Grover's Algorithm Offers No Quantum Advantage

E.M. Stoudenmire[1] and Xavier Waintal[2]

[1]*Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA*
[2]*PHELIQS, Université Grenoble Alpes, CEA, Grenoble INP, IRIG, Grenoble 38000, France*
(Dated: March 21, 2023)

Grover's algorithm is one of the primary algorithms offered as evidence that quantum computers can provide an advantage over classical computers. It involves an "oracle" (external quantum subroutine) which must be specified for a given application and whose internal structure is not part of the formal scaling of the quantum speedup guaranteed by the algorithm. Grover's algorithm also requires exponentially many steps to succeed, raising the question of its implementation on near-term, non-error-corrected hardware and indeed even on error-corrected quantum computers. In this work, we construct a quantum inspired algorithm, executable on a classical computer, that performs Grovers' task in a *linear* number of call to the oracle — an exponentially smaller number than Grover's algorithm — and demonstrate this algorithm explicitly for boolean satisfiability problems (3-SAT). Our finding implies that there is no *a priori* theoretical quantum speed-up associated with Grover's algorithm. We critically examine the possibility of a practical speed-up, a possibility that depends on the nature of the quantum circuit associated with the oracle. We argue that the unfavorable scaling of the success probability of Grover's algorithm, which in the presence of noise decays as the exponential of the exponential of the number of qubits, makes a practical speedup unrealistic even under extremely optimistic assumptions on both hardware quality and availability.

I. INTRODUCTION

Two classes of algorithms dominate the landscape of possible applications for quantum computing. The first class computes a non-trivial result then extracts this result using the quantum Fourier transform. This class includes the seminal Shor's algorithm for integer factorization [1–3] as well as the quantum phase estimation algorithm [3] proposed for solving quantum chemistry problems and several other algorithms [4]. Some of these algorithms, in particular Shor's, offer an exponential speedup over any known classical methods, though only for a handful of rather specific applications.

The second class includes Grover's algorithm (GA) and its generalizations, such as amplitude amplification [5–7]. Grover's algorithm promises a less spectacular quadratic speedup, but in return enjoys wide popularity due to its many possible use cases. Theoretically, quite a large number of problems could be accelerated by simply replacing the critical part of the classical algorithm by a call to a Grover's routine implemented on a quantum computer. It is also appealing that the quadratic speedup of Grover's algorithm can be put on firm mathematical grounds and is provably optimal under certain assumptions [8], in contrast to Shor's where the speedup is only conjectured. This quadratic speedup is very convenient theoretically as it does not require any knowledge of the oracle which encodes the problem into the quantum algorithm. The class of problems for which Grover's algorithm can be applied include instances of NP-complete problems which are extremely challenging computationally. Applications where Grover's algorithm is the main subroutine range from optimizing functions [9] for e.g. analyzing high energy physics data [10] to solving various graph problems [11] to option pricing [12] to pattern recognition (finding a string in a text) [13] to various

form of machine learning (including supervised learning [14], perceptrons [15], active learning agents [15] and reinforcement learning [16]). Specific versions have been implemented on small quantum processors with up to $n \leq 5$ qubits [17–20] but the success probability is still low for the largest systems.

Grover's algorithm solves the problem of inverting an unknown function. Given a function $y = f(b)$ where b can take $N = 2^n$ values, Grover's algorithm finds the value $b = f^{-1}(y)$ for a given y in only $\sqrt{N}$ calls to the quantum "oracle" function which implements $f(b)$ while a naive exhaustive search would require about N calls to the corresponding classical function. Grover's theoretical speedup is of a peculiar nature: it is an abstract speedup that considers the oracle as a black box function and counts the computational cost solely in terms of the number of calls to the oracle. In any actual implementation, however, the oracle must be realized as a specific quantum circuit, hence the internal structure of $f(b)$ must be exposed. In particular the computational cost of one call to the $f(b)$ oracle may (almost always will) depend on N in some way.

In this work, we argue that there is no generic theoretical speed-up in Grover's algorithm. To this end, we construct a quantum inspired Grover algorithm (QiGA) that solves Grover's problem on a classical computer. In contrast to the usual assumptions made to discuss quantum advantage in this context, the input of QiGA is not the classical oracle $f(b)$ but the quantum oracle, which is the same quantum circuit that would be given to the quantum computer. The comparison with a quantum computer is fairer: both computers are given the same input. We find that QiGA solves the problem in at most $O(\log N)$ calls and in many cases a single call to the *quantum* oracle. In other words, if one is capable of simulating the quantum circuit of the oracle once, then one can

obtain the same result that would require a quantum computer to make exponentially many calls to the oracle. Our findings provide a new comparison point for the discussion of any possible advantage.

In the second half of this work we discuss the implications of our work for the possibility of a practical quantum advantage using Grover. We argue that if the oracle is too hard to simulate classically, so that a quantum computer becomes necessary, then generic, practical problems arise. We show Grover's algorithm to be very sensitive to noise with the success probability decaying as the exponential of an exponential. Beyond ruling out near-term implementations for more than a few qubits, such a rapid accumulation of noise would overwhelm any known quantum error correction protocol.

This is why we say Grover's algorithm offers no quantum advantage. By *quantum advantage* we mean for a specific problem of a fixed size, that a quantum algorithm running on actual hardware would reach the solution faster than the best classical strategy. For a given problem, if the oracle is easy to simulate then a quantum computer is not needed, while if it is hard to simulate the quantum implementation will be, as we shall see, beyond the reach of foreseeable hardware.

This article starts with a summary of the main results (Section II) for readers mostly interested in the implications of our work. The rest of the article is split into two parts. In the first (Sections III and IV), we construct the quantum inspired algorithm that mimics Grover's algorithm in a classical computer but takes advantage of its internal structure to require exponentially fewer calls to the oracle. We explicitly demonstrate and test this algorithm using both random and quasi-one-dimensional instances of NP-complete random boolean satisfiability problems (3-SAT). In the second part of the article, we examine the implication of our findings on the possibility of a quantum speed-up (Sections VI and VII). In particular, Section VI establishes the lack of resilience of Grover's algorithm to noise, an important aspect needed for the discussion of a possible quantum advantage.

II. SUMMARY OF THE MAIN RESULTS

The logic behind our main claim is summarized in Fig. 1. The same "oracle" circuit that defines an implementation of GA on a quantum computer is given to a classical simulator which is limited only by entanglement (a tensor network). If this simulator can calculate a *single* oracle circuit with a complexity better than $2^{n/2}$, our QiGA algorithm solves the GA problem parametrically faster than a quantum computer would. Even when the simulator scales as a less favourable exponential, we explicitly demonstrate cases where simulating the oracle can be done in just hours on a desktop computer for $n \lesssim 40$ qubits and likely on a supercomputer for $n \lesssim 80$ qubits. For larger numbers of qubits there must be some instances of NP-complete problems where simulating the
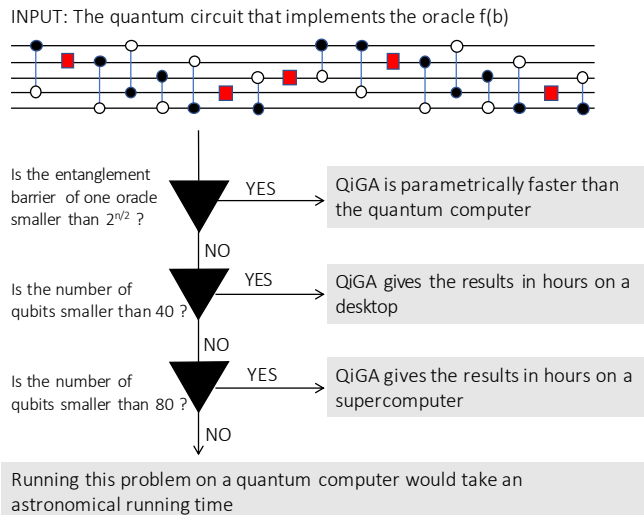

INPUT: The quantum circuit that implements the oracle f(b)

Is the entanglement barrier of one oracle smaller than $2^{n/2}$? — YES → QiGA is parametrically faster than the quantum computer

NO

Is the number of qubits smaller than 40 ? — YES → QiGA gives the results in hours on a desktop

NO

Is the number of qubits smaller than 80 ? — YES → QiGA gives the results in hours on a supercomputer

NO

Running this problem on a quantum computer would take an astronomical running time

FIG. 1. Schematic showing the main implications of our quantum-inspired Grover's algorithm (QiGA) for the possibility of Grover's algorithm offering a genuine quantum advantage.

oracle circuit remains infeasible. Yet for such problems, GA would require at least $2^{80/2} = 2^{40} \approx 10^{12}$ application of the oracle circuit on a quantum computer which translates into an astronomically large time-to-solution even under favourable hardware assumptions.

At the root of this work is the observation that, *in between* the calls to the oracle, the level of entanglement present in the internal state of a quantum computer running GA is extremely low. In between each iteration of GA, the entanglement entropy between any two subgroups of qubits is at most $\log(2)$. In other words, GA strongly leverages the ability of quantum mechanics to create superpositions (quantum parallelism) but it barely uses the possibility to entangle states.

Before describing our "quantum inspired" GA (where we use the resources available in a classical computer in the most efficient way), we start with a "simulation" of GA (where we only use operations implementable on a quantum computer). Figure 2 shows the entanglement entropy of the quantum state during a simulation of Grover's algorithm for $n = 10$ qubits, using the quantum circuit described in Appendix A to implement the oracle and diffusion steps. As claimed, the entanglement in between iterations never exceeds $\log(2)$, a value which is reached when the algorithm is halfway done. The entanglement entropy does become higher during the partial application of the oracle and diffusion operator circuits. The value of this intra-oracle entanglement barrier is strongly problem dependent and will determine the performance of the quantum inspired GA.

