Reducibility of native weighted graphs on Rydberg arrays

J. Kombe[*] and J. D. Pritchard[†]

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

(Received 10 May 2026; revised 25 September 2026; accepted 9 November 2026; published 20 November 2026)

We investigate the classical reducibility of random unit-disk graph (UDG) instances of the maximum independent set and maximum weighted independent set problems, which can be natively realized in Rydberg atom quantum processors. Using state-of-the-art kernelization techniques, we systematically probe how far classical preprocessing can simplify such native optimization problems of varying size and connectivity. While many small or sparse instances can be fully reduced, dense graphs often retain finite irreducible kernels even after extensive reductions. Introducing vertex weights tends to increase reducibility, whereas extending the interaction range in the underlying UDG connectivity suppresses the reduction efficiency. By exploring where classical reductions cease to be effective, we aim to delineate the regime of problem instances that remain computationally demanding—those most relevant for testing and benchmarking near-term quantum optimization hardware. We find that for the remaining finite kernels, quantum execution would require non-native embeddings with substantial resource overheads, suggesting that directly running native instances may be more practical than embedding a reduced kernel.

DOI: 10.1103/vbpq-f9qf

I. INTRODUCTION

Combinatorial optimization problems lie at the core of computer science, mathematics, and operations research, encompassing a vast range of applications from logistics and network design to scheduling, finance, and drug discovery [1]. Since Karp's seminal compilation of 21 NP-complete problems [2], it has been known that many such problems are computationally intractable in the worst case, with solution times that scale exponentially with system size [3,4]. Despite decades of progress in heuristic and approximation algorithms, the intrinsic hardness of combinatorial optimization problems continues to motivate the search for alternative computational paradigms [5–7].

Quantum optimization offers a paradigm of solving certain problem classes through, for example, quantum annealing [8–11], adiabatic evolution [5,6], or quantum-classical hybrid approaches [12,13] such as the variational quantum eigensolver [14–17] or quantum approximate optimization algorithms [18–22]. Problems that efficiently map onto Ising spin models are particularly suitable for these quantum optimization methods, with examples including the maximum (weighted) independent set (MIS/MWIS) and quadratic unconstrained binary optimization (QUBO) problems [23]. In all these approaches, the solution to the problem is encoded into the ground state of quantum many-body Hamiltonian [24,25].

Quantum annealing of Ising spin models has been demonstrated using superconducting circuit systems, which offer tunable couplings and fast annealing timescales, suitable for solving QUBO problems [26,27] or quantum simulation [28]. Open challenges remain around the optimal choice of coupling graphs [29,30], embedding strategies [31–34], and scalability in achieving robust demonstrations of quantum advantage for different classes of optimization problems [35,36].

Among competing architectures, neutral-atom quantum processors based on Rydberg atom arrays have recently emerged as promising candidates for combinatorial optimization tasks [37,38], combining large-scale, programmable qubit registers with tunable long-range interactions, dynamical reconfigurability [39–43], and local addressability [44]. Operating in the Rydberg blockade regime where any pair of atoms separated by less than the blockade radius R_b cannot be simultaneously excited to the Rydberg manifold [45–47], these platforms naturally enforce the independence constraint of certain graph optimization problems, most notably for MIS problems on a restricted class of geometric graphs known as unit disk graphs (UDGs) [48–51]. Furthermore, local control fields [52] realizing local light shifts permit embeddings of

[*]Contact author: johannes.kombe@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.

maximum weighted independent set (MWIS) problems defined on UDGs [44,53].

A fundamental challenge of this approach is that Rydberg interactions generate only UDG connectivity, constrained by the two-dimensional geometry of the planar atomic array, while real-world optimization problems typically correspond to non-UDG graphs with arbitrary topologies. This has motivated a wide range of encoding strategies for embedding arbitrary problems onto UDG graphs with nearest-neighbor or next-nearest neighbor coupling [54–63]. However, while these schemes differ in physical details, they share a common limitation: embedding graphs with N logical variables generally results in embeddings with $\mathcal{O}(N^2)$ physical atoms [61,64], and despite recent attempts to further optimize embedded layouts [65], this quadratic overhead remains a significant obstacle for implementing large-scale combinatorial optimization on near-term neutral-atom platforms.

It is therefore interesting and timely to investigate natively embeddable graphs more closely. To date, most large-scale experimental studies of MIS on Rydberg platforms with several hundred atoms have focused on native unit-disk instances on a King's lattice (a square lattice with edge couplings to nearest neighbors along horizontal and diagonal paths) [50,51]. A subsequent analysis, however, demonstrated that King's lattice UDG MIS instances can be solved to optimality at much larger scales (reaching thousands of nodes) by using classical algorithms on commodity computing hardware, and without relying on instance-specific fine-tuning [66].

Recent works extending these findings introduce hybrid classical-quantum workflows combining graph reduction, embedding, and compilation for Rydberg-based optimization [65,67]. Graph reduction is a classical preprocessing step that acts to simplify a target problem graph by sequential application of reduction rules to create a subset of smaller graph kernels which can be solved more efficiently than direct solution of the initial graph. Highly reducible graphs are deemed "easy" to solve, and in these works the King's lattice MIS instances were identified as being as algorithmically easy and efficiently reduced from thousands of nodes to smaller kernels with tens or hundreds of nodes. However, only a limited set of four basic reduction rules was employed in these studies, leaving open the question of how this "easy-hard" landscape might shift under state-of-the-art, graph-theoretic kernelization techniques [68–70].

In this work, we revisit and extend these studies with the aim of identifying the boundary between "easy" and "hard" instances for native Rydberg graphs, with the workflow schematically depicted in Fig. 1. Specifically, we investigate whether modern, more powerful reduction pipelines can render previously "hard" instances tractable, and how vertex weights (MWIS) and longer-range native

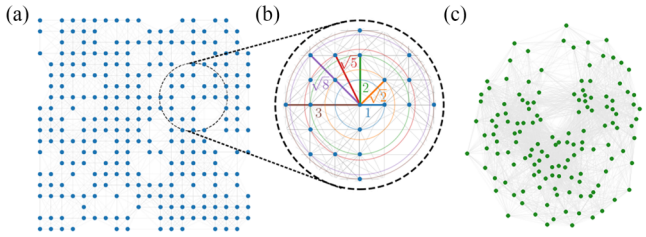


FIG. 1. Algorithmic pipeline of the `LearnAndReduce` framework for graphs native to the Rydberg array system. (a) An initial $L \times L$ grid of atoms is randomly populated to density ρ, after which native UDG edges are added as determined by the blockade radius R_b. The example shows $L = 20$, $\rho = 0.7$, $R_b = 3$. (b) The local vicinity of a node in the graph is shown along with a range of blockade radii, highlighting the emerging UDG connectivity. Scaling the blockade radius from $R_b = 1$ to 3 (see in different colors and associated blockaded circles) increases the maximum number of edge-connected neighbors from 4 to 28. (c) Each UDG graph instance is passed to `LearnAndReduce` which applies classical reductions to create an output kernel with fewer vertices. The kernels are no longer guaranteed to have unit-disk character, and tend to be very dense as demonstrated for the example graph from (a).

connectivities (corresponding to larger blockade radii) affect problem hardness. Rather than addressing the embedding challenge directly, we here focus on the native 2D UDG problem class, assuming either directly realizable graphs or ones that have already undergone embedding using existing schemes discussed above. We utilize a comprehensive open-source reduction framework [71] featuring over 20 reduction rules of varying complexity, and systematically quantify the reduction power and hardness scaling across MIS and MWIS instances for random Rydberg arrays of varying size and connectivity, as set by the blockade radius. By interpreting the reducibility as a quantitative measure of instance hardness, we establish a refined understanding of which native Rydberg problems are classically "easy," and crucially which remain intrinsically challenging—thereby delineating the most promising domains to focus further research efforts and for potential quantum advantage. We emphasize that this is not the same definition of problem hardness as the worst-case notions of complexity theory, but rather pragmatically motivated by and relevant for the context of current near-term quantum devices.

