exponent of $z \sim 1.25$. We are therefore encouraged that the quantum wires will not lead to adiabatic bottlenecks when annealing real graphs, due to the benign closing of their spectral gap.

2. Robustness of a single quantum wire

In order to evaluate how robust the embedding scheme is to shot-to-shot fluctuations in the applied local light shifts, we isolate a single wire module and vary the local detunings randomly. For concreteness, we draw the random weights from a normal distribution centered around the ideal detuning value, and study how these deviations from the ideal implementation can affect the low-energy spectrum of the quantum wire. Let X be a random variable $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$, where $\mu_X = \{\alpha, \beta, c\}$ and $\sigma_X = \{\sigma_\alpha, \sigma_\beta, \sigma_c\}$, respectively. The configurational energies of the wire are thus also random variables, and specifically the distributions of the low-energy states shown in Fig. 2(a) are given by $E_{x_\alpha, x_\beta} \sim \mathcal{N}(\mu_{x_\alpha, x_\beta}, \sigma_{x_\alpha, x_\beta}^2)$, with

$$\mu_{x_\alpha, x_\beta} = -\alpha x_\alpha - \beta x_\beta - c N_{x_\alpha, x_\beta}, \quad (\text{A1})$$

$$\sigma_{x_\alpha, x_\beta}^2 = x_\alpha \sigma_\alpha^2 + x_\beta \sigma_\beta^2 + N_{x_\alpha, x_\beta} \sigma_c^2, \quad (\text{A2})$$

their mean and variance, respectively. Here $x_i \in \{0, 1\}$ indicates the absence (presence) of an excitation on either end of the wire (α or β , respectively), and $N_{x_\alpha, x_\beta} = (L - 2x_\alpha x_\beta)/2$ is the total number of excitations on the ancilla atoms for the considered states. The expectation values of the energies coincide with the ideal wire result as expected, and we note that fluctuations in the detunings

lead to a standard deviation $\sigma = \sqrt{\sigma_\alpha^2 + \sigma_\beta^2 + L\sigma_c^2}$, which scales as $\sigma \sim \sqrt{L}$ with increasing wire length. We also note that this analysis holds both considering only UDG interactions and vdW tails, as in the latter case the long-range interactions are identical in the lowest two energy configurations.

We test empirically the robustness of a single wire to noise in the local light shifts by setting the ground state with $\alpha > \beta$, randomizing the weights of the logical nodes and the wire atoms in the manner discussed above and checking whether the new ground state still corresponds to the $|10\rangle$ state. To speed up the calculations, we only evaluate the energies of the $L + 3$ configurations of the system formed by the logical nodes and the wire that respect the blockade constraint. In Figs. 8(a) and 8(b), we show, respectively, the results corresponding to $\alpha = 0.6, \beta = 0.4$ and $\alpha = 0.8, \beta = 0.2$ for wires of different lengths L and for different values of the standard deviation σ_X of the noise Gaussian, which for each atom is chosen such that σ_X/μ_X is constant across the system. Each data point has been obtained by sampling 100,000 random local detuning realizations. The success probability is larger for the case $\alpha = 0.8, \beta = 0.2$ shown in Fig. 8(b) than for the case $\alpha = 0.6, \beta = 0.4$ of Fig. 8(a). This is to be expected from Eq. (A1), as the expectation value of the gap between the $|10\rangle$ and the rest of states is larger for $\alpha = 0.8, \beta = 0.2$ than for $\alpha = 0.6, \beta = 0.4$. Even though there is a decay in the probability of successfully encoding the solution of the MWIS problem with increasing values of the standard




FIG. 8. Probability of successfully encoding the ground state in the presence of Gaussian fluctuations in the local light shift for wires of different length L as a function of the standard deviation σ_X of the distribution. Plot (a) corresponds to the choice of weights $\alpha = 0.6, \beta = 0.4$, and (b) $\alpha = 0.8, \beta = 0.2$. Each data point has been estimated by sampling across 100,000 random realizations of the local light shifts.

deviation σ_X and the wire length L , for lengths of up to $L = 64$ and standard deviation $\sigma_X/\mu_X = 0.05$, the success probability remains lower bounded by approximately 70% in the case $\alpha = 0.6, \beta = 0.4$ and by approximately 92% for $\alpha = 0.8, \beta = 0.2$. When the wire is embedded in a larger graph, we expect the robustness of the encoding to be enhanced by the additional gaps induced by the structure of the logical problem and the presence of residual long-range vdW interactions.