As we shall see, the low-entanglement states occurring in GA at integer and half integer iteration number have a simple classical representation in terms of a rank-2 "matrix product state" [21–23] (MPS) which can easily be
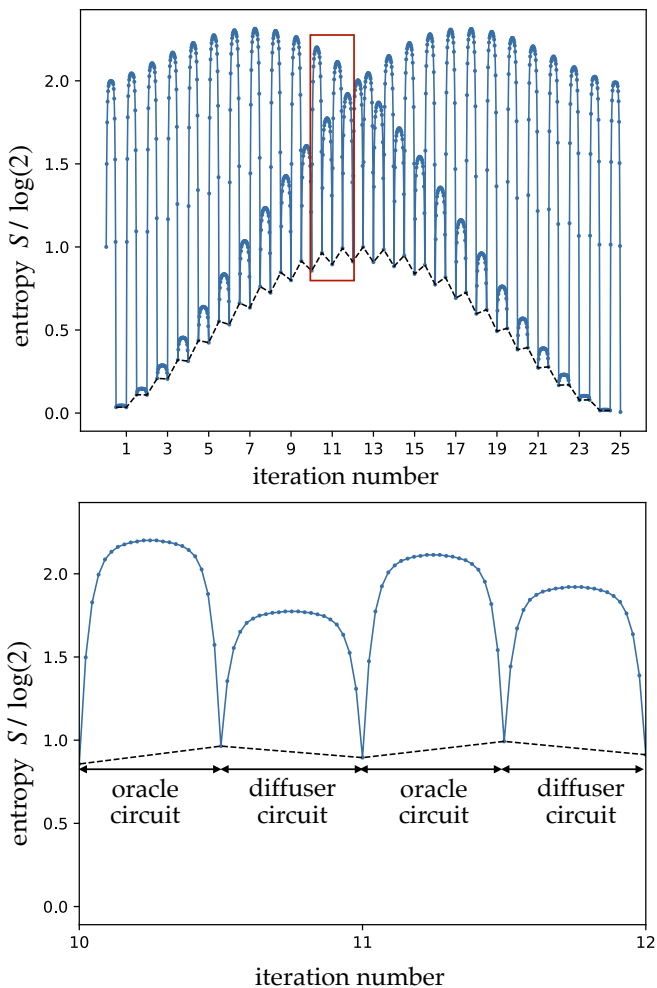
FIG. 2. Entanglement entropy, in units of log(2), of the quantum state during a simulation of Grover's algorithm for $n = 10$ qubits for the optimal number $r = 25$ of iterations. Substeps between each complete iteration show the properties of the state after each layer of the quantum circuits described in Appendix A which implement the oracle and diffusion operators. The entropy is dominated by the oracle step for the first half of the algorithm, then becomes dominated by the diffusion step for the second half of the algorithm. During the simulation the matrix product state (MPS) rank (not shown) goes up and down in a sawtooth shape with maxima of $\chi = 11$ during an iteration and minima of $\chi = 2$ between each iteration. The dashed line along the lower envelope is the theoretical prediction for the entanglement between iterations of the algorithm (the oscillations are a small $1/N$ effect). The lower panel shows a zoom of the region in the box in the upper panel, labeling substeps where the oracle and diffuser circuits act.

kept on a classical computer at a linear cost in the number n of qubits just by storing $2n$ matrices of size 2×2. Similarly, the oracle itself is a rank-2 "matrix product operator" [24, 25] (MPO), another standard object of tensor network theory.

A second, very important observation of this work is

that being able to compute the action of the oracle circuit on the initial product state (the post-oracle state) is always sufficient to complete GA in a *single* call to the oracle. In fact, it is sufficient to even just compute $\sim \log(N)$ *amplitudes* of this state in the $|\pm\rangle$ basis. In this way the quantum inspired algorithm differs from a mere *simulation* of GA that would still require an exponentially large number $\sqrt{N} = 2^{n/2}$ of applications of the oracle circuit. Hence, the initial problem is mapped onto the following question: given a quantum circuit that implements a given oracle, how can one compute its action on an MPS? There exist standard techniques for computing the outputs of quantum circuits as tensor networks [24, 26–28]. We illustrate them in two cases: in the first we take an explicit problem (known as 3-SAT) and show that the explicit knowledge of the oracle's quantum circuit can be exploited to enact the oracle efficiently in many cases. In the second, we consider the generic case of an arbitrary quantum circuit and show that the full post-oracle state can be constructed from $O(n)$ amplitudes of this state using the so-called tensor cross interpolation algorithm.

It remains to discuss the implications of our findings on the possibility of quantum advantage using GA. On a theoretical level our quantum inspired GA requires exponentially fewer number of calls to the oracle. In turn, a classical calculation of the quantum oracle may be exponentially harder than the corresponding task on a quantum computer, depending on the quantum circuit structure and depth. Hence, which of the two approaches (classical or quantum) is faster depends on the actual problem and can not be decided a priori. In particular oracles whose quantum circuits have small depths (and therefore more of chance to be eventually implemented in quantum hardware) can be solved trivially using our quantum inspired approach.

An interesting corollary of the quantum inspired GA algorithms relates to classical complexity theory. It is widely believed that $P \neq NP$, meaning the so-called NP-complete problems are exponentially difficult in the hardest cases. Since NP-complete problems, such as the random 3-SAT examples we consider below, can be simulated with GA, it follows that $P \neq NP$ implies that in at least one instance the entanglement barrier of the 3-SAT oracle must be exponentially high. Otherwise, because the cost of MPS simulations are determined by the entanglement, a single call to the oracle could be simulated in polynomial time and hence the full problem as well, which would imply $P = NP$. Thus there must be a direct connection between the complexity of a classical problem and the entanglement level of the associated quantum circuit.

On a practical level our quantum inspired GA puts very stringent conditions on putative quantum hardware that could be used to implement the GA. In presence of noise or imperfections, the probability of success of GA decreases exponentially with the number of applied gates in the circuit. Since the number of oracle calls, hence the number of gates, also scales exponentially with

n in GA, it follows that the overall success rate decays as an "exponential of an exponential", i.e. very quickly. We estimate the associated constraints on the hardware in terms of qubit number n, noise level, and quantum computer clock frequency and conclude that none of the requirements are realistic, even if fault-tolerant quantum technologies were available.

III. PROBLEM FORMULATION

Grover's algorithm [5, 6] aims to harness the potential of quantum parallelism by gradually extracting the result of a parallel function evaluation through a series of *Grover iterations*. For a problem involving classical states of length N, for which n qubits are needed such that $N = 2^n$, the number of Grover iterations needed to extract the result of the problem scales as $\sqrt{N}$, which is to be compared to worst-case classical strategies such as guessing solutions at random or performing an exhaustive search which have a cost scaling as N.

A. Notation and Problem Setting

Let $b \in \{0, 1, \cdots, N-1\}$ be an integer whose binary representation is $b_{n-1} \cdots b_1 b_0$, that is $b = \sum_{i=0}^{n-1} b_i 2^i$ with $b_i \in \{0, 1\}$. We denote by $|b\rangle = |b_{n-1} \cdots b_1 b_0\rangle$ the corresponding n-qubit state in the computational basis.

Let $f(b)$, be a function that takes a bitstring as an input b and returns

$$f(b) = \begin{cases} 1, & \text{if } b = w \\ 0, & \text{if } b \neq w \end{cases} . \tag{1}$$

Here w is a (unknown) specific bitstring. GA aims to solve the problem of finding the value of w in as few calls to the function f as possible. This problem can be viewed as inverting f, that is, computing $w = f^{-1}(1)$.

GA also assumes one can implement an *oracle operator* U_w such that for states $|b\rangle$ in the computational basis

$$U_w|b\rangle = (-1)^{f(b)}|b\rangle. \tag{2}$$

Since quantum computer can perform classical logic (at the price of adding ancilla qubits to ensure reversibility), a classical algorithm that computes $f(b)$ can be turned into a quantum circuit that implements Eq. (2). The explicit form of U_w reads,

$$U_w|b\rangle = \begin{cases} -|b\rangle, & \text{if } b = w \\ +|b\rangle, & \text{if } b \neq w \end{cases} \tag{3}$$

therefore U_w is equivalent to the operator

$$U_w = 1 - 2|w\rangle\langle w| . \tag{4}$$

However, for any real application of practical interest, one does not know the value of w ahead of time and only knows how to implement U_w as a quantum circuit based on the function f.

GA straightforwardly generalizes to the case of multiple solutions $\{w^\alpha\}_{\alpha=1}^S$ such that $f(w^\alpha) \equiv 1$. One defines the oracle as

$$U_w = 1 - 2\sum_{\alpha=1}^{S} |w^\alpha\rangle\langle w^\alpha|. \tag{5}$$

Each solution w^α has a binary representation $w_{n-1}^\alpha \cdot w_2^\alpha w_1^\alpha w_0^\alpha$. In this article, we focus on the case where the problem has a fixed number of solutions S (or more generically where S grows at most polynomially with n). For problems that have an exponential number of solutions, our algorithm would have to be revisited, but we conjecture that easy classical solutions exist in that case. For each qubit, we define the two states $|+\rangle$ and $|-\rangle$ as,

$$|\pm\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}} \tag{6}$$

and the equal weight superposition state $|s\rangle$ as,

$$|s\rangle = |+++\cdots+\rangle \tag{7}$$

$$= \frac{1}{\sqrt{2^n}} \sum_{x_{n-1}\cdots x_0 \in \{0,1\}^n} |x_{n-1}\cdots x_0\rangle . \tag{8}$$

Last, GA requires a second operator, the diffusion operator U_s that has a structure similar to the oracle but with respect to the known state $|s\rangle$:

$$U_s = 1 - 2|s\rangle\langle s| . \tag{9}$$

B. Definition of the Grover Algorithm

Given an oracle U_w, GA proceeds as follows

1. initiate the qubits in state $|000\cdots 0\rangle$

2. apply a Hadamard gate on each qubit to obtain $|s\rangle$

3. apply the oracle operator U_w

4. apply the diffusion operator U_s

5. repeat steps 3 and 4 each q times

6. measure the qubits in the computational basis and find $|w\rangle$ with a probability very close to one

The optimal number of steps of GA can be estimated to be about $q \approx r$ with $r \equiv \frac{\pi}{4}\sqrt{N} = \frac{\pi}{4}2^{n/2}$. In the case where there are multiple solutions, the measurement at the end produces one of the w^α with uniform probability. GA has an appealing geometrical interpretation [29]: U_w and U_s are mirror operators with respect to the hyper-planes perpendicular to w and s. It follows that the product $U_s U_w$ is a rotation inside the $(|s\rangle, |w\rangle)$ plane that gradually rotates the state from $|s\rangle$ to $|w\rangle$.

