









Efficient algorithms to solve atom reconfiguration problems. I. Redistribution-reconfiguration algorithm

Barry Cimring ¹, Remy El Sabeh ², Marc Bacvanski ³, Stephanie Maaz ⁴, Izzat El Hajj ², Naomi Nishimura ⁴,
Amer E. Mouawad ^{2,4,5} and Alexandre Cooper ^{1,*}

¹*Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

²*Department of Computer Science, American University of Beirut, Beirut 1107 2020, Lebanon*

³*Khoury College of Computer Science, Northeastern University, Boston, Massachusetts 02115, USA*

⁴*David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada*

⁵*AG Theoretische Informatik, University of Bremen, 28359 Bremen, Germany*



(Received 13 October 2022; revised 25 April 2023; accepted 6 July 2023; published 4 August 2023)

We propose the redistribution-reconfiguration (red-rec) algorithm to prepare large configurations of atoms using programmable arrays of optical traps. Red-rec exploits simple heuristics and exact subroutines to solve atom reconfiguration problems on grids. It admits a fast and efficient implementation suitable for real-time operation. We numerically quantify its performance using realistic physical parameters and operational constraints, in both the absence and presence of loss. Red-rec enables assembling large configurations of atoms with high mean success probability. Fast preparation times are achieved by harnessing parallel control operations that actuate multiple traps simultaneously. Faster preparation times are achieved by rejecting configurations of atoms containing fewer atoms than a given threshold. However, the number of traps required to prepare a compact-centered configuration of atoms on a grid with a probability of 0.5 scales as the $3/2$ power of the number of desired atoms. This finding highlights some of the challenges associated with scaling up configurations of atoms beyond tens of thousands of atoms.

DOI: [10.1103/PhysRevA.108.023107](https://doi.org/10.1103/PhysRevA.108.023107)

I. INTRODUCTION

Quantum many-body systems formed by configurations of individual quantum particles, such as neutral atoms, charged ions, and molecules, offer valuable features for quantum information processing (QIP) [1–5]. Operated as programmable quantum simulators [6–13], these systems realize lattice spin models, which can also be mapped onto other physical models relevant for applications. Realizing these models provides access to their static and dynamic properties in regimes inaccessible to classical simulators.

Realizing QIP protocols often requires initializing the system in a deterministic state. For neutral atoms, this requirement implies assembling a configuration of individual atoms according to some predefined spatial geometry. Programmable arrays of optical traps address this requirement by enabling the assembly and dynamic rearrangement of individual atoms in optical traps.

Finding sequences of displacement operations to prepare specific configurations of atoms from arbitrary ones requires solving atom reconfiguration problems [14,15]. These problems are hard combinatorial optimization problems, each with a potentially exponentially large number of valid solutions. These problems are even harder when considering that atoms might be lost during control operations.

Exact and approximation algorithms that admit efficient implementations are known only for the simplest cases [15,16]. These algorithms might not scale well with the problem size, preventing an implementation fast enough for real-time operation. For example, in the absence of loss, exact algorithms exist for minimizing the total distance traveled by all atoms. The Hungarian algorithm [17–19] is such an exact algorithm with a complexity of $O(N_t^3)$, where N_t is the number of traps.

Heuristic algorithms [14,20–24] address the limitations of exact and approximation algorithms. They trade off guarantees of optimal performance for operational simplicity and computational speed. Moreover, they exploit *ad hoc* principles derived from intuition and experience to define specific control subroutines. For example, parallel control operations might be exploited to speed up operational runtime [22,24].

In this paper we introduce the redistribution-reconfiguration (red-rec) algorithm. Red-rec is a simple and fast heuristic algorithm designed for preparing two-dimensional (2D) configurations of atoms with lattice geometries. We numerically benchmark its performance against exact and approximation algorithms, in both the presence and absence of loss. We specifically focus on the problem of preparing centered-compact configurations of atoms on grids. We show that red-rec can readily be used to prepare configurations of a few thousand atoms with high success probability at a fast rate. We also describe an approach to reduce the mean wait time between the preparation of two successful configurations. This speedup is

*alexandre.cooper@uwaterloo.ca

achieved by rejecting configurations containing fewer atoms than a given threshold. However, our evaluation exposes important challenges in scaling up configurations of atoms to larger sizes. These results support the main three goals of this paper: (i) introducing a hardware-agnostic formalism for reconfiguration problems to facilitate the shared development of reconfiguration algorithms, (ii) providing a systematic approach to quantify the performance of reconfiguration algorithms in the absence and in the presence of loss, and (iii) describing the red-rec algorithm for preparing compact-centered configurations of atoms on grids and other lattice geometries.

The presentation of our results proceeds as follows. We first introduce the theory of atom reconfiguration problems (Sec. II), as well as a benchmarking method to numerically evaluate the performance of reconfiguration algorithms (Sec. III). We then describe the red-rec algorithm (Sec. IV) and quantify its performance at preparing centered-compact configurations of atoms in 1D chains (Sec. V) and 2D grids (Sec. VI), in both the absence and presence of loss. We then introduce and characterize an approach to decrease the mean wait time between two successful configurations of atoms (Sec. VII). We conclude by reviewing key results and outlining opportunities to further improve the red-rec algorithm (Sec. VIII). In the Appendixes we review exact and approximation algorithms (Appendix A), benchmark their performance against the red-rec algorithm in the absence of loss (Appendix B), and prove the correctness of the red-rec algorithm (Appendix C).

II. ATOM RECONFIGURATION PROBLEMS

To facilitate the shared development of reconfiguration algorithms, we begin our discussion with formal definitions of fundamental concepts. These definitions aim to reduce ambiguity in nomenclature, e.g., distinguishing configurations of atoms from arrays of optical traps.

We refer to an atom at $\vec{x} \in \mathbb{R}^3$ by $a(\vec{x}, \tilde{p})$, where $\tilde{p} \in [0, 1]$ is the probability of detecting the atom by performing a perfect measurement at \vec{x} . A perfect measurement at $\vec{x} \in \mathbb{R}^3$, $\pi(\vec{x})$, detects an atom at \vec{x} if and only if there is an atom at \vec{x} . If a perfect measurement detects an atom at \vec{x} , then the probability \tilde{p} of $a(\vec{x}, \tilde{p})$ is updated from a probabilistic value $\tilde{p} \in [0, 1]$ to a deterministic value $p \in \{0, 1\}$. An atom $a(\vec{x}, p)$ with $p = 0$ is said to have been lost and can be disregarded.

A configuration of atoms $\tilde{\mathcal{C}} = \{a_j(\vec{x}_j, \tilde{p}_j)\}_{j=1}^{N_a}$ is defined as a collection of N_a atoms that can each be detected by a collection of perfect measurements performed at the location of each atom $\Pi = \{\pi_j(\vec{x}_j)\}_{j=1}^{N_a}$. A deterministic configuration of atoms \mathcal{C} is a collection of atoms whose detection probability is unity, i.e., $p_j = 1$ for $1 \leq j \leq N_a$.

We denote a trap at $\vec{x} \in \mathbb{R}^3$ by $t(\vec{x}, \tilde{p})$, where $\tilde{p} \in [0, 1]$ is the probability that the trap t contains an atom $a(\vec{x}, \tilde{p})$. We set the probability \tilde{p} that t contains an atom to the probability of detecting $a(\vec{x}, \tilde{p})$ at \vec{x} .

We define a trap array $\mathcal{A}(V, S) = \{t(\vec{x}_j, \tilde{p}_j)\}_{j=1}^{N_t}$ as a collection of $N_t = |V|$ traps. Here $V = \{\vec{x}_j\}_{j=1}^{N_t}$ specifies the geometry of the trap array and $S = \{\tilde{p}_j\}_{j=1}^{N_t}$ specifies the occupation state of the trap array. The trap array \mathcal{A} contains

the configuration of atoms $\tilde{\mathcal{C}} = \{a(x_j, \tilde{p}_j) \mid x_j \in V(\mathcal{A}), \tilde{p}_j > 0\}_{j=1}^{N_t}$ that can be mapped onto a deterministic configuration of atoms \mathcal{C} by performing perfect measurements at the location of each trap of the array.

We distinguish a dynamic trap array from a static trap array depending on whether the location and physical properties of the traps can change over time. The set of all dynamic traps is chosen to contain the static trap array. In this way, an atom in a static trap can be displaced to any other static trap by actuating the dynamic trap array. We refer to the set of static traps containing the target configuration of atoms as the target region of the array. We refer to the set of all remaining static traps as the storage region.

In the remainder of this paper we focus our attention on trap arrays located in a restricted field of view, which is defined as the restricted plane $[-L/2, L/2]^2 \subset \mathbb{R}^2$ whose square side length is equal to $L \in \mathbb{R}$. This restricted plane replaces \mathbb{R}^3 in the previous definitions. Typical geometries V in \mathbb{R}^2 include (a) Bravais lattices (oblique, rectangular, centered rectangular, square, or hexagonal), which are specified by their origin and generator vectors; (b) sublattices, which are specified by their origin, generator vectors of the parent lattice, and coordinate numbers (indicating which elements in the parent lattice are present); and (c) arbitrary geometries, which are specified by a list of spatial coordinates.

We generally refer to a column as the primary generator vector of the static trap array along which the number of traps is the largest. When the generator vectors of the lattice are not orthogonal, columns and rows can be understood as distinct primary and secondary generator vectors of the array.

A. Reconfiguration protocols

Atom reconfiguration problems are hidden stochastic problems: The presence of an atom in each trap is the realization of a random variable (stochastic) whose value remains unknown until a measurement is performed (hidden). These problems differ from deterministic reconfiguration problems [25–28], which have been studied in fundamental, graph-theoretic contexts [29,30] and operational, application-specific contexts [31,32].

An atom reconfiguration problem seeks a sequence of control operations to prepare a predetermined target configuration of atoms \mathcal{C}_T from an arbitrary one. An atom reconfiguration problem is said to be solvable if \mathcal{C}_T is reachable in a finite number of control steps. Furthermore, it is efficiently solvable if it is solvable and the number of control steps in a solution is polynomial in the size of the static trap array (the number of static traps N_t). An optimal solution with respect to a given cost function is a solution that minimizes the cost function, e.g., the total number of displacement operations.

