

Number Partitioning With Grover's Algorithm in Central Spin Systems

Galit Anikeeva,^{1,‡} Ognjen Marković^{1,*}, Victoria Borish^{2,‡} Jacob A. Hines^{2,‡}
Shankari V. Rajagopal^{1,‡} Eric S. Cooper,¹ Avikar Periwal,¹ Amir Safavi-Naeini² Emily J. Davis,¹
and Monika Schleier-Smith^{1,†}

¹*Department of Physics, Stanford University, Stanford, California 94305, USA*

²*Department of Applied Physics, Stanford University, Stanford, California 94305, USA*

 (Received 11 September 2020; revised 23 February 2021; accepted 23 March 2021; published 13 May 2021)

Numerous conceptually important quantum algorithms rely on a blackbox device known as an oracle, which is typically difficult to construct without knowing the answer to the problem that the algorithm is intended to solve. A notable example is Grover's search algorithm. Here we propose a Grover search for solutions to a class of NP-complete decision problems known as subset sum problems, including the special case of number partitioning. Each problem instance is encoded in the couplings of a set of qubits to a central spin or boson, which enables a realization of the oracle without knowledge of the solution. The algorithm provides a quantum speedup across a known phase transition in the computational complexity of the partition problem, and we identify signatures of the phase transition in the simulated performance. Whereas the naive implementation of our algorithm requires a spectral resolution that scales exponentially with system size for NP-complete problems, we also present a recursive algorithm that enables scalability. We propose and analyze implementation schemes with cold atoms, including Rydberg-atom and cavity-QED platforms.

DOI: [10.1103/PRXQuantum.2.020319](https://doi.org/10.1103/PRXQuantum.2.020319)

I. INTRODUCTION

Many quantum algorithms that offer a provable speedup over their best classical counterparts rely on the ability to query an oracle: a black box that knows the answer to the problem that the quantum computer is to solve. A paradigmatic example is Grover's search algorithm [1,2], which theoretically speeds up the time to search through an unstructured database of N entries, requiring only $O(\sqrt{N})$ queries of the oracle rather than the classical $O(N)$ queries. By extension, Grover's algorithm can in principle speed up the search for solutions to a wide range of decision problems, including NP-complete problems [3] such as boolean satisfiability, the clique problem, and the number partitioning problem [4,5], with applications from cryptography to finance [6–9].

Formally, any instance of a search or decision problem is represented by an oracle function $f(x)$ that acts on a string

x of n bits and returns either 0 (failure) or 1 (success). The search aims to find a value X such that $f(X) = 1$, while the decision problem asks whether such an X exists at all. In experimental demonstrations to date of Grover's search [10–22], implementing the oracle—a unitary operation controlled by $f(x)$ —requires knowing the solution(s) X . To obtain a true benefit from a quantum algorithm involving an oracle, one requires a physical system that directly encodes the function f in a manner that is agnostic to the solution [23].

In this paper, we propose a genuine application of Grover's algorithm to solving the NP-complete number partitioning problem: *Given n objects with integer weights, does there exist a bipartition that balances a scale?* Our approach can be implemented in physical systems that take the form of either a central spin or central boson model, featuring n qubits interacting with an ancilla spin or photon that plays the role of the oracle. Crucially, the decision problem is encoded in the couplings of the qubits to the ancilla, allowing the oracle to be implemented without *a priori* knowledge of the solution. Numerical simulations of the quantum algorithm illustrate physical manifestations of a known phase transition in the computational complexity of number partitioning, including an exponential scaling of the spectral resolution required to solve hard problem instances. A recursive variant of our algorithm avoids this exponential resource requirement, providing improved

*ognjenm@stanford.edu

†schleier@stanford.edu

‡These authors contributed equally to this work.

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.

scalability. By analyzing proposed implementations with Rydberg atoms and in cavity-QED systems, we show that a speedup is attainable in near-term experiments.

II. ALGORITHM AND IMPLEMENTATION

We begin with a brief review of Grover's algorithm [Fig. 1(a)]. The algorithm starts by initializing a collection of $n = \log_2 N$ qubits in an equal superposition

$$|\psi_0\rangle = \frac{(|0\rangle + |1\rangle)^{\otimes n}}{2^{n/2}} = \sum_x c_{x,0} |x\rangle \quad (1)$$

of all possible standard basis states labeled by n -bit numbers x , with $c_{x,0} = 1/\sqrt{N}$. The objective is to amplify the amplitude c_X of the solution state(s) $|X\rangle$. To this end, the oracle U first marks the solution(s) by applying a π phase shift ($c_X \rightarrow e^{i\pi} c_X$) for all X with $f(X) = 1$. The marked states are then amplified by inversion about the average: $c_x \rightarrow \bar{c} - (c_x - \bar{c})$ for all x , where $\bar{c} = \sum_x c_x/N$. This inversion operation V is accomplished by combining single-qubit Hadamard gates with an n -qubit controlled phase gate that is similar to the oracle but less technically demanding (see Appendix B), or can alternatively be replaced by single-qubit rotations only [24]. Thus, we focus on the challenge of realizing the oracle.

We will show a natural physical incarnation of the oracle for a class of decision problems known as subset sum problems [25], focusing on the special case of number

partitioning. We specify each problem instance by a list of n weights $w_i \in (0, 1]$ of finite bit depth k , and search for a partition into two sublists of equal total weight. To encode the partition problem using n qubits representing the objects with weights w_i , we let each qubit state indicate which subset ($|0\rangle$ or $|1\rangle$) an object is in [Fig. 1(b)], so that the weighted collective spin

$$S_z \equiv \frac{1}{2} \sum_i w_i \sigma_i^z \quad (2)$$

represents the imbalance between the subsets. Implementing the oracle then requires applying a π phase shift to any n -qubit basis state $|x\rangle$ satisfying $S_z|x\rangle = 0$.

The quantum oracle thus requires implementing a collective phase gate $U = e^{i\pi f(x)} = e^{i\pi \delta(S_z)}$, where $\delta(\cdot)$ denotes the Kronecker delta function. To design a physical implementation of this gate, it is helpful to define a generalized oracle $U_\gamma = e^{i\Phi_\gamma(S_z)}$ in terms of an S_z -dependent phase shift

$$\Phi_\gamma(S_z) = 2 \arctan(2S_z/\gamma) + \pi, \quad (3)$$

which steps from zero to 2π as a function of S_z and provides a π phase shift at $S_z = 0$ [Fig. 1(c)]. The ideal oracle is obtained in the limiting case $U \equiv U_{\gamma \rightarrow 0}$ of an infinitely steep phase step.

The collective phase gate U_γ can be enabled by coupling the qubits to an ancilla, which may take the form of an auxiliary qubit or a bosonic mode. We consider either a central spin model

$$H_q = J_{\max} I_z S_z \quad (4)$$

featuring an ancilla qubit represented by a spin-1/2 operator I_z , or a central boson model

$$H_c = J_{\max} c^\dagger c S_z \quad (5)$$

featuring a cavity mode with annihilation operator c . In both cases, the ancilla couples to n system spins in the starlike graph of Fig. 1(d), and hence to the weighted collective spin S_z . The maximum coupling between a system spin and the ancilla is parameterized by J_{\max} .

For concreteness, we describe representative implementations of the central boson and central spin models with cold atoms [Fig. 1(d)]. The system spins are encoded in two internal states $|0\rangle$, $|1\rangle$ and coupled to either a cavity mode [26–33] or an auxiliary atom that can be excited to a Rydberg state [34–46]. Each coupling $w_i J_{\max}$ represents the energy shift of the $|0\rangle \rightarrow |1\rangle$ transition in atom i when either a photon enters the cavity or the auxiliary atom is excited. In the cavity implementation, the photon imparts an ac Stark shift [26–31]. In the Rydberg implementation, the excited ancilla suppresses an ac Stark shift induced by

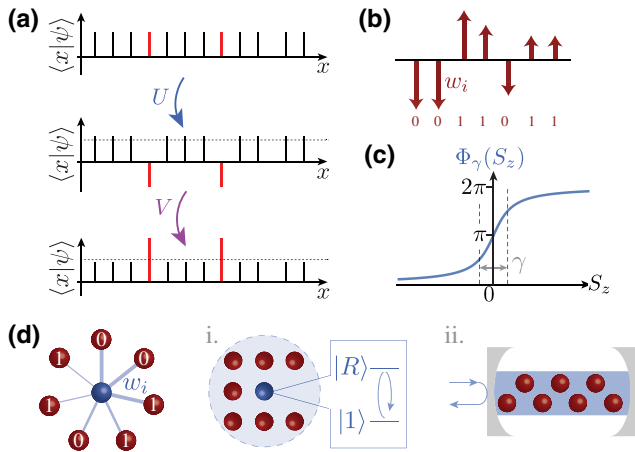


FIG. 1. (a) Sketch of Grover's algorithm, showing the amplitude of each basis state $|x\rangle$ in the system state $|\psi\rangle$. One iteration consists of the oracle U marking the solution states (red) with a π phase shift, followed by inversion about the average V . (b) Number partitioning: a set of weighted spins is partitioned, if possible, into two sets of equal total weight. (c) Phase shift $\Phi_\gamma(S_z)$ applied by generalized oracle with step width γ . (d) The weights w_i are encoded by couplings of system spins (red) to an ancilla (blue), which can be either (i) a central spin (e.g., Rydberg atom); or (ii) a bosonic mode (e.g., cavity).

classical control fields that couple the system atoms' state $|1\rangle$ to a Rydberg state. In both cases, the weights w_i can be programmed via the atomic positions or control fields. The net effect of the couplings $w_i J_{\max}$ on the ancilla is a frequency shift $J_{\max} S_z$ that depends on the weighted collective spin S_z .

The S_z -dependent resonant frequency of the ancilla is crucial to enabling the oracle. In the central boson model, the oracle relies on the phase response of a driven harmonic oscillator. For a one-sided cavity of linewidth κ , the output field is phase shifted by π for a resonant drive compared with the off-resonant case. Having the drive field consist of a single photon that is resonant if and only if $S_z = 0$ yields precisely the oracle operation U_γ , with a phase step of dimensionless width $\gamma = \kappa/J_{\max}$, where we set $\hbar = 1$. In the central spin model, the oracle U_γ is implemented by attempting to drive a 2π rotation of the ancilla with a field that is resonant if the weighted spin S_z is zero. For a suitably shaped drive pulse, the ancilla atom ends up in its initial state irrespective of S_z [47], and the entire system acquires a π geometric phase shift only when $S_z = 0$. The width $\gamma = \kappa/J_{\max}$ of the phase step is now set by the bandwidth $\kappa = 2\pi/\tau$ of the pulse with temporal width τ .

To examine the performance of the generalized oracle, we first introduce a convenient visualization of Grover's algorithm [48]. We define the solution space $\mathcal{A} = \{|X\rangle : S_z|X\rangle = 0\}$ as the set of states that solve the partition problem and let

$$|A\rangle = \frac{1}{\sqrt{N_{\mathcal{A}}}} \sum_{|X\rangle \in \mathcal{A}} |X\rangle \quad (6)$$

denote the equal superposition of all solutions (assuming their existence) where $N_{\mathcal{A}}$ is the number of solutions. We additionally define an orthogonal state

$$|B\rangle \propto |\psi_0\rangle - |A\rangle\langle A|\psi_0\rangle, \quad (7)$$

where $|\psi_0\rangle$ is the initial state of Eq. (1). The states $|A\rangle$ and $|B\rangle$ span an $SU(2)$ subspace that can be visualized on a Bloch sphere with $|A\rangle$ and $|B\rangle$ as poles.

Grover's algorithm ideally takes place entirely within this subspace of the full 2^n -dimensional Hilbert space, iteratively rotating the initial state $|\psi_0\rangle$ towards the solution state $|A\rangle$. Each iteration

$$|\psi_{T+1}\rangle = VU|\psi_T\rangle, \quad (8)$$

comprises the oracle U and inversion about the average V . The net effect of these two operations is a rotation about the $\hat{\mathcal{Y}}$ axis [Fig. 2(a)]. For $N_{\mathcal{A}}/N \ll 1$, a near-unity success probability is achieved after an optimal number of iterations

$$T \approx (\pi/4)\sqrt{N/N_{\mathcal{A}}}. \quad (9)$$

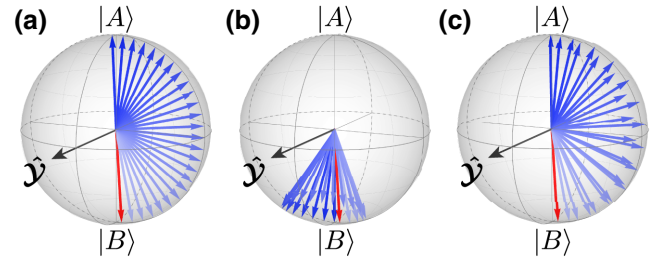


FIG. 2. Visualization of Grover's algorithm and generalized oracle. (a) Grover's algorithm with ideal oracle for $N = 2^{10}$ and $N_{\mathcal{A}} = 1$. Over repeated iterations (blue), the state $|\psi_0\rangle$ (red) approaches the solution state $|A\rangle$. (b) Grover's algorithm with naive application of the generalized oracle for $\epsilon = 0.25$. (c) The spin-echo sequence compensates for the imperfection of the oracle, allowing similar performance to the ideal case.

The generalized oracle with a nonzero step width introduces an error that, to lowest order, is correctable by spin echo. To visualize how, we consider a simplified scenario where there exist only two possible values of the phase $\Phi_\gamma \in \{\epsilon, \pi\}$. For nonzero ϵ , the combination of the generalized oracle and inversion about the average induces the state to rotate about a tilted axis [Fig. 2(b)]. To mitigate accumulation of error, we alternate between applying the oracle U_γ and its Hermitian conjugate U_γ^\dagger . A pair of two Grover iterations then takes the form

$$|\psi_{T+2}\rangle = VU_\gamma^\dagger VU_\gamma |\psi_T\rangle, \quad (10)$$

where $U_\gamma^\dagger = \mathcal{R}_\pi^\dagger(\hat{\mathbf{x}})U_\gamma\mathcal{R}_\pi(\hat{\mathbf{x}})$ is implemented by a spin-echo sequence involving two global π rotations $\mathcal{R}_\pi(\hat{\mathbf{x}})$ about the individual qubits' $\hat{\mathbf{x}}$ axes. The result is the trajectory shown in Fig. 2(c), which achieves similar performance to the ideal oracle in Fig. 2(a).

Even with spin echo, the step width will ultimately limit the resolution of the generalized oracle: selectively amplifying only spin configurations with $S_z = 0$ requires a narrow step. Further, producing a narrow step requires a long coherence time, so that dissipation will place physical limits on the performance of the algorithm. We elaborate on both of these considerations in Secs. III and IV. First, however, we examine the application of Grover's algorithm to number partitioning using a phase step narrow enough to resolve even the least significant bit of the weights.

III. SPEEDUP IN NUMBER PARTITIONING

To analyze the performance for number partitioning, we generate sets of n random k -bit weights and postselect for instances where at least one perfect partition exists. For

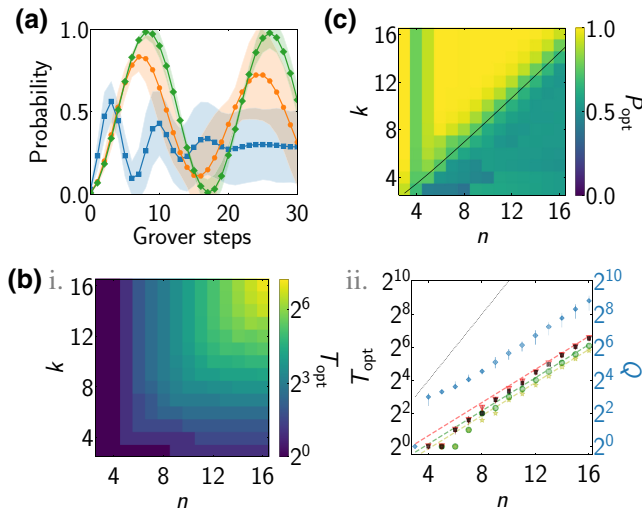


FIG. 3. Number partitioning with generalized oracle of step width $\gamma = 2^{-k}$. (a) Success probability P_T for $n = 8$, $k = 4, 8, 12$ (blue squares, orange circles, and green diamonds). Shading indicates standard deviation over 5000 instances of the weights. (b.i) Optimal number of iterations T_{opt} versus (n, k) . (b.ii) T_{opt} for $n = k$, grouped by number of solutions $N_{\mathcal{A}} = 2, 4, 6$ (red triangles, green circles, and yellow stars) and compared with asymptotic theory $T_{\text{opt}} \propto \sqrt{N}$ (dashed lines). Black squares show average over all instances. Dotted gray line indicates linear scaling $T_{\text{opt}} \propto N$. Blue diamonds show median speedup $[Q]_{0.5}$ for $n = k$, with error bars indicating interquartile range. (c) Probability P_{opt} versus (n, k) . Black line shows critical bit depth $k_c(n)$.

each such instance, we calculate the success probability

$$P_T = \sum_{|X\rangle \in \mathcal{A}} |\langle X | \psi_T \rangle|^2 \quad (11)$$

as a function of the number T of calls to the oracle, applied with spin echo [Eq. (10)]. Figure 3(a) shows examples of P_T for $n = 8$ spins, bit depths $k = 4, 8, 12$, and a step width $\gamma = 2^{-k}$ just narrow enough to resolve changes in the least significant bit of S_z . As expected from the Bloch-sphere picture, the success probability oscillates as a function of T . The maximum probability and the time to reach it combine to determine the effectiveness of the algorithm.

