

Received 4 June 2024; revised 6 October 2024; accepted 15 October 2024; date of publication 22 October 2024;
date of current version 20 November 2024.

Digital Object Identifier 10.1109/TQE.2024.3484650

Fixed-Point Grover Adaptive Search for Quadratic Binary Optimization Problems

ÁKOS NAGY¹, JAIME PARK², CINDY ZHANG³, ATITHI ACHARYA⁴,
AND ALEX KHAN⁵

¹BEIT Canada, Toronto, ON M5G 1L7, Canada

²Vanderbilt University, Nashville, TN 37240, USA

³Cindy Zhang resides in Piscataway, NJ 08854, USA

⁴Rutgers, The State University of New Jersey, New Brunswick, NJ 08901, USA

⁵National Quantum Laboratory, University of Maryland, College Park, MD 20740, USA

Corresponding author: Ákos Nagy (e-mail: contact@akosnagy.com).

ABSTRACT In this article, we study a Grover-type method for quadratic unconstrained binary optimization (QUBO) problems. For an n -dimensional QUBO problem with m nonzero terms, we construct a marker oracle for such problems with a tunable parameter, $\Lambda \in [1, m] \cap \mathbb{Z}$. At $d \in \mathbb{Z}_+$ precision, the oracle uses $O(n + \Lambda d)$ qubits and has total depth of $O(\frac{m}{\Lambda} \log_2(n) + \log_2(d))$ and a non-Clifford depth of $O(\frac{m}{\Lambda})$. Moreover, each qubit is required to be connected to at most $O(\log_2(\Lambda + d))$ other qubits. In the case of a maximum graph cuts, as $d = 2 \lceil \log_2(n) \rceil$ always suffices, the depth of the marker oracle can be made as shallow as $O(\log_2(n))$. For all values of Λ , the non-Clifford gate count of these oracles is strictly lower (at least by a factor of ~ 2) than previous constructions. Furthermore, we introduce a novel fixed-point Grover adaptive search for QUBO problems, using our oracle design and a hybrid fixed-point Grover search, motivated by the works of Boyer et al. (1988) and Li et al. (2019). This method has better performance guarantees than previous Grover adaptive search methods. Some of our results are novel and useful for any method based on the fixed-point Grover search. Finally, we give a heuristic argument that, with high probability and in $O(\frac{\log_2(n)}{\sqrt{\epsilon}})$ time, this adaptive method finds a configuration that is among the best $\epsilon 2^n$ ones.

INDEX TERMS Quantum computing, quantum algorithm, optimization.

I. INTRODUCTION

Grover [12] introduced a quantum search algorithm that can be viewed as an unstructured search algorithm whose query complexity scales with the inverse square root of the ratio of the number of marked items to all items. One drawback of Grover's algorithm is the "soufflé problem," that is, the fact that one needs to know the exact ratio of marked items to compute the correct number of queries prescribed by the algorithm. Not enough or too many queries can "under/overcook" the quantum state. This problem was partially remedied by the fixed-point Grover search (FPGS) of Yoder et al. [25]. The FPGS is similar to Grover's algorithm in many ways: both are quantum amplitude amplification methods, both have the same query complexity, and they can typically be implemented using the same oracles. Geometrically, the difference is that while Grover's algorithm uses reflections in the real plane spanned by the initial and target states, the FPGS uses phase gates in the complexification of that subspace. The benefit of using FPGS is that one no longer needs to know the exact number of marked states; it is enough

to only have a lower bound for that number. An important caveat is that the lower this bound is, the more queries are required, but one cannot overcook the quantum state. The query complexity of FPGS scales with the inverse square root of *lower bound used for the ratio of marked states*. Thus, if this lower bound is close to the exact ratio, this latter fact can superficially be viewed as retaining the "quadratic speedup." However, that is somewhat misleading, as finding lower bounds might still be hard. In order to circumvent this issue, many adaptive methods have been designed that retain the quadratic speedup; the most important ones for our work are [2] and [14]. In particular, we refine the ideas of Li et al. [14] to get an adaptive search method that, to our knowledge, has the best current performance guarantee in terms of total expected oracle calls and can be applied to arbitrary boolean quantum search.

A quadratic unconstrained binary optimization (QUBO) problem consists of a (real-valued) quadratic polynomial on n Boolean variables, and thus, it can be described, uniquely, up to an overall additive constant, by an n -by- n

upper-triangular real matrix. Such problems provide some of the most interesting NP-complete problems, such as cluster analysis, maximum graph cuts, and the Ising model. Algorithms to find approximate maxima of QUBOs have been extensively studied for a long time via classical methods. Recently, classical–quantum hybrid methods have been proposed, where the parameters of the quantum algorithm are classically optimized. The two main types of such algorithms are of quantum approximate optimization algorithm (QAOA) [7], [11], [16], [23] and Grover’s algorithm [9], [10], [21]. Compared to QAOA, Grover-type algorithms have a certain set of advantages and drawbacks. The promise of QAOA is that, using an easily implementable low-depth circuit, one can prepare a quantum state whose dominant components in the computational basis correspond to high value configurations of the QUBO problem. On the other hand, Grover-type algorithms tend to have more complex circuits, and the amplification is only sensitive to whether or not a configuration is above a given threshold, which is a hyperparameter of the method. However, what makes Grover-type methods still appealing is the fact that they usually require much less classical optimizations.