3. Robustness of the crossing gadget

To complete the analysis of the robustness of the different building blocks involved in our architecture, we also test the performance of the crossing gadget in the presence of random shot-to-shot fluctuations in the values of the weights and compare it with that of the quantum wires. As before, we sample the weights $\{\alpha, \beta, \gamma, \delta, c\}$ of the crossing gadget shown in Fig. 2(d) from normal distributions $\mathcal{N}(\mu_X, \sigma_X^2)$ such that the ratio σ_X/μ_X is constant across the system, and we study how likely is the crossing gadget to encode the correct ground state as a function of the standard deviation σ_X . In order to make a fair comparison with the robustness of the wires, we have picked different examples of logical weight configurations and compared with wires of length $L = 2$ and $L = 4$ with logical weights such that the gap of the logical states of the wires is the same as for the states of the gadget. The reason to choose these wire lengths is that the crossing gadget has four ancillary atoms but only one of them is excited in the ground state, so it makes sense to compare the robustness both with a wire with the same number of ancillary atoms ($L = 4$) and the same number of ancillary excitations ($L = 2$).

The results of this analysis are summarized in Fig. 9. We have chosen different values of the weights intervening in the crossing gadget such that in all cases $c = \alpha + \beta + \gamma + \delta = 1$, and that the gap between the logical ground and first excited states takes the values $\Delta E = 0.3$ [Fig. 9(a)], $\Delta E = 0.2$ [Fig. 9(b)], and $\Delta E = 0.1$ [Fig. 9(c)]. We have checked independently that the robustness of the gadget depends only on the logical gap and not the particular choice of the weights. For each of these crossing gadget weight instances, we have chosen values of the wire logical weights such that $c = \alpha + \beta = 1$ and that the gap matches that of the logical states of the crossing gadget. For energy gaps $\Delta E = 0.3, \Delta E = 0.2$, we observe that the crossing gadget is as robust to fluctuations in the weights as a wire of length $L = 2$. For $\Delta E = 0.1$, the robustness of the crossing gadgets deviates slightly from that of the $L = 2$ wire for higher values of the variance, but it still remains higher than the robustness of wire of length $L = 4$ with the same spectral gap. Overall, this analysis allows us to conclude that the crossing gadget has a robustness to random shot-to-shot fluctuations in the




FIG. 9. Probability of successfully encoding the ground state in the presence of Gaussian fluctuations in the local light shift for a crossing gadget and wires of length $L = 2, 4$ as a function of the standard deviation σ_X of the distribution. Plot (a) corresponds to the choice of weights $\{\alpha = \gamma = 0.4, \beta = \delta = 0.1\}$ for the gadget and $\{\alpha = 0.65, \beta = 0.35\}$ for the wires, (b) corresponds to the choice of weights $\{\alpha = 0.4, \gamma = 0.3, \beta = 0.2, \delta = 0.1\}$ for the gadget and $\{\alpha = 0.6, \beta = 0.4\}$ for the wires, and (c) corresponds to the choice of weights $\{\alpha = \gamma = 0.3, \beta = \delta = 0.2\}$ for the gadget and $\{\alpha = 0.55, \beta = 0.45\}$ for the wires. Each data point has been estimated by sampling across 100,000 random realizations of the weights.

values of the detunings comparable to that of wires of length at most $L = 4$.

APPENDIX B: EXPERIMENTAL SETUP

Our experimental system is shown schematically in Fig. 10 and described in detail in Ref. [27]. An array of cesium (Cs) atoms is confined in holographically generated optical tweezers (1064 nm). To implement MWIS annealing, we combine global annealing protocols with the ability to implement arbitrary local light shifts using an auxiliary spatial light modulator. Once the problem is mapped to UDG-MWIS, atoms are arranged spatially to represent the graph problem, where edge-connected qubits are placed within a blockade radius of each other to ensure the strong interactions prevent simultaneous excitation of both qubits.