The solution of an atom reconfiguration problem is a reconfiguration protocol $\mathcal{R} = (\Pi_0, \mathbf{T}_1, \Pi_1, \dots, \mathbf{T}_K, \Pi_K)$. This reconfiguration protocol is defined as an initial measurement Π_0 followed by a sequence of K reconfiguration cycles, each cycle comprising a control sequence \mathbf{T}_k and a perfect measurement Π_k . The flowchart of a typical reconfiguration protocol is presented in Fig. 1.

The k th control sequence is a sequence of elementary control operations $\mathbf{T}_k = (T_{k,1}, T_{k,2}, \dots, T_{k,L_k})$, where $T_{k,l}$ is the

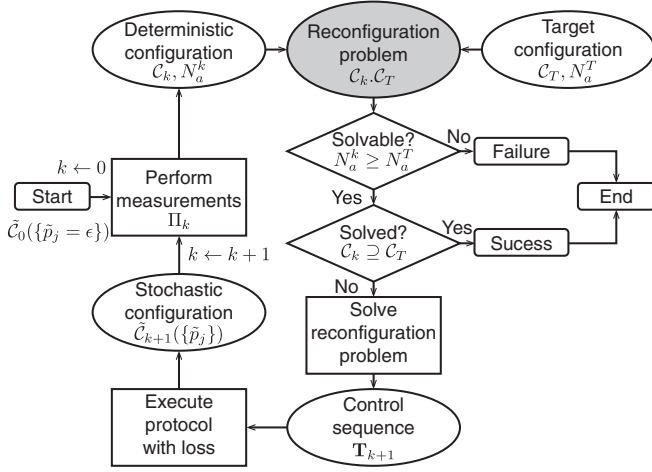


FIG. 1. Flowchart of a typical reconfiguration protocol. The reconfiguration protocol starts by performing measurements on an initial configuration of atoms \tilde{C}_0 . In simulation, \tilde{C}_0 is obtained by adding an atom in each trap of the static trap array and equating the corruption of each atom \tilde{p}_j to the loading efficiency ϵ . The measurements return a deterministic configuration of atoms. A reconfiguration problem is formulated using the measured and desired target configuration of atoms C_T . The protocol first checks if the problem is solvable. If there are fewer atoms than desired, the protocol declares failure and terminates. Otherwise, the protocol checks if the problem is already solved. If the measured configuration contains the desired configuration, the protocol declares success and ends. Otherwise, if the problem is solvable but not yet solved, the protocol solves the reconfiguration problem using the specified reconfiguration algorithm. The algorithm returns a control sequence T_{k+1} to be executed on a real or simulated system. The execution of the protocol results in atom loss. Each atom in the resulting configuration of atoms $\tilde{C}_{k+1}(\{\tilde{p}_j\})$ is corrupted, i.e., the j th atom might have been lost with probability \tilde{p}_j . Starting with a collection of measurements, the protocol then undergoes another reconfiguration cycle until it ends in failure or success.

l th elementary control operation of the k th reconfiguration cycle and L_k is the number of elementary control operations required to execute the k th control sequence. Each elementary control operation $T_{k,l}$ is chosen from a set of six elementary control operations $\{T_{\alpha,v}^{\pm,0}\}$, defined as follows: (i) transfer operations T_{α}^{\pm} , which include extraction operations T_{α}^{+} and implantation operations T_{α}^{-} that transfer atoms from (to) the static trap array to (from) the dynamic trap array; (ii) displacement operations $T_{v_j}^{\pm}$, which displace a dynamic trap by an elementary step along the generator vector \vec{k}_j of the static trap array, either forward $T_{v_j}^{+}$ or backward $T_{v_j}^{-}$; and (iii) no-op operations $\{T_{\alpha}^0, T_v^0\}$, which leave some traps undisturbed while transfer and displacement operations are performed on other traps. Explicitly defining no-op operations facilitates accounting for the loss of idle atoms in both static and dynamic traps.

The k th measurement step reveals the configuration of atoms contained in the static trap array, effectively projecting \tilde{C}_k into the deterministic configuration C_k of N_a^k atoms. If the number of detected atoms is less than the number of desired atoms $N_a^k < N_a^T$, then the reconfiguration problem is said to

be unsolvable. In this case, the protocol is said to have failed and the execution of the protocol is aborted. If the number of detected atoms is greater than or equal to the number of desired atoms $N_a^k \geq N_a^T$ or, more generally, greater than a given threshold $N_a^k \geq N_a^{\text{thresh}} \geq N_a^T$ (see Sec. VII), then the reconfiguration problem is said to be solvable. If the problem is solvable and if the measured configuration contains the target configuration $C_k \supseteq C_T$, then the problem is said to have been solved in $K = k$ reconfiguration cycles and the execution of the protocol is terminated. If the problem is solvable but has not been solved yet, then the execution of the protocol continues until the target configuration has been reached or is no longer reachable because the number of remaining atoms is no longer sufficient. After solving the problem, the surplus atoms located outside the target region are discarded, e.g., by turning off the corresponding static traps outside the target region.

B. Control operations

Control operations might update the state S of both the static trap array $\mathcal{A}_s(V_s, S_s)$ and the dynamic trap array $\mathcal{A}_d(V_d, S_d)$, as well as the geometry V_d of $\mathcal{A}_d(V_d, S_d)$. Transfer operations $T_{\alpha}^{\pm}(\vec{x}_i)$ update the state of the traps at \vec{x}_i in both the static and dynamic trap arrays. Given a static trap $t_s(\vec{x}, \tilde{p})$ at $\vec{x} = \vec{x}_i$ containing an atom with probability \tilde{p} and an overlaid dynamic trap $t_d(\vec{x}, 0)$ at \vec{x} containing no atom, an extraction operation transfers an atom from the static trap to the dynamic trap. An extraction operation thus updates the static and dynamic traps from $[t_s(\vec{x}, \tilde{p}), t_d(\vec{x}, 0)]$ to $[t_s(\vec{x}, 0), t_d(\vec{x}, \tilde{p})]$, where $\tilde{p}' = \tilde{p}p_{\alpha} \exp(-t_{\alpha}/t_{\text{trap}})$ is the updated state probability (see the next paragraph for relevant definitions).

Similarly, an implantation operation transfers an atom from a dynamic trap to a static trap, updating the static and dynamic traps from $[t_s(\vec{x}, 0), t_d(\vec{x}, \tilde{p})]$ to $[t_s(\vec{x}, \tilde{p}'), t_d(\vec{x}, 0)]$. A transfer operation acting on a pair of empty traps returns empty traps, as no atom can be spontaneously created. In contrast, a transfer operation acting on a pair of occupied traps returns a pair of empty traps due to collisional loss. These definitions assume that the static and dynamic traps are overlaid; otherwise, implanting an atom in free space leads to its loss. In the same vein, elementary displacement operations $T_{v_j}^{\pm}$ acting on the dynamic trap $t_d(\vec{x}, \tilde{p})$ at \vec{x} will displace the trap and the atom it contains from \vec{x} to $\vec{x} \pm \vec{k}_j$, where \vec{k}_j is the generator vector of the static trap array, e.g., $\vec{k}_j \in \{\vec{k}_x, \vec{k}_y\}$ with $\vec{k}_x \perp \vec{k}_y$ on the grid.

Control operations acting on different individual traps may be executed in parallel by batching them into a sequence of control operations acting on multiple traps simultaneously. The batching rules used to batch control operations are defined by operational constraints, e.g., the maximum number of dynamic traps, and the set of allowed operations, e.g., whether noncontiguous blocks of atoms can be simultaneously displaced without affecting other trapped atoms.

In the presence of loss, each elementary control operation acting on a trap is imperfect, possibly inducing the loss of the atom it contains. We define the survival probabilities for transfer and displacement operations as p_{α} and p_v , respectively. To account for the finite trapping lifetime, we multiply the survival probability by the exponential decay factor $\exp(-t/t_{\text{trap}})$, where $t \in \{t_{\alpha}, t_v\}$ is chosen as the time

duration of each operation and t_{trap} is the trapping lifetime. We thus refer to the corruption of an atom as the cumulative probability of it being lost given its control history. More formally, $\tilde{p} = p_{\alpha}^{N_{\alpha}} p_v^{N_v} \exp(-t_{\text{tot}}/t_{\text{trap}})$, where N_{α} and N_v are the numbers of transfer and displacement operations performed on the atoms, respectively, and $t_{\text{tot}} = N_{\alpha}^{\Sigma} t_{\alpha} + N_v^{\Sigma} t_v$ is the total duration of the reconfiguration cycle computed from the total number $N_{\alpha,v}^{\Sigma}$ of batched control operations. Our calculation ensures that we do not double count operations performed in parallel.

We choose \tilde{p} to denote both the detection probability and the corruption, as the two quantities are equivalent for perfect measurements. The corruption thus corresponds to the probability of identifying an atom as lost upon performing a perfect measurement. These definitions imply a definition for the survival probability for no-op operations performed on certain atoms while others experience transfer or displacement operations $p_{0,\alpha,v}^{\alpha,v} = \exp(-t_{\alpha,v}/t_{\text{trap}})$. For simplicity, we assume that the no-op survival probability is the same for both static and dynamic traps. We also assume that control errors are limited to loss, that is, we ignore the probability of an atom being left behind in its original trap during a transfer operation. Such an error can however be modeled as a decrease in survival probability.

Finally, to simplify their experimental implementation, we restrict displacements to elementary displacement operations executed in constant time along the generator vectors of the static trap array. Furthermore, to ensure compatibility with acousto-optic deflectors generating dynamic trap arrays, we restrict parallel control operations to chains of atoms. These chains correspond to horizontal subrows or vertical subcolumns in the grid of the static trap array. More elaborate strategies for parallel control operations might achieve greater performance. These strategies include simultaneously displacing contiguous or noncontiguous blocks of atoms or displacing atoms along arbitrary trajectories at variable speed.