As a single figure of merit, we calculate the total number of calls to the oracle required to reach a specified (near-unity) success probability \mathcal{P} . For a search procedure with fixed success probability P per trial, the number of trials M needed to reach a probability $\mathcal{P} = 1 - \varepsilon$ of finding a solution is

$$M(P, \varepsilon) = \frac{\ln(\varepsilon)}{\ln(1 - P)}. \quad (12)$$

Thus, reaching the target error ε with Grover's algorithm requires querying the oracle a total of $T_{\text{total}} = M(P_T, \varepsilon)T$

times. To minimize this quantity, we first calculate its median value as a function of T over many instances of weights at a given (n, k, γ) . We then define T_{opt} as the number of Grover iterations that minimizes the median total number of queries T_{total} . Note that T_{opt} is independent of the target error ε .

Figure 3(b.i) shows the optimal number of queries T_{opt} as a function of the number of spins n and bit depth k , at fixed step width $\gamma = 2^{-k}$. The scaling of T_{opt} with n is shown in Fig. 3(b.ii) for a cut at $n = k$ (black squares), where the number of perfect partitions is typically of order one [4]. We additionally plot T_{opt} for instances of the weights postselected according to the number of solutions $N_{\mathcal{A}} = 2, 4, 6$ (red triangles, green circles, and yellow stars). In each case, the optimal number of iterations approaches the prediction of Eq. (9) (dashed lines) at large $N = 2^n$, scaling as $T_{\text{opt}} \propto \sqrt{N}$. Quantifying the resulting speedup requires additionally examining P_{opt} , the success probability after T_{opt} iterations [Fig. 3(c)].

The dependence of success probability P_{opt} on (n, k) reflects a known phase transition in the computational complexity of the number partitioning problem [4,49]. For small bit depth $k \lesssim n$ (the “easy” phase), there typically exist many perfect partitions. For large bit depth $k \gtrsim n$ (the “hard” phase), perfect partitions are rare and thus—even when postselecting for their existence—the probability of finding them by random guessing is exponentially small in n . By contrast, in our quantum search [Fig. 3(c)], the success probability P_{opt} is everywhere of order unity and highest in the “hard” phase, since Grover's algorithm is most effective when solutions are few. The phase boundary lies at a critical bit depth [4]

$$k_c(n) \equiv n - \frac{1}{2} \log_2 \left(\frac{n\pi}{6} \right), \quad (13)$$

shown by the black curve in Fig. 3(c), where the average number of perfect partitions is $\langle N_{\mathcal{A}} \rangle \sim \sqrt{6/(\pi n)} 2^{n-k} = 1$ [50].

We quantify the advantage of the algorithm by calculating the limited quantum speedup Q , defined as in Ref. [51] by comparing the quantum search with an algorithmically similar classical search. The most analogous classical algorithm is a memoryless search, which at each trial samples (with replacement) a random partition with success probability $P_0 = N_{\mathcal{A}}/N$. The number of memoryless search trials needed to reach a target success probability \mathcal{P} also follows from Eq. (12). For each problem instance, we define speedup Q as the ratio of memoryless trials to total Grover iterations:

$$Q = \frac{1}{T_{\text{opt}}} \frac{\ln(1 - P_{\text{opt}})}{\ln(1 - N_{\mathcal{A}}/N)}. \quad (14)$$

This speedup is independent of the target error ε , thanks to the algorithmic similarity of the two memoryless search

algorithms, as further discussed in Appendix C. Figure 3(b.ii) shows the median speedup $[Q]_{0.5}$, where $[Q]_q$ denotes the q th quantile over problem instances. We observe the expected scaling $Q \propto \sqrt{N}$ of the speedup in query complexity.

A caveat is that physical limitations might preclude successfully implementing the algorithm in cases requiring a narrow step width γ . We have so far assumed a step width $\gamma = 2^{-k}$, motivated by the intuition that γ sets a capture range of S_z values amplified by Grover's algorithm. To test this intuition, we plot the normalized probability distribution $\tilde{P}(S_z) \equiv P(S_z)/P(S_z = 0)$ after $T_{\text{opt}}(\gamma)$ Grover iterations as a function of step width γ [Fig. 4(a.i)], for $n = k = 6$ without postselecting on the existence of perfect partitions. Consistent with our expectation, the width of the distribution is approximately set by the step width γ . An analytic derivation of this capture range is given in Appendix E.

To capture only true solutions $S_z = 0$, the step width γ should be smaller than the smallest nonzero $|S_z|$ value. In the easy regime $k \lesssim n$, a step width $\gamma \lesssim 2^{-k}$ is required

to distinguish $S_z = 0$ from $S_z = \pm 2^{-k}$. However, with increasing k , the typical size of the smallest residue approaches a finite value $|S_z| \approx 2^{-k_c}$ [4]. Thus, the critical bit depth $k_c(n)$ in Eq. (13) represents the resolution required to discriminate the smallest typical residue $|S_z|$ in the large- k limit. For arbitrary (n, k) , we can choose the oracle to have resolution

$$-\log_2 \gamma_c = \min(k_c, k) \approx \min(n, k), \quad (15)$$

coarser than we have so far assumed in the hard regime. We verify Eq. (15) by plotting the resolution $-\log_2 \gamma$ required to reach a fixed success probability P_{opt} , averaging over all pairs (n, k) with $3 \leq n, k \leq 16$, in Fig. 4(a.ii). For each of three different values of $P_{\text{opt}} = 0.4, 0.6, 0.8$, the required step width γ is within a constant factor of γ_c .

We plot the quantum speedup for this less stringent choice of step width γ_c in Fig. 4(b.i). The speedup exhibits a maximum along the phase boundary $k = k_c(n)$ (solid black curve). In Fig. 4(b.ii), we examine the scaling of the speedup along an approximation to this curve chosen

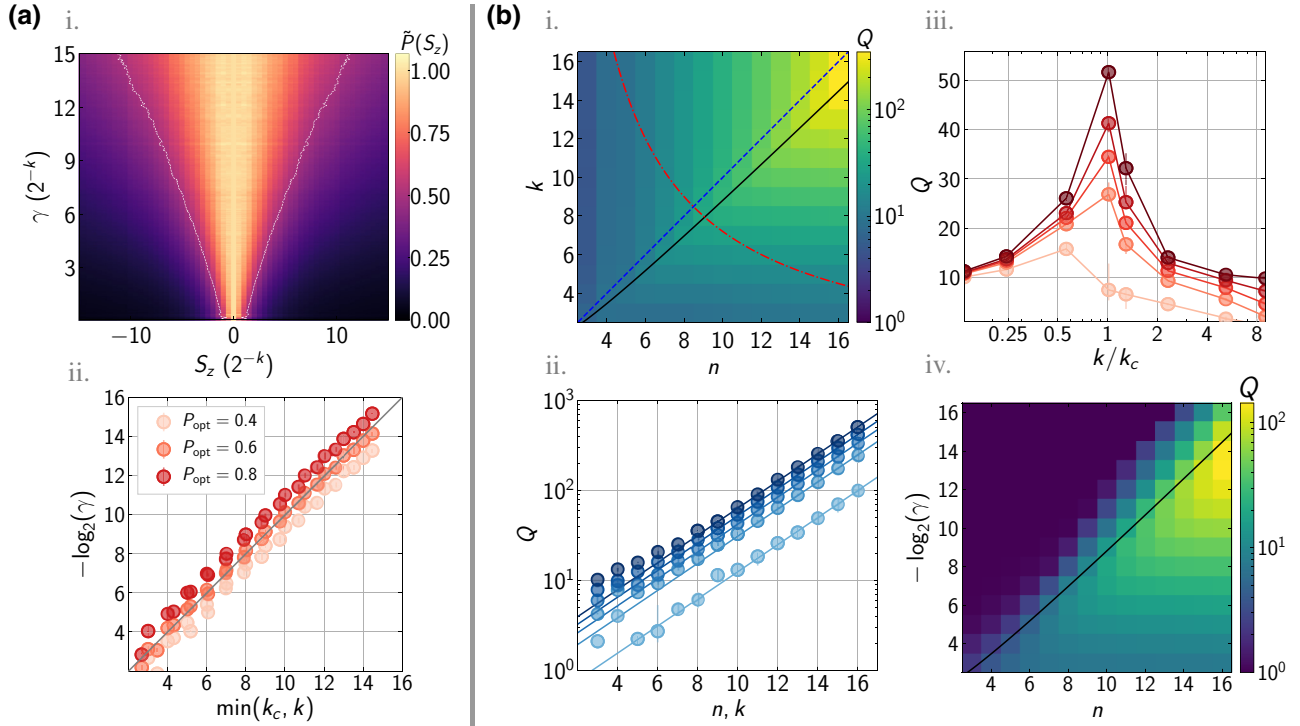


FIG. 4. (a) Step width as capture range. (a.i) Normalized probability distribution $\tilde{P}(S_z)$ versus γ , for $n = k = 6$ with no postselection. White lines indicate contours of $\tilde{P}(S_z) = 0.5$. (a.ii) Step width γ required to obtain $P_{\text{opt}} = [0.4, 0.6, 0.8]$, shaded from lighter to darker. Results are plotted versus $\min(k_c, k)$, with markers showing average over all (n, k) with $3 \leq n, k \leq 16$. Gray line shows $\gamma_c = 2^{-\min(k_c, k)}$. (b) Quantum speedup in the decision and optimization problems. (b.i) Median speedup $[Q]_{0.5}$ versus (n, k) for the decision problem at step width γ_c . Lines denote $k = k_c$ (solid black), $k = n$ (dotted blue), and fixed problem size $nk = 72$ (dashed red). (b.ii) Cuts of Q along $n = k$ for different quantiles $q = [0.01, 0.25, 0.5, 0.75, 0.99]$, shaded from lightest to darkest. Lines denote \sqrt{N} scaling. (b.iii) Cuts of Q at $nk = 72$ [red dashed line in (b.i)] for different quantiles, as in (b.ii). Lines are a guide to the eye. (b.iv) Median speedup $[Q]_{0.5}$ versus n and $k_{\text{eff}} = -\log_2 \gamma$ for the optimization problem with machine-precision weights, approximating the large- k limit. Black line shows $-\log_2 \gamma = k_c$.

to ensure integer values of (n, k) , namely, the $n = k$ cut (dotted blue line). We plot the speedup $[Q]_q$ versus N for different quantiles q (blue circles) and find good agreement with an asymptotic scaling $Q \propto \sqrt{N}$ (solid lines) for all quantiles. Thus, the generalized oracle with the critical step width γ_c suffices to achieve an $O(\sqrt{N})$ speedup, the same scaling that is achieved by the ideal oracle and proven to be optimal for an unstructured search [52–54].

The phase transition in computational complexity manifests in a sharp peak in the speedup at the phase boundary $k = k_c(n)$. We observe this peak in Fig. 4(b.iii) along a cut of fixed problem size nk , i.e., fixing the total number of bits encoding the set of n weights. The peak in the speedup reflects the known result that the hardest problem instances are not deep in the hard regime, but rather near the phase transition [4,55]. In particular, the hardest problems are those with the largest ratio N/N_A of the size of the search space to the number of solutions, after postselecting for the existence for solutions. This ratio reaches a maximum near the phase boundary, explaining the peak in $Q \propto \sqrt{N}/N_A$.

Even in the experimentally relevant case where the weights are not restricted to a finite bit depth, the resolution of the oracle sets an effective bit depth $k_{\text{eff}} = -\log_2 \gamma$ that can reveal the complexity phase transition. For real-numbered weights $w_i \in (0, 1]$, we consider the optimization problem of minimizing $|S_z|$, defining the success probability P_{opt} as that of finding the system in a configuration of minimal $|S_z|$ after an optimal number of Grover iterations. We plot the median speedup $[Q(n, \gamma)]_{0.5}$ in Fig. 4(b.iv). As a function of k_{eff} at fixed n , the speedup first rises to a maximum at $k_{\text{eff}} \approx k_c$ before declining precipitously for $k_{\text{eff}} > k_c$ due to the narrowness of the capture range, providing a striking signature of the complexity phase transition.

IV. EFFECTS OF DISSIPATION

A key challenge for experimental implementations is that producing a narrow phase step requires a long coherence time. Specifically, at fixed interaction strength J_{max} , the step width γ determines the physical time $\kappa^{-1} \sim 1/(\gamma J_{\text{max}})$ to implement the oracle operation U_γ . Even a single error occurring during this time thwarts the amplification process. For concreteness, we consider an error model in which the excited ancilla decays—or, equivalently, the ancilla photon is lost—at rate Γ_a . In terms of the interaction-to-decay ratio $\rho \equiv J_{\text{max}}/\Gamma_a$, the error rate per query of the oracle is then approximately $\Gamma_a/\kappa = 1/(\rho\gamma)$. Thus, on average $T_{\text{max}} \sim \rho\gamma$ Grover iterations can be implemented before incurring an error. For $\rho\gamma_c \lesssim T_{\text{opt}}$, the algorithm must be run at an increased step width $\gamma > \gamma_c$ that reduces the speedup.

Figure 5(a) shows the speedup calculated at finite interaction-to-decay ratio $\rho = 10^3$. We model the decay by modifying the frequency shift of the ancilla's excited

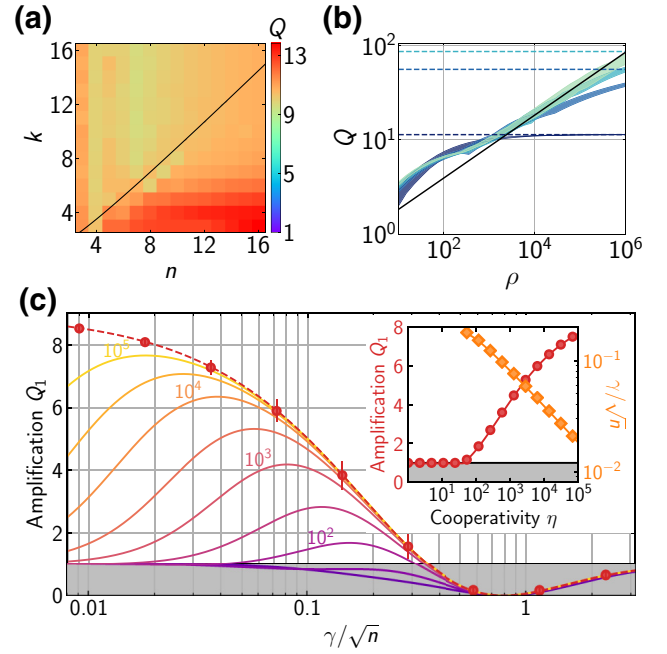


FIG. 5. Effects of dissipation. (a) Median speedup $[Q]_{0.5}$ versus (n, k) in the presence of decay with interaction-to-decay ratio $\rho = 10^3$. Solid black line denotes $k = k_c$. (b) Median speedup $[Q]_{0.5}$ versus ρ for $n = k = (4, 6, 8, 10)$ denoted by dark blue to light green shaded lines. The shading denotes the interquartile range. Black solid line denotes scaling $Q \sim \rho^{1/3}$. Dashed lines show maximum achievable Q for each system size n . (c) Amplification versus step width γ for $T = 1$. Solid curves show average amplification in large- N limit for finite cavity cooperativity $\eta = 10^1, 3 \times 10^1, 10^2, 3 \times 10^2, \dots, 10^5$ (purple to yellow) and for unitary evolution (red dashed). Dark red circles show simulated amplification at $n = k = 12$ with no dissipation; error bars denote standard deviation. Inset shows optimal amplification (red circles) and step width (orange diamonds) versus η for $n = k = 12$, matching the prediction for large N (solid curves).

state (Sec. II) with an imaginary component, $J_{\text{max}}S_z + i\Gamma_a/2$, as detailed in Appendix G. We choose the step width γ for each (n, k) to maximize the speedup, accounting for a reduction in success probability due to the chance of ancilla decay. While the speedup no longer achieves $O(\sqrt{N})$ scaling, we preserve an advantage $Q \approx 10$ compared with the classical search. The dependence of the speedup on interaction-to-decay ratio ρ is shown in Fig. 5(b) for $n = k$ at different system sizes n . The speedup scales as $Q \sim \rho^{1/3}$, consistent with an analytic model derived in Appendix G, before saturating to the value expected for the ideal Grover's algorithm.