Atoms are initially prepared in the hyperfine ground state $|g\rangle = |4, 0\rangle$ and coherently transferred to the Rydberg




FIG. 10. Overview of experimental setup. The tweezer trap array is formed using an SLM to generate holographic tweezers using 1064 nm light, with detection performed using an sCMOS camera. Rydberg excitation is performed using global beams at 459 and 1039 nm, with uniform intensity created using a top-hat beam shaper and an SLM, respectively. Local light shifts are generated using an 800 nm beam imaged off an SLM to enable site-dependent weights.

state $|r\rangle = |80s_{1/2}, m_j = +1/2\rangle$ using a two-photon excitation scheme via the $7p_{1/2}$ intermediate state. The excitation is driven by counterpropagating laser beams at 459 and 1039 nm, each shaped into flat-top intensity profiles to enhance uniformity across the atomic array. The 459 nm beam is flattened using a commercial beam-shaping optic (Osela DTH-1D-250 μm -40 μm -3.75 mm), while the 1039 nm beam is shaped using an SLM [45].

The excitation lasers are detuned from the intermediate state by $\Delta_{\text{int}}/2\pi = 1$ GHz, enabling Rydberg coupling with a measured two-photon Rabi frequency of $\Omega/2\pi = 1.0(5)$ MHz. Fine control of the global detuning is achieved through frequency modulation of the 459 nm beam using an AOM in a double-pass configuration. For the $80s_{1/2}$ Rydberg level, the van der Waals interaction strength is characterized by a coefficient $C_6 = -3376 \text{ GHz } \mu\text{m}^6$ [46], which corresponds to a resonant blockade radius of $r_B = \sqrt[6]{|C_6|/\Omega} = 12.1 \text{ } \mu\text{m}$. To suppress atom loss and decoherence, the optical tweezers are turned off throughout the annealing sequence.

An additional light-shift potential introduces a site-dependent offset given by $\delta_{\text{ac}}^i = w_i \delta_{\text{ac}}$, where w_i denotes the relative weight assigned to site i . The global term δ_{ac} is governed by the total intensity of the light-shift laser, which is dynamically modulated using an AOM before reaching the SLM. This configuration enables the implementation of continuous annealing ramps while

maintaining fixed relative weights throughout the experiment. These ac Stark shifts are realized using 800 nm light from a Ti:Sapph laser. Arrays of focused light-shift spots are projected onto atoms with a measured $1/e^2$ waist of $3.0(2) \text{ } \mu\text{m}$.

For each instance of the UDG-MWIS problem, the intensity-weighted light-shift patterns are generated using a modified weighted-Gerchberg-Saxton phase retrieval algorithm, allowing precise control over the relative amplitudes of the individual sites [47]. Calibration of the resulting Stark shifts is performed via spectroscopy on the $|50s_{1/2}, m_j = +1/2\rangle$ Rydberg state, which is chosen to suppress interatomic interaction effects due to its lower principal quantum number. Measured shifts are used in a closed-loop feedback routine to iteratively update the SLM phase pattern until the rms error in the target intensity distribution converges to a value below a certain threshold, typically $< 7\%$.

To quantify the effect of atom losses, we perform a series of baseline survival measurements. The first measurement is to run the system with no 800 nm or Rydberg lasers active, from which we find a single atom survival of $99(1)\%$. When adding the 800 nm light and running annealing ramps with no Rydberg light, this reduces to $p = 98.5(1)\%$. Repeating the ramp with both 800 nm light and either the 459 or 1039 nm beams on, we see no change in atom loss, recovering $p = 98.5(2)\%$ averaged over graph arrays with up to $N = 28$ atoms.

APPENDIX C: GRAPH WEIGHTS

Figure 11 presents the weights used for all the UDG-MWIS problems investigated in the main text.

APPENDIX D: CLASSICAL POSTPROCESSING

These experiments have identified challenges in scaling up to larger problem sizes, namely minimizing the effects of ground state loss, which act as a false positive for Rydberg excitation (and hence a higher chance of identifying configurations as high-cost blockade violations), and verification of interaction strengths as a function of array position to ensure the array spacings of the encoded problems are within the blockade radius.

In the context of a many-body quantum simulation experiment, these detrimental effects are significant; however, in the context of using the atom array as an annealer for classical optimization, it remains possible to take output bitstrings featuring edge violations, and perform classical postprocessing to recover valid graph solutions. To explore the impact of this, we implement the same postprocessing algorithms introduced by Ref. [16] in the context of solving MIS problems, and apply these retroactively to our MWIS datasets.