The remainder of this paper is organized as follows. In Sec. II, we review the background of the MIS and MWIS problems, and summarize the main graph reduction and kernelization techniques relevant to this study. Section III presents our results and discussion, analyzing reduction performance, hardness trends, and the effects of weights and blockade radius on instance complexity. Finally, we conclude in Sec. IV with a summary of our findings and an outlook on future directions.

II. BACKGROUND

A. MIS/MWIS

The MIS problem is a paradigmatic task in combinatorial optimization. Given a graph $G = (V, E)$ comprised of a set of vertices V, and a collection of edges E between pairs of vertices, an *independent set* is a subset $S \subseteq V$ such that no two vertices in S share an edge. The MIS problem seeks an independent set of maximum cardinality. More generally, in the MWIS problem, each vertex $i \in V$ carries a weight w_i, and the objective is to find an independent set S maximizing the weight $\sum_{i \in S} w_i$.

Both MIS and MWIS can be formulated as a classical minimization over binary variables $x_i \in \{0, 1\}$, where $x_i = 1$ indicates that vertex i is included in the candidate solution. The classical cost function is given by

$$C(\mathbf{x}) = -\sum_{i \in V} w_i x_i + U \sum_{(i,j) \in E} x_i x_j, \tag{1}$$

where the first term rewards selecting high, positive-weight vertices, while the second assigns a penalty $U > \max_i w_i > 0$ for any adjacent pair enforcing the independence constraint. For the unweighted MIS problem $w_i = 1$, and $U > 1$. Finding the optimal assignment $\mathbf{x}^* = \min_{\mathbf{x}} C(\mathbf{x})$ is NP-hard in general, even for the unweighted case [2,23,72,73], and can be natively realized on UDGs using Rydberg atom tweezer arrays [48,50]. Consequently, MIS and MWIS serve as canonical benchmarks for evaluating the performance of both classical and emerging neutral-atom quantum optimization approaches.

B. Classical graph reductions for MIS and MWIS

State-of-the-art classical approaches to the maximum (weighted) independent set (MIS/MWIS) problem make extensive use of *data reduction* or *kernelization* rules, which systematically reduce a graph instance while provably preserving optimality. Many reduction rules have been developed for MIS and MWIS problems and are employed in state-of-the-art algorithms [68–71,74–96] as a preprocessing step before running an exact and heuristic solver on the reduced graph, and subsequently lifted the optimal solution for the original graph from the solution of the reduced kernel.

1. Reduction philosophy

Given a graph $G = (V, E)$ with n_G vertices and vertex weights $w\colon V \to \mathbb{R}_{\geq 0}$, a reduction rule identifies a set of vertices whose membership in an optimal solution is fully determined by its local structure. Removing these vertices and adjusting the objective value (in the weighted case) yields a smaller problem instance $K = (V', E')$ with n_K vertices, termed the *kernel*. Applying all admissible reductions until convergence can often reduce real-world instances dramatically, and the kernel can then be solved by exact branching, branch-and-reduce, or local-search strategies, after which the removed parts of the solution are reconstructed.

We introduce the reducibility ξ which serves as a practical proxy for problem hardness, defined as

$$\xi = 1 - n_K/n_G, \tag{2}$$

where $n_{K/G}$ denotes the number of vertices in the original (G) and reduced kernel (K) graphs, respectively. Highly reducible graphs ($\xi \approx 1$) are deemed "easy" to solve, whereas graphs with $\xi \gtrsim 0$ remain "hard."

2. Learned reduction: KaMIS's `LearnAndReduce`

In Appendix A, we provide examples of some of the rules relevant for MIS and MWIS reductions. Although many reduction rules are inexpensive, several of the most powerful ones—such as advanced neighborhood-dominance checks, generalized twin rules, or deep folding cascades—can be costly to evaluate. The `LearnAndReduce` framework [71] incorporates machine learning to accelerate these expensive reductions. The approach trains a graph neural network (GNN) to predict the likelihood that a reduction rule will apply at a vertex. During kernelization, for each vertex v the GNN outputs a probability $p(v)$ indicating how promising v is for a computationally heavy rule. The solver then evaluates those costly reductions *only* on vertices with high $p(v)$, effectively steering the kernelization process. After reductions saturate, the resulting kernel is passed to a high-performance parallel local-search algorithm (e.g., Concurrent Hybrid Iterated Local Search [71,97]) integrated within KaMIS. This hybrid of classical kernelization with learned heuristics constitutes one of the strongest state-of-the-art classical frameworks for large MIS and MWIS instances, enabling reductions that would otherwise be too computationally expensive to deploy at scale.

III. RESULTS AND DISCUSSION

In the following, we present detailed results of numerical experiments employing `LearnAndReduce` to solve randomly generated UDG instances. We systematically vary the underlying array size L and graph connectivity by tuning the blockade radius R_b and analyze the reduction statistics of the remaining finite kernels, see Fig. 1 for a schematic illustration. Our work considers blockade radii up to $R_b = 3$ (throughout this manuscript we work in units of the natural lattice spacing, and set $a = 1$), initially exploring the impact of the improved reduction algorithms on MIS problems while comparing to earlier works with simplified reduction rules [65–67]. Crucially, we then extend the scope of our study by looking at weighted MWIS problems.

A. MIS

We begin by studying native, randomly generated MIS problem instances on square arrays of varying size $L \in \{10, 20, 30, 40, 50\}$ and connectivity $1 \leq R_b \leq 3$. The generated UDG instances span a wide range of connectivity patterns with the neighborhood of a generic node graph with node density $\rho = 1$ varying in size between 4 and 28, respectively.

Figure 2(a) shows an example of a $L = 20$ square array across different node densities $\rho \in \{0.7, 0.8, 0.9, 1.0\}$. For each problem instance (L, ρ, R_b) we generate $\#G = 1000$ random graph instances (except for $\rho = 1.0$, where only 1 exists). We record the number of finite kernels ($\#K$), shown in each box of Figs. 2(a) and 2(c), the corresponding reduction factors ξ, and plot the kernel ratio $r_k = \#K/\#G$ as a heatmap. We observe a sharp transition in the number of finite kernels as a function of blockade radius across all densities shown here. For $R_b \leq \sqrt{2}$ (nearest, and next-nearest neighbor connectivity) every single randomly generated graph instance was solved to optimality resulting in an empty kernel ($\xi = 1$), indicating that graphs at this size, density, and connectivity do *not* pose a challenge for existing state or the art classical reduction techniques. We have also confirmed that `LearnAndReduce` was able to fully solve less dense graphs ($\rho < 0.7$) generated in this fashion across all studied connectivities (not shown here). Conversely, for high densities $\rho \geq 0.7$ (with the exception of $R_b = \sqrt{8}$), all other generated problem instances ($R_b \geq 2$) retained a finite kernel, demonstrating an easy-hard transition from the perspective of the `LearnAndReduce` algorithm. A blockade radius of $R_b = \sqrt{8}$ appears to be the outlier and is surprisingly easy to be reduced by standard techniques. This observation is in agreement with a time-to-solution (TTS) analysis of an exact B&B (branch and bound) solver for native UDG graphs [66].