An interaction-to-decay ratio $\rho \gtrsim 10^3$ is experimentally accessible in an implementation of the central spin model using Rydberg atoms, as detailed in Appendix 1. In this implementation, the dominant dissipative process is decay of the ancilla from the Rydberg state, whereas decay of the system spins is suppressed by coupling to their Rydberg

states off-resonantly [38–40,42]. In terms of the maximum attainable Rabi frequency Ω_{\max} of this coupling, the interaction-to-decay ratio is limited to $\rho < \Omega_{\max}/(2\sqrt{n}\Gamma)$, which permits values of order $\rho \sim 10^3$ for realistic laser powers and high-lying Rydberg states.

At lower interaction-to-decay ratios, the optimum speedup is obtained by performing only a single Grover iteration. In the absence of dissipation, this single-cycle speedup Q_1 is identical to the amplification factor P_1/P_0 , assuming $P_{0,1} \ll 1$. Figure 5(c) shows $Q_1 = P_1/P_0$ as a function of step width γ for $n = k = 12$ with no dissipation (red circles), corroborating an analytical model derived in Appendix E in the large- N limit (dashed curve). The model shows that the gain is set by γ/\sqrt{n} , which parameterizes the ratio of the step width to the width of the initial S_z distribution, and saturates at a maximum value $Q_1 = 9$ for $\gamma/\sqrt{n} \ll 1$. Ancilla decay reduces the amplification Q_1 below this ideal curve, becoming significant for interaction-to-decay ratios $\rho \lesssim 1/\gamma$. The optimum speedups in Fig. 5(b) are obtained from a single amplification cycle for interaction-to-decay ratios $\rho \lesssim 10^2$.

A single amplification cycle could be performed in near-term realizations of the central boson model with atoms in a cavity (Appendix H2), by driving with a weak coherent field and heralding on the detection of a photon. The coherence of the atom-cavity coupling is quantified by the cooperativity $\eta = 4g^2/\kappa\Gamma_e$, where g is the vacuum Rabi frequency and (κ, Γ_e) are the linewidths of the cavity and an atomic excited state to which it couples. The resulting interaction-to-decay ratio scales as $\rho \propto \eta\gamma/n$, reflecting the fact that decreasing the dimensionless step width $\gamma = \kappa/J_{\max}$ comes at the cost of increasing the photon loss probability by atomic scattering. Achieving amplification requires reaching a step width $\gamma < \sqrt{n}/12$ narrower than the initial S_z distribution while keeping $\rho\gamma > 1$ to avoid photon absorption, and hence requires strong coupling $\eta \gg 1$.

The full dependence of amplification Q_1 on step width γ and cooperativity η is shown by the solid curves in Fig. 5(c). Notably, the amplification at an optimal step width [Fig. 5(c) inset] is independent of the number of spins n , depending only on the cooperativity η . A state-of-the-art optical cavity with demonstrated cooperativity $\eta \sim 200$ [56] thus allows for amplifying solutions to the partition problem at scalable system size. Stronger amplification could be attained by coupling Rydberg atoms or superatoms [57,58] to a high-cooperativity millimeter-wave cavity [29,59,60]. For the parameters of Ref. [29], the cooperativity $\eta = 4 \times 10^8$ is no longer the limiting factor. Instead, finite lifetime Γ^{-1} of the Rydberg states places a limit $\rho < g/(n^{3/2}\Gamma) = 5 \times 10^3/n^{3/2}$ on the interaction-to-decay ratio, which permits near-maximal Q_1 for up to $n \sim 30$ atoms. Rydberg-based implementations might be further enhanced by inhibition of spontaneous emission [61,62].

V. SCALABLE ALGORITHM

The requirement of an exponentially fine resolution of the oracle poses challenges for scalability in the simple application of Grover's algorithm presented so far. Specifically, we showed in Sec. III that the required step width $\gamma \sim 2^{-k}$ becomes exponentially small with increasing system size $n \sim k$ for the hardest problem instances. As a result, if we scale the system-ancilla couplings such that the energy grows extensively with system size by fixing the maximum coupling J_{\max} , then the time required for each query of the oracle grows as $2^k \sim 2^n$. Alternatively, to keep the query time fixed, the energy must be chosen to grow exponentially with increasing system size.

The exponential resource requirement can be avoided by a more sophisticated version of our algorithm that operates at a fixed resolution $\gamma \sim 2^{-m}$ of the oracle for arbitrary (n, k) . This scalable algorithm begins by identifying candidate solutions of the number partitioning problem by searching for spin configurations in which the m least significant bits of S_z are zero. To this end, we first perform Grover amplification with each coupling J_i set to a value given by the m least significant bits of the weights. We thereby amplify only spin configurations for which $2^k S_z$ is a multiple of 2^m , thus producing a superposition state with a sparser distribution of S_z values than the initial state $|\psi_0\rangle$ (Fig. 6). We subsequently consider increasing numbers ℓm of bits in successive layers $\ell = 1, 2, 3, \dots$ of the algorithm, setting couplings

$$J_{i,\ell} = \frac{J_{\max} \text{mod}(2^k w_i, 2^{\ell m})}{2^{\ell m}}, \quad (16)$$

while keeping the resolution of the oracle fixed.

This scalable algorithm retains the benefit of an efficient encoding in a central spin system, but does place additional technical demands compared with our standard algorithm. First, the system-ancilla couplings must be changed between layers of the algorithm (Appendix F1), a capability that is naturally present in our proposed implementation schemes. A second new ingredient is a *modular oracle* that can detect the imbalance $2^k S_z$ modulo a specified power of 2 (Appendix F2). This modular oracle can readily be implemented by applying a multifrequency drive to the ancilla. Finally, successive layers ℓ of the algorithm require increasingly complex operators V_ℓ to invert about the average amplitude of previously amplified states. In fact, as we explain in Appendix F3, the inversion step in layer ℓ involves repeating the entire algorithm up through layer $\ell - 1$. For this reason, we call our scalable algorithm the *recursive algorithm*.

We describe and analyze the recursive algorithm in detail in Appendix F, showing that it solves the number partitioning problem in $O(2^{n/2+cn/m})$ queries, where $c = \log_2(\pi/2)$. Thus, in the limit of a high but fixed resolution of $m \gg 1$ bits, we recover the ideal Grover speedup.

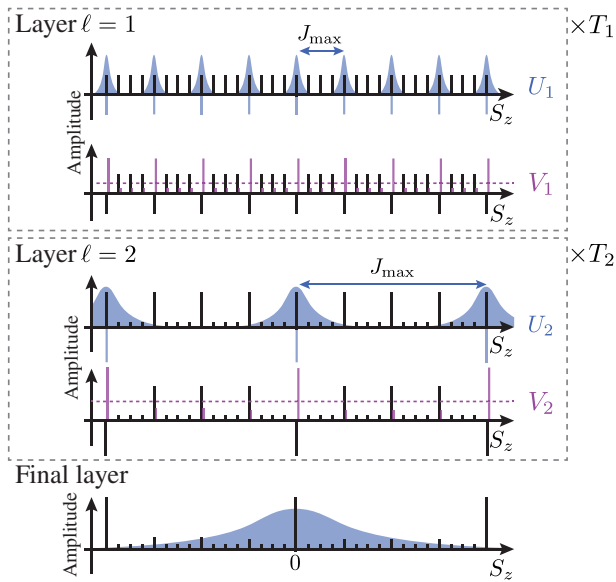


FIG. 6. Sketch of the scalable algorithm, showing amplitudes of basis states versus S_z . Each layer ℓ of the algorithm consists of T_ℓ amplification cycles, each comprising the modular oracle U_ℓ and recursive inversion operator V_ℓ . The modular oracle acts on states spaced in energy by J_{\max} with a spectral resolution γJ_{\max} illustrated by the shaded blue curves. The operator V_ℓ inverts the amplitudes of the states amplified by the previous layer of the algorithm about their average (dashed purple line). Implementing V_ℓ requires recursion to the lower layers of the algorithm. In the final layer $\ell = k/m$, the standard nonmodular oracle is used.

Importantly, we now attain this speedup not only in query complexity but also in the physical time to implement the algorithm in a scalable manner, in the sense that the total interaction energy required to encode the problem grows linearly with the problem size.

A comparison of the recursive algorithm with the standard algorithm of Secs. II–IV is shown in Fig. 7. We simulate both algorithms with the same 1000 instances of weights to examine the scaling of their physical runtimes, which is proportional to T_{total}/γ for a fixed maximum system-ancilla coupling strength J_{\max} [Fig. 7(a)]. The physical runtime of the standard algorithm with $\gamma = \gamma_c$ scales as $O(2^{1.5n})$ due to the exponential narrowing of the step width $\gamma \sim 2^{-k}$ with system size $n = k$, whereas the scaling of the recursive algorithm for $m = 6$ and $\gamma = 2^{-m-1}$ is consistent with the theoretical prediction $O(2^{0.5n+0.65n/m})$ derived in Appendix F. Thus, the recursive algorithm exhibits a scalable quantum speedup.

The performance of the recursive algorithm in realistic implementations with finite interaction-to-decay ratio shows an advantage over the standard algorithm. In Fig. 7(b), we plot the speedup Q versus $n = k$ for different interaction-to-decay ratios, with the step width γ chosen to minimize the total number of Grover queries T_{total} . In the recursive algorithm, the number of amplification cycles per layer of the algorithm (Appendix F 3) is additionally

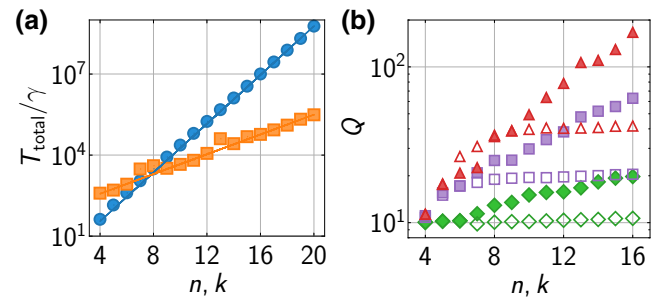


FIG. 7. Comparison of the recursive and standard algorithms. (a) Physical time T_{total}/γ to implement all T_{total} queries of the oracle at fixed J_{\max} , plotted versus $n = k$. Blue circles show the standard algorithm with $\gamma = \gamma_c$. Orange squares show recursive algorithm with $m = 6$ and $\gamma = 2^{-m-1}$. Lines denote scalings $2^{\alpha n}$, with $\alpha = 1.5$ for the standard algorithm and $\alpha = 0.5 + 0.65/m$ for the recursive algorithm. (b) Speedup Q versus $n = k$ for standard algorithm (open markers) and recursive algorithm (solid markers) at interaction-to-decay ratios $\rho = 10^3$ (green diamonds), $\rho = 10^4$ (purple squares), and $\rho = 10^5$ (red triangles).

optimized to minimize T_{total} . Whereas the speedup of the standard algorithm plateaus with increasing system size because dissipation limits the resolution of the oracle, the recursive algorithm achieves a higher speedup because it is designed to operate at fixed resolution of the oracle.

VI. DISCUSSION AND OUTLOOK

In this paper, we describe practical implementations of Grover’s algorithm for the number partitioning problem, relying on a natural encoding in spin systems with a star-like coupling graph. The problem offers an ideal setting for examining the physical manifestations of computational complexity, thanks to a well-understood phase diagram including easy and NP-hard regimes. Numerical simulations of our quantum algorithm show clear signatures of the complexity phase transition, yet even in the hard phase we are able to find an advantage over an analogous classical search.

Specifically, we compare our quantum algorithm to a probabilistic classical search, with query complexity $O(2^n)$ equivalent to that of a brute-force search (see Appendix C). While there exist classical algorithms that match [63,64] and surpass [65,66] our algorithm’s query complexity, they do so at the expense of exponential memory requirements [67,68]. To the best of our knowledge, the leading classical algorithm of polynomial space complexity is that by Esser and May, which achieves a time complexity of $O(2^{0.645n})$ [69]. Our proposed implementation achieves an improved $O(2^{0.5n})$ runtime while remaining hardware efficient, underscoring the significance of attaining a Grover speedup. Further, the possibility of using our algorithm as a subroutine in more sophisticated classical algorithms

[64,66,70] opens several directions for future work [5,71–74].

In quantifying speedup, we define the runtime of the quantum algorithm in terms of the query complexity, i.e., the number of queries to the oracle. An additional consideration is the physical time required to implement a single query. In our standard algorithm, for a fixed maximum pairwise interaction strength J_{\max} , the spectral resolution γJ_{\max} required of the oracle results in a query time that scales as $\gamma^{-1} \sim 2^{k_c} \approx 2^n$ along the phase boundary $k = k_c$. This exponential scaling highlights the importance of considering not only query complexity, but also the time required to implement the oracle given physical limitations of the hardware (the finite interaction energy). At finite interaction-to-decay ratio, this scaling limits the speedup of the standard algorithm in our simulations, whereas the recursive algorithm achieves higher performance limited only by the increase in optimal number of queries with system size.

In near-term experiments, despite fragility to dissipation, even the standard algorithm could produce a speedup in few-qubit systems in the hard regime, and in scalable systems in the easy phase. In the hard regime and along the phase boundary, achieving the ideal performance at scalable system size is precluded by the exponential decrease of the energy gap with n . If instead we vary n at fixed bit depth k , the time to implement each query saturates to a fixed value set by $\gamma^{-1} \sim 2^k$ as we cross the transition into the easy regime, allowing the ideal performance to be maintained at fixed interaction-to-decay ratio. Irrespective of k , if we fix the duration of each query, the standard algorithm samples from a probability distribution $P(S_z)$ of fixed effective temperature set by $J_{\max}\gamma$, which may enable extensions to Boltzmann sampling [75].

Our hardware-efficient approach to implementing the Grover oracle enables near-term realizations in cold-atom systems, as well as comparisons with alternative proposed methods for solving NP-hard problems in similar platforms [76,77]. Our approach also generalizes to other platforms, including trapped ions [20], or superconducting qubits coupled to phononic [78,79] or microwave [80] resonators. The algorithms presented here might be further optimized by a variational approach that adapts the resolution of the oracle and the number of queries over multiple trials [81]. Grover amplification could also be applied to engineer entangled states, e.g., to produce squeezed or Dicke states by amplifying a particular S_z value. For more versatile quantum control, arbitrary superpositions of Dicke states might be amplified by shaping the drive pulse [27,28,82].

ACKNOWLEDGMENTS

This work is supported by the ONR under Grant No. N00014-17-1-2279 and the AFOSR under Grant No. FA9550-20-1-0059. O.M. acknowledges support from the

ARO under Grant No. W911NF-16-1-0490. J.A.H., M.S.-S., and A.S.-N. acknowledge support from the DOE Office of Science, Office of Basic Energy Sciences, under Grant No. DE-SC0019174. E.S.C. and A.P. are supported by the NSF under Grant No. PHY-1753021, the NSF GRFP (E.S.C.), and the NDSEG Fellowship (A.P.). We thank V. Vuletić and T. Zhang for helpful discussions.