C. On the level of entanglement inside Grover's algorithm

The type of classical calculation we will consider involves representing the quantum state as a tensor network, specifically a matrix product state (MPS). An MPS compresses a quantum state by factoring it into a network of smaller tensors contracted in a one-dimensional chain like structure. For states having low to moderate entanglement, the MPS rank or dimension χ of the "bond" indices connecting the MPS tensors can be chosen relatively small while representing the state to very high or even perfect accuracy. Quantum states such as GHZ or W states are exactly MPS of rank $\chi = 2$ and product states such as the initial state $|s\rangle$ of Grover's algorithm are $\chi = 1$ MPS. An important fact we will use below is that any state which is a sum of P product states can be explicitly written as an MPS of rank $\chi = P$ [24].

In the context of GA, one finds that after any application of U_w or U_s, the internal state $|\Psi\rangle$ of the quantum computer lies in a superposition of $|s\rangle$ and $|w\rangle$ [29],

$$|\Psi\rangle = \alpha|s\rangle + \beta|w\rangle \tag{10}$$

with $|\alpha|^2 + |\beta|^2 = 1$, i.e. in the superposition of two unentangled states ($1 + S$ states in the general case). It follows that $|\Psi\rangle$ can be cast into the form,

$$|\Psi\rangle = \begin{bmatrix} \alpha & \beta \end{bmatrix} \begin{bmatrix} |w_1\rangle & 0 \\ 0 & |+\rangle \end{bmatrix} \begin{bmatrix} |w_2\rangle & 0 \\ 0 & |+\rangle \end{bmatrix}$$
$$\cdots \begin{bmatrix} |w_{n-1}\rangle & 0 \\ 0 & |+\rangle \end{bmatrix} \begin{bmatrix} |w_n\rangle & 0 \\ 0 & |+\rangle \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \tag{11}$$

which is manifestly a rank $\chi = 2$ MPS or $\chi = 1 + S$ in the general case. Such a state is minimally entangled and can easily be kept inside the memory of a classical computer at a linear cost in the number of qubits. In other words, while Grover's algorithm takes advantage of quantum parallelism (i.e. superposition), it uses very little entanglement for most of the algorithm. The only possible exception is while the oracle circuit has only been partially applied.

IV. A QUANTUM INSPIRED ALGORITHM FOR SIMULATING GROVER'S ALGORITHM IN A SINGLE CALL TO THE ORACLE

We now detail the different steps of our quantum inspired Grover's algorithm (QiGA). Although we use MPS and MPO technology for both QiGA and mere simulations of GA, we emphasize that the goals are very different. In the first, we aim at solving the Grover problem with as few computations as possible while in the latter we want to mimic what would happen in a (possibly noisy) actual quantum computer.

A. A low rank Matrix Product Operator (MPO) for Grover's oracle

A crucial observation that makes QiGA possible is that the oracle U_w can be cast into the form of a rank-2 MPO (rank $1+S$ in the case of multiple solutions). The explicit form of this MPO is

$$U_w = \begin{bmatrix} 1 & 1 \end{bmatrix} \left(\prod_{i=1}^{n} M_i \right) \begin{bmatrix} 1 \\ -2 \end{bmatrix} \tag{12}$$

with

$$M_i = \begin{bmatrix} I_i & 0 & 0 & \cdots \\ 0 & |w_i^1\rangle\langle w_i^1| & 0 & \cdots \\ 0 & 0 & |w_i^2\rangle\langle w_i^2| & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ \cdots & 0 & 0 & |w_i^S\rangle\langle w_i^S| \end{bmatrix} \tag{13}$$

where I_i is the 2×2 identity matrix acting on qubit i and $|w_i^\alpha\rangle\langle w_i^\alpha|$ the projector on the bitstring i of solution α. We emphasize that this MPO *exists* but its construction is not necessarily easy since, by definition, one does not have access to the solutions w^α. A similar MPO can be written for the diffusion operator U_s with the replacement of M_i by

$$M_i' = \begin{bmatrix} I_i & 0 \\ 0 & |+\rangle\langle+| \end{bmatrix} \tag{14}$$

in Eq.(12).

B. Solving the problem in a single call to the oracle

Assuming one has access to an explicit form of U_w, such as the MPO form above, a product of a small number of MPOs, or an efficiently simulable quantum circuit, one can construct the MPS state $|\Psi_w\rangle = U_w|s\rangle$. Using the definition of the oracle Eq. (5) and the fact that $\langle w^\alpha|s\rangle = \frac{1}{\sqrt{2^n}}$, we obtain

$$|\Psi_w\rangle = |s\rangle - \frac{2}{\sqrt{2^n}} \sum_{i=1}^{S} |w_i\rangle . \tag{15}$$

This expression is explicitly the sum of $1 + S$ product states, thus the state $|\Psi_w\rangle = U_w|s\rangle$ is exactly an MPS of bond dimension $\chi = 1 + S$.

The form of $|\Psi_w\rangle$ as a sum of product states in Eq. (15) immediately presents a classical strategy to extract the solution states $\{|w_i\rangle\}$ in a single step: one simply subtracts the state $|s\rangle$. Subtracting a product state such as $|s\rangle$ from an MPS is a straightforward operation with a cost scaling as $n\chi^3 \propto \log N$. For example, in the case of $n = 100$ qubits and $\chi = 50$ the subtraction takes about 0.2 seconds.

It is important to note that this subtraction operation has no quantum equivalent. This can be seen easily with an argument analogous to the no-cloning theorem: if there existed a unitary matrix that maps $|\Psi\rangle \otimes |s\rangle$ to $(|\Psi\rangle - |s\rangle) \otimes |\Phi\rangle$ for all $|\Psi\rangle$, then this matrix would map $|s\rangle \otimes |s\rangle$ to the null vector which contradicts the assumption that the matrix is unitary. It follows that our algorithm cannot be used as a "classical preconditioner" for amplitude amplification [7]. See the associated discussion in section VII A.

In summary the different steps of QiGA are:

1. Use classical simulation techniques to compute $|\Psi_w\rangle = U_w|s\rangle$ as an MPS of rank $\chi = 1 + S$

2. Compute $|\tilde{W}\rangle = |\Psi_w\rangle - |s\rangle$ and normalize it to obtain $|W\rangle$, an MPS of rank $\chi = S$

3. Sample from $|W\rangle$ to obtain the states $|w^\alpha\rangle$ with uniform probability, using the fact that perfect sampling of MPS can be performed with a cost $n\chi^2$ [30, 31]

If S is small enough, the states $|w^\alpha\rangle$ can also be fully enumerated.

One can modify the classical approach described above not only to sample individual solutions w^α but even to count the number of solutions. To do so, one acts with U_w on the *unnormalized* state $\sum_b |b\rangle$. Then the squared norm of the resulting state gives the number of solutions.

C. Obtaining the MPS using tensor cross interpolation

The problem is therefore reduced to the construction of the explicit MPS form of $|\Psi_w\rangle = U_w|s\rangle$ which is known to have a finite rank $\chi = 1 + S$.

Suppose for a specific bitstring $|b\rangle$ one is capable of classically calculating the amplitude $\langle b|\Psi_w\rangle$. Such a simulation is known as a closed simulation and is much easier [32] than a so-called open simulation which provides the full state $|\Psi_w\rangle$. We will comment in section VII on the practical limitations to these kind of simulations.

There has been recent progress in algorithms able to construct a MPS representation of the state $|\Psi_w\rangle$ in $O(\chi^2 n)$ calls to a routine that calculates $\langle b|\Psi_w\rangle$. Here, we make use of tensor cross interpolation [33–37] following the implementation described in [37]. The advantage of tensor cross interpolation is that it is agnostic to the details of the quantum circuit U_w and only requires an external classical subroutine that computes $\langle b|\Psi_w\rangle$. The algorithm requests the value of $\langle b|\Psi_w\rangle$ for certain values of b using an active learning algorithm. It follows that it is directly compatible with the most advanced methods that have been developed to calculate amplitudes of quantum circuits (including those that leverage on the high level of parallelism available in supercomputers).

Before we can use tensor cross interpolation effectively, a small adjustment must be performed. A direct calculation of $\langle b|\Psi_w\rangle$ provides

$$\langle b|\Psi_w\rangle = \frac{1}{2^{n/2}}\left[1 - 2\sum_{\alpha=1}^{S}\delta_{b,w^\alpha}\right] \quad . \qquad (16)$$

For a random bitstring b, one has $\langle b|\Psi_w\rangle = 1/\sqrt{2^n}$ since it is exponentially unlikely that b matches one of the solutions w^α. It follows that the tensor cross interpolation algorithm would fail to reconstruct the MPS as its exploration of the $\langle b|\Psi_w\rangle$ function would have an exponentially low probability of finding the relevant part (second half of the right hand side of Eq.(16)). Another way to see the same problem is to write $\langle b|\Psi_w\rangle$ in terms of the calls to the function $f(b)$. It reads

$$\langle b|\Psi_w\rangle = \frac{1}{2^{n/2}}(-1)^{f(b)} \qquad (17)$$

i.e. the amplitudes can be calculated in a single call to $f(b)$. Hence, if $|\Psi_w\rangle$ MPS could be reconstructed from $O(n) = O(\log N)$ calls to $\langle b|\Psi_w\rangle$, it would mean that the original problem could be solved in $O(n)$ calls to the function $f(b)$ hence that NP-complete problems could be solved in polynomial time, an unlikely situation.