This article explores a Grover-type method for QUBO problems. Gilliam et al. [9] proposed a marker oracle design for QUBO problems: given an instance of an integer QUBO problem, say, by n -by- n upper-triangular integer matrix, Q , and a threshold $y \in \mathbb{Z}$, their oracle marks states $|x\rangle$ such that $x^T Q x > y$. This is precisely the ingredient needed to run Grover-type search algorithms to find configurations with values above the threshold. Furthermore, in the same paper, Gilliam et al. propose a *Grover adaptive search* for QUBO problems, based on the adaptive Grover search of Boyer et al. [2]. We explain their construction for the marker oracle in Section II-A and their Grover adaptive search in Section IV. This article was inspired by the results of Gilliam et al., and our key results can be viewed as improvements on both their oracle design and adaptive method. Most notably, our method has improved depth and total gate counts for both for Clifford and non-Clifford gates. Furthermore, we provide a variant of the oracle with a tunable number of ancillas, which uses the same (amount and type of) non-Clifford gates and has reduced the total circuit depth, while having only an $O(1)$ -fold increase in the Clifford gate count. In the shallowest case, the depth is $O(\log_2(d))$ —in particular, independent of n —and uses at most $\frac{n^2+n+2}{2}d$ ancillas. While we focus on unconstrained problems, constraints can readily be handled, as in [9, Fig. 7]. Finally, our methods can be combined with other instance-specific circuit optimization techniques; for example [21].

Organization of this article: The rest of this article is organized as follows. In Section II, we introduce QUBO problems and present our design for QUBO quantum dictionary encoders. In Section III, we introduce the FPGS and, using our oracle, apply it to QUBO problems. Section IV defines and studies the adaptive version of the FPGS. In Section V, we test and compare our method to the Grover adaptive search

for QUBO of Gilliam et al. [9]. Finally, Section VI concludes this article.

II. QUANTUM DICTIONARIES AND QUBO PROBLEMS

Given a function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$, the associated (unconstrained) binary optimization problem is the task of finding an element $x \in \mathbb{F}_2^n$ such that $f(x)$ is maximal. Many interesting binary optimization problems are quadratic, such as finding maximum graph cuts or the Max 2-SAT problems, and thus, most of the contemporary research centers around QUBO problems.

The first main contribution of this article is an oracle design for QUBO problems. More concretely, we define and construct (projective) encoder oracles for quadratic functions $f : \mathbb{F}_2^n \rightarrow [-2^{d-1}, 2^{d-1}] \cap \mathbb{Z} \cong \mathbb{F}_2^d$ below. These oracles have applications in, for example, Grover-type algorithms and threshold QAOA [3], [11]. While designs for such oracles have already existed, cf. [9], ours has better circuit depth and size, both for Clifford and non-Clifford gates. Furthermore, we construct a variant with quadratically many ancillas, as well as with logarithmic depth. A generalization of these oracles is given in [18], which can be applied to polynomials of arbitrary degrees.

In this article, we use two common labelings for the computational basis. First, given a bit string $x = (x_0, x_1, \dots, x_{n-1}) \in \mathbb{F}_2^n$, we write $|x\rangle := |x_0\rangle|x_1\rangle \cdots |x_{n-1}\rangle$. Similarly, if $\bar{y} \in \mathbb{Z}$ and $d \in \mathbb{Z}_+$, then let $y = (y_0, y_1, \dots, y_{d-1}) \in \mathbb{F}_2^d$ be such that $\bar{y} \equiv \sum_{j=0}^{d-1} y_j 2^{d-1-j}$ and $|\bar{y}\rangle_d := |y\rangle$, that is, the state is the \bar{y} th(mod 2^d) computational basis element in the usual ordering.

With the above notation, a *quantum dictionary encoder*, $U_{f,d}$, corresponding to a function, $f : \mathbb{F}_2^n \rightarrow \mathbb{Z}$, is any oracle on $n + d$ qubits defined via

$$U_{f,d}|x\rangle|y\rangle_d := |x\rangle|y + f(x)\rangle_d.$$

More generally, $U_{f,d}$ is a *projective* quantum dictionary encoder, if

$$U_{f,d}|x\rangle|y\rangle_d := e^{i\alpha_{f,x,y}}|x\rangle|y + f(x)\rangle_d$$

for some $\alpha_{f,x,y} \in \mathbb{R}$.

Remark 1: In many quantum algorithms, including ours, quantum dictionary encoders appear in pairs in the form

$$\cdots \text{---} \boxed{U_{f,d}} \text{---} \boxed{D} \text{---} \boxed{U_{f,d}^\dagger} \text{---} \cdots$$

where D is diagonal in the computational basis. Thus, using a projective encoder yields the same outcome, since

$$\begin{aligned} U_{f,d}^\dagger D U_{f,d} |x\rangle|y\rangle_d &= e^{i\alpha} U_{f,d}^\dagger D |x\rangle|y + f(x)\rangle_d \\ &= \eta U_{f,d}^\dagger (e^{i\alpha} |x\rangle|y + f(x)\rangle_d) \\ &= \eta |x\rangle|y\rangle_d \end{aligned}$$

where η is the appropriate eigenvalue of D .

A. QUBO ENCODER

Let QFT_d denote the d -qubit quantum Fourier transform (QFT). In order to construct $U_{f,d}$ for a quadratic f , we use Draper’s QFT-based shift operator [6], which we summarize in the following theorem.