III. METHODS FOR QUANTIFYING PERFORMANCE

We now describe our approach to quantify the performance of reconfiguration algorithms. We categorize the performance into three distinct types: (a) Operational performance measures how fast and with which success probability a target configuration of atoms can be prepared, (b) runtime performance measures how fast a specific implementation of the algorithm can be executed on a processor, and (c) algorithmic performance measures the complexity of the algorithm and its scaling with relevant parameters. In the following evaluation we specifically focus on the operational performance.

A. Operational performance in the absence of loss:

Baseline success probability

In the absence of loss, we quantify operational performance by the total number of transfer and displacement operations. We further quantify operational performance in terms of the baseline success probability. The baseline success probability is defined as the probability that the initial configuration of atoms contains at least as many atoms as needed to prepare the target configuration of atoms, $p_0 = P(N_a^0 \geq N_a^T)$.

The number of atoms in the initial configuration (N_a^0) follows a binomial distribution $N_a^0 \sim B(N_t, \epsilon)$. The binomial distribution is determined by the size of the static trap array N_t and the mean loading efficiency ϵ .

The baseline success probability is thus given by the complementary cumulative distribution function of the binomial distribution $p_0 = B_{\text{CCDF}}(N_t, \epsilon; N_a^T)$, which has a known analytical expression. The baseline success probability provides an exact upper bound on the mean success probability in the presence of loss. It also provides an estimate of the minimum number of static traps required to prepare a given target configuration of atoms with a given success probability.

Quantifying performance in the absence of loss decouples the analysis from the realization of random processes and the choice of specific experimental parameters like the duration of control operations. It also provides the means to establish an upper bound on the operational performance in the presence of loss.

Furthermore, quantifying performance in the absence of loss enables a quantitative comparison against exact and approximation algorithms (see Appendix A). Here we benchmark the operational performance of red-rec against two algorithms: an exact assignment algorithm and a three-approximation algorithm. The exact assignment algorithm exactly minimizes the total number of displacement operations. The three-approximation (3-approx) algorithm seeks to minimize the total number of transfer operations; it provides a solution where the total number of transfer operations is at most three times larger than the optimal value (see Appendix B).

B. Operational performance in the presence of loss:

Mean success probability

In the presence of loss, we quantify the operational performance in terms of the mean success probability \bar{p} . The mean success probability cannot exceed the baseline success probability, i.e., $\bar{p} \leq p_0$. To calculate the mean success probability, we average the probability of successfully preparing a target configuration of atoms over an ensemble of randomly sampled initial configurations of atoms and randomly sampled realizations of loss processes.

The mean success probability finds practical applications in remote-access scenarios for which the user is different from the operator. The user first specifies the desired configuration of atoms. The operator then confirms whether this configuration of atoms is achievable. If achievable, the operator informs the user with which probability the configuration of atoms will be prepared. The mean success probability depends on the number of reconfiguration protocols executed and the mean wait time between two successful protocols; it thus correlates with total execution cost.

The performance of different shared-access platforms can thus be compared by computing the mean success probability for a given target configuration or the largest configuration achievable at a fixed mean success probability. Reporting the largest configuration that can be prepared without indicating the mean number of trials required to prepare that configuration is otherwise misleading, as obtaining a configuration of that size might rely on very unlikely sampling realizations.

C. Operational performance benchmarking module

We numerically quantify the operational performance of reconfiguration algorithms using a custom-built benchmarking module that exploits the Monte Carlo sampling method implemented in PYTHON.

The module takes as input a reconfiguration problem, a set of physical parameters, and a set of sampling parameters. The reconfiguration problem defines the static trap array, including its geometry and dimensions, and the target atom configuration. The physical parameters include the loading efficiency and the survival probability for displacement and transfer operations. The sampling parameters include the number of simulation repetitions, the number of atoms in the initial configuration (if fixed), and the name of the algorithm to be used.

First, the module constructs a configuration of N_a^0 atoms. If N_a^0 is not fixed, the module samples N_a^0 from the binomial distribution defined by the loading efficiency ϵ and total number of traps N_t . Given N_a^0 , the module then randomly assigns an atom to each of N_a^0 randomly chosen traps. It does so, for example, by randomly selecting N_a^0 integers in a list of N_t trap indices.

Second, given the initial and target configurations, the module solves the reconfiguration problem using the specified algorithm. The solution is a sequence of control operations. The module performs these operations on a simulated system to compute the numbers of transfer, displacement, and idle operations for each atom. Given the history of control operations for each atom and the survival probability for each control operation, the module computes the corruption of each atom.

Third, given the corruption of each atom, the module generates a configuration of atoms that would be obtained following a perfect measurement. This configuration is computed by sampling a number μ from the uniform distribution on $[0,1]$ for each atom. If $\mu > \tilde{p}$, the atom is lost, i.e., the state of the trap containing the atom is set to 0. If $\mu \leq \tilde{p}$, the atom is kept in the trap and the state of the trap containing the atom is set to 1.

These configuration sampling and problem solving steps continue until the problem has been solved or is no longer solvable. Each data point is typically averaged over 1000 initial configurations.

D. Typical experimental parameters and limitations

To quantify the expected performance under realistic experimental conditions, we numerically compute the mean success probability using realistic physical parameters. We choose the duration of elementary transfer and displacement operations to be $t_\alpha = 15 \mu\text{s}$ and $t_v = 67 \mu\text{s}$, respectively, and the survival probability to be $p_v = p_\alpha = 0.985$ [22]. The survival probability typically decreases with the displacement distance; however, this decrease can typically be mitigated by increasing the displacement time [33]. We conservatively choose the trapping lifetime to be $t_{\text{trap}} = 60 \text{ s}$, although it could be further enhanced up to at least 6000 s using cryogenics [34].

We further choose the mean loading efficiency to be $\epsilon = 0.60$, which is typical for ^{87}Rb atoms loaded in the collisional blockade regime [35]. An efficiency $\epsilon \geq 0.60$ has also been achieved using enhanced-loading techniques [36–42].

To quantify performance under enhanced loading, we choose an increased loading efficiency of $\epsilon = 0.90$. We assume that the loading efficiency is the same for all traps in the static trap array, as can be achieved using closed-loop optimization routines [43].

To quantify the size of the largest configuration of atoms that could be realistically prepared given typical experimental limitations, we restrict the accessible range of trap parameters. An important bottleneck in scaling up atom configurations beyond a few thousand atoms is the limit on the number of optical traps. Limits may be due to restrictions in optical power, diffraction efficiency of active diffractive optical devices like spatial light modulators and acousto-optic deflectors, and transmission efficiency of optical elements. As a realistic, near-term operational limit on the number of optical traps, we choose the maximum size of the static trap array to be $N_t^{\text{max}} = 2048$ traps and set the stretch value to twice that amount, i.e., $N_t^{\text{max}} = 4096$.

A trap array with such a number of traps could be realized with approximately 5 W in laser power at the source given 1 mW per optical trap in the focal plane of the optical system and a total power delivery efficiency of 0.4 from the laser source. This power value is typical for a Ti:sapphire laser operating at 813 nm that is pumped with an 18-W pump laser at 532 nm. Further increase in the number of traps might be achieved by exploiting advances in laser technology. Examples include pumping Ti:sapphire lasers at a higher power, using multiple laser sources, either coherently combined or not (e.g., if operating at a magic trapping wavelength is not crucial), or using high-power fiber lasers and fiber amplifiers.

Another bottleneck is restrictions in the field of view (FOV) of the high-resolution microscopy system used to produce optical traps and image individual atoms. We upper bound the width and height of the trap arrays by $N_t^{y,\text{max}} = N_t^{x,\text{max}} = 100$ traps. Such bounds result, for example, from a grid with an intertrap spacing of 3.0 μm in an optical system with a FOV of 300 μm . An important challenge in increasing the FOV of the optical system is avoiding aberration. Aberration distorts optical traps and their loading efficiency [43]. Significant advances have been made towards solving this problem. A research group in molecular biology has recently reported on the design of a mesolens for operation in the wavelength range of 400–750 nm with a FOV of 6 mm, a lateral resolution of 0.6 μm , and a numerical aperture of 0.5 [44].

IV. THE RED-REC ALGORITHM

In this section we describe the various steps of the red-rec algorithm (see Algorithm 1 and Fig. 2). Before solving the problem, red-rec checks whether the problem is solvable. Given a measured configuration of N_a atoms, where we drop the cycle index k to improve clarity, red-rec compares the number of measured atoms N_a to the number of desired atoms N_a^T . The problem is solvable if $N_a - N_a^T \geq 0$. Additionally, red-rec checks if the atom surplus exceeds a user-defined threshold value N_a^{thresh} , i.e., $N_a - N_a^T \geq N_a^{\text{thresh}}$. If the problem is unsolvable or the threshold is not met, red-rec aborts and declares failure.

If the problem is solvable and the threshold is met, red-rec proceeds to compute the atom imbalance for each column.