The classical postprocessing step acts to improve the quality of the final solutions with respect to the raw




FIG. 11. Weighted graphs. (a) Graph of Fig. 4(b). (b) QUBO problem with weighted nodes and connections of Fig. 5(a). (c) QUBO with all-to-all connectivity of Fig. 6(b).

measured strings. It provides a simple mechanism to, first, ensure that all strings become valid independent sets (by removing edge violations caused by neighboring Rydberg states through vertex reduction) and, second, maximize the allowed number of excitations (thereby lowering the cost) through vertex addition. The vertex reduction step can resolve errors caused by failures of adiabaticity (meaning preparation of higher energy states which violate the IS condition) and atom losses, while the vertex addition step can reduce errors due to imperfect state preparation or imperfect Rydberg detection that result in higher ground state detection than expected.

1. Vertex reduction

For experimental bitstrings featuring edge (or equivalently blockade) violations, whereby pairs of edge-connected vertices are both selected, we can recover valid independent set bitstrings by applying *vertex reduction*. This algorithm is implemented by first counting the number of edge violations on each vertex, and selecting the vertex with the highest number to be flipped from 1 to 0. Where many vertices all have the maximum number of edge violations, the vertex to be flipped is chosen at random. This process is repeated until all edge violations are removed, ensuring the bitstring represents an IS. An example is shown in Fig. 12, where starting with the worst possible bitstring in (a) (all vertices selected), the edge violations are sequentially removed to obtain a valid IS bitstring shown in (b).

2. Vertex addition

Following the application of vertex reduction to the MWIS bitstrings we recover a valid IS but it may not be maximal, meaning it may be possible to select one or more additional vertices to recover a lower cost solution. A *vertex addition* algorithm is implemented using a constant overhead greedy algorithm, whereby for a fixed number of iterations (in this case 10 as used in Ref. [16]) the vertices are selected in a random order, and if a vertex is in state 0 it




FIG. 12. Postprocessing algorithms. (a) Input bitstring data containing edge (blockade) violations. (b) Vertex reduction algorithm removes edge violations and recovers a valid independent set (IS) bitstring. (c) The vertex reduced bitstring may be an IS but have the option of adding additional vertices. Vertex addition uses a constant overhead greedy algorithm to randomly select vertices to convert from 0 to 1 if no edge violation occurs, and then returns the lowest cost solutions.

is transformed to 1 if it would not cause an edge violation. The lowest cost solution is then returned. An example is shown in Fig. 12(c) where the algorithm has added an additional vertex to achieve a lower cost.

APPENDIX E: VERTEX ADDITION DATA FOR THE BLOCK-WISE QUBO PROBLEMS

Figures 13(c)–13(g) present the histograms obtained after postprocessing the data from Figs. 5(c)–5(g), illustrating how classical postprocessing improves the relative frequency of configurations corresponding to the MWIS solution.




FIG. 13. Block-wise decomposition of a QUBO problem with postprocessing. (a) A QUBO problem consisting of eight nodes and ten edges is shown. (b) This problem is mapped to a UDG formulation of the MWIS, enabling native implementation on a Rydberg quantum processing unit. The QUBO graph can be partitioned into blocks of varying sizes by selectively removing edges. (c)–(e) Annealing results after vertex reduction and addition for each individual block derived from (a), shown from left to right. Block (e) exhibits two degenerate ground state solutions. Data shown for 870, 763, and 749 defect-free realizations, respectively, out of 1000 trials. (f) MWIS results corresponding to the QUBO problem composed of the two leftmost blocks (664/1000 defect-free realizations). (g) MWIS results for the full three-block QUBO problem for 669 defect-free realizations. In (c)–(g), the ground-state configurations are highlighted in red.




FIG. 14. Performance comparison between QPU-generated and randomly selected bitstring inputs. (a) Mean ground state relative frequencies computed over 100 independent runs of classical postprocessing, using 728 bitstrings obtained from the QPU and 728 randomly sampled bitstrings. (b) Aggregated histograms from 100 repetitions of the numerical experiment, showing the distribution of relative frequencies and costs for the most frequent states for random (top) and QPU-derived bitstrings (bottom).

APPENDIX F: QPU AS A SAMPLER

As a way to evaluate the QPU's performance as a sampler despite noise, we compare its output to randomly generated bitstrings for the graph of Fig. 6(b). We collect 728 QPU-derived bitstrings from a fixed annealing ramp and generate an equal number of random bitstring samples. Each set is then used as input for 100 independent runs of the classical postprocessing algorithm. As illustrated in Fig. 14(a), the QPU samples exhibit a notably higher mean ground state relative frequency. The QPU's advantage is reinforced by the histograms in Fig. 14(b), which show a clear bias toward ground state configurations in the QPU data.