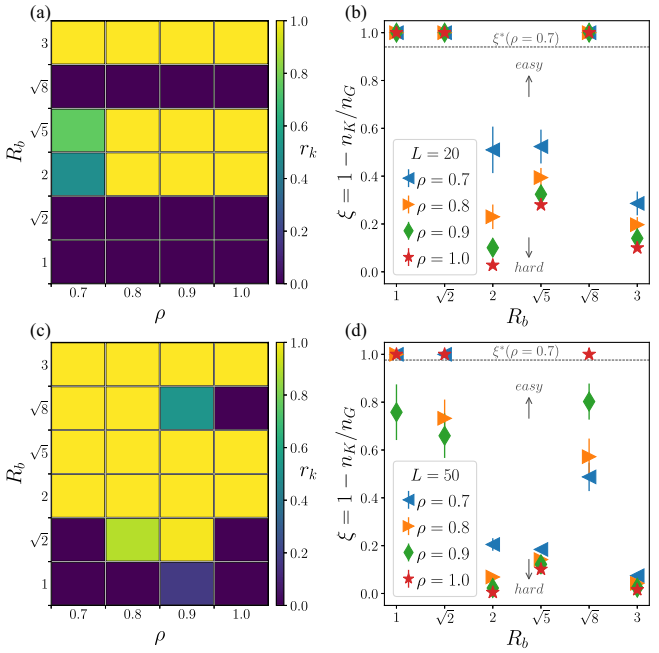


FIG. 2. (a) Results for the $L = 20$ MIS problems, and (b) their corresponding (mean) reduction factors with error bars denoting the standard deviation. (c) The bottom row shows the same results as the top for the largest considered array of size $L = 50$, with corresponding reduction factors shown in (d). Note that for $\rho = 1.0$ there is only a single geometric node arrangement ($\#G = 1$) but for all other parameters, $\rho < 1.0$, we generated $\#G = 1000$ random graph realizations for our statistical ensemble. The numbers in the left subplots indicate the number of finite kernels $\#K$ from the ensemble. The dashed line in the right subplots marks $\xi^* = 1 - 1/\sqrt{N}$, the threshold above which it is more resource efficient to embed the reduced kernel—here chosen for $\rho = 0.7$ (higher densities would result in an even larger threshold ξ^*).

To better understand the structure of these graphs, Fig. 2(b) shows the distribution of the reduction factor ξ across the *finite* kernels shown in (a). Excluding the fully reduced graphs for $R_b \leq \sqrt{2}$ and $R_b = \sqrt{8}$, we note that reduction factors are generally relatively low, $\xi < 0.5$, indicating these graphs are already not trivial.

In particular we observe that the kernelization techniques show a pronounced peak in ξ for the intermediate regime $2 < R_b < 3$, but appear to struggle for blockade radii of 2 and 3. In [66] the TTS for an exact sweeping line algorithm, and a commercial B&B solver was examined as the graph connectivity was tuned. It was found that problem hardness (as measured by the TTS) can be systematically tuned over several orders of magnitude, with the correct behavior in the asymptotic cases of 0-regular ($R_b \to 0$) and complete ($R_b \to \sqrt{2}L$) graphs, as well as pronounced peaks in the TTS at integer blockade radii of $R_b = \{2, 3, 4\}$. Here we corroborate that $R_b = 2, 3$ are not only challenging connectivities for exact solvers, but also for classical reduction algorithms, while $R_b = \{\sqrt{5}, \sqrt{8}\}$ present "easier" problem instances for these classical techniques, with similar behavior across all densities shown here, and an overall trend toward lower ξ values as ρ increases.

In line with the findings in [66] we find that dense, large graphs ($\rho \geq 0.7$ and $L \gtrsim 40$) become less tractable, and result in kernels with much lower reduction factors, even for the King's lattice connectivity at $R_b = \sqrt{2}$ and $R_b = \sqrt{8}$. Figure 2(c) shows the same analysis for the largest studied array of $L = 50$. The vast majority of graphs did yield a finite kernel with, e.g., the largest mean reduction factor for $\rho = 0.9$ (at $R_b = \sqrt{5}$) dropping by a factor of 3 from $\bar{\xi} \sim 0.35$ ($L = 20$) down to $\bar{\xi} \sim 0.12$ for $L = 50$. The other blockade ranges show even smaller reductions, and we again find very similar behavior for the different fillings ρ in the array. Even for $R_b = \sqrt{8}$ and $\rho < 1$ did we find graphs which are not fully reducible,

where arrays with a lower density appear to be harder from a kernelization perspective.

Finally, as noted above, using the full reduction suite of `LearnAndReduce` (in particular the most powerful techniques) means that the resulting graph kernels are no longer guaranteed to be UDG native. This brings us back to the embedding problem outlined at the beginning of this article. Current embedding schemes still require a quadratic resource overhead for embedding non-UDG instances into UDG arrays. This implies that a kernel with a reduction factor ξ requires $\sim(1-\xi)^2N^2$ nodes to be embedded natively on a neutral-atom platform. From a pure resource perspective this leads us to conclude that there is a natural trade-off between embedding a reduced kernel, and solving the original native graph directly, with the former only being favorable for large reductions $\xi \gtrsim \xi^* = 1 - 1/\sqrt{N}$. For a dense completely filled ($\rho = 1$) graph of modest size $L = 10$, this already suggests that we would require reductions in excess of 90% before embedding the resulting kernels becomes beneficial. None of the investigated kernels have exhibited such large reductions, which leads us to identify larger, more connected problem instances as target problem areas where quantum algorithms can be useful, and which may lead to a potential quantum speedup for future experiments with Rydberg atom arrays.

B. MWIS

We now turn to the MWIS problem on native, randomly generated UDG instances. We here focus on the largest array size $L = 50$, as its unweighted MIS problem presented the richest finite kernel landscape. We also confirmed that we obtain qualitatively very similar results for the smaller graphs. For each geometric instance (L, ρ, and R_b) we explore the impact weighting has on the reduction distribution by assigning random integer vertex weights sampled from a uniform distribution $w_i \in [1, W]$, where $W \in \{10, 100, 1000\}$. For each triple (ρ, R_b, W) we randomly select 100 different geometries (from the previous MIS data) and, for every such geometric graph instance produce 100 independent weight realizations. As before, we record the resulting reduction factors ξ. We have studied the effect the *range* of the weighting might have on the reduction factors, but found qualitatively very similar results for all cases (see Appendix B for details). In the following, unless otherwise stated, we present results for $W = 10$.

Figure 3(a) illustrates typical UDG-MWIS instances for $L = 50$ as a function of blockade radius and node density. One notable observation is that adding weights can make the instances *more* amenable to classical reductions: while every single unweighted problem instance for $R_b \geq 2$ (with the exception of $R_b = \sqrt{8}$) was deemed "hard" (i.e., had a finite kernel), here we see that certain weighted instances are solved to optimality (empty kernel).

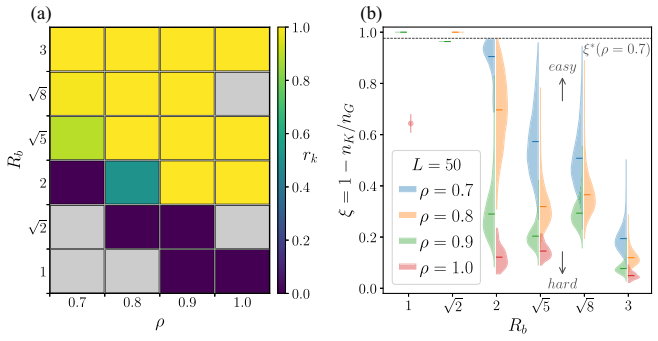


FIG. 3. Reducibility of native MWIS graphs. (a) Results for $L = 50$, where data are generated by randomly sampling 100 graph geometries for each combination of (ρ, R_b) (one for $\rho = 1$), and for each then generating a further 100 unique vertex weightings using uniform random weights with a range of $W = 10$. For certain parameters the unweighted MIS problem was fully reducible, in which case we did not run any weighted problems and highlighted this with a gray, blank square. (b) Reduction factor distribution as half-violins for all node densities ρ and blockade radii R_b. The dashed line marks $\xi^* = 1 - 1/\sqrt{N}$, the threshold above which it is more resource efficient to embed the reduced kernel—here chosen for $\rho = 0.7$ (higher densities would result in an even larger threshold ξ^*).