APPENDIX A: GENERALIZATION TO SUBSET SUM PROBLEM

The number partitioning problem is a special case of the more general class of decision problems known as subset sum problems. These problems answer the question: given a set of n objects with positive weights $w_i \in (0, 1]$ of finite bit depth k , does there exist a subset $\mathcal{X} \subset \{w_i\}$ of total weight $\sum_{\mathcal{X}} w_j = W_*$ for a specified value W_* ? The entire class of problems are naturally implemented with the two experimental realizations that we present in detail in Appendices H 1–H 2.

For general subset sum problems, implementing the oracle requires applying a π phase shift to the system of qubits if and only if the total weight of the qubits in state $|1\rangle$ is a specified target weight W_* , i.e., if the system is in an eigenstate of

$$W_1 \equiv \sum_i w_i |1\rangle_i \langle 1|_i, \quad (\text{A1})$$

with eigenvalue W_* . Experimentally, the target weight is set by the frequency of a field that drives the ancilla. For the special case of the partition problem, the target weight is set to $W_* = \sum_i w_i/2$, and the condition in Eq. (A1) then reduces to the condition $S_z|x\rangle = 0$ of the main paper. More generally, the oracle phase shift in terms of W_1 is given by

$$\Phi_\gamma(W_1) = 2 \arctan [2(W_* - W_1)/\gamma] + \pi. \quad (\text{A2})$$

APPENDIX B: INVERSION ABOUT THE AVERAGE

The operator V that performs inversion about the average, also known as the diffusion operator, requires a multi-qubit controlled phase gate similar to the Grover oracle. In particular, the operator

$$V = 2|\psi_0\rangle\langle\psi_0| - \mathbb{1} = H_n R H_n \quad (\text{B1})$$

can be decomposed into two n -qubit Hadamard transforms H_n and a multiqubit controlled phase gate R [2]. The operation H_n is performed by applying a single-qubit Hadamard gate to each qubit. The operator

$$R = 2|0\rangle\langle 0| - \mathbb{1} \quad (\text{B2})$$

is a diagonal matrix in the basis of spin configurations $|x\rangle$, with matrix elements $R_{00} = 1$ and $R_{xx} = -1$ for $x \neq 0$.




FIG. 8. Comparison of speedups Q between two diffusion operator methods: Q versus the interaction-to-decay ratio ρ for the perfect diffusion operator (full lines) and the diffusion operator using the generalized oracle (dashed lines) for $n = k = (4, 6, 8, 10)$, shaded from darkest to lightest.

Thus, R applies a phase shift of π to all basis states except for $|0\rangle$.

The multiqubit controlled phase gate R can be implemented by adapting the protocol used for the generalized oracle. Specifically, the phase gate R is equivalent, up to a global phase, to the Grover oracle for a subset sum problem [Eq. (A1)] with target weight zero. A generalized version R_γ can be implemented by setting all of the weights to the maximum value $w_i = 1$, and simultaneously choosing the detuning to set the target weight $W_* = 0$. We expect the resulting generalized diffusion operator $H_n R_\gamma H_n$ to produce the desired amplification for a relatively broad diffusion step width, requiring only $\gamma < 1$. Notably, the step width permissible for diffusion is much broader than that required for the oracle, allowing inversion about the average to occur with negligible added dissipation even at finite interaction-to-decay ratio ρ .

We verify that added dissipation due to the generalized diffusion operator has negligible effect by examining the quantum speedup. In Fig. 8, we compare the achievable quantum speedup between the perfect diffusion operator and the generalized diffusion operator, in the latter case including effects of decay during diffusion as well as the nonzero step width. The speedup is reduced by at most 23% over a wide range of ρ values, thanks to the less stringent requirement on the step width during the generalized diffusion transform compared with the oracle. Thus, for simplicity, we directly apply the ideal diffusion operator V in the calculations presented in Figs. 2–5 of the main paper.

It is also possible to replace the diffusion operator with only single-qubit rotations, e.g., a global transverse field as in Ref. [24]. While a detailed analysis of this alternative is beyond the scope of the present work, we simulate the application of a transverse field for a time $t = \pi/n$ in lieu of inversion about the average, finding success probabilities approximately half as large as those achieved with the multiqubit diffusion operator. The transverse field thus enables a technically convenient scheme in which the only multiqubit gate is the oracle.

APPENDIX C: CLASSICAL SEARCH ALGORITHMS

In the main text, we evaluate our implementations of Grover’s algorithm by comparing them to the most analogous classical algorithm, memoryless search. We begin this section by quantifying that relationship, discussing the expected and worst-case performance for each method. We then consider increasingly more complex classical number partitioning algorithms and identify their benefits and drawbacks. This allows us to consider how our algorithm compares with the *best* classical algorithms, and indicates prospects for more sophisticated versions of our quantum algorithm.

Both Grover’s algorithm and the classical memoryless search have a probability of success p that is the same for every trial. For such search algorithms, the number of trials M to obtain a solution is a random variable with expected value $E[M] = 1/p$. For Grover’s algorithm the success probability is $P(T_{\text{opt}}) = P_{\text{opt}}$, as defined in the main text, so the expected number of Grover readout measurements M_G is

$$E[M_G] = \frac{1}{P_{\text{opt}}}. \quad (\text{C1})$$

For memoryless search with $N = 2^n$ possible partitions and N_A exact solutions, $p = N_A/N$. The expected number of memoryless trials M_M is then

$$E[M_M] = \frac{N}{N_A}. \quad (\text{C2})$$

Incidentally, when $T = 0$, Grover’s algorithm reduces to measuring an equal superposition of configuration states. The success probability is then $P_0 = N_A/N$, equivalent to memoryless search.

We also consider the worst-case performance of both algorithms. This is equivalent to the number of queries required to reach $\mathcal{P} = 1 - \varepsilon$ probability of having found a solution, in the limit $\varepsilon \rightarrow 0$. For both algorithms, even after an arbitrarily large number of queries, there remains an exponentially small probability that a perfect partition exists but has not been found. We quantify this worst-case




FIG. 9. Speedup scaling over all (n, k) , shown by plotting median total Grover iterations versus memoryless search trials required to reach $\mathcal{P} = 0.99$ probability of success. Each point represents a particular (n, k) , where n and k each take values over the range $[3, 16]$. Black and red lines denote linear and square root dependences, respectively.

performance when $N \gg N_A$ by allowing ε to remain finite, so that the \mathcal{P} quantile of $M_G T_{\text{opt}}$ can be written as

$$[M_G T_{\text{opt}}]_{\mathcal{P}} \in O \left[\ln \left(\frac{1}{\varepsilon} \right) \times \frac{T_{\text{opt}}}{P_{\text{opt}}} \right], \quad (\text{C3})$$

and the \mathcal{P} quantile of M_M is

$$[M_M]_{\mathcal{P}} \in O \left[\ln \left(\frac{1}{\varepsilon} \right) \times \frac{N}{N_A} \right]. \quad (\text{C4})$$

Figure 9 shows the relative median scaling of M_M and $M_G T_{\text{opt}}$, each calculated according to Eq. (12), with Grover’s algorithm showing the expected \sqrt{N} speedup.

While memoryless search follows the same probability distribution as Grover readout measurements, it is not as efficient as linear search through an unsorted list. The expected number of trials M_L for linear search is

$$E[M_L] = \frac{N + 1}{N_A + 1}. \quad (\text{C5})$$

For $N, N_A \gg 1$, the expected trial scaling of both memoryless and linear search algorithms is $O(N/N_A)$. The largest difference occurs with postselection in the hard regime, where $E[N_A] \approx 2$ and memoryless search is expected to take 1.5 times as many trials as linear search.

For the linear search, the worst-case performance is $N - N_A$. More generally, we can take ε arbitrarily close

to 0 such that $[M_L]_{\mathcal{P}}$ converges to $N - N_A$, while retaining the scaling of $[M_M]_{\mathcal{P}}$ in Eq. (C4). Thus, while both algorithms are both worst-case linear in N , worst-case memoryless search requires $O[\ln(1/\varepsilon)/N_A]$ times as many queries as worst-case linear search in the hard regime. Further, because both algorithms are unstructured, they do not need to precalculate a potentially exponential number of values before performing queries. Thus their memory scaling is $O(n)$, set by the number of values to be partitioned.

Improving upon memoryless and linear search requires us to consider a variety of structured search algorithms, which can be grouped based on the difficulty of the problem instance they aim to solve. An instance’s difficulty is related to its density, defined for a set of integer weights $\mathbf{a} = (a_1, \dots, a_n)$ as the ratio $d = n / \log_2(\max_i a_i)$ of the number of weights to the number of bits needed to represent the largest weight [83,84]. Thus, $d < 1$ corresponds to the “hard phase” and $d > 1$ to the “easy phase” [4]. In the easy phase there are typically many perfect partitions and a problem instance can generally be solved efficiently by various classical methods, with the best based on the Karmarkar-Karp differencing algorithm [50,70]. In the hard phase, classical algorithms have been demonstrated to solve “almost all” problems of density $d < d_c < 1$ in polynomial time, with subsequently published algorithms pushing d_c closer to 1 [83–86]. In such “low-density attack” algorithms, the number partitioning problem is reduced to the shortest vector problem, for which there exist algorithms that produce good approximations in polynomial time.

Indeed, the hardest instances of the number partitioning problem are not deep into the hard phase, but near the phase transition at a density close to 1 [4,55]. For such instances, classical algorithms with the best known time complexity are subject to a space-time trade-off; improvements in runtime come at the cost of exponential memory requirements [67,68]. However, Esser and May devised a classical algorithm that achieves a time complexity of $O(2^{0.645n})$ while maintaining polynomial space complexity [69]. This algorithm offers a compelling comparison to our proposed Grover implementations, as each algorithm is hardware efficient in its use of memory or qubits. With our Grover implementation requiring $O(2^{0.5n})$ queries, a direct comparison would yield a speedup $O(2^{0.145n})$.

Finally, an interesting open question is whether one can design hardware-efficient quantum algorithms that exploit the problem structure of number partitioning. Answers to this question would build on recent work that combined quantum and classical methods to produce hybrid algorithms with exponential time, memory, and qubit trade-offs [5,71–74]. One avenue to explore is the use of our Grover search as a subroutine in a differencing algorithm, in which a pair of large weights w_i, w_j is replaced by their difference to reduce the size of the search space. Such differencing

could be performed either classically (representing $w_i - w_j$ by a single spin) or quantumly (representing $w_i - w_j$ by an entangled state $|01\rangle + |10\rangle$ of two spins). While classical differencing has the potential benefit of reducing the dynamic range of the weights, quantum differencing generalizes to initializing the system in a superposition state that reflects classically computed probabilities of finding certain pairs of spins on opposite sides of a perfect partition.

APPENDIX D: NUMERICAL METHODS

The simulations of Grover's algorithm are performed numerically, by matrix multiplication according to Eq. (10). Each simulation for a specific problem size is performed on an ensemble of lists of randomly selected weights. To postselect on the existence of solutions, for each list of weights we first use the classical complete Karmarkar-Karp differencing algorithm [70] to search for solutions, and simulate the quantum algorithm only for instances with solutions. The number of problem instances in an ensemble, after postselection where applicable, ranges from 1000 to 5000 for all datasets except that used for Fig. 4(a.i), in which each probability distribution $P(S_z)$ is determined from 5×10^4 instances.

To find a sufficiently narrow step width γ to reach a specified success probability P_{opt} in Fig. 4(a.ii), we generate an ensemble of weights and numerically optimize γ using the Nelder-Mead algorithm to reach the specified value P_{opt} . To find the optimal step width γ in the presence of decay [Figs. 5 and 7(b)], we similarly optimize γ to minimize the median total number of Grover iterations using a gradient-descent algorithm.

APPENDIX E: CAPTURE RANGE AND AMPLIFICATION

The interpretation of the step width γ as a capture range for S_z values is illustrated in Fig. 4(a.i) of the main text, where we plot the amplification factor after T_{opt} Grover iterations. Here, we additionally present an analytic derivation of the amplification factor after a single Grover iteration. Specifically, for a given spin configuration $|x\rangle$, we show that the amplification factor after the first Grover cycle is of the Lorentzian form

$$\left| \frac{c_{x,1}}{c_{x,0}} \right|^2 = \frac{A}{1 + (2S_z/\gamma)^2} + B, \quad (\text{E1})$$

with width γ set by the width of the phase step. While the amplitude A and offset B depend on the set of weights, we analytically derive their values averaged over instances of the weights to determine the amplification factor at $S_z = 0$ as a function of step width.

We first consider the combined effect of the generalized oracle and inversion about the average on a generic state

$$|\psi_T\rangle = \sum_x c_{x,T} |x\rangle. \quad (\text{E2})$$

The state $|\psi_{T+1}\rangle = VU_\gamma |\psi_T\rangle$ is characterized by coefficients

$$c_{x,T+1} = -e^{i\Phi_\gamma(x)} c_{x,T} + \frac{2}{N} \sum_{x'} e^{i\Phi_\gamma(x')} c_{x',T}. \quad (\text{E3})$$

Equation (E3) simplifies for the case of $T = 0$, where all coefficients $c_{x,0} = 1/\sqrt{N}$ are equal. Thus, after the first Grover iteration, we have

$$\frac{c_{x,1}}{c_{x,0}} = -e^{i\Phi_\gamma(x)} + \frac{2}{N} \sum_{x'} e^{i\Phi_\gamma(x')}. \quad (\text{E4})$$

In terms of phasors $\chi(x) = e^{i\Phi_\gamma(x)}$ and the average phasor $\bar{\chi} = \sum_x \chi(x)/N$, the gain in probability of finding the system in state $|x\rangle$ is then given by

$$G(x) \equiv \left| \frac{c_{x,1}}{c_{x,0}} \right|^2 = 4|\bar{\chi}|^2 - 4\text{Re}[\chi(x)\bar{\chi}] + |\chi(x)|^2. \quad (\text{E5})$$

We now proceed to account for the specific functional form $\Phi_\gamma(x) = 2 \arctan(2S_z/\gamma) + \pi$ of the oracle's phase response. Defining $\mu(x) \equiv 2S_z(x)/\gamma$ as the weighted spin normalized by the step width, we have

$$\chi = \frac{\mu^2 - 1}{\mu^2 + 1} - i \frac{2\mu}{1 + \mu^2}. \quad (\text{E6})$$

Furthermore, since for each spin configuration $|x\rangle$ with weighted spin S_z there exists a complementary spin configuration with weighted spin $-S_z$, the average phasor $\bar{\chi}$ is always real. Equation (E5) then reduces to

$$G(\mu) = (1 - 2\bar{\chi})^2 + \frac{8\bar{\chi}}{1 + \mu^2}. \quad (\text{E7})$$

This result is of the Lorentzian form in Eq. (E1), with amplitude $A = 8\bar{\chi}$ and offset $B = (1 - 2\bar{\chi})^2$. The gain in the first Grover cycle for a solution state ($\mu = 0$) is bounded above by $G_{\text{max}} = 9$, which is achieved if $\bar{\chi} = 1$ and approached in the limit where the number of solutions is small and the step is narrow.

For illustration, we examine a single iteration of Grover's algorithm applied to number partitioning with $n = 12$ random weights of bit depth $k = 12$. Figure 10 shows the amplification factor G averaged over all spin configurations $|x\rangle$ with the same value of the weighted spin S_z , as a function of step width γ . Cuts at fixed γ are well




FIG. 10. Top: Amplification factor G after one Grover iteration for a single representative instance of weights at $n=k=12$. Fits with the Lorentzian model of Eq. (E1) are shown as solid lines. Offset B is the sole free fit parameter, which is related to A by Eq. (E7). Bottom: Initial S_z probability density distribution for $n = k = 12$, averaged over 10^4 instances (yellow points with shading). Expected initial distribution, a Gaussian with $\sigma_{S_z} = 1$, shown as a solid orange line.

fit by the Lorentzian form in Eq. (E7) with $\mu = 2S_z/\gamma$, confirming that the step width γ sets the capture range for amplification. The peak amplification $G_0 \equiv G(0)$ remains near its maximum possible value $G_{\max} = 9$ until the width γ grows to roughly σ_{S_z}/G_{\max} , where σ_{S_z} denotes the width of the initial S_z distribution, which we plot for comparison in the bottom panel of Fig. 10.