- [1] V. T. Paschos, *Applications of Combinatorial Optimization* (John Wiley & Sons, New York, 2014), Vol. 3.
- [2] A. Abbas *et al.*, Challenges and opportunities in quantum optimization, *Nat. Rev. Phys.* **6**, 718 (2024).
- [3] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, *Rep. Progress Phys.* **83**, 054401 (2020).

- [4] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, Quantum annealing: An overview, *Philos. Trans. R. Soc. A* **381**, 20210417 (2023).
- [5] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [6] N. Astrakhantsev, G. Mazzola, I. Tavernelli, and G. Carleo, Phenomenological theory of variational quantum ground-state preparation, *Phys. Rev. Res.* **5**, 033225 (2023).
- [7] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [8] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature (London)* **549**, 242 (2017).
- [9] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum chemistry calculations on a trapped-ion quantum simulator, *Phys. Rev. X* **8**, 031022 (2018).
- [10] C. Kokail, C. Maier, R. van Bijnen, T. Brydges, M. K. Joshi, P. Jurcevic, C. A. Muschik, P. Silvi, R. Blatt, C. F. Roos, and P. Zoller, Self-verifying variational quantum simulation of lattice models, *Nature (London)* **569**, 355 (2019).
- [11] A. Lucas, Ising formulations of many np problems, *Front. Phys.* **2**, 5 (2014).
- [12] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).
- [13] M. Morgado and S. Whitlock, Quantum simulation and computing with rydberg-interacting qubits, *AVS Quantum Sci.* **3**, 023501 (2021).
- [14] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, Quantum optimization for maximum independent set using Rydberg atom arrays, [arXiv:1808.10816](https://arxiv.org/abs/1808.10816).
- [15] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, Computational complexity of the Rydberg blockade in two dimensions, [arXiv:1809.04954](https://arxiv.org/abs/1809.04954).
- [16] S. Ebadi *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, *Science* **376**, 1209 (2022).
- [17] K. Kim, M. Kim, J. Park, A. Byun, and J. Ahn, Quantum computing dataset of maximum independent set problem on king lattice of over hundred Rydberg atoms, *Sci. Data* **11**, 111 (2024).
- [18] R. S. Andrist, M. J. A. Schuetz, P. Minssen, R. Yalovetzky, S. Chakrabarti, D. Herman, N. Kumar, G. Salton, R. Shaydulin, Y. Sun, M. Pistoia, and H. G. Katzgraber, Hardness of the maximum-independent-set problem on unit-disk graphs and prospects for quantum speedups, *Phys. Rev. Res.* **5**, 043277 (2023).
- [19] M. Kim, K. Kim, J. Hwang, E.-G. Moon, and J. Ahn, Rydberg quantum wires for maximum independent set problems, *Nat. Phys.* **18**, 755 (2022).
- [20] A. Byun, M. Kim, and J. Ahn, Finding the maximum independent sets of platonic graphs using rydberg atoms, *PRX Quantum* **3**, 030305 (2022).
- [21] C. Dalyac, L.-P. Henry, M. Kim, J. Ahn, and L. Henriët, Exploring the impact of graph locality for the resolution of mis with neutral atom devices, [arXiv:2306.13373](https://arxiv.org/abs/2306.13373).
- [22] A. Byun, J. Jung, K. Kim, M. Kim, S. Jeong, H. Jeong, and J. Ahn, Rydberg-atom graphs for quadratic unconstrained binary optimization problems, *Adv. Quantum Technol.* **7**, 2300398 (2024).
- [23] A. Byun, S. Jeong, and J. Ahn, Programming higher-order interactions of rydberg atoms, *Phys. Rev. A* **110**, 042612 (2024).
- [24] S. Jeong, M. Kim, M. Hhan, J. Y. Park, and J. Ahn, Quantum programming of the satisfiability problem with rydberg atom graphs, *Phys. Rev. Res.* **5**, 043037 (2023).
- [25] M. Lanthaler, K. Ender, C. Dłaska, and W. Lechner, Quantum optimization with globally driven neutral atom arrays, [arXiv:2410.03902](https://arxiv.org/abs/2410.03902).
- [26] K. Goswami, R. Mukherjee, H. Ott, and P. Schmelcher, Solving optimization problems with local light-shift encoding on Rydberg quantum annealers, *Phys. Rev. Res.* **6**, 023031 (2024).
- [27] A. G. de Oliveira, E. Diamond-Hitchcock, D. M. Walker, M. T. Wells-Pestell, G. Pelegrí, C. J. Picken, G. P. A. Malcolm, A. J. Daley, J. Bass, and J. D. Pritchard, Demonstration of weighted-graph optimization on a rydberg-atom array using local light shifts, *PRX Quantum* **6**, 010301 (2025).
- [28] M.-T. Nguyen, J.-G. Liu, J. Wurtz, M. D. Lukin, S.-T. Wang, and H. Pichler, Quantum optimization with arbitrary connectivity using Rydberg atom arrays, *PRX Quantum* **4**, 010316 (2023).
- [29] L. Bombieri, Z. Zeng, R. Tricarico, R. Lin, S. Notarnicola, M. Cain, M. D. Lukin, and H. Pichler, Quantum adiabatic optimization with Rydberg arrays: Localization phenomena and encoding strategies, *PRX Quantum* **6**, 020306 (2025).
- [30] W. Lechner, P. Hauke, and P. Zoller, A quantum annealing architecture with all-to-all connectivity from local interactions, *Sci. Adv.* **1**, e1500838 (2015).
- [31] M. Lanthaler, C. Dłaska, K. Ender, and W. Lechner, Rydberg-blockade-based parity quantum optimization, *Phys. Rev. Lett.* **130**, 220601 (2023).
- [32] J. Park, S. Jeong, M. Kim, K. Kim, A. Byun, L. Vignoli, L.-P. Henry, L. Henriët, and J. Ahn, Rydberg-atom experiment for the integer factorization problem, *Phys. Rev. Res.* **6**, 023241 (2024).
- [33] M. J. A. Schuetz, R. S. Andrist, G. Salton, R. Yalovetzky, R. Raymond, Y. Sun, A. Acharya, S. Chakrabarti, M. Pistoia, and H. G. Katzgraber, Quantum compilation toolkit for Rydberg atom arrays with implications for problem hardness and quantum speedups, *Phys. Rev. Res.* **7**, 033107 (2025).
- [34] X. Qiu, P. Zoller, and X. Li, Programmable quantum annealing architectures with ising quantum wires, *PRX Quantum* **1**, 020311 (2020).
- [35] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, The unconstrained binary quadratic programming problem: A survey, *J. Comb. Optimiz.* **28**, 58 (2014).
- [36] QuEra Computing, UnitDiskMapping.jl (2025), <https://github.com/QuEraComputing/UnitDiskMapping.jl>.
- [37] M. J. A. Schuetz, R. Yalovetzky, R. S. Andrist, G. Salton, Y. Sun, R. Raymond, S. Chakrabarti, A. Acharya,