To solve this issue, we turn to the $|\pm\rangle$ basis and calculate $\langle \beta|\Psi_w\rangle$ instead of $\langle b|\Psi_w\rangle$ where $|\beta\rangle$ is a product of $|\pm\rangle$ states (e.g. $|+-+-...-+\rangle$). Denoting the binary representation of β as $\beta_0, \beta_1...\beta_n$ with $\beta_i = 0$ for a state $|+\rangle$ and $\beta_i = 1$ for a state $|-\rangle$, we find

$$\langle \beta|\Psi_w\rangle = \delta_{\beta,0} - \frac{2}{2^n}\sum_{\alpha=1}^{S}(-1)^{\sum_{i=0}^{n-1}\beta_i w_i^\alpha} \qquad (18)$$

This form is directly suitable for tensor cross interpolation since information about the solutions w^α is now present for any bitstring β. We emphasize that QiGA itself knows nothing about the solution w^α and only uses the amplitudes $\langle \beta|\Psi_w\rangle$. Calculating these amplitudes $\langle \beta|\Psi_w\rangle$ has the same complexity as calculating $\langle b|\Psi_w\rangle$ since the two quantum circuits only differ by a layer of Hadamard gates at the end. Similarly, when the MPS is known in the $|\beta\rangle$ basis, it is simply a matter of applying local Hadamard gates to get it back in the $|b\rangle$ basis. We have checked in explicit numerical calculations that our implementation of tensor cross interpolation can reconstruct the MPS in $O(n)$ calls to the $\langle \beta|\Psi_w\rangle$ subroutine up to at least $n = 1000$.

In terms of call the the function $f(b)$, the amplitudes $\langle \beta|\Psi_w\rangle$ take the form,

$$\langle \beta|\Psi_w\rangle = \frac{1}{2^n}\sum_{b=0}^{2^n-1}(-1)^{f(b)+\sum_{i=0}^{n-1}b_i\beta_i} \qquad (19)$$

which takes $O(2^n)$ calls to the classical function, if one does not take advantage of its quantum circuit form. Hence we recover the expected $O(N)$ classical scaling

to solve the search problem if one does not use insights about the function f. When using the quantum circuit to compute the oracle amplitudes, the QiGA complexity will depend on the entanglement barrier present in a single application of the quantum oracle as illustrated in Fig. 2.

We emphasize that the approach outlined above is only feasible for a closed, classical simulation of the oracle circuit; it cannot be attempted on a quantum computer. Indeed, a quantum computer only provides bitstrings β distributed according to the probability $|\langle\beta|\Psi_w\rangle|^2$ but it does not provide the actual value $\langle\beta|\Psi_w\rangle$ (nor can one choose the value of β).

V. ILLUSTRATION WITH AN EXPLICIT CONSTRUCTION FOR THE 3-SAT PROBLEM

To illustrate our quantum inspired GA with a practical application, we have implemented a simulation of the oracle for the 3-SAT boolean satisfiability problem. 3-SAT is an NP-complete problem and finding fast, possibly heuristic, algorithms for solving it is the subject of active research, with applications including cryptanalysis [38, 39], industrial operations research [40], and computational biology [41].

In a SAT problem, the function $f(b)$ is given by a set of p clauses that must all be satisfied for b to be a solution. In 3-SAT, each clause δ is constructed out of 3 variables b_{i_δ}, b_{j_δ} and b_{k_δ} from the binary representation of b. $f(b)$ takes the form,

$$f(b) = (\tilde{b}_{i_1} \vee \tilde{b}_{j_1} \vee \tilde{b}_{k_1}) \wedge (\tilde{b}_{i_2} \vee \tilde{b}_{j_2} \vee \tilde{b}_{k_2}) \wedge \ldots$$
$$\ldots \wedge (\tilde{b}_{i_p} \vee \tilde{b}_{j_p} \vee \tilde{b}_{k_p}) \quad (20)$$

where $\vee$ means logical "or", $\wedge$ logical "and", and $\tilde{b}_a = b_a$ or $\tilde{b}_a = \neg b_a$ (not b_a) depending on the clause.

The problems we consider are defined by choosing the ratio $\alpha = p/n$ of clauses to variables (or qubits) to be fixed, usually between $4 < \alpha < 5$ since in this range the number of satisfying solutions S becomes small. Otherwise the choice of which variables enter into each clause and whether a variable is negated is made with uniform probability. Below we will consider totally random SAT clauses in section V B then clauses with quasi-one-dimensional structure in section V C.

A. Tensor Network SAT Approach

To explicitly implement the Grover's oracle operator for 3-SAT and construct the post-oracle state $|\Psi_w\rangle$, first prepare the state of the system to be

$$|+\rangle_1|+\rangle_2|+\rangle_3 \cdots |+\rangle_n|1\rangle_A = |s\rangle|1\rangle_A \quad (21)$$

where the extra qubit in the $|1\rangle_A$ state acts as an ancilla whose role is to record which states of the previous n qubits either satisfy ($|1\rangle_A$) or do not satisfy ($|0\rangle_A$) all of the SAT clauses applied so far.

Next, each SAT clause C such as $C = (b_3 \vee \neg b_7 \vee b_8)$ is mapped to an operator by noting that there is only one bitstring which *fails to satisfy* the clause. In the example above, this bitstring is $0_3, 1_7, 0_8$. Using this bitstring, one defines an operator

$$O_C = P_3^0 \otimes P_7^1 \otimes P_8^0 \otimes P_A^0 + (1 - P_3^0 \otimes P_7^1 \otimes P_8^0) \otimes 1_A . \quad (22)$$

which projects the ancilla qubit into the $|0\rangle_A$ state for any state containing the unsatisfying bitstring. Otherwise it leaves the ancilla unchanged. Here $P_i^0 = |0\rangle\langle0|$ and $P_i^1 = |1\rangle\langle1|$ are projectors onto the $|0\rangle$ or $|1\rangle$ states for qubit i.

In our classical implementation, the operator O_C can be applied in its above form using straightforward tensor network methods. We used the approach of implementing each O_C as an MPO and applying these MPOs to the quantum state represented as an MPS. As an MPO, O_C has a rank $\chi = 3$, which can be understood from the fact that when one expands all the terms it is the sum of three product operators [24].

After applying the O_C operators for every SAT clause, the state of the system becomes

$$\frac{1}{\sqrt{2^n}} \sum_{i=1}^{S} |w_i\rangle|1\rangle_A + \frac{1}{\sqrt{2^n}} \sum_{j=1}^{2^n - S} |\tilde{w}_j\rangle|0\rangle_A \quad (23)$$

where the $\{w_i\}$ are the satisfying bitstrings and the $\{\tilde{w}_j\}$ are the unsatisfying ones. To convert this state to a conventional Grover's post-oracle state Eq. (15), one can perform a post-selected, or forced, measurement of the ancilla qubit to be in the $|-\rangle = H|1\rangle$ state. (Note that for a tensor network such a forced measurement always succeeds on the first try and can be done by just acting with a single-qubit projection operator.) After dropping the now unentangled ancilla qubit, the state will take the form of Eq. (15). If one is only interested in solving the Grover problem rather than constructing the post-oracle state, one simply projects the ancilla of Eq. (23) onto the state $|1\rangle_A$.

B. Random SAT Experiments

We have tested this classical oracle implementation on fully random SAT clauses (variables b_{i_p} chosen by drawing i_p randomly from $1, 2, .., n$ and with each variable negated with probability $1/2$) for up to $n = 40$ variables or qubits, using the ITensor software [42, 43] running on a single workstation with four Xeon 3.6 GHz processors and 256 Gb of RAM. For all experiments we used a ratio of clauses to variables of $p/n = 4.2$. The results shown in Table I are for various experiments over a range of n and different random instances for the same n with different numbers S of satisfying bitstrings. We report both the

n	S	$\chi_{\max}$	time
30	4	467	21 s
32	2	954	1.8 minutes
34	48	1162	3.2 minutes
36	16	1994	8.3 minutes
38	8	5867	1.6 hours
40	0	1402	4.2 minutes
40	28	2926	21 minutes
40	161	5690	1.65 hours
40	174	10374	6.5 hours

TABLE I. Maximum MPS ranks $\chi_{\max}$ and execution times to compute the post-oracle state corresponding to random 3-SAT problem instances involving n variables (n qubits). The table also shows the number of satisfying assignments or bit-strings S for each problem instance.

maximum MPS rank $\chi_{\max}$, which was the largest rank encountered during application of the O_C operators, and the total time required to apply all of the operators and construct the post-oracle state.

After each post-oracle state was prepared, its quality was checked by projecting (post-selecting) the ancilla qubit into the state $|1\rangle_A$ then sampling 5000 bitstrings from the other n qubits to verify that all samples satisfied the SAT clauses. To count the number S of satisfying bitstrings (#SAT problem) we applied the MPOs O_C to an *unnormalized* state with each qubit (except the ancilla) initialized to $(|0\rangle + |1\rangle)$. Afterward, we projected the ancilla into the $|1\rangle_A$ state and computed the norm of the resulting state which is equal to S. For smaller systems these counts were verified to be correct by checking with exact enumeration. The resulting counts S are shown in the second column of Table I,

These results indicate the post-Grover's-oracle state can be prepared classically for typical 3-SAT instances for at least $n = 40$ qubits on a single workstation. For problems of this size, the optimal number of iterations of Grover's algorithm would be $r = 823,500$ in contrast to the single application of the oracle used here. The largest MPS rank encountered across the experiments we performed was $\chi = 10,374$ which is a small fraction (1%) of the maximum possible rank $2^{40/2} \approx 10^6$. The entanglement barrier in the random 3-SAT problem is not only relatively modest for finite problem sizes, but was typically observed to be lower for the case of fewer satisfying solutions S. Hence QiGA appears to perform better on problems with few solutions.

It is important to note that the approach above can be straightforwardly parallelized through a "slicing" approach, similar to other recent high-performance quantum circuit simulations [44–46]. Instead of initializing all the n input qubits to the $|+\rangle$ state, a subset of p of these qubits can be initialized to either the $|0\rangle$ or $|1\rangle$ states. By running 2^p separate calculations for each setting of these computational-basis qubits one can obtain the same final result by combining the solutions obtained from 2^p com-

puters working with no communication overhead. When trying this in practice, we found that the total computational effort (sum of running times of each of the parallel computers) was comparable to the original (serial) running time, while the maximum time for any one computer was significantly less than the original time. Because the oracle acts classically on the computational basis qubits, the maximum entanglement during the oracle application is generally much lower for each sliced input so that the parallel approach results in a faster time to solution.