The amplification G_0 of solution states depends to lowest order only on the ratio of γ to the width $\sigma_{S_z} \propto \sqrt{n}$ of the S_z distribution. To calculate the dependence of G_0 on γ/\sqrt{n} from Eq. (E7), we express $\bar{\chi}$ in terms of the number of partitions $g(\mu)$ with a given value of the imbalance μ :

$$\bar{\chi} = \frac{1}{N} \sum_{\mu} g(\mu) \frac{\mu^2 - 1}{\mu^2 + 1}. \quad (\text{E8})$$

Here, we have used the relation $g(\mu) = g(-\mu)$ to eliminate the term that is odd in μ . Assuming a large number $N \gg 1$ of spin configurations, we approximate the average multiplicity $\langle g(\mu) \rangle$ over many instances of the weights using a normal distribution

$$p(\mu) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\mu^2/(2\sigma^2)} \quad (\text{E9})$$

of standard deviation $\sigma = w_{\text{rms}}\sqrt{n}/\gamma$, where $w_{\text{rms}} \equiv \sqrt{\langle w_i^2 \rangle} = 1/\sqrt{3}$ for weights chosen from a uniform distribution on $(0, 1]$. In terms of $p(\mu)d\mu \approx \langle g(\mu) \rangle/N$, we have

$$\begin{aligned} \langle \bar{\chi} \rangle &= \int_{-\infty}^{\infty} p(\mu) \frac{\mu^2 - 1}{\mu^2 + 1} d\mu \\ &= 1 - \frac{\sqrt{2\pi} e^{1/(2\sigma^2)}}{\sigma} \text{erfc} \left(\frac{1}{\sqrt{2}\sigma} \right), \end{aligned} \quad (\text{E10})$$

where erfc is the complementary error function. The average amplification over many instances of the weights is given in terms of $\bar{\chi}$ by

$$\langle G_0 \rangle = 1 + 4\langle \bar{\chi} \rangle + 4\langle \bar{\chi}^2 \rangle \geq 1 + 4\langle \bar{\chi} \rangle + 4\langle \bar{\chi} \rangle^2. \quad (\text{E11})$$

This bound is tight in the large- N limit, where the variance in $\bar{\chi}$ over different instances of the weights is small. We plot the lower bound in Eq. (E11) as the dashed red curve in Fig. 5(c). There, we denote the amplification as $Q_1 \equiv G_0$ to emphasize its equivalence to the quantum speedup for a single Grover cycle. We compare our model with the amplification calculated at $n = k$ for $n = 12$, in each case averaging over 10^3 instances of the weights with postselection. We observe excellent agreement between the model and the simulation.

APPENDIX F: SCALABLE ALGORITHM

In Sec. V, we outline a recursive version of our algorithm that allows for operating at a fixed resolution $\gamma \sim 2^{-m}$ of the oracle for arbitrary problem size. The essence of our approach is to consider only the ℓm least significant bits of the weights for successive values $\ell = 1, 2, 3, \dots$. These truncated weights suffice to amplify, in each layer ℓ of the algorithm, candidate solutions satisfying the condition $\text{mod}(2^k S_z, 2^{\ell m}) = 0$. In the final layer of the algorithm, the fixed m -bit resolution of the oracle suffices to identify only true solutions satisfying $S_z = 0$, thanks to the preamplification of a sparse distribution of S_z values in prior layers.

In this appendix, we elaborate on the details of the scalable algorithm, including the encoding of the weights, the modular oracle, and the recursive implementation of inversion about the average. Finally, we derive the asymptotic scaling of the query complexity and present numerical simulations corroborating our analysis.

1. Encoding the weights

Key to our approach is the ability to dynamically change the mapping from weights w_i to system-ancilla couplings J_i between successive queries of the oracle. We define the following set of mappings from weights to system-ancilla

couplings:

$$J_{i,\ell} = \frac{J_{\max} \text{mod}(2^k w_i, 2^{\ell m})}{2^{\ell m}}, \quad (\text{F1})$$

with $\ell = 1, 2, \dots, k/m$. Here, we scale the couplings to a fixed maximum value J_{\max} as usual, but for a given value ℓ we use only the ℓm least significant bits of the weights. [For $\ell > 1$, we are keeping more bits than the oracle can actually resolve, to avoid subtleties of accounting for carry bits that can add up. It should never be necessary to program the weights with a resolution of more than $m + \log_2(n)$ bits, but the higher precision assumed in Eq. (F1) also does no harm.]

2. Grover amplification with modular oracle

In each layer ℓ of our algorithm, our objective is to amplify spin configurations satisfying the condition $\text{mod}(2^k S_z, 2^{\ell m}) = 0$, i.e., spin configurations for which the ℓm least significant bits of the imbalance S_z are zero. To check this condition for a given value of ℓ , we need only to know the ℓm least significant bits of each weight, so we use the couplings $J_{i,\ell}$ defined in Eq. (F1) for each layer ℓ . We further define

$$\begin{aligned} S_{z,\ell} &= 2^{-\ell m} \sum_{i=1}^n \text{mod}(2^k w_i, 2^{\ell m}) \sigma_i^z / 2 \\ &= \frac{1}{J_{\max}} \sum_{i=1}^n J_{i,\ell} \sigma_i^z / 2, \end{aligned} \quad (\text{F2})$$

representing the imbalance at layer ℓ using only the ℓm least significant bits of the weights. We wish to design the oracle to produce a π phase shift if $\text{mod}(2^{\ell m} S_{z,\ell}, 2^{\ell m}) = 0$, which is equivalent to the ancilla resonance energy shift being an integer multiple of J_{\max} .

This modular oracle can be implemented in the central spin or central boson model by subjecting the ancilla to a multifrequency drive field, consisting of a comb with spacing J_{\max} . Since $S_{z,\ell} \leq n$, there are approximately n different possible values $2^{\ell m} S_{z,\ell}$ that are equivalent to zero modulo $2^{\ell m}$, and correspondingly only approximately n drive frequencies are needed. Each tooth of the comb of drive fields has a spectral width γJ_{\max} , which we choose to be independent of ℓ , with a value $\gamma \sim 2^{-m} \ll 1$. The separation of scales between the width and the spacing of the teeth ensures that the phase response of the oracle is well approximated as

$$\Phi_\ell \approx 2 \arctan \left[\frac{2 \text{mod}(2^k S_z, 2^{\ell m}, -2^{\ell m-1})}{2^{\ell m} \gamma} \right] + \pi, \quad (\text{F3})$$

where $\text{mod}(\cdot, d, b)$ denotes the modulo with divisor d and offset b . In terms of the phase shift Φ_ℓ at layer ℓ , we define

the modular oracle $U_\ell = \exp(i\Phi_\ell)$. In the final layer of the algorithm, where $\ell m = k$, we want to amplify only the partitions with $S_z = 0$ without taking the modulus, so we apply our usual oracle $U_{k/m} = \exp(i\Phi)$ with resolution γ .

3. Recursive algorithm

The first layer of our algorithm consists simply of applying the modular oracle in alternation with the diffusion operator $V = H_n R H_n$, where H_n is the n -qubit Hadamard and R is a multiqubit controlled phase gate (Appendix B). For an imperfect oracle, we apply the usual spin-echo sequence to produce a state

$$|\psi_1\rangle = (V U_1^\dagger V U_1)^{T_1/2} |\psi_0\rangle, \quad (\text{F4})$$

where $|\psi_0\rangle = H_n |0\rangle^{\otimes n}$ is the equal superposition of all spin configurations. We assume the number of amplification cycles T_1 to be even for notational simplicity, but Eq. (F4) can also be generalized to allow an odd number of cycles. Since only a fraction 2^{-m} of the spin configurations satisfy the condition $\text{mod}(2^m S_{z,1}, 2^m) = 0$, we expect to require approximately $T_\ell \approx (\pi/4) 2^{m/2}$ amplification cycles in the first layer $\ell = 1$. Upon completion of this layer of the algorithm, the state $|\psi_1\rangle \equiv \mathcal{G}_1 |\psi_0\rangle$ is approximately an equal superposition of all spin configurations that are candidate solutions to the number partitioning problem based on the m least significant bits of S_z .

Naively one might expect subsequent layers of our algorithm to be analogous to Eq. (F4) with the replacement $U_1 \rightarrow U_\ell$. However, an important subtlety is that the diffusion operator V must be modified so that at layer ℓ it rotates by π about the state $|\psi_{\ell-1}\rangle$, i.e.,

$$V_\ell = 2|\psi_{\ell-1}\rangle\langle\psi_{\ell-1}| - \mathbb{1}. \quad (\text{F5})$$

In particular, it is important to rotate about $|\psi_{\ell-1}\rangle$ —as opposed to $|\psi_0\rangle$ —so that the amplitude of the solution states is inverted about the average amplitude in the sparse superposition of S_z values produced in the preceding layer, while ignoring the near-zero amplitudes of the spin configurations that have already been suppressed.

To understand how to implement the generalized diffusion operator V_ℓ , we first recall how our usual diffusion operator is constructed [Eqs. (B1)–(B2)]. We can perform a π rotation about any state $|\psi\rangle$ by a combination of (1) the operator R that rotates about the state $|0\rangle \equiv |0\rangle^{\otimes n}$ and (2) a unitary operator \mathcal{O} that transforms $|\psi\rangle$ to $|0\rangle$. In terms of these ingredients, the rotation about $|\psi\rangle$ is implemented by applying the compound operator $\mathcal{O}^\dagger R \mathcal{O}$. For the usual Grover's algorithm, the n -qubit Hadamard $\mathcal{O} = H_n = \mathcal{O}^\dagger$ is the operator that transforms $|\psi\rangle$ to $|0\rangle$ and back.

To construct the diffusion operator V_ℓ for any layer of our algorithm, we thus require an operator \mathcal{O} that transforms the state $|\psi_{\ell-1}\rangle$ to state $|0\rangle$. Conveniently, we know

exactly how to perform this transformation for arbitrary ℓ —by applying Grover's algorithm all the way up to layer $\ell - 1$. If we define the Grover operator at level ℓ as

$$\mathcal{G}_\ell \equiv \left(V_\ell U_\ell^\dagger V_\ell U_\ell \right)^{T_\ell/2}, \quad (\text{F6})$$

such that $|\psi_\ell\rangle = \mathcal{G}_\ell |\psi_{\ell-1}\rangle$, then the operator $\mathcal{O}^\dagger = \left(\prod_{\ell'=1}^{\ell-1} \mathcal{G}_{\ell'} \right) H_n$ transforms $|0\rangle^{\otimes n}$ to $|\psi\rangle_{\ell-1}$. Thus, the diffusion operator needed in layer ℓ of the algorithm is

$$V_\ell = \left(\prod_{\ell'=1}^{\ell-1} \mathcal{G}_{\ell'} \right) H_n R H_n \left(\prod_{\ell'=1}^{\ell-1} \mathcal{G}_{\ell'} \right)^\dagger. \quad (\text{F7})$$

Note that Eq. (F7) correctly reduces to $V_1 = V$ for the first layer of our algorithm.

4. Query complexity

Due to the recursive nature of the algorithm, the query complexity grows exponentially with k and hence with n in the hard regime. This should not surprise us, since Grover's algorithm cannot produce an exponential speedup. The key performance metric, then, is the coefficient α in the exponent of the $O(2^{\alpha n})$ query complexity.

The query complexity is given by

$$T_{\text{tot}} = \sum_{\ell=1}^{k/m} T_\ell \tau_\ell, \quad (\text{F8})$$

where T_ℓ is the number of amplification cycles at layer ℓ and τ_ℓ is the number of calls to the oracle required in each amplification cycle, including the queries involved in implementing the diffusion operator V_ℓ for $\ell > 1$. We expect to need $T_\ell \approx (\pi/4)2^{m/2}$ amplification cycles at each layer except the final one, by the same argument given above for $\ell = 1$. The final layer takes a factor of \sqrt{n} more steps, but this factor will only introduce a subexponential correction to the query complexity so we can ignore it in the following analysis. The number of calls to the oracle in each amplification cycle of the ℓ th layer is

$$\tau_\ell = 1 + \sum_{\ell'=1}^{\ell-1} 2T_{\ell'} \tau_{\ell'}, \quad (\text{F9})$$

based on Eq. (F7). Put another way, we have

$$\begin{aligned} \tau_\ell &= \tau_{\ell-1} (1 + 2T_{\ell-1}) \\ &\approx \tau_{\ell-1} \left[1 + 2^{m/2} \left(\frac{\pi}{2} \right) \right]. \end{aligned} \quad (\text{F10})$$

Since the first layer requires only $\tau_1 = 1$ call to the oracle per amplification cycle, for general ℓ we have

$$\tau_\ell = \left[1 + 2^{m/2} \left(\frac{\pi}{2} \right) \right]^{\ell-1}, \quad (\text{F11})$$

as can readily be verified by induction.

The total number of calls to the oracle given by Eq. (F8) thus takes the form of a finite geometric series. Evaluating the geometric series yields

$$\begin{aligned} T_{\text{tot}} &= \left(\frac{\pi}{4} \right) 2^{m/2} \left(\frac{[1 + 2^{m/2} (\pi/2)]^{k/m} - 1}{2^{m/2} (\pi/2)} \right) \\ &\approx 2^{k/2-1} (\pi/2)^{k/m} \\ &= 2^{k(1/2+c/m)-1}, \end{aligned} \quad (\text{F12})$$

where $c = \log_2(\pi/2) \approx 0.65$. For $n \approx k$, we obtain a query complexity $O(2^{\alpha n})$ with $\alpha = 0.5 + 0.65/m$. We thus need $m \approx 5$ bits of resolution to outperform the best scalable classical algorithm [69].

While the query complexity of the recursive algorithm is at best (i.e., for large m) the same as that of the standard algorithm, the recursive algorithm offers the benefit that the actual runtime in a scalable implementation with fixed J_{max} is directly proportional to the query complexity, and thus exhibits a Grover speedup. We thus eliminate the exponential overhead that is present in the simplest algorithm, and we do so without compromising on hardware efficiency.

5. Simulation

A representative comparison of the standard algorithm and the recursive algorithm is shown in Fig. 11(a), where we simulate a single instance of number partitioning with $n = k = 12$. The instance is selected to have exactly one pair of solutions. The recursive algorithm is performed with $m = 4$ bits of resolution, and the amplification steps per layer $(T_1, T_2, T_3) = (2, 3, 2)$ are chosen to maximize the probability at each layer. The resolution of the oracle is set to $\gamma = 2^{-m-1}$ for the recursive algorithm (orange squares), compared with $\gamma = 2^{-k}$ for the standard algorithm (blue circles). While the recursive algorithm required approximately 3 times as many queries as the standard algorithm, the physical time per query at fixed J_{max} is a factor of $2^{k-m-1} = 2^7$ times longer for the standard algorithm than for the recursive algorithm. Thus, in this example the recursive algorithm produces a significant reduction in runtime for a fixed maximum system-ancilla coupling.