Theorem 1: Let $\tilde{U}_{f,d}$ be defined as

$$\tilde{U}_{f,d}|x\rangle|y\rangle_d = \exp\left(\frac{2\pi i}{2^d} f(x)y\right) |x\rangle|y\rangle_d. \quad (1)$$

Then

$$U_{f,d} = (\mathbb{1} \otimes \text{QFT}_d^\dagger) \circ \tilde{U}_{f,d} \circ (\mathbb{1} \otimes \text{QFT}_d). \quad (2)$$

is a quantum dictionary encoder for f .

Theorem 1 allows the computation of $f(x)$ to be done in the “phase space” and reduces the construction of $U_{f,d}$ to the construction of $\tilde{U}_{f,d}$, which we do later.

Before we present the design of $\tilde{U}_{f,d}$, let us make a few definitions. Let now $f : \mathbb{F}_2^n \rightarrow [-2^{d-1}, 2^{d-1}] \cap \mathbb{Z}$ be quadratic, and write

$$f(x) = f(0) + \sum_{j,k=1}^n Q_{jk}x_jx_k$$

where $(Q_{jk})_{j,k \in [n]}$ is a real symmetric matrix. Let \oplus be the addition modulo 2 (that is, logical XOR). Then, we have

$$x_jx_k = \frac{1}{2}(x_j + x_k - x_j \oplus x_k)$$

and thus, if we set $q_j := \sum_k Q_{jk}$, then

$$f(x) = f(0) + \sum_{j=0}^{n-1} q_jx_j - \sum_{j=0}^{n-2} \sum_{k=j+1}^{n-1} Q_{jk}x_j \oplus x_k. \quad (3)$$

For $k \in \mathbb{R}$, let us define the d -qubit operators

$$\mathcal{P}_d(k) := \otimes_{j=0}^{d-1} \begin{pmatrix} 1 & 0 \\ 0 & e^{\frac{\pi i}{2^j}k} \end{pmatrix}$$

or, equivalently

$$\mathcal{P}_d(k)|y\rangle_d := e^{\frac{2\pi i}{2^d}ky}|y\rangle_d.$$

Let $\tilde{U}_\emptyset := \mathbb{1} \otimes \mathcal{P}_d(q_\emptyset)$, where

$$q_\emptyset := f(0) + \frac{1}{4}\text{tr}(Q) + \frac{1}{4}\text{sum}(Q)$$


and thus

$$\tilde{U}_\emptyset|x\rangle|y\rangle_d = e^{i\alpha_\emptyset y}|x\rangle|y\rangle_d$$

where

$$\alpha_{j,y} = \frac{2\pi}{2^d}q_\emptyset y, \quad (4)$$


Let the gates \tilde{U}_j be



where

$$\alpha_{j,y} = \frac{2\pi}{2^d}q_j \left(x_jy - \frac{x_j}{2}(2^d - 1) - \frac{y}{2}\right) \quad (5)$$


and the gates \tilde{U}_{jk} be



where

$$\alpha_{j,k,y} = \frac{2\pi}{2^d}Q_{jk} \left(- (x_j \oplus x_k)y + \frac{x_j \oplus x_k}{2}(2^d - 1) + \frac{y}{2}\right). \quad (6)$$

Note that the *fanout*, or *multitarget* CNOT gates



can be implemented using $\approx 2d$ many CNOT gates with $\approx 2 \log_2(d)$ depth; cf. [15, Appendix B.1]. With these definitions, let

$$\tilde{U}_{f,d} := \tilde{U}_\emptyset \circ \prod_{j=0}^{n-1} \tilde{U}_j \circ \prod_{j=0}^{n-2} \prod_{k=j+1}^{n-1} \tilde{U}_{jk}.$$

Theorem 2: For each $x \in \mathbb{F}_2^n$ and $y \in \mathbb{F}_2^d$, we have that

$$\tilde{U}_{f,d}|x\rangle|y\rangle_d = \exp\left(\frac{2\pi i}{2^d} f(x)y + i\alpha_{f,x,y}\right) |x\rangle|y\rangle_d$$

with $\alpha_{f,x,y} := \frac{\pi(1-2^d)}{2^d}(f(x) - f(0))$. Thus, using (2) and (3) and the aforementioned equation, we get that

$$U_{f,d} := (\mathbb{1} \otimes \text{QFT}_d^\dagger) \circ (\tilde{U}_{f,d}) \circ (\mathbb{1} \otimes \text{QFT}_d).$$

is a projective quantum dictionary encoder for f , with the same *garbage phases* $\alpha_{f,x,y}$.

Proof: The verifications of (4)–(6) is straightforward, and they yield

$$\tilde{U}_{f,d}|x\rangle|y\rangle_d = \exp\left(\frac{2\pi i}{2^d} \varphi_{x,y}\right) |x\rangle|y\rangle_d$$

where

$$\begin{aligned} \varphi_{x,y} := & q_\emptyset y + \sum_{j=0}^{n-1} q_jx_jy - \sum_{j=0}^{n-2} \sum_{k=j+1}^{n-1} Q_{jk}x_j \oplus x_ky \\ & - \frac{2^d - 1}{2} \left(\sum_{j=0}^{n-1} q_jx_j + \sum_{j=0}^{n-2} \sum_{k=j+1}^{n-1} Q_{jk}x_j \oplus x_k \right) \end{aligned}$$