Furthermore, Fig. 3(b) shows that we find larger mean reduction factors compared with the MIS case. This behavior is consistent with the interpretation that adding weights can partially lift the degeneracies present in the MIS problem, thereby providing the reduction machinery with more structure to exploit.

These conclusions are further corroborated by known approximability theory results for the MIS and MWIS problem; UD-MWIS is in the polynomial-time approximation scheme (PTAS) class [98–102], while the general MIS problem is APX-complete [72,103,104]. Thus MIS is believed not to admit a PTAS algorithm, and cannot be in the PTAS class (unless P = NP) [105]. Likewise following the quantum PCP conjecture [106], MIS is believed to remain in the quantum equivalent of the APX class even with quantum algorithms.

It should be noted that weighting in turn could lead to *smaller* spectral gaps generically for the logical problems, which would on the other hand lead to adiabatic bottlenecks in quantum annealing procedures. Nevertheless, even in the weighted setting we still observe predominantly modest reduction factors, with $\xi \lesssim 0.5$ for many instances, indicating that these problems remain nontrivial in the majority of cases, and only very few outliers come close to the reduction target ξ^* above which embedding of the reduced kernel might be favorable over solving the original native graph instance.

To quantify when MWIS remains challenging, Fig. 3(b) shows the ξ distribution for the finite kernels across the different blockade radii and for all node densities ρ. Generally we find that weight-induced simplifications

shift the distribution toward higher values of ξ as compared with the unweighted MIS data in Fig. 2. In particular graphs with a relatively local structure (e.g., $R_b < 2$) are solved to near optimality even for graphs with few defects ($\rho = 0.8$ or 0.9). In contrast, for the longer-range connectivity $R_b = \{\sqrt{5}, 3\}$, we find that kernelization performance sharply deteriorates: the distribution becomes more skewed toward lower values of ξ, exhibiting a long tail corresponding to the minority of instances where reductions are more sizable (ξ larger than its mean $\bar{\xi}$). For $R_b = \sqrt{8}$, which was efficiently reducible for small MIS instances, we observe that weighting acts to reduce the mean ξ relative to MIS (e.g., for $\rho = 0.9$ the mean drops from 0.85 to 0.3) while still proving more reducible than instances with $R_b = \{\sqrt{5}, 3\}$. The overwhelming majority of weighted graphs therefore remain difficult to reduce. This mirrors the behavior seen in MIS, suggesting that here too longer-range UDG structures are intrinsically more resistant to classical reductions regardless of the weight scale. Since weights and geometry both influence reducibility, it is natural to ask whether MWIS hardness is primarily set by the underlying connectivity pattern or by the choice of weights. While a full geometry-resolved study is beyond the scope of this work, we note that for a fixed geometry the spread of reduction factors across different weight realizations can be substantial, particularly for moderate blockade radii. This indicates that both sources contribute in a nontrivial way: geometry determines a baseline difficulty, whereas weights modulate how effectively reduction rules can break symmetries or eliminate ambiguous local configurations.

We conclude this section by studying how the reduction factors behave as a function of the size $n_G = |G|$ of the graphs, shown in Fig. 4. At moderate connectivity ($R_b = 2$) ξ displays a very curious behavior for the different node densities. For $\rho = 0.7$ the reduction factor dips to a minimum for $20 \leq L \leq 30$ but remains very high for very small or large graphs. Slightly denser graphs do not exhibit the resurgence of ξ for large arrays but plateau to a steady value of $\xi \sim 0.7$. Increasing the node density further inverts the observed behavior completely. The smallest $L = 10$ graph has the smallest reduction factor, which dramatically increases when the array size increases to 20, after which ξ gradually decreases with system size as one would naively expect. Higher connectivity graphs $R_b = \sqrt{5}$ and $R_b = \sqrt{8}$ follow this more consistently. Smaller graphs exhibit higher reductions, which then gradually decrease with increasing system size. Further we can nicely observe that higher density graphs present more of a challenge to these reduction techniques since, e.g., $\rho = 0.7$ (blue diamonds) reduction factors are consistently larger than higher values of ρ.

Finally, as in the MIS case, invoking the full reduction suite of `LearnAndReduce` means that resulting MWIS

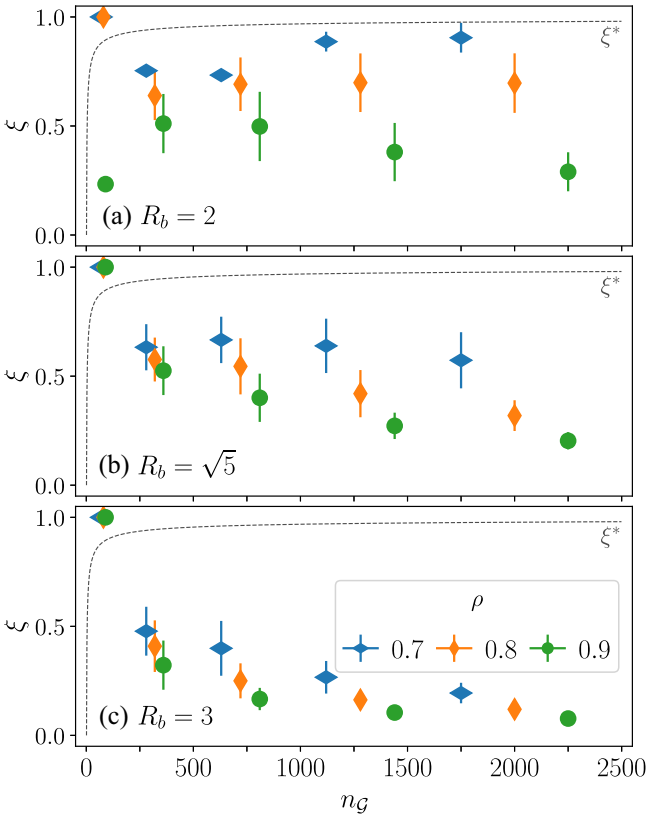


FIG. 4. Scaling of the reduction factor ξ with the graph size $n_G = |G| = \rho L^2$ for three different node densities ρ, and connectivities R_b (a)–(c) as depicted in the legend. The gray dashed line indicates the critical reduction ξ^* beyond which an embedding of the kernel is preferable (from an atom number resource perspective) compared with solving the original graph.

kernels are generally no longer UDG-native. For $L = 50$, this corresponds to an even larger reduction target of $\xi^* \sim 0.98$. Across all MWIS instances we examined—excluding the fully reducible cases—no kernels achieved such dramatic reductions. This leads us to the same conclusions as in the unweighted MIS case: in most practical situations it may be preferable to run the *original* native MWIS instance on a Rydberg atom processor rather than embedding a reduced kernel. Consequently, large, dense, and longer-range UDG-MWIS instances emerge as further candidates for probing the computational frontier of near-term neutral-atom hardware.

We note that for the discussed results we have used `LearnAndReduce` with the GNN filter enabled. However, we find no change in conclusion when implementing the complete set of reduction methods. This confirms that a lack of reducibility is not a consequence of the GNN filter having been trained on an unsuitable class of graphs, but instead is due to native, high-connectivity graphs being intrinsically harder to reduce (see Appendix C for details).

IV. CONCLUSION

In this work, we have examined the extent to which state-of-the-art classical reduction techniques can simplify native UDG MIS and MWIS problem instances relevant to Rydberg-atom quantum hardware. Our analysis reveals several clear trends.