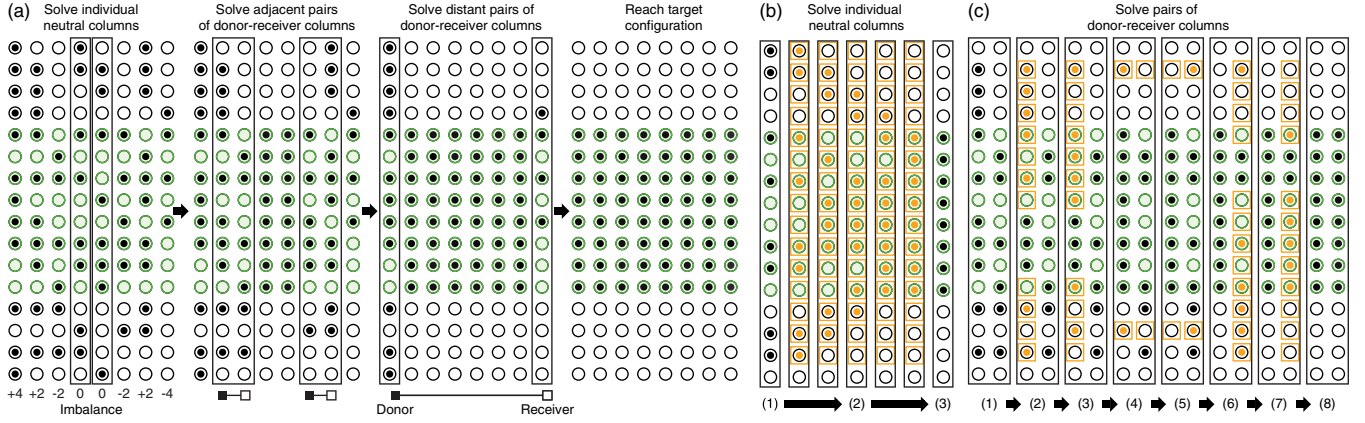


FIG. 2. Single realization of the red-rec algorithm in the absence of loss. (a) Sequence of control operations to prepare a target configuration of $N_a^T = 8 \times 8$ atoms (green shaded disks) in a static trap array of $N_t = 8$ (width of the trap array) \times 16 (height of the trap array) = 128 traps (circles) from an arbitrary configuration of $N_a^0 = 64$ atoms (black dots). After computing the imbalance of each column as the difference between the number of detected and desired atoms, the algorithm sequentially reconfigures individual neutral columns, neighboring pairs of donor-receiver columns, and then distant pairs of donor-receiver columns. (b) An individual neutral column is solved by extracting atoms from the static trap array into the dynamic trap array (orange squares), displacing all atoms to their desired target locations while avoiding collisions, and implanting the atoms back into the static trap array. (c) A pair of donor-receiver columns is solved using a sequence of control operations that [steps (1)–(3)] reconfigures the donor column while displacing atoms to be exchanged to open rows, [steps (4)–(5)] redistributes surplus atoms from the donor column to the receiver column in a single parallel displacement step, and [steps (6)–(8)] reconfigures the receiver column.

The imbalance of the j th column, denoted by ΔN_a^j , is calculated as $N_a^j - N_a^{T,j}$, where N_a^j and $N_a^{T,j}$ represent the numbers

Algorithm 1. The red-rec algorithm.

Require: A trap array $\mathcal{A}(V, S)$ with a grid geometry, the number of rows r , the number of columns c , the number of measured atoms N_a , the number of desired atoms N_a^T , and a user-defined threshold value N_a^{thresh} ,

if $N_a - N_a^T < N_a^{\text{thresh}}$ **then**
 return “failure”
end if

for $j \leftarrow 1$ **to** c **do**
 compute the atom imbalance for column j in \mathcal{A}
 if column j in \mathcal{A} is a neutral column **then**
 solve column j in \mathcal{A} using the exact 1D algorithm
 end if
end for

while there exists a column with negative imbalance in \mathcal{A} **do**
 select the closest pair of donor-receiver columns in \mathcal{A}
 (tie-breaking rules are applied whenever necessary)
 solve the selected pair in \mathcal{A} following six steps:
 1. donor extraction
 2. donor displacement
 3. donor implantation
 4. donor-column redistribution
 5. receiver extraction
 6. receiver displacement and implantation
end while

for $j \leftarrow 1$ **to** c **do**
 if column j in \mathcal{A} is not solved **then**
 solve column j in \mathcal{A} using the exact 1D algorithm
 end if
end for

of atoms in the j th column of the measured and desired configurations, respectively. Here we use the shorthand notation N_a^j to represent $N_a^{k,j}$, where $N_a^{k,j}$ denotes the number of atoms in the j th column during the k th reconfiguration cycle. Based on the imbalance, each column is labeled as (a) a donor, if $\Delta N_a^j > 0$; (b) neutral, if $\Delta N_a^j = 0$; or (c) a receiver, if $\Delta N_a^j < 0$.

After labeling the columns, red-rec breaks down the 2D reconfiguration problem into simpler reconfiguration problems involving individual columns and pairs of columns.

First, red-rec solves each neutral column using an exact reconfiguration algorithm on chains, referred to as the exact 1D algorithm (see Appendix A), which runs in time linear in the number of traps. When applied to an individual column, the algorithm minimizes the total number of displacement operations and ensures that each atom is extracted and implanted at most once. However, in the presence of loss, it does not guarantee the overall minimization of displacement and transfer operations across multiple reconfiguration cycles.

Solving neutral columns facilitates the redistribution of atoms among distant pairs of donor and receiver columns. Doing so creates distribution rows in which no obstructing atoms are present between the donor and receiver columns.

Second, red-rec solves pairs of donor-receiver columns. It begins with adjacent pairs where the donor column can fully satisfy the receiver column, i.e., $\Delta N_{\text{donor}} + \Delta N_{\text{receiver}} \geq 0$. Next, it solves adjacent pairs of columns. Finally, it solves distant pairs, starting with the pairs that can exchange the highest number of atoms in a single redistribution sequence.

Pairs of donor-receiver columns are solved using one or more redistribution sequences [Fig. 2(c)]. To realize a redistribution sequence, the atoms in the donor column are categorized into three groups.

(i) Reconfigured atoms populate the target region of the donor column. They are selected in a way that minimizes the total number of displacement operations. The identification of reconfigured atoms involves utilizing the exact 1D algorithm (see Appendix A).

(ii) Redistributed atoms populate the target region of the receiver column. They are chosen from the storage region of the donor column using a heuristic. The heuristic selects atoms starting from those located farthest away from the target region. These atoms are assigned to the traps in the receiver column that are situated on the nearest distribution row. The selection of redistributed atoms aims to have approximately equal numbers of atoms above and below the target region.

(iii) Idle atoms remain undisturbed in the donor column and do not undergo any displacement.

The redistribution sequence then involves executing the following six steps.

(1) *Donor extraction.* The redistributed and reconfigured atoms are extracted from the donor column [Fig. 2(c), step (2)].

(2) *Donor displacement.* The reconfigured atoms are displaced to traps in the target region of the donor column. The redistributed atoms are displaced to traps located on distribution rows in the storage region of the donor column [Fig. 2(c), step (3)].

(3) *Donor implantation.* The reconfigured atoms in the target region are implanted into the static traps [Fig. 2(c), step (4)]. Meanwhile, the redistributed atoms remain extracted in the dynamic traps.

(4) *Donor-receiver redistribution.* The extracted redistributed atoms are simultaneously displaced to the receiver column [Fig. 2(c), step (5)].

(5) *Receiver extraction.* The atoms in the receiver column are extracted [Fig. 2(c), step (6)].

(6) *Receiver displacement and implantation.* The receiver column is solved using the exact 1D algorithm [Fig. 2(c), step (7)].

This approach ensures that the redistributed atoms are only extracted and implanted once, which helps minimize the overall number of transfer operations.

If not all redistributed atoms can be redistributed from the donor column to the receiver column in a single redistribution sequence, then the redistribution sequence is repeated until the donor column has donated all its surplus atoms. Such a situation could occur, for example, if the number of distribution rows is less than the number of atoms to be distributed. If the receiver column has not been fully solved, then it participates again in the donor-column pairing.

Pairs of donor-receiver columns are thus solved using redistribution sequences until all columns achieve non-negative imbalances. In Appendix C we show that this step requires solving a finite number of donor-receiver column pairs, which implies the correctness of the red-rec algorithm.

Finally, red-rec solves all columns that have not yet been solved. The termination of the second step guarantees that all such columns have positive imbalance, meaning they can be individually solved.

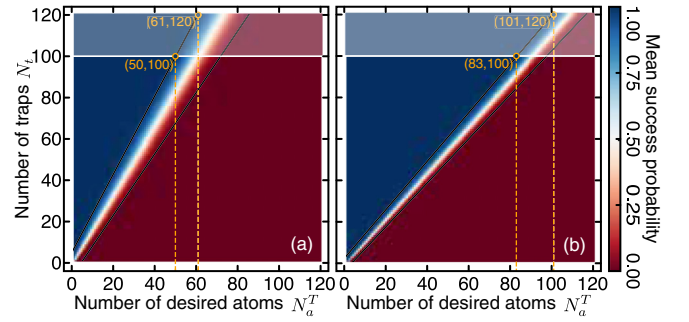


FIG. 3. Baseline success probability for 1D chains: probability of successfully preparing a centered-compact configuration of N_a^T atoms in a 1D chain of N_t traps for a loading efficiency of (a) $\epsilon = 0.60$ and (b) $\epsilon = 0.90$ in the absence of loss. A transition region separates regions of near-certain failure ($\bar{p} \leq 0.02$, red) and near-certain success ($\bar{p} \geq 0.98$, blue). The maximum number of traps is set to $N_t^{\max} = 100$ (white shaded region). The transition midline at $p_0 = 0.5$ is realized for $N_t = N_a^T / \epsilon$ traps, showing that, in the absence of loss, the number of needed traps scales linearly with the number of desired atoms.

V. PERFORMANCE FOR 1D CHAINS

The red-rec algorithm solves atom reconfiguration problems on grids by reducing them to a sequence of reconfiguration problems on 1D chains. These problems on chains can be efficiently solved using the exact 1D algorithm. The performance of the red-rec algorithm is thus intrinsically tied to the performance of the exact 1D algorithm.

We thus start our analysis by numerically evaluating the performance of the exact 1D algorithm (see Sec. III C for its detailed description). We perform the analysis in both the absence and presence of loss. For this analysis, we compute the baseline success probability for preparing a compact chain of N_a^T atoms at the center of a chain of N_t static traps for various values of N_a^T , N_t , and $\epsilon \in \{0.6, 0.9\}$.

In the absence of loss, the baseline success probability surfaces (Fig. 3) exhibit a sharp transition between a region of near-certain failure ($p_0 \leq 0.02$) and a region of near-certain success ($p_0 \geq 0.98$). We note that the values delineating the regions have been chosen arbitrarily without implications for the key results of this paper.

The transition region is centered around the isoprobability line $p_0 = 0.5$ realized for $N_t^0 = \eta_0 N_a^T = N_a^T / \epsilon$, i.e., for a baseline overhead factor $\eta_0 = 1/\epsilon$. The largest centered-compact configuration of atoms that can be prepared in chains of 100 (120) static traps with near-certain success thus contains $N_a^T = 50$ (61) atoms for $\epsilon = 0.60$ and 83 (101) atoms for $\epsilon = 0.90$.