Note that our implementation of a SAT solver as a Grover oracle operator is nevertheless far slower than the most efficient classical SAT solving algorithms, some of which also happen to use tensor networks [47] and can solve typical $n = 40$ instances of 3-SAT in under one second.

C. Quasi-One-Dimensional SAT Experiments

In this section, we design instances of the 3-SAT problem where the QiGA approach has *provably linear scaling* in the number of qubits n, that is a logarithmic scaling with problem size $\log(N)$. This is in sharp contrast to the 2^n scaling of an unstructured classical problem. The goal of the construction and associated experiments we perform below is to illustrate two points. First, it shows that there *are* classes of problems for which QiGA is always exponentially faster than GA. Second, the underlying structure that makes the problem easier for QiGA *need not* be known a priori: QiGA discovers this structure and takes advantage of it automatically.

We consider a quasi-1D case that involves grouping variables into blocks of size B along a 1D path, with the first block $(1, 2, \ldots, B)$, second block $(B + 1, \ldots, 2B)$, etc. The SAT problem is then defined by two sets of SAT clauses required to be satisfied altogether:

1. N_{intra} fully random SAT clauses where variables in each clause only act within each block

2. L_{inter} layers of random SAT clauses where variables span across two neighboring blocks

The cases we consider will allow N_{intra} to be any size while L_{inter} is fixed to a small value such as $L_{\text{inter}} = 1, 2, 3$.

The proof of linear scaling consists of bounding the cost of applying the constraints of clauses in sets (1) and (2) above separately. We will use a similar approach as in Section V A above, with the slight modification that we will project out any unsatisfying bitstrings for a clause $C_p = (b_{i_p} \vee b_{j_p} \vee b_{k_p})$ by acting with an operator

$$O_{C_p} = (1 - P_{i_p}^{b_{i_p}} P_{j_p}^{b_{j_p}} P_{k_p}^{b_{k_p}}) \qquad (24)$$

that sets the state containing the single bitstring not satisfying C_p to zero. There is no ancilla qubit in this approach. The process of applying the MPOs O_{C_p} is depicted in Fig. 3 and explained in detail below.

FIG. 3. Schematic of the process of applying the operators (MPOs) O_{C_p} enforcing the SAT clauses defining the quasi-1D SAT problem. The qubits are organized into blocks of size B and a number of random, intra-block SAT clauses are enforced which act only inside each block. Then a fixed number of layers of inter-block clauses are enforced which act between neighboring blocks.

Starting from the product superposition state

$$|s\rangle = |+\rangle_1 |+\rangle_2 |+\rangle_3 \cdots |+\rangle_n \qquad (25)$$

and applying the operators for the set (1) of intra-block clauses, the cost scales as $(N_{\text{intra}} \ n \ 2^{3B/2})$ in the worst case. To see why, note that the state after applying the operators will be a product of MPS for each block, and the maximum bond dimension of an MPS with B sites is $\chi = 2^{B/2}$.. Algorithms for applying operators to MPS of this size scale as $B\chi^3 = B\,2^{3B/2}$. One has to apply each of the operators O_{C_p} and there are N_{intra} of these. Finally there are n/B blocks. Multiplying each of these costs gives the above scaling. Thus the cost of enforcing the set (1) clauses scales only linearly with number of qubits n.

Next one enforces the L_{inter} layers inter-block clauses in set (2). For each layer of such clauses, one can group the clauses into two sets, one acting across blocks 1 and 2, then 3 and 4, then 5 and 6, etc. and the next set acting across blocks 2 and 3, 4 and 5, 6 and 7, etc. The cost of applying the first layer scales at most as $n \ 2^{3B/2}$ and doubles the bond dimension between blocks in the worst case, since the bond dimension of the O_{C_p} MPOs is 2. The cost of applying the next layer will then scale at most as $(n \ 8 \ 2^{3B/2})$, the extra factor of $8 = 2^3$ coming from the doubling of bond dimension due to the first layer. In general the cost of enforcing the $(2L_{\text{inter}})$ layers of inter-block clauses will be $(n \ 2^{2L_{\text{inter}}-1} \ 2^{3B/2})$.

Therefore the overall cost of enforcing all of the 1D 3-SAT clauses scales as

$$(n \ N_{\text{intra}} \ 2^{3B/2}) + (n \ 2^{2L_{\text{inter}}-1} 2^{3B/2}) \qquad (26)$$

which is manifestly linear in n, assuming B and L_{inter} are held constant, and that N_{intra} depends only on B (is chosen proportional to B).

In practice, when implementing the above method on test systems of $n = 40$ qubits, using a fixed block size of $B = 10$, taking $L_{\text{inter}} = 2$ and choosing $N_{\text{intra}} = 3.7 \cdot B$, we find that all the clauses can be enforced in a running time of under 1 second. The maximum MPS bond dimension observed is $\chi = 22$. Systems with $n = 60$ qubits are just as easy for the same block size and number of

n	S	$\chi_{\max}$	time
40	99	21	0.973s
40	50	22	0.973s
40	108	16	0.989s
40	0	19	0.985s
60	4530	22	2.00s
60	0	19	1.98s
60	17920	19	1.96s

TABLE II. Maximum MPS ranks $\chi_{\max}$ and execution times to compute the post-oracle state corresponding to blocked 1D 3-SAT problem instances involving n variables (n qubits). In all cases the block size was chosen as $B = 10$ and $N_{\text{intra}} = 37$ random clauses were applied within each block while $N_{\text{inter}} = 2$ layers of clauses were applied between neighboring blocks. The table also shows the number of satisfying assignments or bitstrings S for each problem instance.

clauses per block, with similar maximum bond dimension and running times just under 2 seconds. See Table II for detailed results.

VI. SCALING OF ERRORS IN GROVER'S ALGORITHM

We now turn to the second part of this work where we discuss the possibility of quantum advantage for GA. Before we can do that, we need to study the sensitivity of GA to the presence of imperfections such as gate errors or decoherence. There has been some previous literature on this subject [48]. The advantage of our tensor network approach is the capability to study relatively large system from which we can extract the actual scaling of the errors in GA.

At the end of GA, the quantum computer is supposed to be in the state $|w\rangle$ so that upon measuring the qubits one finds the solution w almost certainly (up to negligible exponentially small $1/N$ corrections). However, due to imperfection in the gates or environmental coupling and decoherence the system will be in some other state instead, described by the density matrix ρ. The probability to get the correct output w when one performs the measurement of the qubits is $F = \langle w|\rho|w\rangle$. F is known as the *fidelity* of the calculation which matches the success probability of the algorithm for GA. Generically, the fidelity decays exponentially with the number of gates applied in the calculation as well as with any idle time (decoherence). This exponential behavior has been observed ubiquitously in experiments and was established on a large scale in the seminal "quantum supremacy" experiment of Google [49]. Denoting by ϵ the typical error per gate and N_g the total number of gates used, we therefore expect the decay rate to be $F \approx e^{-\epsilon N_g}$ with $N_g = r(N_o + N_d)$ where N_o (resp. N_d) is the number of gates in the quantum circuit of the oracle (resp. diffusion operator). In other words, the success probability
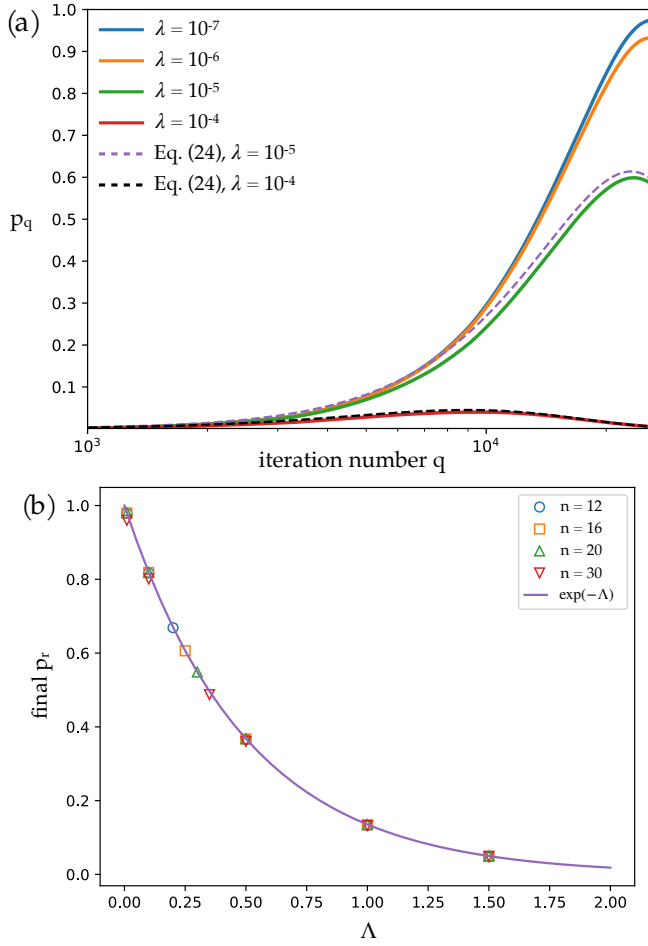
FIG. 4. (a) Probability of success versus number of iterations of Grover's algorithm for $n = 30$ qubits and different levels of noise λ. Also shown are theoretical fits to the cases $\lambda = 10^{-5}, 10^{-4}$. For $n = 30$ the theoretically optimal number of iterations is $r = 25,736$. The final success probability reaches a reasonable value for noise $\lambda < 5 \times 10^{-5}$, but once the noise becomes as large as $\lambda = 10^4$, the success probability reaches a maximum of only 0.04 at the 10,000th iteration then falls to about 0.006 by the final iteration. (b) Final probability of success after the optimal number of iterations r as a function of total noise $\Lambda = \lambda r$ where λ is the amount of depolarizing noise per iteration.

of GA is expected to follow a rather unfavorable *double exponential* decay with n,

$$F \approx \exp\left[-\frac{\pi\epsilon}{4}(\sqrt{2})^n (N_o + N_d)\right]. \qquad (27)$$

To substantiate the scaling of GA with noise and number of qubits, we have performed classical simulations using an MPO state representation capable of representing mixed states. For these simulations we also implement the oracle and diffusion operators as MPO tensor networks as described in Section IV A. As discussed earlier, without noise the MPS bond dimension χ of the state in between the oracle and diffusion steps (so at every step of these simulations) is $\chi = 2$, leading to an MPO-mixed-

state representation of only $\chi = 4$ so that each step can be performed rather quickly and it is possible to run the whole algorithm up to about $n = 30$ qubits in under an hour. Adding noise to the simulations only modifies the bond dimension of the state very slightly and observed bond dimensions always remain less than $\chi \lesssim 10$.