First, even for moderate array sizes, many dense UDG-MIS instances remain *not* fully reducible: sizable irreducible kernels persist despite the extensive suite of reductions employed. Introducing vertex weights, however, tends to increase reducibility, with some instances that admit finite MIS kernels becoming fully reducible once weights are added, indicating that weights can partially lift degeneracies that otherwise obscure the action of reduction rules.

Second, nontrivial kernels arise predominantly in moderately large and dense systems, typically for $L \gtrsim 20$ and $\rho \gtrsim 0.8$. Within this regime, the structure of the underlying UDG connectivity plays a decisive role. Longer-range connectivities—corresponding to larger blockade radii—systematically degrade the performance of classical preprocessing, leading to broad distributions of low reduction factors and signaling that these graphs remain intrinsically difficult for kernelization.

For instances that do retain finite kernels, executing these kernels on a Rydberg quantum processor necessitates embedding them into a native UDG layout. Current embedding schemes incur a quadratic resource overhead, implying that only kernels with reduction factors $\xi \gtrsim 1 - 1/\sqrt{N}$ are advantageous from a purely resource-based perspective. Even for a modest $N = 100$ node graph, this would require $\xi > 0.9$, a level of reduction essentially absent from our data (excluding the fully reducible cases). This indicates that, in the overwhelming majority of practically relevant situations, direct execution of the *original* native graph may be preferable to embedding a reduced kernel.

Taken together, these results identify a clear region of parameter space—dense, moderately large arrays with extended UDG connectivity—which provides a natural target for assessing the computational capabilities of near-term neutral-atom quantum processors. To connect these parameter regimes to real systems, typical platforms offer atom spacing around $a = 4$ μm, and two-photon Rabi frequencies $\Omega/2\pi \sim 2$ MHz [107]. The blockade radius is given by $R_b = \sqrt[6]{|C_6|/\Omega}$, where C_6 is the dispersion coefficient for the Rydberg interaction strength which scales as n^{11} for $nS_{1/2}$ Rydberg states. To achieve $R_b > 3a$ requires $|C_6| > 6 \times 10^6$ MHz μm^6, corresponding to choosing $n \geq 83$ in rubidium, making this high-connectivity regime experimentally feasible [108].

Our findings open several compelling avenues for follow-up work. A significant next step is to determine how exact classical solvers (ideally in a co-designed pipeline together with existing reduction algorithms) perform on these native graphs, and to pinpoint when they begin to struggle. Likewise, it remains to be tested whether quantum annealing or adiabatic strategies remain viable at these problem sizes, particularly beyond King's lattice connectivity. If an embedded finite kernel becomes a weighted King's lattice instance, and if such MWIS problems indeed remain tractable for classical methods, then embedding—despite its overhead—could become competitive in certain regimes. A detailed comparison of time-to-solution across reductions, exact solvers, and quantum hardware will be essential for establishing which classes of UDG-MIS and UDG-MWIS problems constitute meaningful benchmarks for future experimental demonstrations. Finally, to exploit the intrinsic advantage offered by the higher native connectivity, it is essential to develop embedding strategies that can map real-world graph problems onto MIS or MWIS instances with $R_b > \sqrt{2}$.

ACKNOWLEDGMENTS

We thank Ernestine Großmann, Kenneth Langedal, Arthur La Rooij, Gerard Pelegrí, Paul Warburton, Martin J. A. Schuetz, Ruben S. Andrist, and Helmut G. Katzgraber for useful discussions. Particular thanks to Ernestine Großmann and Kenneth Langedal for their help in setting up and running `LearnAndReduce`. This work was supported by EPSRC through Grants No. EP/Z53318X/1 and No. EP/T005386/1.

DATA AVAILABILITY

The data presented in this work are available at [109].

CONFLICTS OF INTEREST

Authors declare no conflict of interest.

APPENDIX A: EXAMPLES OF FUNDAMENTAL REDUCTION RULES

In the following section, we provide examples of some of the rules relevant for MIS and MWIS reductions.

1. Examples of fundamental reduction rules

In the following, we give a brief description of a subset of reduction rules used in `LearnAndReduce`. A complete list and details can be found in Table 1 of [97].

a. Isolated vertex. If a vertex v has an open neighborhood (simply called the neighborhood) $N(v) = \{u | (v,u) \in E\} = \emptyset$, then v must belong to every MIS/MWIS solution. Thus one includes v in the MIS/MWIS, removes it from G to form K, and records a weight offset of w_v.

b. Simplicial vertex. A vertex v is *simplicial* when $N(v)$ induces a clique. For MWIS, if $w_v \geq \max_{u \in N(v)} w_u$, then selecting v dominates any choice of its neighbors, and one includes v in the MWIS, removes the closed neighborhood

$N[v] = \{v\} \cup N(v)$, and records a weight offset of w_v. Otherwise, v can be removed [70,74].

c. Pendant (degree one) vertex. If a vertex v has exactly one neighbor u, then an optimal solution must include v whenever $w_v \geq w_u$. In the unweighted case, the pendant vertex v always dominates the neighbor u. Accordingly, v is selected, and both v and u are removed. If $w_v < w_u$, we form K by removing v, updating $w_u \to w_u - w_v$, and recording the weight offset w_v. It holds that $u \in \mathrm{MWIS}(G)$ iff $u \in \mathrm{MWIS}(K)$ [110].

d. Twin reduction. Let $u, v \in V$ have equal, independent neighborhoods $N(u) = N(v) = \{p, q, r\}$. Any optimal solution contains either both u, v or neither. If $w(\{u, v\}) \geq w(\{p, q, r\})$, include u, v and reduce graph to $K = G - N[\{u, v\}]$, with an offset $w_u + w_v$. If $w(\{u, v\}) < w(\{p, q, r\})$ but $w(\{u, v\}) > w(\{p, q, r\}) - \min\{w_x | x \in \{p, q, r\}\}$, fold $\{u, v, p, q, r\}$ into new vertex v' with $N(v') = N(\{p, q, r\})$, and $w_{v'} = w(\{p, q, r\}) - w(\{u, v\})$. The new graph is given by $K = G((V \cup \{v'\}) \backslash N[\{u, v\}])$ with an offset of $w(\{u, v\})$. To reconstruct the solution for G, one selects $\{p, q, r\}$ if v' is in the solution to K, and $\{u, v\}$ otherwise. Note, that the twin rule can be generalized to larger neighborhoods, and ones which are not an independent set [70].

e. Unconfined reduction. Let $S \subseteq V$ be an independent set, and assume S is contained in every MWIS of G. A vertex $x \in N(S)$ is a child of S if $w_x \geq w(S \cap N(x))$. It is an extending child if there exists $y \in N(x) \backslash N[S]$ such that $w_x \geq w(S \cap N(x)) + w(I_y)$, where I_y is an MWIS of $G[(N(x) \backslash \{y\}) \backslash N[S]]$. The vertex y is called a satellite, and it can be shown that every MWIS contains the satellites from each extending child x of S. Now, iteratively extend $S = \{v\}$ by adding satellites of extending children. If at some step there exists a child x with $w_x \geq w(S \cap N(x)) + w(G[N(x) \backslash N[S]])$, then v is unconfined, and can be removed $K = G - v$. If no further satellites exist, S confines v, and every MWIS containing v, also contains S and one cannot remove v [97].

f. Vertex folding. In the unweighted MIS problem [111], if a vertex v has exactly two nonadjacent neighbors u and w, then every optimal solution includes either v or both u and w. Thus $\{u, v, w\}$ can be *folded* into a single supervertex v'. For the weighted case, let $N(v)$ be independent. If $w(N(v)) > w(v)$ but $w(N(v)) - \min_{u \in N(v)}\{w_u\} < w_v$, fold v and $N(v)$ into new vertex v' with $w_{v'} = w(N(v)) - w_v$. To reconstruct the MWIS on G one selects $N(v)$ if v' is in the solution to K, and v otherwise [70].