In the presence of loss, the mean success probability surface (Fig. 4) exhibits the same two regions of near-certain success and near-certain failure as in the lossless case. There is a sharp transition region whose boundary is determined by the loading efficiency and loss parameters. In contrast to the lossless case, however, the transition curve is nonlinear. The nonlinearity indicates that an increasingly large overhead factor is required to prepare increasingly large chains of atoms with near-certain success. We find that the overhead ratio for $\bar{p} \approx 0.50$ is well approximated by

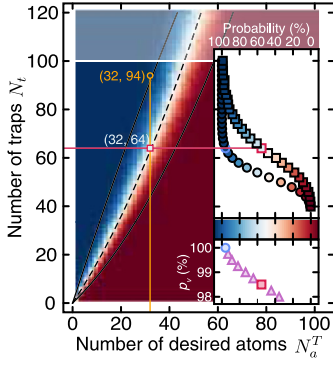


FIG. 4. Mean success probability for 1D chains: mean probability of successfully preparing a centered-compact configuration of N_a^T atoms in a 1D chain of N_t traps for a loading efficiency of $\epsilon = 0.6$ in the presence of loss. The number of traps at the transition midline (dashed line) scales quadratically with the number of desired atoms $N_t \sim O((N_a^T)^2)$. The inset shows the mean success probability of preparing a chain of $N_a^T = 32$ atoms in the absence (disks) and presence (squares) of loss for varying success probability during control operations (bottom). At least $N_t = 64$ traps are required to prepare a compact chain of $N_a^T = 32$ atoms with $\bar{p} \geq 0.5$.

$\eta(N_a^T) \approx \eta_0 + \eta_1 N_a^T$, where $\eta_0 = 1.50(3)$ and $\eta_1 = 0.014(3)$. This statement implies that $N_t = \eta_0 N_a^T + \eta_1 (N_a^T)^2 \sim O((N_a^T)^2)$ traps are required to successfully prepare a chain of N_a^T atoms. The largest centered-compact configuration of atoms that can be prepared with near-certain success in a chain of 100 (120) static traps thus contains $N_a^T = 34$ (42) atoms in the presence of loss, as opposed to $N_a^T = 50$ (61) atoms in the absence of loss.

Because of its relevance for preparing 2D configurations of $32 \times 32 = 1024$ atoms, we further analyze the problem of preparing a chain of $N_a^T = 32$ atoms in an array of $N_t = 64$ traps. The mean success probability is equal to $\bar{p} \approx 0.5$ and increases to unity as the survival probability during control operations increases to unity (inset of Fig. 4). Focusing on successful and unsuccessful reconfiguration protocols (Fig. 5), we observe that the mean number of control cycles is smaller for unsuccessful protocols than for successful protocols. Consequently, protocols are quicker to fail than to succeed: Half of unsuccessful protocols fail after one cycle, whereas half of the successful protocols require at least two cycles to succeed [Fig. 5(a)].

We also observe that successful protocols have more atoms on average in their initial configurations than unsuccessful protocols [Fig. 5(b)]. The surplus atoms are used in late reconfiguration cycles to replace atoms lost in early reconfiguration cycles. This information is also contained in the probability distribution functions for the numbers of control cycles and the numbers of atoms in the initial configuration for successful and unsuccessful protocols [Figs. 5(c) and 5(d)].

These observations indicate that rejecting configurations containing less than a certain number of loaded atoms might prevent the execution of protocols unlikely to succeed (see Sec. VII). For example, the probability that a control protocol is successful given that N_a^0 atoms have been loaded is greater than 0.5 for $N_a^0 \geq 37$; configurations containing fewer atoms could be rejected (see Sec. VII).

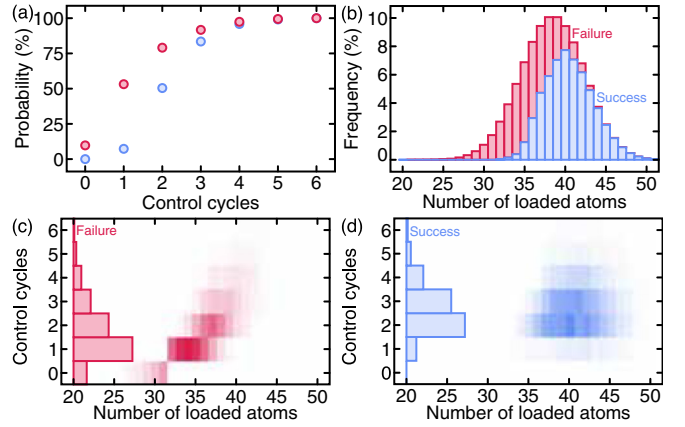


FIG. 5. Control cycles for 1D chains. (a) Cumulative distribution function of the number of control cycles performed during unsuccessful (red) and successful (blue) protocols for preparing a configuration of $N_a^T = 32$ atoms in a chain of $N_t = 64$ traps. Protocols fail faster than they succeed: Half of unsuccessful protocols have failed after one cycle, whereas half of the successful protocols require at least two cycles to succeed. (b) Stacked distributions of the number of atoms loaded in the static trap array. Successful protocols start with more atoms on average. (c) and (d) Distribution of (c) the number of control cycles and (d) the number of atoms loaded in the static array. Integrating over control cycles returns the stacked distributions of (b). Unsuccessful protocols fail in a few control cycles because the number of loaded atoms is not enough to compensate for loss.

VI. PERFORMANCE FOR 2D GRIDS

We now quantify the performance of the red-rec algorithm at solving reconfiguration problems on 2D grids. We specifically focus on the problem of preparing a configuration of $N_a^T = \sqrt{N_a^T} \times \sqrt{N_a^T}$ atoms at the center of a square lattice of $N_t = N_t^x \times N_t^y$ static traps. Here we choose the width of the trap array to be equal to the width of the target configuration, i.e., $N_t^x = \sqrt{N_a^T}$. We further choose the height of the trap array to contain a larger number of traps $N_t^y = \eta \sqrt{N_a^T}$, where $\eta \geq 1$. Figure 2 provides an example of such a grid with rectangular dimensions.

In the absence of loss, the baseline success probability surfaces exhibit a region of near-certain failure ($\bar{p} \leq 0.02$) and a region of near-certain success ($\bar{p} \geq 0.98$) (Fig. 6). The transition region is much sharper than in the 1D case (Fig. 3). In square-lattice trap arrays containing at most $N_t = 2048$ (4096) optical traps, the largest centered-compact configuration of atoms that can be prepared with near-certain success contains $N_a^T = 34^2 = 1156$ ($42^2 = 1756$) atoms [Fig. 6(a)]. Increasing the loading efficiency to $\epsilon = 0.90$ increases the configuration size to $N_a^T = 48^2 = 2304$ ($60^2 = 3600$) atoms [Fig. 6(b)]. Assuming height limits of 100 (120) traps, the largest configuration of atoms that can be prepared with near-certain success contains $N_a^T = 58^2 = 3364$ ($70^2 = 4900$) atoms for $\epsilon = 0.60$ and $N_a^T = 88^2 = 7744$ ($106^2 = 11236$) atoms for $\epsilon = 0.90$. Given a loading efficiency of $\epsilon = 0.60$, preparing a configuration of $100 \times 100 = 10\,000$ atoms with near-certain success would thus require at least $100 \times 330 = 33\,000$ static traps, a formidable operational challenge.

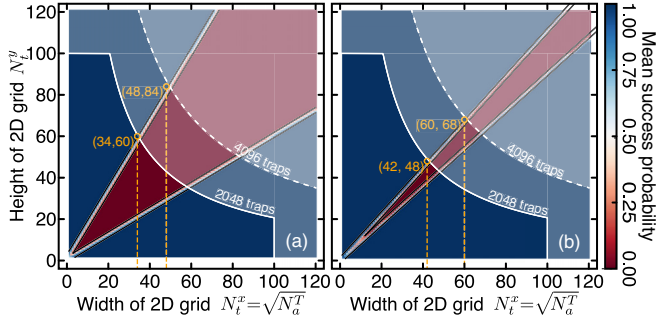


FIG. 6. Baseline success probability for 2D grids: probability of successfully preparing a centered-compact configuration of $N_a^T = \sqrt{N_x^T} \times \sqrt{N_y^T}$ atoms in a square-lattice array of $N_t^x \times N_t^y$ traps, where $N_t^x = \sqrt{N_a^T}$ in the absence of loss. Increasing the loading efficiency from (a) $\epsilon = 0.60$ to (b) $\epsilon = 0.90$ narrows down the failure region (red). The largest achievable configuration is ultimately restricted by the maximum number of optical traps along each direction ($N_t^{x,y} \leq 100$) and in total [$N_t \leq 2048$ (4096)]. As for 1D chains, the transition midline at $p_0 = 0.5$ is realized for $N_t = N_a^T / \epsilon$ traps ($N_t^x = \sqrt{N_a^T} / \epsilon$). In the absence of loss, the number of needed traps scales linearly with the number of desired atoms.

In the presence of loss, the near-success region is achieved for a larger numbers of traps (Fig. 7) than in the lossless case. These additional traps are required to load the surplus atoms needed to compensate for loss during control operations. The transition curve evaluated at $\bar{p} \approx 0.5$ shows that the height of the trap array scales approximately quadratically with the width of the trap array, $N_t^y \sim O((N_t^x)^2) = O(N_a^T)$, and thus linearly with the number of atoms. Choosing $N_t^y \approx \eta_0 + \eta_1 N_a^T$, we find that $\eta_0 = 0.01564$ and $\eta_1 = 1.579$ from linear regression. This relationship implies that the number of traps required to assemble a configuration of N_a^T atoms scales as $N_t = N_t^x \times N_t^y \sim O((N_a^T)^{3/2})$, in contrast to the linear scaling observed in the absence of loss. This result is consistent with the red-rec algorithm solving ap-




FIG. 7. Mean success probability for 2D grids: mean probability of successfully preparing a configuration of $N_a^T = \sqrt{N_x^T} \times \sqrt{N_y^T}$ atoms in a grid of $N_t^x \times N_t^y$ traps in the presence of loss. The loading efficiency is set to $\epsilon = 0.60$. The midline curve of the transition region scales quadratically with the width $N_t^x = \sqrt{N_a^T}$ of the grid, $N_t^y \sim O((N_t^x)^2) = O(N_a^T)$, and thus linearly with the number of traps. In the presence of loss, the number of traps needed thus scales as the $3/2$ power of the number of desired atoms $N_t \sim O((N_a^T)^{3/2})$.