$$\begin{aligned}
 & -\frac{1}{2} \left(\sum_{j=0}^{n-1} q_j y - \sum_{j=0}^{n-2} \sum_{k=j+1}^{n-1} Q_{jk} y \right) \\
 & = (q_\emptyset + f(x) - f(0))y \\
 & \quad - \frac{2^d - 1}{2} (f(x) - f(0)) \\
 & \quad - \frac{1}{2} \left(\text{sum}(Q) - \frac{1}{2} (\text{sum}(Q) - \text{tr}(Q)) \right) y \\
 & = f(x)y - \frac{2^d - 1}{2} (f(x) - f(0))
 \end{aligned}$$

which concludes the proof. \square

Remark 2: We have so far restricted our discussion to integer-valued QUBO problems. However, the definition of $\mathcal{P}_d(k)$ makes sense for any $k \in \mathbb{R}$, not just for integers. When $k \notin \mathbb{Z}$, and we let $[k] \in \mathbb{Z}$ and $\tilde{k} \in (0, 1)$ be its integer and fractional parts, respectively, then we get

$$\text{QFT}_d^\dagger \circ \mathcal{P}_d(k) \circ \text{QFT}_d |y\rangle_d = \sum_{z=0}^{2^d-1} \frac{e^{2\pi\tilde{k}i} - 1}{e^{2\pi i 2^d(y-z+k)} - 1} |z\rangle_d.$$

This yields a Fejér-type distribution on the bit strings with the probability of measuring either $y + [k]$ or $y + [k] + 1$ being at least $\frac{8}{\pi^2} \approx 81\%$; cf. [9, Appendix B.2].

Furthermore, since \mathbb{F}_2^n is a finite set, small perturbations of the coefficient matrix do not change the order of the values, meaning that small “roundings” of the coefficients yield equivalent optimization problems.

Thus, the same oracle design can be, in some cases, used for real-valued QUBO problems.

B. USING ANCILLAS TO REDUCE DEPTH

Let now m be the number of nonzero terms in (3) and let $1 \leq \Lambda \leq m$ be an integer. We claim, without proof, that a circuit for $U_{f,d}$ can be constructed (with the same garbage phases as in Theorem 2) using $(\Lambda - 1)d$ ancillas, where the R_Z depth in general is $\lceil \frac{m}{\Lambda} \rceil$ (assuming that Λd many R_Z gates can be run in parallel) and the total depth is $O(\frac{m}{\Lambda} \max(\{1, \log_2(\frac{\Lambda}{n})\}) + \log_2(d))$. We only present the construction in detail in the most extreme case, when m is maximal, that is, when $m = 1 + n + \binom{n}{2} = \frac{n^2+n+2}{2}$. We also focus on the shallowest case, which is also the case with the most ancillas, that is $\Lambda = m$. Versions for more sparse functions and different values of Λ can then easily be constructed, and we address this issue after our proof for the $m = \Lambda = \frac{n^2+n+2}{2}$ case, when the R_Z depth is 1 and the total depth is $O(\log_2(n) + \log_2(d))$.

Let us assume now that $m = \Lambda = \frac{n^2+n+2}{2}$ and start with a quantum circuit on $n + md$ qubits. Our goal is still the realization of $\tilde{U}_{f,d}$, as in (1), on the first $n + d$ qubits. We assume that the last $(m - 1)d$ qubits are all initialized in zero. Let U_{QUBO} be defined as

$$\begin{aligned}
 & U_{\text{QUBO}} |x\rangle |y\rangle_d |0\rangle_{(m-1)d} \\
 & := |x\rangle |y\rangle_d \bigotimes_{j=0}^{n-1} \bigotimes_{a=0}^{d-1} \left(|y_a \oplus x_j\rangle \bigotimes_{k=j+1}^{n-1} |y_a \oplus x_j \oplus x_k\rangle \right).
 \end{aligned}$$

Note that U_{QUBO} is independent of f .

Theorem 3: U_{QUBO} can be constructed using only CNOT gates with depth $O(\log_2(n) + \log_2(d))$.

Proof: Note that each y_a needs to be added to m qubits, and no two of them need to be added to the same one; thus, using again the fanout gates of [15, Appendix B.1], this can be achieved using $\sim 2md$ CNOT gates with a depth of $\sim 2 \log_2(m) \lesssim 4 \log_2(n)$. After this, each x_j needs to be added to nd terms, albeit some of them overlap, but we can still achieve this with two parallel sets of fanout gates; first (for each j) let $|x_j\rangle$ control a fanout gate that targets the qubits we want to end up in the state either $|y_a \oplus x_j\rangle$ or $|y_a \oplus x_j \oplus x_k\rangle$, where $j < k$, and in the second set, let $|x_j\rangle$ control a fanout gate that targets the qubits we want to end up in the state $|y_a \oplus x_j \oplus x_k\rangle$, where $j > k$. In each round, the fanout gates have no overlapping controls or targets, thus can be run parallel, and the maximum depth in each is $\approx 2 \log_2(n)$. Thus, the total depth is $\lesssim (4 + 2 + 2) \log_2(n) = 8 \log_2(n)$. \square

Now, let

$$\mathbb{P}_f := \mathbb{1} \otimes \mathcal{P}_d(q_\emptyset) \bigotimes_{j=0}^{n-1} \left(\mathcal{P}_d\left(-\frac{q_j}{2}\right) \bigotimes_{k=j+1}^{n-1} \mathcal{P}_d\left(\frac{Q_{jk}}{2}\right) \right).$$

Note that \mathbb{P}_f contains the exact same number and type of R_Z rotations as the design in Section II-A, but now they all run in parallel.