APPENDIX B: GRAPH WEIGHTING

In investigating the efficacy of graph reduction methods on weighted graphs, we implement random vertex weightings chosen uniformly from $[1, W]$, with weight ranges $W = \{10, 100, 1000\}$. Figure 5 shows the results for the reduction of graphs with array size $L = 50$ and weight ranges $W = 100$ (a) and $W = 1000$ (b). Comparing with the data above in Fig. 3(b) for $W = 10$, we find no significant change in the performance of the `LearnAndReduce` reduction algorithm. We are thus confident that our conclusions above based on weighted data for $W = 10$ extend to all graphs with a uniform weighting $W > 1$.

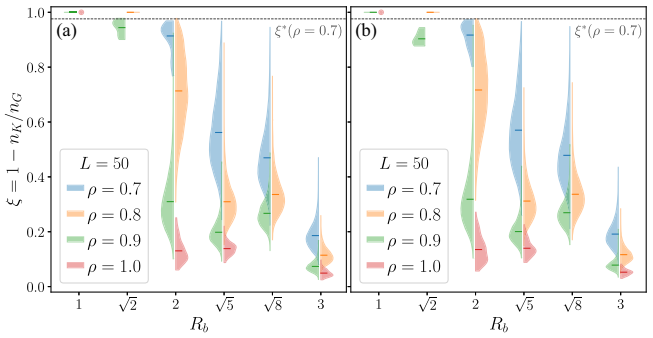


FIG. 5. Comparing the effect of different weight ranges for (a) $W = 100$, and (b) $W = 1000$ on the reductions. The results do not show any significant differences, suggesting that the size of the weightings W does not affect the overall resulting kernel distribution.

APPENDIX C: GNN FILTERING AND RUNTIME

While the focus of this work is on exploring regimes in which native Rydberg graphs are robust to reduction methods, it is relevant to consider both the impact of the GNN methods employed by the `LearnAndReduce` algorithm and its runtime. For the presented results above, `LearnAndReduce` was run with the GNN filter enabled attempting to identify promising nodes and guide further reduction techniques. To verify that the corresponding reduction efficiency is not significantly affected by the training set used for the GNN filter, Fig. 6(a) shows the reduction efficiency with (left half-violin, solid line) and without (right half-violin, dashed line) the GNN filter enabled for some of the hardest graph configurations we observed ($L = 50$, $\rho = 0.9$, $R_b = 3$, and with weight ranges $W = 1$ for MIS and $W = 10$ for MWIS instances). The left subplot (a) shows that the kernel size, using the complete set of reductions, is only marginally smaller (i.e., a ξ distribution shifted to slightly larger values) than that achieved using the GNN selected methods. Comparing the runtimes, however, in Fig. 6(b) shows that the GNN offers an average speedup of around 20%, with the runtime for MIS graphs on average around 0.3 s for MIS and 0.4 s for MWIS configurations with GNN enabled. When considering quantum advantage, time to solution is a key metric. The `LearnAndReduce` method is fast compared with the current runtime of neutral-atom platforms with clock rates of a few Hz [50,107]; however, the development of continuously loaded platforms [43] and fast readout techniques using single-photon detector arrays [112] provide a pathway to clock speeds >1 kHz.

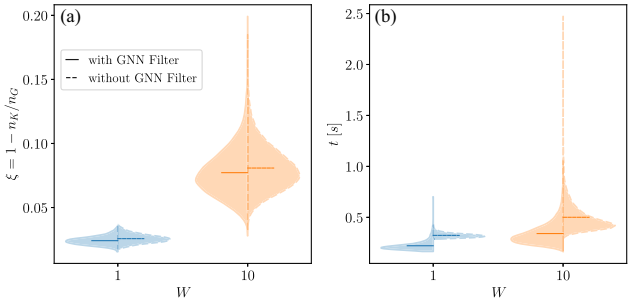


FIG. 6. Comparing the reduction distribution and runtime of the `LearnAndReduce` algorithm when employing the GNN (left half-violin, solid line), or applying the reduction rules exhaustively (right half-violin, dashed line). The overall distribution of reduction values in (a) is shifted to slightly larger values, indicating that employing the full reduction suite exhaustively can bring about smaller kernels on average. The performance gain is, however, almost negligible, and in reality must be balanced with a noticeable increase in the runtime (excluding I/O timings) as shown in (b). Data shown for $L = 50$, $\rho = 0.9$, and $R_b = 3$ corresponding to some of the least reducible graph examples found above.

[1] V. T. Paschos, *Applications of Combinatorial Optimization* (Wiley, Hoboken, NJ, USA, 2014), Vol. 3.

[2] R. M. Karp, Reducibility among combinatorial problems, in *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations*, edited by R. E. Miller-J. W. Thatcher-J. D. Bohlinger (Springer US, Boston, MA, 1972), pp. 85–103.

[3] C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Courier Corporation, Mineola, New York, 1998).

[4] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms* (Springer, Berlin, Heidelberg, 2018) ISBN 978-0-486-40258-1.

[5] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, arXiv: quant-ph/0001106.

[6] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem, Science **292**, 472 (2001).

[7] F. Glover, G. Kochenberger, R. Hennig, and Y. Du, Quantum bridge analytics I: A tutorial on formulating and using QUBO models, Ann. Oper. Res. **314**, 141 (2022).

[8] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, Phys. Rev. E **58**, 5355 (1998).

[9] A. Das and B. K. Chakrabarti, Colloquium: Quantum annealing and analog quantum computation, Rev. Mod. Phys. **80**, 1061 (2008).

[10] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, Perspectives of quantum annealing: Methods and implementations, Rep. Prog. Phys. **83**, 054401 (2020).

[11] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, Quantum annealing: An overview, Phil. Trans. R. Soc. A **381**, 20210417 (2023).

[12] 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).

[13] N. Astrakhantsev, G. Mazzola, I. Tavernelli, and G. Carleo, Phenomenological theory of variational quantum ground-state preparation, Phys. Rev. Res. **5**, 033225 (2023).

[14] 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).

[15] 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).

[16] 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).

[17] 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).

[18] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.

[19] D. Wecker, M. B. Hastings, and M. Troyer, Progress towards practical quantum variational algorithms, Phys. Rev. A **92**, 042303 (2015).

[20] J. Wurtz and P. Love, Maxcut quantum approximate optimization algorithm performance guarantees for $p > 1$, Phys. Rev. A **103**, 042612 (2021).

[21] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, The quantum approximate optimization algorithm and the Sherrington-Kirkpatrick model at infinite size, Quantum **6**, 759 (2022).

[22] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, A review on quantum approximate optimization algorithm and its variants, Phys. Rep. **1068**, 1 (2024).

[23] A. Lucas, Ising formulations of many NP problems, Front. Phys. **2**, 5 (2014).

[24] H. Zhang, K. Boothby, and A. Kamenev, Cyclic quantum annealing: Searching for deep low-energy states in 5000-qubit spin glass, Sci. Rep. **14**, 30784 (2024).

[25] L. F. Pérez Armas, S. Creemers, and S. Deleplanque, Solving the resource constrained project scheduling problem with quantum annealing, Sci. Rep. **14**, 16784 (2024).

[26] M. W. Johnson *et al.*, Quantum annealing with manufactured spins, Nature (London) **473**, 194 (2011).

[27] D. De Santis, S. Tirone, S. Marmi, and V. Giovannetti, Optimized QUBO formulation methods for quantum computing, Quantum Sci. Technol. **11**, 015056 (2026).