To compare the time complexity of the algorithms, we simulate both algorithms for a range of problem sizes (n, k) with 1000 instances of weights [Figs. 11(b) and 11(c)]. For each layer ℓ of the recursive algorithm with $m = 6$ and $\gamma = 2^{-m-1}$, the number of amplification cycles T_ℓ is optimized

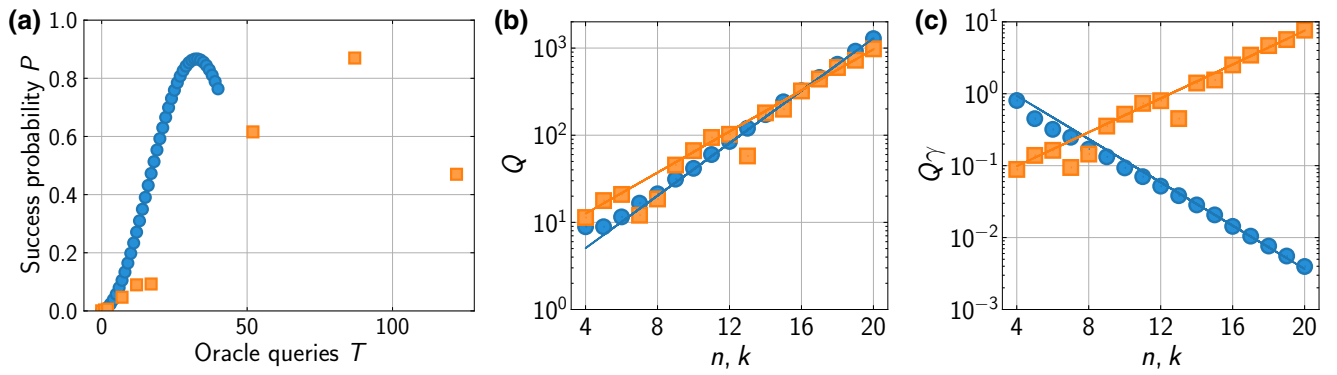


FIG. 11. Comparison of the standard (blue circles) and recursive (orange squares) algorithms. (a) Success probability P versus number of oracle queries T for a single instance of number partitioning at $n = k = 12$. Recursive algorithm is performed with $m = 4$ and $\gamma = 2^{-m-1}$. The physical time per query (at fixed J_{\max}) is longer by a factor of $2^{k-m-1} = 2^7$ for the standard algorithm than for the recursive algorithm. (b) Median speedup $[Q]_{0.5}$ in query complexity, plotted versus $n = k$ for the standard algorithm with $\gamma = \gamma_c$ and the recursive algorithm with $m = 6$ and $\gamma = 2^{-m-1}$. Lines denote the $2^{0.5n}$ scaling for the standard algorithm and $2^{(0.5-0.65/m)n}$ scaling for the recursive algorithm. (c) Speedup $[Q]_{0.5} \gamma$ in physical runtime at fixed J_{\max} , plotted vs. $n = k$ for the standard algorithm with $\gamma = \gamma_c$ and the recursive algorithm with $m = 6$ and $\gamma = 2^{-m-1}$. Lines denote the $2^{-0.5n}$ scaling for the standard algorithm and $2^{(0.5-0.65/m)n}$ scaling for the recursive algorithm.

to minimize the median total number of Grover oracle queries T_{total} . We expect the total number of Grover queries T_{total} in the recursive algorithm to follow the query complexity derived in the preceding section (Appendix F 4). For $m = 6$, the expected scaling is $T_{\text{total}} \in O(2^{0.61n})$, leading to a $Q \in O(2^{0.39n})$ scaling of the speedup with system size, which is confirmed by the simulation in Fig. 11(b). While the speedup Q in query complexity for the recursive algorithm at finite bit depth m is slightly lower than that of the standard algorithm, the benefit of the recursive algorithm becomes apparent when we plot the speedup $Q\gamma$ in physical runtime at fixed J_{\max} [Fig. 11(c)]. The growth in $Q\gamma$ with system size in the recursive algorithm confirms its scalability.

APPENDIX G: EFFECTS OF DECOHERENCE

Two forms of decoherence that can limit the performance of our algorithm in realistic implementations are decay of the ancilla and decay of the system spins. In this section, we first provide an analytic estimate of the scaling of the quantum speedup with a generic interaction-to-decay ratio in the standard algorithm (Appendix G 1). We then describe how we calculate the speedup in the numerical simulations of Fig. 5, focusing on decay of the ancilla, which is the dominant decay channel in the near-term experimental implementations proposed and analyzed in Appendix H.

1. Quantum speedup in presence of decay

Decay during the generalized Grover's oracle limits the maximum achievable quantum speedup. Here, we analytically derive the scaling of optimal quantum speedup with the interaction-to-decay ratio for the standard algorithm

presented in Secs. II–III. The speedup is maximized at a step width γ_{opt} set by a competition between the reduction in capture range at narrower step widths, which ideally increases the success probability, and the accompanying increase in decay. Figure 12 shows the optimal step width and the optimal number of Grover iterations T_{opt} that produce the speedup shown in Fig. 5(b) of the main text. At small interaction-to-decay ratios, it is optimal to use a single amplification cycle with a wide phase step, while at larger interaction-to-decay ratios, the optimal step width is narrower, allowing for a performance closer to that of the ideal Grover's algorithm.

To estimate the optimal step width, we observe that the number of partitions $N_{\mathcal{A}}^{\text{eff}}(\gamma)$ within the capture range




FIG. 12. (a) Ratio of optimal step width γ to critical step width γ_c versus the interaction-to-decay ratio ρ for $n = k = (4, 6, 8, 10)$ denoted by markers shaded from darkest to lightest. (b) Optimal number of Grover iterations T_{opt} versus the interaction-to-decay ratio ρ for $n = k = (4, 6, 8, 10)$ denoted by markers shaded from darkest to lightest. Dashed lines represent the number of iterations T_{opt} at which the speedup is maximized in the ideal Grover's algorithm.

$|S_z| \lesssim \gamma$ sets the behavior of the generalized Grover's algorithm in roughly the same way as the number of perfect partitions N_A sets the behavior of the ideal Grover's algorithm. With increasing step width, in the absence of dissipation, the number of iterations required to maximize the success probability decreases as

$$T_{\text{opt}}^* \approx \frac{\pi}{4} \sqrt{\frac{N}{N_A^{\text{eff}}}}, \quad (\text{G1})$$

in analogy to Eq. (9). (In defining T_{opt}^* to maximize the success probability, we are choosing a slightly different definition from that of T_{opt} in the main text.) For large step widths, where we capture a larger number N_A^{eff} of spin configurations than the actual number of solutions N_A , we can approximate $N_A^{\text{eff}} \approx \gamma N \sqrt{\frac{6}{\pi n}}$ from the theoretical distribution of total weights in the partition problem [87,88]. Thus, in terms of the step width γ , we have

$$T_{\text{opt}}^* = \frac{\pi}{4\gamma^{1/2}} \left(\frac{\pi n}{6}\right)^{1/4}. \quad (\text{G2})$$

The decrease in the optimal number of iterations T_{opt}^* with increasing step width comes at the cost of a reduced success probability $P_{\text{opt}} \approx N_A/N_A^{\text{eff}}$, even before accounting for dissipation. Thus, employing a narrower step for a larger number of iterations T is preferable unless decay results in an appreciable reduction in P_{opt} . To estimate the optimal number of Grover iterations at finite interaction-to-decay ratio ρ , we first determine the maximum number T_C of iterations that can be performed with a given probability e^{-C} of incurring no error. Here, C is a constant that we choose to optimize the speedup. The error rate per iteration is $D/(\rho\gamma)$, where D is an order-unity factor that is derived in Appendix G 2 for the case of the first amplification step and, more generally, can be obtained from a fit to numerical data. We thus estimate the maximum number of iterations as $T_C \approx C\rho\gamma/D$.

We expect the optimum number of iterations in the presence of decay to be given by $T_{\text{opt}}^* = T_C$ for some order-unity value C . Combining the expression for T_C and the relationship between T_{opt}^* and γ_{opt} [Eq. (G2)], the optimal step width is then

$$\gamma_{\text{opt}} = \left(\frac{\pi D}{4C\rho}\right)^{2/3} \left(\frac{\pi n}{6}\right)^{1/6}. \quad (\text{G3})$$

To estimate the speedup Q_{opt} , we approximate P_{opt} in the presence of dissipation as $P_{\text{opt}} \approx e^{-C} N_A/N_A^{\text{eff}}$. The speedup Q_{opt} is then given by

$$Q_{\text{opt}} = \frac{\log(1 - P_{\text{opt}})}{T_{\text{opt}}^* \log(1 - P_0)} \approx \frac{P_{\text{opt}}}{T_{\text{opt}}^* P_0}, \quad (\text{G4})$$

where $P_0 = N_A/N$ and we assume $P_{\text{opt}} \ll 1$ and $P_0 \ll 1$. Finally, collecting the expressions, we find

$$\begin{aligned} Q_{\text{opt}} &= \frac{e^{-C}}{T_{\text{opt}}^* \gamma_{\text{opt}}} \sqrt{\frac{\pi n}{6}} \\ &= \left(\frac{4}{\pi}\right)^{4/3} \left(\frac{\pi n}{6}\right)^{1/6} e^{-C} \left(\frac{C\rho}{D}\right)^{1/3}. \end{aligned} \quad (\text{G5})$$

The scaling of the optimal speedup as a function of interaction-to-decay ratio is given by $Q_{\text{opt}} \sim \rho^{1/3}$. For high values of ρ , the optimal speedup will start to saturate to the quantum speedup of the ideal Grover's algorithm. This saturation occurs when the optimal step width becomes smaller than the smallest nonzero $|S_z|$ values, which for $n = k$ is at $\gamma_{\text{opt}} \approx \sqrt{n}/N$, with $N = 2^n$. Thus, the interaction-to-decay ratio where the speedup starts to saturate scales as $\rho \sim N^{3/2}/\sqrt{n}$. This scaling exemplifies the fact that reaching the ultimate quantum speedup allowed by Grover's algorithm requires exponentially increasing the interaction-to-decay ratio with problem size.

The numerical results of the generalized Grover's algorithm with ancilla decay in Fig. 5(b) are well described by the model of Eq. (G5) with constants $C = 1/3$ and $D = 1.2$. This equation is applicable in a region between $100 \lesssim \rho \lesssim N^{3/2}/\sqrt{n}$. The upper limit of this regime of validity comes from the saturation of the speedup to the ideal Grover's algorithm limit, while the lower limit is reached when $T_{\text{opt}}^* = 1$.

2. Generalized oracle with ancilla decay

The effect of ancilla decoherence during the generalized Grover's oracle can be modeled as an imaginary term in the oracle phase shift [Eq. (3)]. A particular system spin configuration $|x\rangle$ will shift the ancilla excited state from resonance by $\Delta_x = (W_* - W_1)J_{\text{max}}$, where W_1 and W_* are the actual and target weights in the subset sum problem as defined in Appendix A. To include the effect of ancilla decoherence, we make a substitution $\Delta_x \rightarrow \Delta_x + i\Gamma_a/2$, where Γ_a is the linewidth of the ancilla excited state [89]. Thus, the oracle phase shift applied to the spin configuration $|x\rangle$ is given by

$$\begin{aligned} \Phi_\gamma(W_1) &= 2 \arctan [2(W_* - W_1)/\gamma + i\Gamma_a/(J_{\text{max}}\gamma)] + \pi \\ &= 2 \arctan (\mu + ir) + \pi. \end{aligned} \quad (\text{G6})$$

Here, $\mu = 2(W_* - W_1)/\gamma$ in an analogy to the definition in Appendix E and

$$r \equiv \frac{\Gamma_a}{J_{\text{max}}\gamma} = \frac{1}{\rho\gamma} \quad (\text{G7})$$

parameterizes the decay rate per query of the oracle, assuming the decay is dominated by the ancilla decay.

The effect of the oracle on the amplitudes of the spin states is given by $\chi(W_1) = \exp[i\Phi_\gamma(W_1)]$. Using Eq. (G6) we derive

$$\chi(W_1) = -\frac{1 + i\mu - r}{1 - i\mu + r}. \quad (\text{G8})$$

This full form of the oracle including dissipation modifies the single-cycle amplification formula given in Appendix E. To see how, we rewrite χ in terms of its real and imaginary components,

$$\chi = \frac{\mu^2 + r^2 - 1}{\mu^2 + (r+1)^2} - i\frac{2\mu}{(1+r)^2 + \mu^2}, \quad (\text{G9})$$

where we use the fact that both r and μ are real.

The expression for the amplification in Eq. (E5) now reduces to

$$G(\mu) = 4\bar{\chi}(\bar{\chi} - 1) + \frac{(1-r)^2}{(1+r)^2} + \frac{8\bar{\chi}(1+r)}{(1+r)^2 + \mu^2}. \quad (\text{G10})$$

As before, $\bar{\chi}$ is real and thus depends only on the real components of χ , weighted by the density of states $g(\mu)$:

$$\bar{\chi} = \frac{1}{N} \sum_{\mu} g(\mu) \frac{\mu^2 + r^2 - 1}{\mu^2 + (r+1)^2}. \quad (\text{G11})$$

Taking the continuum limit and using the probability distribution $p(\mu)$ derived in Appendix E yields the updated expectation value,

$$\begin{aligned} \langle \bar{\chi} \rangle &= \int_{-\infty}^{\infty} p(\mu) \frac{\mu^2 + r^2 - 1}{\mu^2 + (r+1)^2} d\mu \\ &= 1 - \frac{\sqrt{2\pi} e^{(1+r)^2/(2\sigma^2)}}{\sigma} \operatorname{erfc}\left(\frac{1+r}{\sqrt{2}\sigma}\right). \end{aligned} \quad (\text{G12})$$

From Eqs. (G10) and (G12) we compute the average amplification over many instances:

$$\langle G_0 \rangle \geq \frac{(1-r)^2}{(1+r)^2} + \left(\frac{8}{1+r} - 4\right) \langle \bar{\chi} \rangle + 4\langle \bar{\chi} \rangle^2. \quad (\text{G13})$$

The amplification in Eq. (G13) is a lower bound both due to the substitution of $\langle \bar{\chi} \rangle^2$ for $\langle \bar{\chi}^2 \rangle$ and due to the small additional probability, which we elsewhere neglect, that the spins end up in a solution state following a dissipation event. The inequality becomes exact in the limit of large N and low dissipation $r \ll 1$. To estimate the reduction in amplification due to dissipation in this limit, we assume a phase step sufficiently narrow that $\langle \bar{\chi} \rangle \approx 1$. Expanding Eq. (G13) to lowest order in r then yields

$$\langle G_0 \rangle \approx 9(1 - 4r/3). \quad (\text{G14})$$

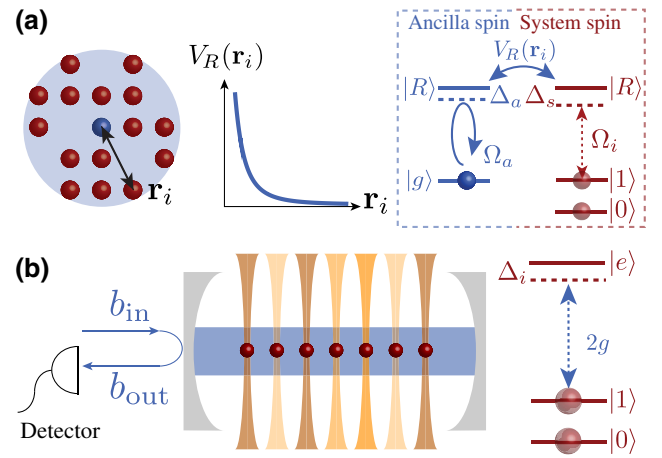


FIG. 13. (a) Central spin model realized by Rydberg-dressed atoms (red) interacting with ancilla qubit encoded on a ground-to-Rydberg transition (blue). (b) Central boson model realized by driving one-sided cavity coupled to system spins and heralding on photodetection.

APPENDIX H: EXPERIMENTAL IMPLEMENTATIONS

1. Central spin model with Rydberg atoms

As a central spin system for encoding subset sum problems, we consider an array of atoms that can be optically coupled to Rydberg states to controllably turn on the interaction Hamiltonian H_q [Eq. (4)]. The implementation is illustrated in Fig. 13(a). The spins of the system atoms are encoded in two ground states $|0\rangle, |1\rangle$. The ancilla qubit is encoded using a ground state $|g\rangle$ and a Rydberg state $|R\rangle$, in terms of which we define the spin raising operator $I_+ = |R\rangle\langle g|$ and lowering operator $I_- = |g\rangle\langle R|$. The system is initialized with the ancilla in state $|g\rangle$ and the system spins in state $|\psi_0\rangle$.

To turn on the system-ancilla interactions, the system atoms are individually addressed by control fields that off-resonantly couple state $|1\rangle$ of the i th atom to the Rydberg state $|R\rangle$ with Rabi frequency Ω_i and detuning $|\Delta_s| \gg \Omega_i$. In this regime, the lowest-order effect of the light on the atomic states is an ac Stark shift given by $\Omega_i^2/(4\Delta_s)$. Thus we can write the interaction Hamiltonian as

$$H_R = |R\rangle\langle R| \sum_i J_i |1\rangle\langle 1|_i, \quad (\text{H1})$$

where

$$J_i = \frac{\Omega_i^2}{4} \left[\frac{1}{\Delta_s - V_R(\mathbf{r}_i)} - \frac{1}{\Delta_s} \right] \quad (\text{H2})$$

and $V_R(\mathbf{r}_i)$ is the Rydberg pair potential between the i th system atom and the ancilla. We choose V_R and Δ_s to have opposite signs. If the ancilla is in the Rydberg

state, the interaction energy V_R then increases the detuning $|\Delta_s - V_R|$, thereby suppressing the ac Stark shift of atom i by an amount J_i . The result [Eq. (H1)] is equivalent to the central spin model in Eq. (4) up to overall energy shifts, with weights $w_i = J_i/J_{\max}$, where J_{\max} is the largest of the system-ancilla couplings J_i .