Lemma 1: Let $U_{f,d} := U_{\text{QUBO}} \circ \mathbb{P}_f \circ U_{\text{QUBO}}$. Then, for each $x \in \mathbb{F}_2^n$ and $y \in F_2^d$, we have that

$$U_{f,d} |x\rangle |y\rangle_d |0\rangle_{(m-1)d} = e^{i\alpha_{f,x,y}} |x\rangle |y + f(x)\rangle_d |0\rangle_{(m-1)d}$$

where $\alpha_{f,x,y}$ is the same as in Theorem 2.

Proof: Analogous to the proof of Theorem 2. \square

Remark 3: Note that the garbage phases can also be eliminated by applying phase gates with angles $\frac{2\pi}{2^d} q_i$ to the qubits $|x_j\rangle$ and controlled phase gates with angles $-\frac{2\pi}{2^d} Q_{jk}$ to the qubits $|x_j\rangle |x_k\rangle$ (for all relevant values of j and k). However, for our current purposes, a projective encoder is already sufficient, as explained in Section II.

III. FPGS FOR QUBO

FPGS [25] is a variant of Grover’s search algorithm that retains the original version’s query complexity while not suffering from the soufflé problem (more details on this are provided later in this article). In this section, we introduce

the algorithm of Yoder et al. [25] and use our oracle design from Section II-A to implement FPGS for QUBO problems. Furthermore, in Section III-A, we study the problem of optimal parameter choice for FPGS, and in Section III-B, we analyze the time complexity of FPGS for QUBO problems using our oracle design.

As opposed to Grover’s algorithm, where the only input is the set of *good/*marked/*target* configurations, $\mathcal{M} \subset \mathbb{F}_2^n$, the FPGS algorithm requires an additional one that can be chosen to be either the target probability or the query complexity. We assume that the target probability, $\mathbb{P}_{\text{target}} \in (0, 1)$, is given, that is, we want the algorithm to produce a state that, when measured in the computational basis, yields a marked state with probability at least $\mathbb{P}_{\text{target}}$.

The price of this flexibility (and of the elimination of half of the soufflé problem) is that one needs to implement not only a pair of oracles but also two families of them, called $S_s(\alpha)$ and $S_t(\beta)$ (where the subscripts refer to the *start* and *target* states, and α and β are real parameters), with the following properties. Let

$$|s\rangle := \frac{1}{\sqrt{2^n}} \sum_{x \in \mathbb{F}_2^n} |x\rangle$$

$$|t\rangle := \frac{1}{\sqrt{|\mathcal{M}|}} \sum_{x \in \mathcal{M}} |x\rangle.$$

Let $\lambda := |\langle s|t\rangle|^2 = \frac{|\mathcal{M}|}{2^n} \in (0, 1)$, and

$$|\bar{s}\rangle := \frac{\sqrt{\lambda}|s\rangle - |t\rangle}{\sqrt{1+\lambda}} \quad |\bar{t}\rangle := \frac{-|s\rangle + \sqrt{\lambda}|t\rangle}{\sqrt{1+\lambda}}.$$

Then, both $\{|s\rangle, |\bar{s}\rangle\}$ and $\{|t\rangle, |\bar{t}\rangle\}$ are orthonormal bases of the same complex plane. Now, we need $S_s(\alpha)$ and $S_t(\beta)$ to be unitaries such that on this plane, their actions are

$$S_s(\alpha)(A|s\rangle + B|\bar{s}\rangle) = e^{i\alpha}A|s\rangle + B|\bar{s}\rangle$$

$$S_t(\beta)(D|t\rangle + C|\bar{t}\rangle) = e^{i\beta}C|t\rangle + D|\bar{t}\rangle.$$

Let $G(\alpha, \beta) := S_s(\alpha)S_t(\beta)$.

Once in possession of such oracles and a target probability $\mathbb{P}_{\text{target}} \in (0, 1)$, the FPGS of [25] can be summarized as follows: for each $\mu \in (0, \lambda]$ and $\delta \in (0, \sqrt{1 - \mathbb{P}_{\text{target}}}]$, there is $l = l(\mu, \delta) \in \mathbb{Z}_+$, such that if, for all $j \in \{1, 2, \dots, l\}$, we set

$$\alpha_j := 2 \operatorname{arccot} \left(\tan \left(\frac{2\pi j}{2l+1} \right) \tanh \left(\frac{\operatorname{arccosh}(1/\delta)}{2l+1} \right) \right)$$

$$G_j := S_s(\alpha_j) S_t(\alpha_{l+1-j}).$$

Then

$$\mathbb{P}_{\text{success}} := |\langle t|G_l G_{l-1} \cdots G_1 |s\rangle|^2 \geq \mathbb{P}_{\text{target}}.$$

In fact, l and $\mathbb{P}_{\text{success}}$ can be computed via

$$l = \min \left\{ l' \mid \sqrt{1 - \lambda} T_{1/2l'+1}(1/\delta) > 1 \right\} \quad (7)$$

$$\mathbb{P}_{\text{success}} = 1 - \left(T_{2l+1} \left(\sqrt{1 - \lambda} T_{1/2l+1}(1/\delta) \right) \right)^2 \delta^2 \quad (8)$$

where T_l is the Chebyshev polynomial of the first kind. Now

$$\mathbb{P}_{\text{success}} \geq 1 - \delta^2 \geq \mathbb{P}_{\text{target}}.$$

The proof of the aforementioned statements can be found in [25].