To model the effects of noise, we apply a depolarizing noise channel $\Delta_\lambda(\rho) = (1 - \lambda)\rho + \frac{\lambda}{2^n}I$ in between Grover iterations, that is once per iteration. By not including any noise during the application of the oracle or the diffuser our simulations are being very generous in favor of the success of GA. The noise per iteration λ relates simply to the noise per gate $\lambda = \epsilon(N_o + N_d)$.

The results of our simulations for $n = 30$ qubits and various levels of noise λ are shown in Fig. 4(a). As long as the noise $\lambda \lesssim 5 \times 10^{-5}$ the final probability of success reaches a reasonable value after the optimal number of iterations, which for $n = 30$ is $r = 25,736$. However, the probability of success after r iterations falls below 1% once the noise becomes larger than about 10^{-4}. Note that non-zero noise leads to a maximum in the success probability at an earlier iteration $q^* < r$. We analyze the height and location of this maximum further below.

Due to the transparent nature of GA and the depolarizing noise model, one can actually work out an exact result for the state after q steps of the algorithm. By a recursive procedure, one finds:

$$\rho_q = (1 - \lambda)^q |\Psi_q\rangle\langle\Psi_q| + (1 - (1 - \lambda)^q)I\frac{1}{2^n} \qquad (28)$$

where $|\Psi_q\rangle$ is the ideal pure state after q noiseless Grover iterations. Using a well-known result for the probability of success of the noiseless GA after q iterations [50]

$$|\langle w|\Psi_q\rangle|^2 = \sin^2((2q + 1)\theta) \qquad (29)$$

$$\theta = \arcsin\left(\frac{1}{\sqrt{2^n}}\right) \qquad (30)$$

it follows that the probability of success after iteration q is given by

$$p_q = \langle w|\rho_q|w\rangle = (1 - \lambda)^q \sin^2((2q + 1)\theta) \quad . \qquad (31)$$

Since $q \gg 1$, ignoring exponentially small corrections, we have:

$$p_q = e^{-\lambda q} \sin^2(2q\theta) \quad . \qquad (32)$$

We show in Fig. 4 that this fit works well, though there is a slight disagreement which we attribute not to the fit, but to a loss of precision in the early iterations of the numerical simulations due to the use of double-precision floating point numbers and the very small signal of GA through the early iterations [51].

Interestingly, the form of p_q Eq. (32) means that if one defines the noise level in terms of a parameter Λ such that $\lambda = \Lambda/r$ then the final success probability is

$$p_r = e^{-\Lambda} \qquad (33)$$

regardless of the number of qubits, using the fact that $\sin^2(2r\theta) \approx 1$ up to exponentially small corrections. One can interpret $\Lambda = \lambda \cdot r$ as the total accumulated noise throughout a complete run of GA. In Fig. 4 we show how the final success probability p_r observed in noisy simulations fits very well to $e^{-\Lambda}$.

In the presence of large noise, due to a maximum that develops in the fidelity curve, it is advantageous for GA to stop the iteration at a certain smaller value of $q = q^*$. An explicit calculation of the optimum of Eq.(32) provides

$$\tan[2(r - q^*)\theta] = \frac{\Lambda}{\pi} \tag{34}$$

from which we arrive at the optimum success probability,

$$p_{\text{success}} = \frac{e^{\frac{2\Lambda}{\pi} \arctan(\frac{\Lambda}{\pi})}}{1 + (\Lambda/\pi)^2} e^{-\Lambda}. \tag{35}$$

This formula behaves as $p_{\text{success}} \approx e^{-\Lambda}$ for small Λ (approaching 1.0 as $\Lambda \to 0$) and behaves as $p_{\text{success}} \approx e^{-2} \cdot (\pi/\Lambda)^2$ for large Λ.

Because p_{success} depends only on Λ, an important conclusion is that for the GA success probability to scale successfully to large values of n, the total noise Λ must be held independent of n. The noise per iteration λ must therefore scale as

$$\lambda = \frac{\Lambda}{r} \propto \frac{\Lambda}{(\sqrt{2})^n} \tag{36}$$

showing the noise per iteration must be reduced exponentially with n.

VII.   ON THE POSSIBILITY OF QUANTUM ADVANTAGE IN GROVER'S ALGORITHM

There are two kinds of quantum advantages. The theoretical one, i.e. the possibility that in an idealized world a perfect quantum computer could perform parametrically better than a classical one for a given task. And the practical one, i.e. the possibility that an actual device does something useful faster than a classical machine.

With respect to the first question our QiGA merely reverses the charge of the proof: we show that there is no theoretical quantum advantage unless proven otherwise and quantum advantage has to be decided in a case-by-case manner.

With respect to the second kind of quantum advantage involving an actual machine, the existence of QiGA and the demands for implementing GA place drastic bounds on the hardware needed which we will argue are simply out of reach. When discussing hardware, there is long list of specifications that we could consider including heat management, control electronics, classical data bandwidth for e.g. syndrom analysis, device variability, power consumption...Here we limit ourselves to discussing three aspects: the total qubit count, the error budget per gate and the time to solution.

A.   Absence of generic theoretical quantum advantage

QiGA implies that if the quantum circuit for the oracle can be simulated classically, then $O(n)$ such calculations are sufficient to solve the problem while a quantum computer needs $O(2^{n/2})$ calls to the oracle. An immediate consequence is that no theoretical quantum advantage can be claimed generically, i.e. irrespectively of the nature of the underlying quantum circuit for the oracle. This is an important point to make with respect to the large literature which assumes, implicitly or explicitly, the existence of a quantum speed-up every time a GA replaces its classical counterpart [52].

If the complexity for calculating one amplitude of the oracle is smaller than $(\sqrt{2})^n$, then QiGA is parametrically faster than its quantum counterpart. Constructing an oracle whose classical simulation is provably harder than $(\sqrt{2})^n$ can most likely be done. Indeed, in the large depth limit classical simulations of quantum circuits have a generic complexity of 2^n. Yet, we are not aware of such a construction for an actual oracle (i.e. a circuit whose output amplitudes are only ± 1). Conversely, there are clear cases where classical algorithms win. For instance, if the oracle can be simulated with a fixed depth, then the problem can be solved in linear time using MPS technology while GA would require an exponential time. The Quasi-1D SAT is another example.

We emphasize that our work does not contradict previous work that formally proves that the quantum speed-up of GA is optimal [8]. While the proof is certainly valid technically, its usefulness for a given problem requires the best known classical strategy to scale as 2^n (i.e. worst-case classical scaling) for that problem. But for any specific problem, if Grover's can be applied there must exist an explicit circuit for the oracle. So there is always at least some structure to the problem: the structure of the oracle circuit. One can always try to simulate this oracle circuit by applying it to a tensor network. Until such a "simulability check" has been done, the applicability of the proof remains in doubt because the problem might not satisfy the proof's assumptions. In other words, one must be very careful with using the concept of an "oracle" which, however appealing theoretically, might not be relevant to practical applications.

A corollary of the existence of QiGA is that the quantum circuit associated to the oracle of an NP complete problem must be exponentially difficult to simulate in the general case, i.e. presents an exponentially high entanglement barrier. Indeed, otherwise, one could simulate it in polynomial time which would prove $P = NP$ a statement widely believed to be untrue. Hence QiGA provides a direct connection between classical complexity theory and the entanglement level of a quantum circuit.

Lastly, we would like to discuss the relation of this work to amplitude amplification [7], a direct generalisation of GA. In some cases, there exist fast classical heuristic algorithms that can solve NP-hard problems faster than GA,

though still scaling exponentially. For instance, there exist very fast classical algorithms for the 3-SAT problem which we considered earlier in our QiGA demonstrations (incidentally, among the best are tensor network approaches [47]). Amplitude amplification is claimed to recover the quadratic Grover speedup over such fast classical algorithms by combining these algorithms with a slight modification of GA. Below, we show that QiGA applies in this context as well. We again argue that the question of whether the oracle can be simulated classically is a crucial one.

A classical heuristic algorithm takes the form of a function $b = G(r)$ that proposes a bitstring b as a possible solution. Here the variable r captures the probabilistic nature of the heuristic, e.g. it can be the seed of the pseudo random number generator used in the heuristic. In a good heuristic, r need to span a much smaller number of values than b. For instance, in the context of 3-SAT, the Schoning algorithm [53] solves 3-SAT (with a complexity $(4/3)^n$) by an iterative process over bitstrings where at each stage an unsatisfied clause is chosen randomly and then a bit of the clause is flipped in order for the clause to become satisfied. To transform this heuristic into a quantum algorithm, amplitude amplification does not search for the bitstring b that satisfies $f(b) = 1$ but instead uses GA to search for the value of r that satisfies $f(G(r)) = 1$ (see the discussion around Theorem 6 of [7]), i.e. one needs to use GA with the oracle $U_w|r\rangle = (-1)^{f(G(r))}|r\rangle$. It follows that our QiGA approach applies directly to the amplitude amplification of these heuristics: one only needs to modify the oracle in the same way it would be modified for the quantum computer. Hence, the question of the existence of a theoretical quantum advantage translates again into the one of an entanglement barrier, here of the new oracle $f(G(r))$. Anticipating on the discussion of the next section, these heuristics make classical problems computable up to very large values of n. For instance for 3-SAT, problems with $n \sim 10,000$ are well within range. It follows that the entrance level for GA for these problems would be much higher than for the problems for which no heuristic is available. As we shall see in the estimates below, this would translate into inaccessible number of required qubits and astronomically large times to solution.