FIG. 8. Largest achievable configuration of atoms: maximum configuration size achievable with near-certain success in the absence (blue disks) and presence (red squares) of loss for a loading efficiency of $\epsilon = 0.6$. Preparing a configuration of $32^2 = 1024$ atoms with near-certain success requires 2304 traps. The inset shows that the maximum configuration size achievable with 2048 traps increases with the survival probability during displacement and transfer operations. We choose $p_v = p_\alpha$. The red square is the data point associated with the specific survival probability used in the performance study. The purple triangles are the data points associated with other typical values for the survival probability.

proximately $N_t^x = \sqrt{N_a^T}$ columns, each column being solved with the 1D exact algorithm, which scales as $N_t^y \sim O(\sqrt{N_a^T})^2$ (see Fig. 4).

The largest configuration of atoms that can be prepared with near-certain success increases with the number of traps (Fig. 8) and the survival probability during control operations. Fixing the total number of traps to be $N_t^{\max} = 2048$ (32×64), the mean success probability for preparing a configuration of 32×32 atoms is $\bar{p} = 0.21(1)$. Increasing the total number of traps to $N_t = 2304$ (32×72) increases the mean success probability to $\bar{p} = 0.993(2)$. Increasing the survival probability during control operations from $p_v = p_\alpha = 0.985$ to $p_v = p_\alpha = 0.990$ increases the mean success probability from $\bar{p} = 0.21(1)$ to $\bar{p} = 0.996(4)$ (inset of Fig. 8). Further increasing the total number of traps to be $N_t^{\max} = 4096$ limits the configuration size for near-certain success to $N_a^T = 40^2 = 1600$ atoms. The results highlight once again the challenges of scaling to large configuration sizes.

To identify the conditions required for red-rec to succeed, we compute the mean total number of displacement and transfer operations performed during each reconfiguration cycle (Fig. 9). We restrict our analysis to the problem of preparing a 2D configuration of $N_a^T = 32 \times 32$ atoms in a static array of $N_t = 32 \times 64$ traps.

First, we observe that successful protocols perform fewer transfer and displacement operations during their first reconfiguration cycle [Fig. 9(a)] and more control operations in all their subsequent reconfiguration cycles than unsuccessful protocols [Fig. 9(b)]. Second, we observe that the relative total number of control operations decreases in later reconfiguration cycles. Specifically, fewer operations lead to less loss and thus fewer atoms to be replaced, *ad infinitum*, so that the late cycles perform fewer control operations to replace fewer lost atoms [Fig. 9(c)]. Third, we observe that the relative numbers of displacement and transfer operations increase with the size of the configuration of atoms [Fig. 9(d)]. We explain the increase in the relative numbers of displacement operations by the fact that, in larger arrays, surplus atoms left in their original positions must travel large distances to replace




FIG. 9. Control operations for 2D grids. (a) Distribution of the total number of transfer and displacement operations during the first reconfiguration cycle of successful (blue) and unsuccessful (red) protocols. The protocols seek to prepare a configuration of $N_a^T = 32 \times 32$ atoms in an array of $N_t = 32 \times 64$ traps. During the first cycle, unsuccessful protocols perform more operations than successful protocols. (b) Distribution of total transfer and displacement operations performed during all the subsequent reconfiguration cycles expressed relative to the first reconfiguration cycle. Successful protocols perform more displacement and transfer operations than unsuccessful protocols during their subsequent reconfiguration cycles. (c) Probability that a successful protocol will perform at least N_c control cycles (shaded disks). Less than 10% (2%) of successful protocols perform more than 15 (19) reconfiguration cycles. The total numbers of transfer (orange) and displacement (green) operations performed during successful protocols, expressed relative to the first cycle, monotonically decrease. (d) Given a fixed overhead factor (here $\eta = 2$ with $N_t^y = \eta N_t^x$), these numbers increase with the desired configuration size: Increasingly more operations are performed in subsequent reconfiguration cycles to compensate for loss.

lost atoms in subsequent cycles. Because more displacement operations result in greater loss, more reconfiguration cycles are required to replace the lost atoms. We explain the increase in the relative numbers of transfer operations by the fact that receiver columns in late reconfiguration cycles are typically paired with multiple donor columns. Each pairing requires executing multiple redistribution sequences, each of which requires extracting and implanting the atoms.

VII. REDUCING MEAN WAIT TIME VIA CONFIGURATION REJECTION

Besides improving the mean success probability of preparing a target configuration of atoms, an important operational requirement is increasing the duty cycle of experiments by reducing the mean wait time between two successful reconfiguration protocols. The mean success probability depends on the number of atoms in the initial configuration. These atoms are needed to prepare the target configuration of atoms and replace atoms lost during the multiple reconfiguration cycles.

We partition the distribution of the number of atoms in the initial configuration partitioned for successful and unsuccessful protocols [Fig. 10(a)]. We identify a first threshold below which a protocol is guaranteed near-certain failure ($\bar{p} \leq 0.02$)

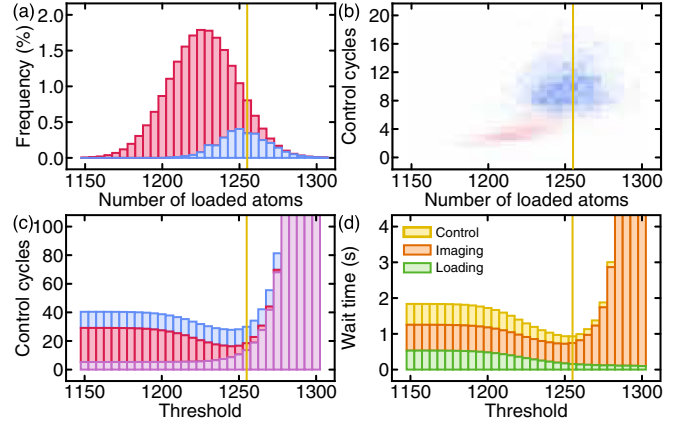


FIG. 10. Reducing mean wait time via configuration rejection. (a) Stacked distributions of the number of atoms loaded into a square-lattice array of $N_t = 32 \times 64 = 2048$ traps for successful (blue) and unsuccessful (red) protocols when preparing a configuration of $N_a^T = 32 \times 32 = 1024$ atoms. The threshold (vertical orange line) indicates the number of atoms N_a^{thresh} below which configurations are rejected. (b) Distribution of the number of control cycles per control protocol and the number of atoms loaded for successful (blue) and unsuccessful (red) protocols. Integrating over control cycles returns the stacked distributions of (a). (c) Number of measurements performed between two successful protocols for various threshold values. Measurements are performed whenever the algorithm fails to load more than N_a^{thresh} atoms (purple) or fails (red) or succeeds (blue) in preparing the target configuration given more than N_a^{thresh} atoms. (d) Stacked wait times between two successful protocols for various threshold values. The time is spent loading the magneto-optical trap (green), imaging configurations of atoms (orange), or performing control operations (yellow). The threshold (vertical orange line) is chosen so as to minimize the total wait time.

and a second threshold below which a protocol is more likely to fail than to succeed. Rejecting configurations containing fewer atoms than either one of those two thresholds would increase the probability of successfully solving the reconfiguration problems for the retained configuration. However, even though the probability of success increases with the number of atoms in the initial configuration, the probability of measuring configurations with such a large number of atoms decreases. Thus, rejecting these configurations monotonically decreases the probability of success. Still, because unsuccessfully solving a problem might require multiple reconfiguration cycles [Fig. 10(b)] and thus more time than sampling a new configuration, there exists an optimal rejection threshold (orange vertical line in Fig. 10) that minimizes the mean wait time between two successful protocols. This threshold is obtained by discarding as many unsuccessful configurations as possible while rejecting as few successful configurations as possible without letting the number of measurements diverge [Fig. 10(c)].

The mean number of measurements needed to sample a configuration of atoms that contains more atoms than the threshold is $1/P(N_a^0 \geq N_a^{\text{thresh}})$, where $P(N_a^0 \geq N_a^{\text{thresh}})$ is given by the cumulative distribution function of the binomial distribution. The mean amount of time spent preparing a configuration of atoms containing more atoms than the threshold

is therefore $t_{\text{MOT}} + t_{\text{image}}/P(N_a^0 \geq N_a^{\text{thresh}})$. Here t_{MOT} is the time needed to cool and trap atoms in a magneto-optical trap (MOT) and t_{image} is the time needed to image the configuration of atoms. We calculate the time elapsed between successful preparations of the target configuration by considering the MOT loading time, the time required to measure a configuration of atoms (imaging time), and the time required to execute the control operations. We take the MOT loading time to be $t_{\text{MOT}} = 100$ ms, the time required to image the static array to be $t_{\text{image}} = 20$ ms, and the execution time of parallel transfer and displacement operations to be $t_\alpha = 15$ μ s and $t_\nu = 67$ μ s, respectively (see Sec. III D). Atoms are loaded into the MOT and then continuously imaged in the static trap array until a configuration with sufficient atoms is measured. The MOT loading time is thus incurred only once for each attempted reconfiguration protocol.

From our numerical analysis, we find that the threshold value at which the minimum wait time occurs is $N_a^{\text{thresh}} = 1255$ atoms. As a result, a total of 87.6% of all configurations are discarded. Among these discarded configurations, 90.7% of them would have otherwise resulted in failure. After thresholding, 61.2% of control protocols result in success. For a configuration of $N_a^T = 32 \times 32$ atoms in a square-lattice array of size $N_t = 32 \times 64$, thresholding reduces the mean time between successful preparations of the target configuration by a factor of approximately 2, from 1.84 s to 0.932 s. We also observe that post-thresholding, the distributions in the mean number of reconfiguration cycles and the control time for successful and unsuccessful protocols are similar [Fig. 10(d)]. These results indicate that achieving success depends mostly on random realizations of atom loss rather than on the number of atoms in the initial configuration.