Remark 4: Grover [13] also introduced a “fixed-point” version of his original method; however, this algorithm did not maintain the quadratic speedup.


Remark 5: The fixed-point property is crucial for our method, because the number of marked states for QUBO problems is unknown. This rules out the direct application of other known methods such as “quantum search with certainty” [24] or “nonboolean” versions of quantum search (see, e.g., [1] and [22]). This can be somewhat remedied by the quantum mean estimation method of [22]; it is unclear if such a combined method would not lose the computational advantages.¹

Furthermore, the proof of a fixed-point version of the nonboolean quantum amplitude amplification (NBQAA) method of [22] is currently prepared for publication by the first author, which will open up the possibility of combining the advantages of NBQAA with our adaptive algorithm for binary optimization.

We implement $S_s(\alpha)$ and $S_t(\beta)$ in a straightforward manner. If U_s is the state preparation oracle, that is, $|s\rangle = U_s|0\rangle$, and $\text{MCP}_n(\alpha)$ is the $(n - 1)$ -controlled phase gate on n qubits, then

$$S_s(\alpha) = U_s \circ X^{\otimes n} \circ \text{MCP}_n(\alpha) \circ X^{\otimes n} \circ U_s^\dagger$$

works. Note that MCP_n can be implemented with one clean ancillas and two multicontrolled NOT gates via



In turn, the n -controlled NOT gates can be implemented with a depth of $O(\log_2(n)^3)$ with no ancillas [4] or at depth $O(\log_2(n))$ and a single ancilla [20].

In general, the implementation of $S_t(\beta)$ is case specific. In our method in Section IV for QUBO problems, we first specify thresholds, $y \in \mathbb{Z}$, and want to mark states, $|x\rangle$, such that $f(x) > y$, that is, the set of marked states is $\mathcal{M}_{f,y} := \{x \in \mathbb{F}_2^n \mid y > f(x)\}$. We achieve this as follows. First replace f with $f_y(x) := y - f(x)$ and repeat one of the constructions of Section II-A to get $U_t := U_{f_y,d}$. Note that $|x\rangle$ is marked exactly when $f_y(x) < 0$, which happens exactly when the first qubit of $|f_y(x)\rangle_d$ is $|1\rangle$ (and otherwise it is $|0\rangle$). Now, to get $S_t(\beta)$, we apply both U_t and U_t^\dagger as follows:

Again, since $U_{f_y,d}$ is used in pairs with $U_{f_y,d}^\dagger$ in a way that the operator between them is diagonal, all garbage phases are uncomputed, and thus, a projective quantum dictionary encoder is sufficient.

¹Based on private communications with Shyamsundar [22].




FIGURE 1. δ -dependence of the expected time to find a marked state.

A. PARAMETER OPTIMIZATION FOR FIXED-POINT GROVER ADAPTIVE SEARCH

Recall that in the FPGS, the parameter δ (or, equivalently, $\mathbb{P}_{\text{target}}$) can be chosen arbitrarily within the interval $(0, 1)$. In this section, we show how to optimize it for our adaptive search.

First of all, just to illustrate our ideas, let us assume that λ is known and, thus, we choose $\mu = \lambda$. In order to run FPGS, we need to still choose δ , which then determines $l =: l_\delta$ and $\mathbb{P}_{\text{success}} =: \mathbb{P}_\delta$. Since the number of FPGS runs needed to succeed follows the geometric probability distribution, the expected amount of time required to find a solution is proportional to $\frac{l_\delta}{\mathbb{P}_\delta}$. Since we know that the query complexity scales as $\frac{1}{\sqrt{\lambda}}$, let us define $\tau_\delta := \sqrt{\lambda} \frac{l_\delta}{\mathbb{P}_\delta}$, and minimize τ_δ in the $\lambda \rightarrow 0^+$ limit. This is a nontrivial problem to solve analytically (as the formulas involve nondifferentiable ceiling functions), but we found that this optimization can be carried out numerically, yielding stable approximative values. An instance of this is given in Fig. 1.

We find that $\delta \approx 0.6049$ minimizes the expected time and yields $\tau \approx 0.8582$. As expected, this is higher than $\frac{\pi}{4} \approx 0.7854$, which corresponds to Grover's algorithm. In fact, Grover's algorithm can be achieved to $\tau \approx 0.69$, as in pointed out in [26, Sect. 1]. Thus, when the ratio of marked states is known, we should not use FPGS, but Grover's algorithm.

However, when λ is not known, Grover's algorithm cannot directly be applied. This was addressed in [2]: The authors propose a randomized algorithm that finds a marked state, even when λ is unknown, and they provide an upper bound $\tau \lesssim \frac{9}{2} = 2.25$. Here, "randomized" means that the number of queries is randomly chosen and τ is still defined as $\sqrt{\lambda}$ times the expected number of queries, as $\lambda \rightarrow 0^+$. This is the method used in [9]. Li et al. [14] proposed a similar method, using FPGS, that they called *deterministic*, in the sense that the number of queries is chosen deterministically (growing exponentially using a specified base), and this number surpasses the critical number of queries needed for FPGS [given in (7)] in $O(\frac{1}{\sqrt{\lambda}})$ time. However, they only prove $\tau \lesssim 5.643$ for their parameters. In the following, we refine their argument to (numerically) achieve $\tau \approx 1.433$. Thus, our method yields a deterministic method with the best known worst case complexities. The caveat is that we do not prove our bound analytically and just provide strong numerical evidence for it down to $\lambda \geq 2^{-40}$.