The oracle is implemented by simultaneously turning on the couplings J_i and attempting to drive a 2π pulse on the $|g\rangle \rightarrow |R\rangle$ transition of the ancilla. The ancilla is driven with a field of Rabi frequency $\Omega_a(t)$, with the pulse shape chosen to ensure that the qubit ends up in its ground state irrespective of whether the pulse is resonant. This condition is satisfied for a pulse shape [47]

$$\Omega_a(t) = \frac{2\pi}{\tau} \operatorname{sech}\left(\frac{\pi t}{\tau}\right), \quad (\text{H3})$$

where τ sets the width of the oracle phase step. In practice, we must restrict the pulse to a finite window $-t_p/2 < t < t_p/2$, where a duration $t_p \gtrsim 3\tau$ suffices to provide a smooth turn on. The detuning Δ_a of the ancilla's control field sets the target weight $W_* = \Delta_a/J_{\max}$ in the subset sum problem [Eq. (A1)]: for configurations of the system spins with weight $W_1 = W_*$ in state $|1\rangle$, the ancilla undergoes a 2π rotation that imparts a geometric phase of π .

More generally, this protocol produces a unitary transformation

$$U_R = \mathcal{T} e^{-i \int_{-t_p/2}^{t_p/2} H(t) dt}, \quad (\text{H4})$$

where we set $\hbar = 1$, \mathcal{T} denotes time ordering, and

$$H(t) = H_R + \Omega_a(t)I_x, \quad (\text{H5})$$

where $I_x = (I_+ + I_-)/2$. For the hyperbolic secant pulse in Eq. (H3), we obtain a W_1 -dependent phase shift $U_R = e^{i\Phi_\gamma}$ where

$$\Phi_\gamma = 2 \arctan [2(W_* - W_1)/\gamma] + \pi, \quad (\text{H6})$$

and the width of the phase step is given by $\gamma = 2\pi/(J_{\max}\tau)$ [90].

Two effects that can limit the performance of the Rydberg implementation are the finite lifetime $1/\Gamma_R$ of the Rydberg state and residual interactions among the system spins. The residual interactions between the system spins are smaller than the system-ancilla couplings by a factor of order $(\Omega_i/\Delta_s)^2$ assuming $|V_R(\mathbf{r}_i)| \gtrsim |\Delta_s|$. If necessary, these interactions can furthermore be cancelled by an echo procedure in which the control fields Ω_i are applied again with the signs of Δ_s and V_R reversed, the latter by tuning the electric field near a Förster resonance [91]. We therefore neglect residual interactions in our analysis and focus on the limits set by Rydberg decay.

To estimate the requirements for implementing Grover's algorithm while keeping the probability of Rydberg decay small, we define the maximum Ω_{\max} of the Rabi frequencies Ω_i and the dressing amplitude $\epsilon = \Omega_{\max}/(2|\Delta_s|)$. Our perturbative analysis of the dressing assumes that $\epsilon^2 < 1/n$, where n is the number of system spins. Let us furthermore assume that the most strongly weighted atom is sufficiently close to the ancilla that $|V_R| \gtrsim |\Delta_s|$, such that its coupling is

$$J_{\max} \approx \Omega_{\max}^2/(4\Delta_s) = \epsilon\Omega_{\max}/2. \quad (\text{H7})$$

During the oracle pulse, the probability of decay for a system atom due to the coupling to the Rydberg state will be $t_p\epsilon^2\Gamma_R$. The worst-case decay probability of the system spins when each spin is in state $|1\rangle$ is $3\pi n\epsilon^2/\rho\gamma$, based on the pulse time $t_p = 3\pi/\gamma J_{\max}$. In addition, the error rate due to ancilla decay during the generalized oracle is approximately $\Gamma_R/J_{\max}\gamma$. In the weak dressing limit $n\epsilon^2 \ll 1$, the decay due to the ancilla dominates over the decay of the dressed system spins.

We now present concrete experimental parameters for implementing the central spin model with cesium atoms. Coupling to high-lying Rydberg states is beneficial as the lifetime scales as the cube of the principal quantum number. By coupling to the $|80P_{3/2}\rangle$ state, we can achieve $\Omega_{\max} \approx 2\pi \times 10$ MHz with realistic laser parameters [42,92]. The Rydberg interaction strength is given by $V_R(r) = -C_6/r^6$, where $C_6 \approx 2\pi \times 7000$ GHz μm^6 for $|80P_{3/2}\rangle$ [93]. For a typical distance between neighboring atoms in an optical tweezer array $r_0 \approx 4$ μm , the interaction shift will be $V_R(r_0) \approx 2\pi \times 1.7$ GHz.

The achievable interaction strength in the Rydberg implementation will depend on system size n , as the weak dressing condition $\epsilon^2 < 1/n$ puts an upper limit on $J_{\max} < \Omega_{\max}/(2\sqrt{n})$. To give a particular example, for a system size $n = 6$, with $n\epsilon^2 = 0.1$ and $\Omega_{\max} = 2\pi \times 10$ MHz, the interaction strength is $J_{\max} \approx 650$ kHz for $\Delta_s \approx 2\pi \times 39$ MHz. The interaction shift $|V_R(r_0)| > |\Delta_s|$ is large enough to extinguish the light shift of the most strongly coupled atom as we assumed in the preceding analysis. For the state $|80P_{3/2}\rangle$ in cesium, $\Gamma_R \approx 2\pi \times 0.5$ kHz, giving the interaction-to-decay ratio $\rho \approx 1200$.

2. Central boson model with atoms in a cavity

As a central boson system for encoding subset sum problems, we consider n spins that are coupled to a cavity of linewidth κ . We require a dispersive atom-light interaction described by a Hamiltonian

$$H = c^\dagger c \sum_i J_i |1\rangle\langle 1|_i. \quad (\text{H8})$$

Here, $J_i = g_i^2/\Delta_i$ is the shift of the cavity resonance when the i th spin is flipped, in terms of the vacuum Rabi frequency g_i and the detuning $\Delta_i \gg \Gamma_e$ of the cavity from resonance with a transition $|1\rangle \rightarrow |e\rangle$ of linewidth Γ_e [Fig. 13(b)].

To implement the oracle, the cavity is driven by a weak, narrow-band coherent field $|\alpha\rangle$ of frequency $\omega = \omega_c + \delta$, where ω_c is the resonance frequency of the bare cavity. The output and input modes

$$b_{\text{out}} = \chi b_{\text{in}} \quad (\text{H9})$$

are related by the cavity response function [94]

$$\chi = -\frac{\kappa/2 + i\delta'}{\kappa/2 - i\delta'}, \quad (\text{H10})$$

where $\delta' = \delta - J_{\text{max}} W_1$, assuming that cavity losses are negligible compared with transmission. The weighted sum W_1 is defined as in Appendix A using weights determined by the couplings of each spin to the cavity, $w_i = J_i/J_{\text{max}}$. The choice of detuning of the drive field from bare cavity resonance δ sets the target weight $W_* = \delta/J_{\text{max}}$ for the subset sum problem. This can be tuned to specifically implement the partition problem (see Appendix A).

More generally, we can also account for a photon loss rate Γ_a , including any absorption by the atoms, by letting

$$\delta' = \delta - J_{\text{max}} W_1 + i\Gamma_a/2. \quad (\text{H11})$$

The magnitude and phase of the cavity response function χ determine, respectively, the probability $|\chi|^2$ of successfully detecting the ancilla photon and the resulting oracle phase shift. On resonance, the magnitude of the response function is

$$|\chi(0)| = \frac{\kappa - \Gamma_a}{\kappa + \Gamma_a}, \quad (\text{H12})$$

which yields a detection probability $|\chi|^2 \approx 1 - 4\Gamma_a/\kappa$ for small Γ_a/κ . The phase shift is given by

$$\Phi(W_1) \equiv \arg[\chi] = 2 \arctan(2\delta'/\kappa) + \pi. \quad (\text{H13})$$

The phase Φ increases from 0 to 2π in a step of characteristic width κ , assuming low losses $\Gamma_a \lesssim \kappa/2$, as a function of the atom-dependent detuning between the drive and cavity resonance. We parameterize the step width by the dimensionless value $\gamma = \kappa/J_{\text{max}}$.

To apply the oracle U_γ , we initialize the system in a product state of the atoms, the vacuum field in the cavity, and a weak, narrow-band coherent state in the input mode:

$$|\Psi\rangle = |\psi_0\rangle|0_c\rangle|\alpha_{b_{\text{in}}}\rangle. \quad (\text{H14})$$

The coherent field leaks through the input mirror into the cavity mode, where the light and atoms interact according

to Eq. (H8), then leaks into the output mode b_{out} . After a time $t \gg 1/(\Delta\omega) \gg 1/\kappa$, where $\Delta\omega$ is the bandwidth of the input field, the state evolves to

$$|\Psi_t\rangle = e^{i\alpha\chi b_{\text{out}}^\dagger} |\psi_0\rangle|0_c\rangle|0_{b_{\text{out}}}\rangle. \quad (\text{H15})$$

The action of $e^{i\alpha\chi b_{\text{out}}^\dagger}$ displaces the vacuum state of the output mode $|0_{b_{\text{out}}}\rangle$ such that the detection of a single photon in the output mode heralds the state

$$\langle 1_{b_{\text{out}}} | \Psi_t \rangle = e^{i\Phi(W_1)} |\psi_0\rangle|0_c\rangle, \quad (\text{H16})$$

thus applying the oracle.

As an alternative to the coherent drive and heralding, an ancilla atom can be used as an intracavity single-photon source. By coupling the ancilla to the cavity via a two-photon transition, with the first leg being a classical field, the cavity can be controllably excited from the vacuum to the single-photon state. The bosonic mode is thus reduced to two levels $|0\rangle_c, |1\rangle_c$ that are coupled by the control field on the ancilla, so that we effectively recover a central spin model. The implementation of the oracle then proceeds much as in Appendix 1, by driving a shaped 2π pulse that returns the ancilla atom to its initial state and the cavity to the vacuum state. The width τ of this pulse now controls the step width $\gamma = 2\pi/(J_{\text{max}}\tau)$, subject to the requirement that the pulse be short compared to the cavity lifetime.

We now proceed to estimate the cavity parameters required to observe Grover amplification [as in Eq. (E7)], as well as the attainable interaction-to-decay ratio. Amplifying the probability of solution states requires a phase step narrower than the initial probability distribution $P(W_1)$, which in turn requires strong atom-light coupling. In particular, we will show that the single-atom cooperativity $\eta = 4g^2/(\kappa\Gamma_e)$ sets an upper bound on the dispersive cavity shift J_{max} achievable at low photon loss rate $\Gamma_a < \kappa$, and hence a lower bound on the dimensionless step width $\gamma = \kappa/J_{\text{max}}$ in the driven cavity.

The lower bound on the step width γ arises because increasing the dispersive coupling J_{max} comes at the cost of increased chance of atomic absorption. In the worst-case scenario where all n atoms are in state $|1\rangle$ in the scheme of Fig. 13(b), atomic absorption produces a photon loss rate

$$\Gamma_a = \Gamma_e \sum_{i=1}^n \frac{g_i^2}{\Delta_i^2} = \Gamma_e J_{\text{max}}^2 \sum_{i=1}^n \frac{w_i^2}{g_i^2} \quad (\text{H17})$$

in terms of the weights w_i . While each weight can be tuned via either the atom-cavity coupling g_i or the detuning Δ_i , the latter is preferable because it allows all atoms to benefit from the maximum cavity cooperativity. Thus we set $g_i \equiv$

g to be maximal for all atoms, reducing Eq. (H17) to

$$\frac{\Gamma_a}{\kappa} = \frac{\Gamma_e \kappa}{\gamma^2 g^2} \sum_{i=1}^n w_i^2 = \frac{4n w_{\text{rms}}^2}{\eta \gamma^2}, \quad (\text{H18})$$

where w_{rms}^2 represents the mean-squared value of weights and is given by $w_{\text{rms}}^2 = 1/3$ for weights drawn from a uniform distribution $w_i \in (0, 1]$. Thus, keeping photon loss small ($\Gamma_a/\kappa \lesssim 1$) requires a step width $\gamma \gtrsim \sqrt{n/\eta}$.

Equation (H18) gives the decay parameter $r = \Gamma_a/\kappa$ necessary to determine the single-cycle amplification in Eq. (G13). Notably, we can re-express the decay parameter in terms of the variance $\sigma^2 = n w_{\text{rms}}^2/\gamma^2$ of the normalized weighted spin $\mu = 2(W_* - W_1)/\gamma$ and the cooperativity:

$$r = \frac{4\sigma^2}{\eta}. \quad (\text{H19})$$

Achieving amplification requires $\sigma^2 > 1$, i.e., the probability distribution of W_1 should be broader than the width γ of the phase step. To achieve this condition at low loss $r < 1$, we require strong coupling $\eta \gg 1$. This requirement is corroborated by plots of the amplification versus step width for various cooperativities in Fig. 5. The maximum achievable single-cycle amplification, shown in Fig. 5(c), becomes larger than 1 for $\eta \gtrsim 50$. This condition can be satisfied in state-of-the-art optical cavities, where the highest cooperativities achieved are $\eta \sim 10^2$ [56,95], at scalable system size n .

Achieving substantial quantum speedups requires operating in the ultrastrong coupling regime $\eta \gg n$ to reach step widths $\gamma \ll 1$. A cooperativity as high as $\eta = 4 \times 10^8$ has been achieved by coupling circular Rydberg atoms to a superconducting millimeter-wave cavity [59], with $(g, \kappa, \Gamma) = 2\pi \times (2.5 \times 10^4, 1.4, 4.4)$ Hz. To access this high cooperativity, both spin states $|0\rangle, |1\rangle$ must be Rydberg states with finite lifetime Γ^{-1} , and the dominant decay channel is then atomic decay rather than photon loss, resulting in an interaction-to-decay ratio $\rho \approx J_{\text{max}}/(n\Gamma)$. The detunings Δ_i should be set to maximize the couplings, up to $J_{\text{max}} = \epsilon g$, where $\epsilon \equiv g/\min(\Delta_i)$ is limited by the requirement $n\epsilon^2 < 1$ to avoid absorption of the photon. Fixing $n\epsilon^2 = 0.1$ allows an interaction-to-decay ratio $\rho \approx 2 \times 10^3/n^{3/2}$ for the parameters of Ref. [59].

-
- [1] L. K. Grover, in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, 1996), p. 212.
- [2] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, *Phys. Rev. Lett.* **79**, 325 (1997).
- [3] R. M. Karp, in *Complexity of Computer Computations* (Springer, Boston, MA, 1972), p. 85.
- [4] S. Mertens, Phase Transition in the Number Partitioning Problem, *Phys. Rev. Lett.* **81**, 4281 (1998).