VIII. CONCLUSION

We introduced the red-rec algorithm to solve atom reconfiguration problems on lattices, specifically focusing on grids. The red-rec algorithm reduces reconfiguration problems on grids to simpler reconfiguration problems on individual columns and pairs of columns. These problems are solved quickly and efficiently using simple algorithms. Individual columns are solved using the exact 1D algorithm; pairs of donor-receiver columns are solved heuristically by redistributing atoms from a donor column to a receiver column.

Although we described the red-rec algorithm over grids, it can be extended to other lattice geometries, as well as other sublattice geometries embedded in the grid. Examples are checkerboard sublattices and oblique lattices with rhomboidal boundaries, including triangular lattices (with two generator vectors of equal norm oriented at a relative angle of $\pi/3$; e.g., see Ref. [24]).

We numerically compared the performance of the red-rec algorithm against exact and approximation algorithms, in both the absence and presence of loss. In the absence of loss, we showed that red-rec performs well at minimizing displacement operations. However, it performs more transfer operations than the value achieved by the 3-approx algorithm (see Fig. 12 in Appendix B). Further reduction in transfer operations might be achieved by allowing for the redistribution of atoms in the target region and solving for blocks of

columns containing more than two columns, e.g., containing two donors and one receiver.

Improved reconfiguration algorithms enable preparing large configurations of atoms for a given mean success probability. In the presence of loss, red-rec prepares configurations of 256 atoms in 512 traps with a mean success probability of $\bar{p} = 0.913(2)$. Moreover, it prepares configurations of 1024 atoms in 2048 traps with a mean success probability of $\bar{p} = 0.21(1)$.

Equivalently, improved algorithms facilitate achieving fast preparation times for a given configuration sizes. Rejecting configurations of atoms containing fewer atoms than a given threshold enables even faster preparation times.

The red-rec algorithm is readily and efficiently implementable. It is also compatible with standard acquisition and control systems. Future work should focus on extending red-rec to other geometries and higher dimensions, implementing it on parallel-processing computing devices, and quantifying its operational performance experimentally.

In addition, further opportunities exist to develop exact and approximation algorithms that simultaneously optimize multiple-objective functions. These algorithms might in turn guide the development of improved exact-heuristic algorithms for atom reconfiguration problems, which might also support other domains of applications.

Finally, we refer the interested reader to our complementary study on the assignment-rerouting-ordering (ARO) algorithm [45]. The ARO algorithm exploits subroutines to improve the performance of typical assignment algorithms. Among others, we provide a constructive algorithm to find a partial ordering of displacement trajectories that guarantee that each atom is displaced at most once.

The source code for the benchmarking module and reconfiguration algorithms will be made available upon reasonable request.

ACKNOWLEDGMENTS

We acknowledge early contributions from H. Xiang, D. R. Chowdhury, J. T. Huang, and H. J. Kim. This work was supported by the Canada First Research Excellence Fund. The work of A.E.M. was supported by the Alexander von Humboldt Foundation and partially supported by the PHC Cedre Project No. 2022 “PLR.” The work of S.M. and N.N. was supported by the Natural Sciences and Research Council of Canada.

APPENDIX A: EXACT AND APPROXIMATION ALGORITHMS

A reconfiguration algorithm provides a deterministic procedure that can be implemented as a sequence of programmable instructions to compute a valid solution to a given reconfiguration problem. Exact algorithms return solutions that are provably optimal, whereas approximation algorithms return approximate solutions with provable guarantees on their distance to an optimal solution. In other words, μ -approximation algorithms return solutions that are at most (at least) μ times the optimal value in the case of a minimization (maximization) problem.




FIG. 11. Minimal example for exact and approximation algorithms. (a) The MWMCM algorithm minimizes the total displacement distance without concern for the number of atoms displaced, here displacing two atoms (black circles), each by an elementary step. (b) The 3-approx algorithm approximately (in this specific example, exactly) minimizes the number of atoms displaced without concern for the total displacement distance, here displacing one atom by four elementary steps.

We have previously studied exact and approximation algorithms to solve reconfiguration problems on graphs from a theoretical standpoint [15]. We now briefly describe four such algorithms: (1) the exact tree algorithm, (2) the exact 1D algorithm, (3) the minimum-weight maximum-cardinality matching (MWMCM) algorithm, and (4) the 3-approx algorithm.

The exact tree algorithm is an exact reconfiguration algorithm on trees, i.e., graphs in which any pair of vertices is connected by exactly one path, which is defined as a finite sequence of edges. The exact tree algorithm returns a sequence of elementary displacement operations whose number is the minimum required to solve the reconfiguration problem. These operations can be further batched into a sequence of parallel control operations that minimize the number of control steps or total reconfiguration time. This algorithm runs in time linear in the number of vertices in the tree [16] and is readily amenable to efficient implementation on computing hardware.

The exact 1D algorithm is an exact reconfiguration algorithm on linear chains. While it is possible to execute the exact tree algorithm on linear chains (given that linear chains are trees), the exact 1D algorithm is computationally more efficient, even if marginally. If the number of atoms is equal to the number of target traps ($\Delta N_a = N_a - N_a^T = 0$), then there exists a unique assignment of atoms to target traps because atoms cannot move past each other. The assignment can be computed in time linear in the number of traps in the chain. In this case, the exact 1D algorithm assigns the farthestmost atoms to the farthestmost target traps, from the edges to the center of the array, in order of their position. Such an assignment minimizes the total number of displacement operations in a single extraction-displacement-implantation (EDI) sequence [Fig. 2(b)]. This EDI sequence simultaneously extracts, displaces, and implants the atoms that are not located in their target traps. In addition, it leaves the atoms that do not need to be displaced in their respective static traps. The exact 1D algorithm is also capable of solving reconfiguration problems where $\Delta N_a > 0$ in time linear in the number of traps in the chain [46].

The MWMCM algorithm [Fig. 11(a)] is an assignment algorithm that can be easily adapted into an exact reconfiguration algorithm valid on any graph. This algorithm returns an optimal solution to the problem of

minimizing the total distance traveled by atoms, i.e., the total number of displacement operations, and it does so in polynomial time. The atoms in the initial and target configurations of atoms, associated with occupied traps in the static trap array, naturally define a weighted bipartite graph upon which this problem may be solved. Here the weight of an edge represents the length of a shortest path in the original graph. This algorithm does not consider transfer operations, i.e., it does not distinguish between moving one atom twice or two atoms once. Indeed, for any pair of initial and final traps, the algorithm chooses any arbitrary path among the set of all of shortest paths, irrespective of the number of atoms on that path. The operational performance of the algorithm could be further improved by choosing the shortest paths that minimize the number of atoms displaced. This algorithm has a running time that is cubic in the number of vertices in the graph [17].

The 3-approx algorithm [Fig. 11(b)], which is valid for any graph, returns a three-approximate solution to the problem of minimizing the total number of atoms displaced. This problem is NP-complete even on grid graphs [16]. Given a graph, the algorithm partitions the graph into a collection of balanced trees, where a balanced tree is a tree containing as many atoms as target traps. It does so by solving the U -Steiner problem [16], removing edges in order to transform the graph into a collection of balanced trees. It then applies the exact tree algorithm to each tree. This algorithm could be further improved with regard to the total number of displacement operations by choosing a weighting function that penalizes convoluted paths through trees. In addition, the function could consider “useful” trees to be those that reduce the number of displacement operations, or more generally, the total weighted distance traveled by all atoms.

Both the MWMCM algorithm, an exact algorithm that minimizes total displacement operations, and the 3-approx algorithm, an approximation algorithm that seeks to minimize transfer operations, target the minimization of a single-objective function. However, they are not optimal for any nontrivial linear combinations of both parameters, i.e., multiple-objective functions. There exist no known optimal or approximation algorithms on grids, graphs, or any other relevant geometries that solve reconfiguration problems for multiple-objective functions. Examples of such functions might minimize a linear combination of the numbers of displacement and transfer operations, in polynomial time. In fact, we know that such algorithms do not exist unless $P = NP$, because minimizing the number of transfer operations alone is NP-complete [16]. Since optimizing for the multiple-objective function is a more general problem, solving it in polynomial time would imply a polynomial-time algorithm for minimizing the number of transfer operations. Our proposed red-rec algorithm heuristically seeks to simultaneously minimize the total number of displacement and transfer operations.

APPENDIX B: BENCHMARKING PERFORMANCE AGAINST EXACT ALGORITHMS

To obtain bounds on the minimum number of displacement and transfer operations required to solve a given atom reconfiguration problem, we further benchmark the red-rec algorithm against the MWMCM and 3-approx algorithms.

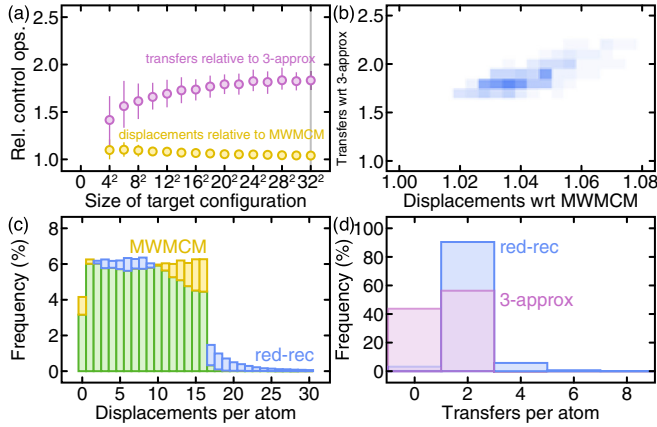


FIG. 12. Relative numbers of control operations in the absence of loss. (a) Relative numbers of transfer and displacement operations performed by the red-rec algorithm expressed relative to the 3-approx and MWMCM algorithms, respectively. The numbers are reported for varying configuration sizes at constant overhead factor $\eta = 2$ with no surplus atoms ($N_a^0 = N_a^T$). The relative number of transfer operations increases with configuration size. (b) Distribution of the number of transfer operations relative to 3-approx and displacement operations relative to MWMCM for preparing a configuration of 32×32 atoms in an array of 32×64 traps. Red-rec extracts and implants significantly more atoms than the 3-approx algorithm. (c) Distribution of the number of displacement operations performed per atom for the red-rec (blue) and MWMCM (yellow) algorithms. Red-rec displaces some atoms more times than the largest number of displacements achieved by MWMCM. (d) Distribution of the number of transfer operations per atom for the red-rec (blue) and 3-approx (purple) algorithms. As opposed to the 3-approx algorithm that extracts and implants each atom at most once, red-rec extracts and implants the same atom more than once.