- R. Shaydulin, M. Pistoia, and H. G. Katzgraber, qReduMIS: A quantum-informed reduction algorithm for the maximum independent set problem, [arXiv:2503.12551](#).
- [38] T. Koch *et al.*, Quantum optimization benchmarking library—The intractable decathlon, [arXiv:2504.03832](#).
- [39] A. Oliveira, J. Kombe, G. Pelegri, P. Schroff, M. Wells-Pestell, D. Walker, A. Daley, and J. Pritchard, Data for “Sparse Graph Optimization using Weighted Quantum Wires in Rydberg Atom Arrays”, [10.15129/e3899a58-be72-4ed8-b033-32c52b1414e7](#) (2026).
- [40] S.R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [41] S.R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [42] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (Amsterdam)* **326**, 96 (2011).
- [43] M. Fishman, S.R. White, and E.M. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* **4** (2022).
- [44] M. Fishman, S.R. White, and E.M. Stoudenmire, Codebase release 0.3 for ITensor, *SciPost Phys. Codebases* **4** (2022).
- [45] P. Schroff, A. La Rooij, E. Haller, and S. Kuhr, Accurate holographic light potentials using pixel crosstalk modeling, *Sci. Rep.* **13**, 3252 (2023).
- [46] N. Sibalić, J.D. Pritchard, C.S. Adams, and K.J. Weatherill, ARC: An open-source library for calculating properties of alkali Rydberg atoms, *Comp. Phys. Comm.* **220**, 319 (2017).
- [47] D. Kim, A. Keesling, A. Omran, H. Levine, H. Bernien, M. Greiner, M.D. Lukin, and D.R. Englund, Large-scale uniform optical focus array generation with a phase spatial light modulator, *Opt. Lett.* **44**, 3178 (2019).