B. Absence of practical quantum advantage on a noisy quantum computer

We now turn to the question of a practical advantage and provide some estimates about the specifications a quantum hardware must fullfill to solve a task better than what one can do classically. We start by estimating the total gate count N_g. The diffusion operator typically requires $2n$ Toffoli gates and the oracle at least the same (a simpler oracle would likely be rather easy to solve classically). Each Toffoli gate must be decomposed into 15 gates (including 6 Control-NOT and 7 T gates). We ar-

rive at a total gate count for GA of $N_g \geq 60\,n\,2^{n/2}$ assuming perfect connectivity (i.e. that the two-qubit gates can be applied between any pairs of qubits). In order for the final success probability to be of order unity (here we choose $p_{\text{success}} = 1/2$) we need $\Lambda \approx 0.8$ which translates into $\epsilon \leq 1/(60\,n\,2^{n/2})$.

It follows that in order to apply GA on just 5 qubits, one needs $\epsilon \leq 5.10^{-4}$ which is much better than any existing hardware. Indeed, the experimental value of the error per gate ϵ has been mostly stable in the last ten years, typically around $\epsilon \approx 0.01$ for the state of the art [49] and slightly better for some systems with few (< 10) qubits. Previous applications of GA for a few qubits used a much smaller gate count in order to retain a large enough fidelity. This is possible for contrived examples where one uses Eq. (3) instead of Eq. (2) i.e. one explicitly uses information about the solution w instead of performing the classical logic of computing $f(b)$. While this is perfectly acceptable for proof of principle experiments, this does not correspond to the full application of GA to actually solve a problem. Going to $n = 40$ which we can easily solve on a desktop computer using QiGA leads to $\epsilon \leq 4.10^{-10}$. Manipulating tens of qubits with precisions better than one part in a billion for a total of billions of gates is in our view completely unrealistic. Using the best available algorithms available on supercomputers (see [32] for a recent review) to perform QiGA we estimate that $n = 80$ problems are accessible on supercomputers (most probably $n \geq 100$). Solving such a problem would require $\epsilon \leq 2.10^{-16}$ as well as a time to solution (assuming very fast 10 ns gates) of more than one year of uninterrupted running time.

C. Absence of practical quantum advantage on a fault tolerant quantum computer

The problem of limited qubit fidelity can in principle be solved by quantum error correction which should allow to lower the effective error level per gate ϵ_L by constructing robust logical qubits out of several physical qubits. However, quantum error correction trades a better effective ϵ_L with a much higher qubit count n_L as well as a much higher time to solution since logical qubits require many physical qubit operations in order to make a single logical one. Hence, we can already anticipate that quantum error correction is unlikely to help given the above already high estimate of the time to solution.

To make quantitative estimates, we focus on the example of the surface code, one of the most robust QEC code with a clear path to implementation [54]. We ignore non-correctable errors [55] for simplicity. We also ignore the so-called syndrome analysis although this would be quite problematic in practice. In the surface code, the error ϵ_L per logical gate on a logical qubit scales as

$$\epsilon_{\text{L}} \propto \epsilon_{\text{ph}} \left(\frac{\epsilon_{\text{ph}}}{\epsilon_{\text{th}}} \right)^{\sqrt{N_c}/2} \tag{37}$$

where N_c is the number of physical qubit per logical qubit, ϵ_{ph} is the error per gate on the physical qubits and ϵ_{th} is the threshold of the code, around $\epsilon_{th} \approx 0.01$ for the surface code [54]. Ignoring logarithmic corrections, we immediately observe that the exponentially large running time which implies that one must have $\epsilon_L(\sqrt{2})^n \leq 1$ translates into

$$N_c \propto n^2. \tag{38}$$

i.e. the number of physical qubits per logical one increases quadratically with n in sharp contrast to e.g. Shor's algorithm where $N_c \propto (\log n)^2$ has a weak logarithmic increase. We can already surmise that, since the estimates for algorithms with exponential speed-up already involve millions, or more realistically billions, of very high quality physical qubits to address $n = 100$, the added complexity in Grover would move these estimates far beyond anything reasonable [56].

To continue, we need to modify our gate count estimate and consider separately the Control-NOT gates and the T gates. In the surface code, Control-NOT are performed through "braiding", a process that requires making a loop of one qubit around the other which costs an extra $\sqrt{n}\sqrt{N_c}$ factor. Also, the time scale is no longer limited by the time it takes to execute one gate but by the time to make a projective measurement of the so-called stabilizers. For instance in superconducting transmon qubit, the former can be as low as 10 ns while the later takes around $1\mu s$. Systems based on atomic physics such as cold ions are typically much slower. Also, the error per gate ϵ is likely to be limited by the measurement errors that are typically much worst than the error per gate. Considering only Control-NOT operations we arrive at a time to solution of 10^5 years for $n = 80$. Applying T gates on a surface code requires the availability of special states that can be obtained through "magic state distillation", a lengthy process that involve several logical qubits for each T gate that one wants to apply. In order for the already rather large time to solution not to be slowed down by the T gates, one would need to incorporate a large number of "T gate factories" on the quantum chip thereby raising the total qubit count dramatically. We need not to elaborate further, it should be clear at this stage that the fate of quantum error correction for the implementation of GA is extremely doubtful.

VIII. CONCLUSION

Grover's algorithm is an elegant intellectual construction. Unfortunately our analysis indicates that it will remain so for the foreseeable future.

We have constructed a quantum inspired version of Grover's algorithm that can solve quantum problem in a single call to the oracle, requiring exponentially fewer steps than Grover's algorithm, provided one is able to compute individual amplitudes of the oracle. We have

also provided specific cases where this "classical advantage" can be realized.

Since our classical algorithm is fully general, it provides a clear benchmark against which one can evaluate the potential speed-up of Grover algorithm both theoretically and practically. While we cannot exclude a *theoretical* quantum scaling advantage for every problem, assuming an idealized quantum implementation, we estimate that practical quantum implementations will be associated with astronomically large computing times. On the other hand, problems for which a quantum implementation may seem necessary could have hidden structure revealed by our classical algorithm in the form of low entanglement barriers on the way to a solution. And even if the entanglement barrier does grow with problem size and produces an exponential scaling, it remains possible this scaling could be better than $2^{n/2}$ for specific classes of problems.

A work which has some overlap with ours is the proposal of Chamon and Mucciolo [57] to use MPS tensor networks as a classical computing platform for solving Grover's algorithm problems. An important technical difference, however, is that their algorithm depends on computing and comparing exponentially small numbers, which could become challenging when working with fixed precision floating-point. Also, while their work discusses the growth of entanglement in Grover oracles, only worst-case bounds are stated.

A separate line of work based on the problem size needed for a quantum/classical scaling crossover has also cast doubt on the usefulness of Grover's algorithm since it only offers at best a quadratic speedup, while the estimated speeds of error-corrected quantum operations are expected to remain much slower than classical operations [58, 59]. This is an entirely distinct argument from ours which further casts doubt on the usefulness of Grover's algorithm as a practical tool.

Beyond the above rather negative results, our quantum inspired algorithm may also lead to more positive results. For instance, while we have focused on exact calculations of the the quantum circuit amplitudes, an interesting possibility would be to construct the MPS from *approximate* calculations of the amplitudes $\langle \beta | \Psi_w \rangle$ using standard MPS compression techniques. It is unclear if the resulting MPS would provide an efficient heuristic for solving the Grover problem.

[1] P. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124–134.

[2] S. Beauregard, Circuit for shor's algorithm using 2n+3 qubits 10.48550/ARXIV.QUANT-PH/0205095 (2002).

[3] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem (1995).

[4] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, Phys. Rev. Lett. **103**, 150502 (2009).

[5] L. K. Grover, Quantum mechanics helps in searching for a needle in a haystack, Phys. Rev. Lett. **79**, 325 (1997).

[6] L. K. Grover, From Schrödinger's equation to the quantum search algorithm, Pramana - J. Phys. **56**, 333 (2001).

[7] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, Quantum amplitude amplification and estimation, Contemporary Mathematics **305**, 53 (2002).

[8] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, Strengths and weaknesses of quantum computing, SIAM Journal on Computing **26**, 1510 (1997), https://doi.org/10.1137/S0097539796300933.

[9] W. P. Baritompa, D. W. Bulger, and G. R. Wood, Grover's quantum algorithm applied to global optimization, SIAM Journal on Optimization **15**, 1170 (2005).

[10] A. Y. Wei, P. Naik, A. W. Harrow, and J. Thaler, Quantum algorithms for jet clustering, Phys. Rev. D **101**, 094015 (2020).

[11] C. Dürr, M. Heiligman, P. HOyer, and M. Mhalla, Quantum query complexity of some graph problems, SIAM Journal on Computing **35**, 1310 (2006).

[12] N. Stamatopoulos, D. J. Egger, Y. Sun, C. Zoufal, R. Iten, N. Shen, and S. Woerner, Option Pricing using Quantum Computers, Quantum **4**, 291 (2020).

[13] H. Ramesh and V. Vinay, String matching in o(n+m) quantum time, Journal of Discrete Algorithms **1**, 103 (2003), combinatorial Algorithms.

[14] E. Aïmeur, G. Brassard, and S. Gambs, Quantum speedup for unsupervised learning, Machine Learning **90**, 261 (2013).

[15] A. Kapoor, N. Wiebe, and K. Svore, Quantum perceptron models, in *Advances in Neural Information Processing Systems*, Vol. 29, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016).

[16] D. Dong, C. Chen, H. Li, and T.-J. Tarn, Quantum reinforcement learning, IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics) **38**, 1207 (2008).