The MWMCM algorithm minimizes the total number of displacement operations and the 3-approx algorithm bounds the total number of transfer operations by at most three times its optimal value (see Appendix A for a description of these two algorithms).

We specifically focus on the problem of preparing a square-compact configuration of $N_a^T = \sqrt{N_a^T} \times \sqrt{N_a^T}$ atoms in the center of rectangular array of $N_t = \sqrt{N_a^T} \times 2\sqrt{N_a^T}$ traps, i.e., with twice as many traps as desired atoms ($\eta = 2$). We set the number of surplus atoms to zero by choosing the number of atoms in the initial configuration to be equal to the number of atoms in the target configuration, i.e., $N_a^0 = N_a^T$.

The red-rec algorithm performs slightly more displacement operations than the optimal value achieved by the MWMCM algorithm. The relative numbers of displacement

operations decrease with increasing system size and reach a value of 1.04(1) for a configuration of $N_a^T = 32^2 = 1024$ atoms [Fig. 12(a)]. The slight excess in displacement operations results from a small fraction of atoms, such as those redistributed across distant pairs of donor-receiver columns, undergoing up to twice as many displacements as the maximum per-atom displacement value measured for MWMCM [Fig. 12(b)]. The number of displacement operations is positively correlated with the number of transfer operations [Fig. 12(c)]. Compared to the 3-approx algorithm, the red-rec algorithm performs 1.5 to 2.0 times more transfer operations. We note that transfer operations include both extraction and implantation operations, so the total number of transfer operations is even by construction, that is, all extracted atoms are implanted back in the static trap array. The relative numbers of transfer operations increase with increasing system size, up to a value of 1.8(1) for a configuration of $N_a^T = 32^2$ atoms [Fig. 12(a)]. Whereas the 3-approx algorithm leaves nearly half of the atoms idle, the red-rec algorithm moves nearly all atoms at least once as it sequentially reconfigures each column at least once with fewer than 7% of the columns getting reconfigured more than twice [Fig. 12(d)].

APPENDIX C: TERMINATION OF THE RED-REC ALGORITHM

We claim that, as long as the problem is solvable, there always exists at least one pair of donor-receiver columns for which at least one atom can be redistributed from the donor to the receiver (where at least one distribution row exists between the columns). This invariant guarantees that the red-rec algorithm indeed terminates. To see why, suppose instead that there are no distribution rows between any of the donor-receiver pairs. Consider the donor-receiver pair that minimizes the distance between the donor column and the receiver column and denote this pair by (c_d, c_r) . By assumption, these two columns cannot be adjacent as otherwise we have at least one distribution row. Assume that distribution rows between adjacent columns are not restricted to the storage region. Moreover, since we select (c_d, c_r) by minimizing the distance between the two columns, there cannot exist any donor column or receiver column between c_d and c_r . Otherwise, we can find a donor-receiver pair consisting of two columns separated by fewer columns than (c_d, c_r) , contradicting our choice of pair. Hence, all columns between c_d and c_r must be neutral columns that have therefore been previously solved. This in turn implies that the storage region between c_d and c_r must be free of atoms and each row in this region can be used as a distribution row.

[1] C. Monroe, *Nature (London)* **416**, 238 (2002).
[2] P. S. Jessen, I. H. Deutsch, and R. Stock, *Quantum Inf. Process.* **3**, 91 (2004).
[3] K.-A. B. Soderberg, N. Gemelke, and C. Chin, *New J. Phys.* **11**, 055022 (2009).
[4] M. Saffman, T. G. Walker, and K. Mølmer, *Rev. Mod. Phys.* **82**, 2313 (2010).

[5] A. M. Kaufman and K.-K. Ni, *Nat. Phys.* **17**, 1324 (2021).
[6] R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).
[7] S. Lloyd, *Science* **273**, 1073 (1996).
[8] C. Gross and I. Bloch, *Science* **357**, 995 (2017).
[9] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu, and Y. Takahashi, *Nat. Rev. Phys.* **2**, 411 (2020).
[10] A. Browaeys and T. Lahaye, *Nat. Phys.* **16**, 132 (2020).

- [11] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, *Rev. Mod. Phys.* **93**, 025001 (2021).
- [12] M. Morgado and S. Whitlock, *AVS Quantum Sci.* **3**, 023501 (2021).
- [13] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, *Nature (London)* **607**, 667 (2022).
- [14] K.-N. Schymik, V. Lienhard, D. Barredo, P. Scholl, H. Williams, A. Browaeys, and T. Lahaye, *Phys. Rev. A* **102**, 063107 (2020).
- [15] A. Cooper, S. Maaz, A. E. Mouawad, and N. Nishimura, in *WALCOM: Algorithms and Computation, Proceedings of the 16th International Conference and Workshops, Jember, 2022*, edited by P. Mutzel, M. S. Rahman, and Slamain, Lecture Notes in Computer Science Vol. 13174 (Springer, Cham, 2022), pp. 263–274.
- [16] G. Călinescu, A. Dumitrescu, and J. Pach, *SIAM J. Discrete Math.* **22**, 124 (2008).
- [17] H. W. Kuhn, *Nav. Res. Logist. Q.* **2**, 83 (1955).
- [18] J. Edmonds and R. M. Karp, *J. ACM* **19**, 248 (1972).
- [19] W. Lee, H. Kim, and J. Ahn, *Phys. Rev. A* **95**, 053424 (2017).
- [20] C. Sheng, J. Hou, X. He, P. Xu, K. Wang, J. Zhuang, X. Li, M. Liu, J. Wang, and M. Zhan, *Phys. Rev. Res.* **3**, 023008 (2021).
- [21] T. Mamee, W. Anukool, N. Thaicharoen, N. Chattrapiban, and P. Sompet, *J. Phys.: Conf. Ser.* **2145**, 012024 (2021).
- [22] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho *et al.*, *Nature (London)* **595**, 227 (2021).
- [23] Z.-J. Tao, L.-G. Yu, P. Xu, J.-Y. Hou, X.-D. He, and M.-S. Zhan, *Chin. Phys. Lett.* **39**, 083701 (2022).
- [24] W. Tian, W. J. Wee, A. Qu, B. J. M. Lim, P. R. Datla, V. P. W. Koh, and H. Loh, *Phys. Rev. Appl.* **19**, 034048 (2023).
- [25] N. Nishimura, *Algorithms* **11**, 52 (2018).
- [26] J. van den Heuvel, in *Surveys in Combinatorics 2013*, edited by S. R. Blackburn, S. Gerke, and M. Wildon (Cambridge University Press, Cambridge, 2013), pp. 127–160.
- [27] T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno, *Theor. Comput. Sci.* **412**, 1054 (2011).
- [28] N. Bousquet, A. E. Mouawad, N. Nishimura, and S. Siebertz, *arXiv:2204.10526*.
- [29] S. Fujita, T. Nakamigawa, and T. Sakuma, *Discrete Appl. Math.* **184**, 139 (2015).
- [30] T. Kato, T. Nakamigawa, and T. Sakuma, *Eur. J. Comb.* **95**, 103325 (2021).
- [31] S. D. Han, N. M. Stiffler, K. E. Bekris, and J. Yu, *IEEE Robot. Autom. Lett.* **3**, 1608 (2018).
- [32] E. D. Demaine, S. P. Fekete, P. Keldenich, H. Meijer, and C. Scheffer, *SIAM J. Comput.* **48**, 1727 (2019).
- [33] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, S. Choi, M. Serbyn, M. Greiner, V. Vuletic, and M. D. Lukin, *Science* **371**, 1355 (2021).
- [34] K.-N. Schymik, S. Pancaldi, F. Nogrette, D. Barredo, J. Paris, A. Browaeys, and T. Lahaye, *Phys. Rev. Appl.* **16**, 034013 (2021).
- [35] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, *Nature (London)* **411**, 1024 (2001).
- [36] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen, *Nat. Phys.* **6**, 951 (2010).
- [37] A. V. Carpentier, Y. H. Fung, P. Sompet, A. J. Hilliard, T. G. Walker, and M. F. Andersen, *Laser Phys. Lett.* **10**, 125501 (2013).
- [38] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, *Phys. Rev. Lett.* **115**, 073003 (2015).
- [39] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu, and C. A. Regal, *Phys. Rev. X* **9**, 011057 (2019).
- [40] M. M. Aliyu, L. Zhao, X. Q. Quek, K. C. Yellapragada, and H. Loh, *Phys. Rev. Res.* **3**, 043059 (2021).
- [41] J. Ang’ong’a, C. Huang, J. P. Covey, and B. Gadway, *Phys. Rev. Res.* **4**, 013240 (2022).
- [42] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman, *Phys. Rev. X* **12**, 021027 (2022).
- [43] K.-N. Schymik, B. Ximenez, E. Bloch, D. Dreon, A. Signoles, F. Nogrette, D. Barredo, A. Browaeys, and T. Lahaye, *Phys. Rev. A* **106**, 022611 (2022).
- [44] G. McConnell, J. Trägårdh, R. Amor, J. Dempster, E. Reid, and W. B. Amos, *eLife* **5**, e18659 (2016).
- [45] R. El Sabeh, J. Bohm, Z. Ding, S. Maaz, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper, following paper, *Phys. Rev. A* **108**, 023108 (2023).
- [46] R. M. Karp and S. R. Li, *Discret. Math.* **13**, 129 (1975).