Remark 6: We assume that queries for FPGS and Grover's algorithm cost the same, which is not always the case: to implement arbitrary phase rotations, one often needs two marker calls (whereas a phase π shift only needs one) and one also needs to consume more magic states in a fault-tolerant (Clifford + T) implementation. This is not the case for our method; however, as in both cases, one needs two marker calls to uncompute the ancillas, and the non-Clifford overhead of implementing arbitrary phases is small compared to that of the marker oracles.


Note that often FPGS queries are assumed to cost twice as much as regular Grover queries, in which case our bound is somewhat weaker than that of [2].

In order to achieve $\tau \approx 1.433$, we mimic the ideas of Li et al. [14]; however, we use sharper bounds. More precisely, our method can be described as follows: Fix $\delta \in (0, 1)$, $\alpha > 1$, and let $\text{FPGS}_{\delta,l}$ be FPGS with parameter δ and l queries. Let $\ell = 1$ and $l_{\text{total}} = 0$. In the first round, run $\text{FPGS}_{\delta,\lceil \ell \rceil}$, add $\lceil \ell \rceil$ to l_{total} , and, if a marked state is found, then stop. If after each round, if no marked states have been found, then replace ℓ with $\alpha \ell$, run $\text{FPGS}_{\delta,\lceil \ell \rceil}$, add $\lceil \ell \rceil$ to l_{total} , and, if a marked state is found, then stop. Again, let $\tau_{\delta,\alpha}$ be the $\sqrt{\lambda}$ time the expected value of l_{total} . We provide evidence for the following conjecture.

Conjecture 1: Let $\delta = 0.4038$ and $\alpha = 1.975$. Then, $\tau \leq 1.434$.

Numerical evidence for Conjecture 1: We can write

$$\begin{aligned} \tau_{\delta,\alpha} &= \sqrt{\lambda} \mathbb{E}[l_{\text{total}}] \\ &= \sqrt{\lambda} \sum_{s=1}^{\infty} \left(\sum_{r=1}^s \lceil \alpha^{r-1} \rceil \right) \mathbb{P} \left(\text{FPGS}_{\delta,\lceil \alpha^{s-1} \rceil} \text{ succeeds first} \right) \\ &= \sqrt{\lambda} \sum_{s=1}^{\infty} \lceil \alpha^{s-1} \rceil \sum_{r=s}^{\infty} \mathbb{P} \left(\text{FPGS}_{\delta,\lceil \alpha^{r-1} \rceil} \text{ succeeds first} \right) \end{aligned}$$


FIGURE 2. Phase portrait of $\tau_{\delta, \alpha}$ at $\lambda = 2^{-40}$.

$$= \sqrt{\lambda} \sum_{s=1}^{\infty} \lceil \alpha^{s-1} \rceil \prod_{r=1}^{s-1} \mathbb{P}(\text{FPGS}_{\delta, \lceil \alpha^{r-1} \rceil} \text{ fails}).$$

Let $l_s := \lceil \alpha^{s-1} \rceil$. Then, by (8), we have

$$\mathbb{P}(\text{FPGS}_{\delta, l_s} \text{ fails}) = \left(T_{2l_s+1} \left(\sqrt{1 - \lambda} T_{1/2l_s+1}(1/\delta) \right) \right)^2 \delta^2.$$

Let

$$\begin{aligned} T_{s, \delta} &:= (2l_s + 1) \left(\sqrt{1 - \lambda} T_{1/2l_s+1}(1/\delta) \right) \\ &= (2l_s + 1) \left(\sqrt{1 - \lambda} \cosh \left(\frac{\text{arccosh}(1/\delta)}{2l_s + 1} \right) \right). \end{aligned}$$

When $l_s < l_{\text{crit}}$, then this yields

$$\mathbb{P}(\text{FPGS}_{\delta, l_s} \text{ fails}) = \cosh^2(\text{arccosh}(T_{s, \delta}))^2 \delta^2.$$

Otherwise

$$\begin{aligned} \mathbb{P}(\text{FPGS}_{\delta, l_s} \text{ fails}) &\leq \mathbb{P}(\text{FPGS}_{\delta, l_{s-1}} \text{ fails}) \delta^2 \\ &\leq \mathbb{P}(\text{FPGS}_{\delta, l_{s_0}} \text{ fails}) \delta^{2(s-s_0)} \end{aligned}$$

where s_0 is the smallest integer such that $l_{s_0} \geq l_{\text{crit}}$. Let

$$Q_s := \prod_{r=1}^{s-1} \mathbb{P}(\text{FPGS}_{\delta, \lceil \alpha^{r-1} \rceil} \text{ fails}).$$