[17] A. Dewes, R. Lauro, F. R. Ong, V. Schmitt, P. Milman, P. Bertet, D. Vion, and D. Esteve, Quantum speeding-up of computation demonstrated in a superconducting two-qubit processor, Phys. Rev. B **85**, 140503 (2012).

[18] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe, Complete 3-qubit grover search on a programmable quantum computer, Nature communications **8**, 1 (2017).

[19] A. Mandviwalla, K. Ohshiro, and B. Ji, Implementing Grover's Algorithm on the IBM Quantum Computers, in *2018 IEEE International Conference on Big Data (Big Data)* (2018) pp. 2531–2537.

[20] B. Pokharel and D. Lidar, Better-than-classical grover search via quantum error detection and suppression (2022).

[21] S. Östlund and S. Rommer, Thermodynamic limit of density matrix renormalization, Phys. Rev. Lett. **75**, 3537 (1995).

[22] G. Vidal, Efficient classical simulation of slightly entangled quantum computations, Phys. Rev. Lett. **91**, 147902 (2003).

[23] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, Quantum Info. Comput. **7**, 401–430 (2007).

[24] I. P. McCulloch, From density-matrix renormalization group to matrix product states, Journal of Statistical Mechanics: Theory and Experiment **2007**, P10014 (2007).

[25] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Matrix product density operators: Simulation of finite-temperature and dissipative systems, Phys. Rev. Lett. **93**, 207204 (2004).

[26] G. M. Crosswhite and D. Bacon, Finite automata for caching in matrix product algorithms, Phys. Rev. A **78**, 012356 (2008).

[27] B. Pirvu, V. Murg, J. I. Cirac, and F. Verstraete, Matrix product operator representations, New Journal of Physics **12**, 025012 (2010).

[28] M. P. Zaletel, R. S. K. Mong, C. Karrasch, J. E. Moore, and F. Pollmann, Time-evolving a matrix product state with long-ranged interactions, Phys. Rev. B **91**, 165112 (2015).

[29] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information (10th Anniversary edition)* (Cambridge University Press, 2016).

[30] A. J. Ferris and G. Vidal, Perfect sampling with unitary tensor networks, Phys. Rev. B **85**, 165146 (2012).

[31] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, New Journal of Physics **12**, 055026 (2010).

[32] T. Ayral, T. Louvet, Y. Zhou, C. Lambert, E. M. Stoudenmire, and X. Waintal, A density-matrix renormalization group algorithm for simulating quantum circuits with a finite fidelity (2022).

[33] I. Oseledets and E. Tyrtyshnikov, Tt-cross approximation for multidimensional arrays, Linear Algebra and its Applications **432**, 70 (2010).

[34] D. Savostyanov and I. Oseledets, Fast adaptive interpolation of multi-dimensional arrays in tensor train format, in *The 2011 International Workshop on Multidimensional (nD) Systems* (IEEE, 2011) pp. 1–8.

[35] D. V. Savostyanov, Quasioptimality of maximum-volume cross interpolation of tensors, Linear Algebra and its Applications **458**, 217 (2014).

[36] S. Dolgov and D. Savostyanov, Parallel cross interpolation for high-precision calculation of high-dimensional integrals, Computer Physics Communications **246**, 106869 (2020).

[37] Y. Núñez Fernández, M. Jeannin, P. T. Dumitrescu, T. Kloss, J. Kaye, O. Parcollet, and X. Waintal, Learning feynman diagrams with tensor trains, Phys. Rev. X **12**, 041018 (2022).

[38] F. Massacci and L. Marraro, Logical cryptanalysis as a sat problem, Journal of Automated Reasoning **24**, 165 (2000).

[39] I. Mironov and L. Zhang, Applications of sat solvers to cryptanalysis of hash functions, in *International Conference on Theory and Applications of Satisfiability Testing* (Springer, 2006) pp. 102–115.

[40] L. Perron and V. Furnon, Or-tools.

[41] F. Corblin, L. Bordeaux, Y. Hamadi, E. Fanchon, and L. Trilling, A sat-based approach to decipher gene regulatory networks, Integrative Post-Genomics, RIAMS, Lyon (2007).

[42] M. Fishman, S. R. White, and E. M. Stoudenmire, The ITensor Software Library for Tensor Network Calculations, SciPost Phys. Codebases , 4 (2022).

[43] M. Fishman, S. R. White, and E. M. Stoudenmire, Codebase release 0.3 for ITensor, SciPost Phys. Codebases , 4 (2022).

[44] J. Chen, F. Zhang, C. Huang, M. Newman, and Y. Shi, Classical simulation of intermediate-size quantum circuits (2018), arXiv:1805.01450 [quant-ph].

[45] J. Gray and S. Kourtis, Hyper-optimized tensor network contraction, Quantum **5**, 410 (2021).

[46] F. Pan, K. Chen, and P. Zhang, Solving the sampling problem of the sycamore quantum circuits, Phys. Rev. Lett. **129**, 090502 (2022).

[47] S. Kourtis, C. Chamon, E. R. Mucciolo, and A. E. Ruckenstein, Fast counting with tensor networks, SciPost Phys. **7**, 060 (2019).

[48] D. Koch, A. Torrance, D. Kinghorn, S. Patel, L. Wessing, and P. M. Alsing, Simulating quantum algorithms using fidelity and coherence time as principle models for error (2019), arxiv:1098.04229.

[49] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[50] K. Zhang and V. E. Korepin, Depth optimization of quantum search algorithms beyond grover's algorithm, Phys. Rev. A **101**, 032346 (2020).

[51] A. SaiToh, A multiprecision C++ library for matrix-product-state simulation of quantum computing: Evaluation of numerical errors, in *Journal of Physics: Conference Series*, Vol. 454 (IOP Publishing, 2013) p. 012064.

[52] A. Montanaro, Quantum algorithms: an overview, npj Quantum Information **2**, 1 (2016).

[53] T. Schoning, A probabilistic algorithm for k-sat and constraint satisfaction problems, in *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)* (IEEE, 1999) pp. 410–414.

[54] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).

[55] X. Waintal, What determines the ultimate precision of a quantum computer, Phys. Rev. A **99**, 042318 (2019).

[56] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Elucidating reaction mechanisms on quantum computers, Proceedings of the National Academy of Sciences **114**, 7555 (2017), https://www.pnas.org/doi/pdf/10.1073/pnas.1619152114.

[57] C. Chamon and E. R. Mucciolo, Virtual parallel computing and a search algorithm using matrix product states, Phys. Rev. Lett. **109**, 030503 (2012).

[58] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven, Focus beyond quadratic speedups for error-corrected quantum advantage, PRX Quantum **2**, 010103 (2021).

[59] M. Troyer, Disentangling hype from reality: Achieving practical quantum advantage., Q2B Practical Quantum Computing Conference, 2020.

[60] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).

[61] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, The complexity of nisq (2022).

[62] D. Gottesman, The heisenberg representation of quantum computers 10.48550/ARXIV.QUANT-PH/9807006 (1998).

[63] A. Garcia-Saez and J. I. Latorre, An exact tensor network for the 3SAT problem 10.48550/ARXIV.1105.3201 (2011).

[64] L. Ma and C. Yang, Low rank approximation in simulations of quantum algorithms, Journal of Computational Science **59**, 101561 (2022).

[65] O. Regev and L. Schiff, Impossibility of a quantum speed-up with a faulty oracle, in *International Colloquium on Automata, Languages, and Programming* (Springer, 2008) pp. 773–781.

[66] M. V. d. Nest, Classical simulation of quantum computation, the gottesman-knill theorem, and slightly beyond 10.48550/ARXIV.0811.0898 (2008).

[67] S. Yang, W. Zi, B. Wu, C. Guo, J. Zhang, and X. Sun, Efficient quantum circuit synthesis for sat-oracle with limited ancillary qubit (2021), arXiv:2101.05430.

Appendix A: Quantum Circuits for Oracle and Diffusion Operators

We discuss here some details of the quantum circuit used in the *simulations* of GA, not to be confused with the QiGA calculations.

For obtaining the entanglement entropy plot Fig. 2, which shows the entanglement not only between Grover iterations but also inside of the substeps of the oracle and diffusion circuits, we used the circuits shown in the figure below. These circuits are for the case where the target bitstring w is known. (The target for the diffusion operator is always known, since it can be implemented as the oracle which targets $|000...0\rangle$ pre- and post-processed by a Hadamard gate on each qubit.) The circuit pattern below uses at most three-qubit gates. This is in contrast to the implementation sometimes seen where the oracle is implemented by a "multi-controlled" gate, which is equivalent to our observation in the main text that the oracle can always in principle be implemented by a rank $\chi = 2$ MPO (for the case of a single targeted bitstring w).

For the specific case of four qubits and a target bitstring $w = 1011$, the oracle circuit pattern used for the Fig. 2 simulation was:

For Grover's algorithm on n qubits, the operator above is initialized by preparing $n+1$ additional ancilla qubits, n in the $|0\rangle$ state and the $n+1$ qubit in the $|-\rangle = H|1\rangle$ state. By using Toffoli gates acting on the upper and lower registers, the ancillas are flipped to indicate that each qubit of the target bitstring has been found (upper control) and that all previous bits have been found (lower control). If so, the next ancilla qubit is flipped to 1.

The notation of the white circle for the upper control of the second Toffoli gate stands for an "anti-control" meaning the gate acts only if that qubit is $|0\rangle$. This kind of control can be viewed as shorthand for:

that is, NOT gates on either side of a regular control.

At the center of the circuit, if the n^{th} ancilla qubit is 1, then the $n+1$ ancilla is acted on by a NOT gate which results in a minus sign ("phase kickback mechanism") for the amplitude of the state with the target bitstring in the upper register. Lastly, the Toffoli gates are acted in reverse order to "uncompute" the ancilla register, restoring all of the ancillas to their initial product state. It is easy to check by inspection that applying the above circuit to $|b\rangle|0000-\rangle_A$ leads to $\pm|b\rangle|0000-\rangle_A$ depending on wither $b = w$ or not.