[28] A. D. King *et al.*, Beyond-classical computation in quantum simulation, Science **388**, 199 (2025).

[29] A. Zaribafiyan, D. J. J. Marchand, and S. S. Changiz Rezaei, Systematic and deterministic graph minor embedding for Cartesian products of graphs, Quantum Inf. Process. **16**, 136 (2017).
[30] E. Pelofske, Comparing three generations of D-wave quantum annealers for minor embedded combinatorial optimization problems, Quantum Sci. Technol. **10**, 025025 (2025).
[31] V. Choi, Minor-embedding in adiabatic quantum computation: I. The parameter setting problem, Quantum Inf. Process. **7**, 193 (2008).
[32] V. Choi, Minor-embedding in adiabatic quantum computation: II. Minor-universal graph design, Quantum Inf. Process. **10**, 343 (2011).
[33] P. Date, R. Patton, C. Schuman, and T. Potok, Efficiently embedding QUBO problems on adiabatic quantum computers, Quantum Inf. Process. **18**, 117 (2019).
[34] A. Ceselli and M. Premoli, On good encodings for quantum annealer and digital optimization solvers, Sci. Rep. **13**, 5628 (2023).
[35] S. Kim, S.-W. Ahn, I.-S. Suh, A. W. Dowling, E. Lee, and T. Luo, Quantum annealing for combinatorial optimization: A benchmarking study, npj Quantum Inf. **11**, 77 (2025).
[36] F. A. Quinton, P. A. S. Myhr, M. Barani, P. Crespo del Granado, and H. Zhang, Quantum annealing applications, challenges and limitations for optimisation problems compared to classical solvers, Sci. Rep. **15**, 12733 (2025).
[37] L. F. Henriet, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, Quantum **4**, 327 (2020).
[38] M. Morgado and S. Whitlock, Quantum simulation and computing with Rydberg-interacting qubits, AVS Quantum Sci. **3**, 023501 (2021).
[39] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, Nature (London) **604**, 451 (2022).
[40] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, Nature (London) **622**, 268 (2023).
[41] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, Logical quantum processor based on reconfigurable atom arrays, Nature (London) **626**, 58 (2024).
[42] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, Nature (London) **647**, 60 (2025).
[43] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. L. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, S. H. Li, X. Lyu, L. M. Peters, D. Bluvstein, T. T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Continuous operation of a coherent 3,000-qubit system, Nature (London) **646**, 1075 (2025).
[44] 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).
[45] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller, Dipole blockade and quantum information processing in mesoscopic atomic ensembles, Phys. Rev. Lett. **87**, 037901 (2001).
[46] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Parallel implementation of high-fidelity multiqubit gates with neutral atoms, Phys. Rev. Lett. **123**, 170503 (2019).
[47] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, Rev. Mod. Phys. **82**, 2313 (2010).
[48] 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.
[49] 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.
[50] S. Ebadi *et al.*, Quantum optimization of maximum independent set using Rydberg atom arrays, Science **376**, 1209 (2022).
[51] 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).
[52] 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).
[53] A. G. de Oliveira, J. Kombe, G. Pelegrí, P. Schroff, M. T. Wells-Pestell, D. M. Walker, A. J. Daley, and J. D. Pritchard, Sparse-graph optimization using weighted quantum wires in Rydberg-atom arrays, Phys. Rev. Appl. **26**, 024007 (2026).
[54] 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).
[55] A. Byun, M. Kim, and J. Ahn, Finding the maximum independent sets of platonic graphs using Rydberg atoms, PRX Quantum **3**, 030305 (2022).
[56] C. Dalyac, L.-P. Henry, M. Kim, J. Ahn, and L. Henriet, Exploring the impact of graph locality for the resolution of MIS with neutral atom devices, arXiv:2306.13373.
[57] 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).
[58] A. Byun, S. Jeong, and J. Ahn, Programming higher-order interactions of Rydberg atoms, Phys. Rev. A **110**, 042612 (2024).
[59] 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).

[60] M. Lanthaler, K. Ender, C. Dlaska, and W. Lechner, Quantum optimization with globally driven neutral atom arrays, arXiv:2410.03902.

[61] 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).

[62] M. Lanthaler, C. Dlaska, K. Ender, and W. Lechner, Rydberg-blockade-based parity quantum optimization, Phys. Rev. Lett. **130**, 220601 (2023).

[63] 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).

[64] J. Park, S. Jeong, M. Kim, K. Kim, A. Byun, L. Vignoli, L.-P. Henry, L. F. Henriet, and J. Ahn, Rydberg-atom experiment for the integer factorization problem, Phys. Rev. Res. **6**, 023241 (2024).

[65] 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, arXiv:2412.14976.

[66] R. S. Andrist, M. J. Schuetz, P. Minssen, R. Yalovetzky, S. Chakrabarti, D. Herman, N. Kumar, G. Salton, R. Shaydulin, Y. Sun *et al.*, Hardness of the maximum-independent-set problem on unit-disk graphs and prospects for quantum speedups, Phys. Rev. Res. **5**, 043277 (2023).

[67] 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.

[68] S. Lamm, P. Sanders, C. Schulz, D. Strash, and R. F. Werneck, Finding near-optimal independent sets at scale, J. Heuristics **23**, 207 (2017).

[69] D. Hespe, C. Schulz, and D. Strash, Scalable kernelization for maximum independent sets, ACM J. Exp. Algorithmics **24**, 1 (2019).

[70] S. Lamm, C. Schulz, D. Strash, R. Williger, and H. Zhang, Exactly solving the maximum weight independent set problem on large real-world graphs, in *2019 Proceedings of the Meeting on Algorithm Engineering and Experiments (ALENEX)* (2019), pp. 144–158.

[71] E. Großmann, K. Langedal, and C. Schulz, Accelerating reductions using graph neural networks and a new concurrent local search for the maximum weight independent set problem, arXiv:2412.14198.

[72] M. R. Garey and D. S. Johnson, "Strong " NP-completeness results: Motivation, examples, and implications, J. ACM **25**, 499 (1978).

[73] D. S. Johnson and M. R. Garey, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (WH Freeman, San Francisco, California, 1979) ISBN 978-0-7167-1044-8.

[74] T. Akiba and Y. Iwata, Branch-and-reduce exponential/FPT algorithms in practice: A case study of vertex cover, Theor. Comput. Sci. **609**, 211 (2016).

[75] D. Hespe, S. Lamm, C. Schulz, and D. Strash, Wegotyoucovered: The winning solver from the pace 2019 challenge, vertex cover track, in *2020 Proceedings of the SIAM Workshop on Combinatorial Scientific Computing (CSC)* (2020), pp. 1–11.

[76] S. Szabo and B. Zaválnij, Combining algorithms for vertex cover and clique search, In Proceedings of the 22nd International Multiconference INFORMATION SOCIETY–IS (2019), **vol. 1**, pp. 71–74.

[77] R. Plachetta and A. van der Grinten, SAT-and-reduce for vertex cover: Accelerating branch-and-reduce by SAT solving, in *2021 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)* (2021), pp. 169–180.

[78] D. Hespe, S. Lamm, and C. Schorr, Targeted branching for the maximum independent set problem, arXiv:2102.01540.

[79] K. Langedal, D. Hespe, and P. Sanders, Targeted branching for the maximum independent set problem using graph neural networks, in *22nd International Symposium on Experimental Algorithms (SEA 2024), Leibniz International Proceedings in Informatics (LIPIcs)*, edited by L. Liberti (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2024), pp. 20:1–20:21.

[80] J. S. Warren and I. V. Hicks, Combinatorial branch-and-bound for the maximum weight independent set problem, Relatório Técnico, Texas A&M University, Citeseer **9**, 17 (2006), https://www.cmor-faculty.rice.edu/~ivhicks/jeff.rev.pdf.