- [5] D. J. Bernstein, S. Jeffery, T. Lange, and A. Meurer, in *International Workshop on Post-Quantum Cryptography* (Springer, Berlin, Heidelberg, 2013), p. 16.
- [6] R. Merkle and M. Hellman, Hiding information and signatures in trapdoor knapsacks, *IEEE Trans. Inf. Theory* **24**, 525 (1978).
- [7] V. Lyubashevsky, A. Palacio, and G. Segev, in *Theory of Cryptography Conference* (Springer, Berlin, Heidelberg, 2010), p. 382.
- [8] H. M. Weingartner and D. N. Ness, Methods for the solution of the multidimensional 0/1 knapsack problem, *Oper. Res.* **15**, 83 (1967).
- [9] M. Gilli, D. Maringer, and E. Schumann, *Numerical Methods and Optimization in Finance* (Academic Press, Cambridge, MA, 2019).
- [10] I. L. Chuang, N. Gershenfeld, and M. Kubinec, Experimental Implementation of Fast Quantum Searching, *Phys. Rev. Lett.* **80**, 3408 (1998).
- [11] J. A. Jones, M. Mosca, and R. H. Hansen, Implementation of a quantum search algorithm on a quantum computer, *Nature* **393**, 344 (1998).
- [12] L. M. Vandersypen, M. Steffen, M. H. Sherwood, C. S. Yannoni, G. Breyta, and I. L. Chuang, Implementation of a three-quantum-bit search algorithm, *Appl. Phys. Lett.* **76**, 646 (2000).
- [13] P. Kwiat, J. Mitchell, P. Schwindt, and A. White, Grover's search algorithm: An optical approach, *J. Mod. Opt.* **47**, 257 (2000).
- [14] J. Ahn, T. Weinacht, and P. Bucksbaum, Information storage and retrieval through quantum phase, *Science* **287**, 463 (2000).
- [15] M. Anwar, D. Blazina, H. Carteret, S. Duckett, and J. Jones, Implementing Grover's quantum search on a para-hydrogen based pure state NMR quantum computer, *Chem. Phys. Lett.* **400**, 94 (2004).
- [16] K.-A. Brickman, P. C. Haljan, P. J. Lee, M. Acton, L. Deslauriers, and C. Monroe, Implementation of Grover's quantum search algorithm in a scalable system, *Phys. Rev. A* **72**, 050306(R) (2005).
- [17] P. Walther, K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer, and A. Zeilinger, Experimental one-way quantum computing, *Nature* **434**, 169 (2005).
- [18] R. Prevedel, P. Walther, F. Tiefenbacher, P. Böhi, R. Kaltenbaek, T. Jennewein, and A. Zeilinger, High-speed linear optics quantum computing using active feed-forward, *Nature* **445**, 65 (2007).
- [19] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, Demonstration of blind quantum computing, *Science* **335**, 303 (2012).
- [20] C. Figgatt, D. Maslov, K. Landsman, N. M. Linke, S. Debnath, and C. Monroe, Complete 3-qubit Grover search on a programmable quantum computer, *Nat. Commun.* **8**, 1918 (2017).
- [21] C. Godfrin, A. Ferhat, R. Ballou, S. Klyatskaya, M. Ruben, W. Wernsdorfer, and F. Balestro, Operating Quantum States in Single Magnetic Molecules: Implementation of Grover's Quantum Algorithm, *Phys. Rev. Lett.* **119**, 187702 (2017).
- [22] Y. Wu, Y. Wang, X. Qin, X. Rong, and J. Du, A programmable two-qubit solid-state quantum processor under ambient conditions, *npj Quantum Inf.* **5**, 1 (2019).

- [23] M. Roget, S. Guillet, P. Arrighi, and G. Di Molfetta, Grover Search as a Naturally Occurring Phenomenon, *Phys. Rev. Lett.* **124**, 180501 (2020).
- [24] Z. Jiang, E. G. Rieffel, and Z. Wang, Near-optimal quantum circuit for Grover’s unstructured search using a transverse field, *Phys. Rev. A* **95**, 062317 (2017).
- [25] C. Moore and S. Mertens, *The Nature of Computation* (OUP, Oxford, 2011).
- [26] L. Jiang, G. K. Brennen, A. V. Gorshkov, K. Hammerer, M. Hafezi, E. Demler, M. D. Lukin, and P. Zoller, Anyonic interferometry and protected memories in atomic spin lattices, *Nat. Phys.* **4**, 482 (2008).
- [27] W. Chen, J. Hu, Y. Duan, B. Braverman, H. Zhang, and V. Vuletić, Carving Complex Many-Atom Entangled States by Single-Photon Detection, *Phys. Rev. Lett.* **115**, 250502 (2015).
- [28] E. J. Davis, Z. Wang, A. H. Safavi-Naeini, and M. H. Schleier-Smith, Painting Nonclassical States of Spin or Motion with Shaped Single Photons, *Phys. Rev. Lett.* **121**, 123602 (2018).
- [29] S. Gleyzes, S. Kuhr, C. Guerlin, J. Bernu, S. Deleglise, U. B. Hoff, M. Brune, J.-M. Raimond, and S. Haroche, Quantum jumps of light recording the birth and death of a photon in a cavity, *Nature* **446**, 297 (2007).
- [30] S. Welte, B. Hacker, S. Daiss, S. Ritter, and G. Rempe, Photon-Mediated Quantum Gate between two Neutral Atoms in an Optical Cavity, *Phys. Rev. X* **8**, 011018 (2018).
- [31] R. McConnell, H. Zhang, J. Hu, S. Ćuk, and V. Vuletić, Entanglement with negative wigner function of almost 3000 atoms heralded by one photon, *Nature* **519**, 439 (2015).
- [32] G. Barontini, L. Hohmann, F. Haas, J. Estève, and J. Reichel, Deterministic generation of multiparticle entanglement by quantum zeno dynamics, *Science* **349**, 1317 (2015).
- [33] E. J. Davis, A. Periwal, E. S. Cooper, G. Bentsen, S. J. Evered, K. Van Kirk, and M. H. Schleier-Smith, Protecting Spin Coherence in a Tunable Heisenberg Model, *Phys. Rev. Lett.* **125**, 060402 (2020).
- [34] M. Saffman and K. Mølmer, Efficient Multiparticle Entanglement via Asymmetric Rydberg Blockade, *Phys. Rev. Lett.* **102**, 240502 (2009).
- [35] K. Mølmer, L. Isenhower, and M. Saffman, Efficient Grover search with Rydberg blockade, *J. Phys. B: At. Mol. Opt. Phys.* **44**, 184016 (2011).
- [36] X. L. Zhang, L. Isenhower, A. T. Gill, T. G. Walker, and M. Saffman, Deterministic entanglement of two neutral atoms via Rydberg blockade, *Phys. Rev. A* **82**, 030306(R) (2010).
- [37] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, Entanglement of two Individual Neutral Atoms Using Rydberg Blockade, *Phys. Rev. Lett.* **104**, 010502 (2010).
- [38] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, Entangling atomic spins with a Rydberg-dressed spin-flip blockade, *Nat. Phys.* **12**, 71 (2016).
- [39] J. Zeiher, R. Van Bijnen, P. Schau, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross, Many-body interferometry of a Rydberg-dressed spin lattice, *Nat. Phys.* **12**, 1095 (2016).
- [40] J. Zeiher, J.-Y. Choi, A. Rubio-Abadal, T. Pohl, R. van Bijnen, I. Bloch, and C. Gross, Coherent Many-Body Spin Dynamics in a Long-Range Interacting Ising Chain, *Phys. Rev. X* **7**, 041063 (2017).
- [41] C. J. Picken, R. Legaie, K. McDonnell, and J. D. Pritchard, Entanglement of neutral-atom qubits with long ground-Rydberg coherence times, *Quantum Sci. Technol.* **4**, 015011 (2018).
- [42] V. Borish, O. Marković, J. A. Hines, S. V. Rajagopal, and M. Schleier-Smith, Transverse-Field Ising Dynamics in a Rydberg-Dressed Atomic gas, *Phys. Rev. Lett.* **124**, 063601 (2020).
- [43] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, High-fidelity entanglement and detection of alkaline-earth Rydberg atoms, *Nat. Phys.* **16**, 857 (2020).
- [44] 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).
- [45] J. T. Young, P. Bienias, R. Belyansky, A. M. Kaufman, and A. V. Gorshkov, Asymmetric blockade and multi-qubit gates via dipole-dipole interactions, arXiv:2006.02486 (2020).
- [46] Y. Ashida, T. Shi, R. Schmidt, H. R. Sadeghpour, J. I. Cirac, and E. Demler, Quantum Rydberg Central Spin Model, *Phys. Rev. Lett.* **123**, 183001 (2019).
- [47] N. Rosen and C. Zener, Double Stern-Gerlach experiment and related collision phenomena, *Phys. Rev.* **40**, 502 (1932).
- [48] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, United Kingdom, 2010).
- [49] C. Borgs, J. Chayes, and B. Pittel, Phase transition and finite-size scaling for the integer partitioning problem, *Random Struct. Algorithms* **19**, 247 (2001).
- [50] M. Mézard and A. Montanari, in *Information, Physics, and Computation* (Oxford University Press, Oxford, 2009), p. 141.
- [51] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Defining and detecting quantum speedup, *Science* **345**, 420 (2014).
- [52] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, Strengths and weaknesses of quantum computing, *SIAM J. Comput.* **26**, 1510 (1997).
- [53] M. Boyer, G. Brassard, P. Hyer, and A. Tapp, Tight bounds on quantum searching, *Fortschr. Phys.* **46**, 493 (1998).
- [54] C. Zalka, Grover’s quantum searching algorithm is optimal, *Phys. Rev. A* **60**, 2746 (1999).
- [55] R. Impagliazzo and M. Naor, Efficient cryptographic schemes provably as secure as subset sum, *J. Cryptol.* **9**, 199 (1996).
- [56] Y. Colombe, T. Steinmetz, G. Dubois, F. Linke, D. Hunger, and J. Reichel, Strong atom–field coupling for Bose–Einstein condensates in an optical cavity on a chip, *Nature* **450**, 272 (2007).

- [57] J. Zeiher, P. Schauß, S. Hild, T. Macrì, I. Bloch, and C. Gross, Microscopic Characterization of Scalable Coherent Rydberg Superatoms, *Phys. Rev. X* **5**, 031015 (2015).
- [58] A. Paris-Mandoki, C. Braun, J. Kumlin, C. Tresp, I. Mirgorodskiy, F. Christaller, H. P. Büchler, and S. Hofferberth, Free-Space Quantum Electrodynamics with a Single Rydberg Superatom, *Phys. Rev. X* **7**, 041010 (2017).
- [59] S. Haroche and J.-M. Raimond, *Exploring the Quantum: Atoms, Cavities, and Photons* (Oxford University Press, Cambridge, United Kingdom, 2006).
- [60] A. Suleymanzade, A. Anferov, M. Stone, R. K. Naik, A. Oriani, J. Simon, and D. Schuster, A tunable high- q millimeter wave cavity for hybrid circuit and cavity QED experiments, *Appl. Phys. Lett.* **116**, 104001 (2020).
- [61] R. G. Hulet, E. S. Hilfer, and D. Kleppner, Inhibited Spontaneous Emission by a Rydberg Atom, *Phys. Rev. Lett.* **55**, 2137 (1985).
- [62] T. L. Nguyen, J. M. Raimond, C. Sayrin, R. Corti nas, T. Cantat-Moltrecht, F. Assemat, I. Dotsenko, S. Gleyzes, S. Haroche, G. Roux, T. Jolicoeur, and M. Brune, Towards Quantum Simulation With Circular Rydberg Atoms, *Phys. Rev. X* **8**, 011032 (2018).
- [63] E. Horowitz and S. Sahni, Computing partitions with applications to the knapsack problem, *J. ACM* **21**, 277 (1974).
- [64] R. Schroepfel and A. Shamir, A $T = O(2^{n/2})$, $S = O(2^{n/4})$ algorithm for certain NP-complete problems, *SIAM J. Comput.* **10**, 456 (1981).
- [65] N. Howgrave-Graham and A. Joux, New generic algorithms for hard knapsacks, Cryptology ePrint Archive, Report 2010/189 (2010).
- [66] A. Becker, J.-S. Coron, and A. Joux, Improved generic algorithms for hard knapsacks, Cryptology ePrint Archive, Report 2011/474 (2011).
- [67] I. Dinur, O. Dunkelmann, N. Keller, and A. Shamir, in *Proceedings of the 32nd Annual Cryptology Conference on Advances in Cryptology – CRYPTO 2012 – 7417* (Springer-Verlag, Berlin, Heidelberg, 2012), p. 719.
- [68] P. Austrin, P. Kaski, M. Koivisto, and J. Määtä, in *Automata, Languages, and Programming*, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2013), p. 45.
- [69] A. Esser and A. May, Low weight discrete logarithms and subset sum in $2^{0.65n}$ with polynomial memory, Cryptology ePrint Archive, Report 2019/931 (2019), <https://eprint.iacr.org/2019/931>.
- [70] R. Korf, A complete anytime algorithm for number partitioning, *Artif. Intell.* **106**, 181 (1998).
- [71] A. Helm and A. May, in *13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018)*, edited by S. Jeffery Leibniz, International Proceedings in Informatics (LIPIcs) Vol. 111 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018), p. 5:1.
- [72] Y. Li and H. Li, Improved quantum algorithm for the random subset sum problem, arXiv:1912.09264 (2019).
- [73] A. Helm and A. May, in *International Conference on Post-Quantum Cryptography* (Springer, Cham, Switzerland, 2020), p. 445.
- [74] X. Bonnetain, R. Bricout, A. Schrottenloher, and Y. Shen, Improved classical and quantum algorithms for subset-sum, arXiv:2002.05276 (2020).
- [75] D. S. Wild, D. Sels, H. Pichler, and M. D. Lukin, Quantum sampling algorithms for near-term devices, arXiv:2005.14059 (2020).
- [76] 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 (2018).
- [77] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, *Phys. Rev. X* **10**, 021067 (2020).
- [78] M. Pechal, P. Arrangoiz-Arriola, and A. H. Safavi-Naeini, Superconducting circuit quantum computing with nanomechanical resonators as storage, *Quantum Sci. Technol.* **4**, 015006 (2018).
- [79] C. T. Hann, C.-L. Zou, Y. Zhang, Y. Chu, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, Hardware-Efficient Quantum Random Access Memory with Hybrid Quantum Acoustic Systems, *Phys. Rev. Lett.* **123**, 250501 (2019).
- [80] R. Naik, N. Leung, S. Chakram, P. Groszkowski, Y. Lu, N. Earnest, D. McKay, J. Koch, and D. Schuster, Random access quantum information processors using multimode circuit quantum electrodynamics, *Nat. Commun.* **8**, 1 (2017).
- [81] M. E. S. Morales, T. Tlyachev, and J. Biamonte, Variational learning of Grover's quantum search algorithm, *Phys. Rev. A* **98**, 062333 (2018).
- [82] T. Keating, C. H. Baldwin, Y.-Y. Jau, J. Lee, G. W. Biedermann, and I. H. Deutsch, Arbitrary Dicke-State Control of Symmetric Rydberg Ensembles, *Phys. Rev. Lett.* **117**, 213601 (2016).
- [83] E. F. Brickell, in *Advances in Cryptology: Proceedings of Crypto 83*, edited by D. Chaum (Springer, Boston, MA, USA, 1984), p. 25.
- [84] J. C. Lagarias and A. M. Odlyzko, Solving low-density subset sum problems, *J. ACM* **32**, 229 (1985).
- [85] M. J. Coster, A. Joux, B. A. LaMacchia, A. M. Odlyzko, C.-P. Schnorr, and J. Stern, Improved low-density subset sum algorithms, *Comput. Complex.* **2**, 111 (1992).
- [86] C. P. Schnorr and M. Euchner, Lattice basis reduction: Improved practical algorithms and solving subset sum problems, *Math. Program.* **66**, 181 (1994).
- [87] S. Mertens, Random Costs in Combinatorial Optimization, *Phys. Rev. Lett.* **84**, 1347 (2000).
- [88] S. Mertens, in *Computational Complexity and Statistical Physics*, edited by A. Percus, G. Istrate, and C. Moore (Oxford University Press, New York, 2006), Chap. 5, p. 125.
- [89] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, in *Atom-Photon Interactions: Basic Processes and Applications*, Wiley-Interscience publication (J. Wiley, Hoboken, NJ, 1992), p. 201.
- [90] R. T. Robiscoe, Extension of the Rosen–Zener solution to the two-level problem, *Phys. Rev. A* **17**, 247 (1978).
- [91] T. Vogt, M. Viteau, J. Zhao, A. Chotia, D. Comparat, and P. Pillet, Dipole Blockade at Förster Resonances in High

- Resolution Laser Excitation of Rydberg States of Cesium Atoms, *Phys. Rev. Lett.* **97**, 083003 (2006).
- [92] A. M. Hankin, Y.-Y. Jau, L. P. Parazzoli, C. W. Chou, D. J. Armstrong, A. J. Landahl, and G. W. Biedermann, Two-atom Rydberg blockade using direct $6S$ to nP excitation, *Phys. Rev. A* **89**, 033416 (2014).
- [93] N. Sibalić, J. D. Pritchard, K. J. Weatherill, and C. S. Adams, Arc: An open-source library for calculating properties of alkali Rydberg atoms, *Comput. Phys. Commun.* **220**, 319 (2017).
- [94] C. W. Gardiner and M. J. Collett, Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation, *Phys. Rev. A* **31**, 3761 (1985).
- [95] M. Wolke, J. Klinner, H. Keßler, and A. Hemmerich, Cavity cooling below the recoil limit, *Science* **337**, 75 (2012).