[81] D. V. Andrade, M. G. C. Resende, and R. F. Werneck, Fast local search for the maximum independent set problem, J. Heuristics **18**, 525 (2012).

[82] A. Gellner, S. Lamm, C. Schulz, D. Strash, and B. Zaválnij, Boosting data reduction for the maximum weight independent set problem using increasing transformations, in *Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, edited by M. Farach-Colton (Society for Industrial and Applied Mathematics (SIAM), 2021), p. 128.

[83] C. Ebenegger, P. Hammer, and D. de Werra, Pseudo-boolean functions and stability of graphs, *Algebraic and Combinatorial Methods in Operations Research*, North-Holland Mathematics Studies, Vol. 95, edited by R. Burkard, R. Cuninghame-Green, and U. Zimmermann (North-Holland, Amsterdam, the Netherlands, 1984), pp. 83–97.

[84] G. Alexe, P. L. Hammer, V. V. Lozin, and D. de Werra, Struction revisited, Discrete Appl. Math. **132**, 27 (2003).

[85] W. Zheng, J. Gu, P. Peng, and J. X. Yu, Efficient weighted independent set computation over large graphs, in *2020 IEEE 36th International Conference on Data Engineering (ICDE)* (2020), pp. 1970–1973.

[86] M. Xiao, S. Huang, and X. Chen, Maximum weighted independent set: Effective reductions and fast algorithms on sparse graphs, Algorithmica **86**, 1293 (2024).

[87] A. Figiel, V. Froese, A. Nichterlein, and R. Niedermeier, There and back again: On applying data reduction rules by undoing others, in *30th Annual European Symposium on Algorithms (ESA 2022), Leibniz International Proceedings in Informatics (LIPIcs)*, edited by S. Chechik,

G. Navarro, E. Rotenberg, and G. Herman (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2022), Vol. 244, pp. 53:1–53:15.

[88] Z. Fang, C.-M. Li, and K. Xu, An exact algorithm based on MaxSAT reasoning for the maximum weight clique problem, J. Artif. Int. Res. **55**, 799 (2016).

[89] S. Held, W. Cook, and E. C. Sewell, Maximum-weight stable sets and safe lower bounds for graph coloring, Math. Program. Comput. **4**, 363 (2012).

[90] H. Jiang, C.-M. Li, and F. Manya, An exact algorithm for the maximum weight clique problem in large graphs, Proc. AAAI Conf. Artif. Intell. **31**, 830 (2017).

[91] S. Butenko, P. Pardalos, I. Sergienko, V. Shylo, and P. Stetsyuk, Finding maximum independent sets in graphs arising from coding theory, in *Proceedings of the 2002 ACM Symposium on Applied Computing, SAC '02* (Association for Computing Machinery, New York, NY, USA, 2002), pp. 542–546.

[92] S. Butenko and S. Trukhanov, Using critical sets to solve the maximum independent set problem, Oper. Res. Lett. **35**, 519 (2007).

[93] D. Strash, On the power of simple reductions for the maximum independent set problem, *Computing and Combinatorics*, edited by T. N. Dinh-M. T. Thai (Springer International Publishing, Cham, 2016), pp. 345–356.

[94] L. Chang, W. Li, and W. Zhang, Computing a near-maximum independent set in linear time by reducing-peeling, in *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17* (Association for Computing Machinery, New York, NY, USA, 2017), pp. 1181–1196.

[95] S. Butenko, P. Pardalos, I. Sergienko, V. Shylo, and P. Stetsyuk, Estimating the size of correcting codes using extremal graph problems, *Optimization: Structure and Applications*, edited by C. Pearce and E. Hunt (Springer New York, New York, NY, 2009), pp. 227–243.

[96] E. Großmann, S. Lamm, C. Schulz, and D. Strash, Finding near-optimal weight independent sets at scale, in *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '23* (Association for Computing Machinery, New York, NY, USA, 2023), pp. 293–302.

[97] E. Großmann, K. Langedal, and C. Schulz, Concurrent iterated local search for the maximum weight independent set problem, in *23rd International Symposium on Experimental Algorithms (SEA 2025), Leibniz International Proceedings in Informatics (LIPIcs)*, edited by P. Mutzel and N. Prezza, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2025), Vol. 338, pp. 22:1–22:18.

[98] T. Erlebach, K. Jansen, and E. Seidel, Polynomial-time approximation schemes for geometric graphs, in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '01* (Society for Industrial and Applied Mathematics, USA, 2001), pp. 671–679.

[99] T. Matsui, Approximation algorithms for maximum independent set problems and fractional coloring problems on unit disk graphs, in *Discrete and Computational Geometry*, edited by J. Akiyama, M. Kano, and M. Urabe (Springer, Berlin, Heidelberg, 2000), pp. 194–200.

[100] G. K. Das, M. De, S. Kolay, S. C. Nandy, and S. Sur-Kolay, Approximation algorithms for maximum independent set of a unit disk graph, Inf. Process. Lett. **115**, 439 (2015).

[101] S. C. Nandy, S. Pandit, and S. Roy, Faster approximation for maximum independent set on unit disk graph, Inf. Process. Lett. **127**, 58 (2017).

[102] M. Bonamy, E. Bonnet, N. Bousquet, P. Charbit, and S. Thomasse, EPTAS for max clique on disks and unit balls, in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, CA, USA, 2018), pp. 568–579.

[103] J. Håstad, Clique is hard to approximate within $n^{1-\epsilon}$, Acta Math. **182**, 105 (1999).

[104] C. Bazgan, B. Escoffier, and V. T. Paschos, Completeness in standard and differential approximation classes: Poly-(D)APX- and (D)PTAS-completeness, Theor. Comput. Sci. **339**, 272 (2005).

[105] C. de Correc, T. Ayral, and C. Bertrand, Approximate combinatorial optimization with Rydberg atoms: The barrier of interpretability, Phys. Rev. A **112**, 042441 (2025).

[106] D. Aharonov, I. Arad, and T. Vidick, Guest column: The quantum PCP conjecture, ACM SIGACT News **44**, 47 (2013).

[107] J. Wurtz, A. Bylinskii, B. Braverman, J. Amato-Grill, S. H. Cantu, F. Huber, A. Lukin, F. Liu, P. Weinberg, J. Long, S.-T. Wang, N. Gemelke, and A. Keesling, Aquila: QuEra's 256-qubit neutral-atom quantum computer, arXiv:2306.11727.

[108] N. Sibalić, J. D. Pritchard, C. S. Adams, and K. J. Weatherill, ARC: An open-source library for calculating properties of alkali Rydberg atoms, Comput. Phys. Commun. **220**, 319 (2017).

[109] J. Kombe and J. PritchardReducibility of native weighted graphs on Rydberg Arrays, 10.15129/69d10cd4-8cc8-4bcc-8329-3275804f5199.

[110] J. Gu, W. Zheng, Y. Cai, and P. Peng, Towards computing a near-maximum weighted independent set on massive graphs, in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD'21* (Association for Computing Machinery, New York, NY, USA, 2021), pp. 467–477.

[111] J. Chen, I. A. Kanj, and W. Jia, Vertex cover: Further observations and further improvements, J. Algorithms **41**, 280 (2001).

[112] X.-Z.-Q. Zeng, C. You, Q.-W. Wang, Z.-F. Li, Y. Ji, D. An, C. Yu, J.-R. Liu, Z.-M. He, J.-R. Gu, Y. Mei, H.-W. Cheng, Y.-C. Zhang, R. Lin, Z. Wu, J. Rui, J. Zhang, M.-C. Chen, Y.-H. Deng, C.-Y. Lu, and J.-W. Pan, Fast nondestructive readout for high-clock-rate atom array quantum processor, arXiv:2608.17189.