Thus, we get, as long as $\alpha \delta^2 < 1$, that

$$\begin{aligned} \tau_{\delta, \alpha} &= \sqrt{\lambda} \sum_{s=1}^{s_0-1} \lceil \alpha^{s-1} \rceil \prod_{r=1}^{s-1} \mathbb{P}(\text{FPGS}_{\delta, \lceil \alpha^{r-1} \rceil} \text{ fails}) \\ &\quad + \sqrt{\lambda} \sum_{s=s_0}^{\infty} \lceil \alpha^{s-1} \rceil \prod_{r=1}^{s-1} \mathbb{P}(\text{FPGS}_{\delta, \lceil \alpha^{r-1} \rceil} \text{ fails}) \end{aligned}$$

$$\begin{aligned} &\leq \overbrace{\sqrt{\lambda} \sum_{s=1}^{s_0-1} \lceil \alpha^{s-1} \rceil \prod_{r=1}^{s-1} \left(\cosh^2(\text{arccosh}(T_{s, \delta}))^2 \delta^2 \right)}^{\mathbb{E}_{\lambda, \delta, \alpha} :=} \\ &\quad + \sqrt{\lambda} Q_{s_0} \alpha^{s_0-1} \sum_{s=s_0}^{\infty} \alpha^{s-s_0} \delta^{2(s-s_0)} \\ &= \mathbb{E}_{\lambda, \delta, \alpha} + \frac{\sqrt{\lambda} Q_{s_0}}{1 - \alpha \delta^2}. \end{aligned}$$

In the range where $\alpha \delta^2 < 1$, $\mathbb{E}_{\lambda, \delta, \alpha}$ is bounded, but it is hard to analytically prove the convergence of $\tau_{\delta, \alpha}$ as $\lambda \rightarrow 0^+$. Nevertheless, numerical simulations suggest that for $\delta = 0.4038$ and $\alpha = 1.975$, $\tau_{\delta, \alpha}$ converges to ≈ 1.433 . The results of a numerical evaluation of the above upper bound for a wide range of parameters is shown in Fig. 2.

Remark 7: Note that while the value of the optimal δ appears sharp in Fig. 2, the graph of the α - τ dependence has many local minima and, furthermore, those minima correspond to lower α s but similar τ s. Thus, when the depth of the circuits is a concern, which will probably be the case for early fault-tolerant quantum computers, then such parameter values might be more beneficial.

Python code implementing the above computation can be found in the repository complementing this article [17].

B. TIME COMPLEXITY

In this section, we analyze the total time complexity of the FPGS for QUBO, using our oracle design and assuming Conjecture 1.

The expected time complexity of the fixed-point Grover adaptive search in Section III-A depends on three quantities.

- 1) The ratio λ is marked states to all states.
- 2) The circuit depth of the diffusion operator, $S_s(\alpha)$. Let us call this complexity C_s .

3) The circuit depth of the operator, $S_t(\beta)$. Let us call this complexity C_t .

The time complexity, assuming Conjecture 1, then is $\approx \frac{1.433}{\sqrt{\lambda}}(C_s + C_t)$.

Now, let us assume that a QUBO problem is given with $S_s(\alpha)$ and $S_t(\beta)$, as defined in Section III.

The complexity of the diffusion operator is $C_s = O(\log_2(n))$.

The complexity, C_t , depends on the number of ancillas used. In the case of the ancilla-free oracle, straightforward computation shows that the depth of $U_{f,d}$ is $O((m+d)\log_2(d))$, where $m = O(n^2)$ is the number of nonzero terms in (3). Since in our applications the ancillas start and end in the all zero state, the initial QFT and the terminal inverse QFT can be replaced by d many parallel Hadamard gates (the “quantum Walsh–Hadamard transform”), which is much shallower and easier to implement. However, the QFT and its inverse in the middle of U_t need to be implemented, which can be done with approximately with a T count of $O(d\log_2(d))$ and T -depth of $O(d)$; cf. [19]. Beyond these two subroutines, the oracle uses only CNOT gates and R_Z rotation. In fault-tolerant implementations, further T gates come the latter. The absolute values of the angles of the R_Z gates are at least $\frac{\pi q_f}{2d-1}$, where q_f is the smallest of the absolute value of the nonzero coefficients appearing in the \mathcal{P}_d gates above. Thus, fault-tolerant implementations of such gates require a T -depth (and count) of $O(d + \log_2(\frac{1}{q_f}))$; cf. [8]. Finally, when ancillas are allowed, the non-Clifford cost is identical, the Clifford cost only grows by an $O(1)$ factor, and the depth is given in Section II-B. Thus, C_t ranges from $O(\log_2(nd))$ to $O(\log_2(d)n^2)$.

Put the above together, we get that the best total time complexity (as long as $d = O(n)$) is $O(\frac{\log_2(n)}{\sqrt{\lambda}})$. Note that this beats random guessing, which is $O(\frac{1}{\lambda})$, as long as $\lambda \ll \frac{1}{\log_2(n)^2}$.

Remark 8: In the case when f is the cut function of a simple, unoriented, and undirected graph, then n is the number vertices, m is the number of edges, Q is the graph Laplacian, and d can be chosen to be $O(\log_2(n))$. Thus, in the shallowest case, one query in the FPGS for maximum graph cuts can be implemented on $n + O(m \log(n))$ qubits, with a depth of $O(\log_2(n))$.

The Edward–Erdős bound for maximum cuts is $\frac{m}{2} + \frac{n-1}{4}$, and the best known classical method [5] for finding such cuts have time complexity $O(n^4)$. It is an interesting and nontrivial question whether FPGS could yield a faster algorithm. In particular, if the ratio of cuts above the Edward–Erdős bound among all, 2^n many, cuts is $\lambda \in (0, 1)$, then for graphs that satisfy

$$\frac{\log_2(n)^2}{n^8} = o(\lambda)$$

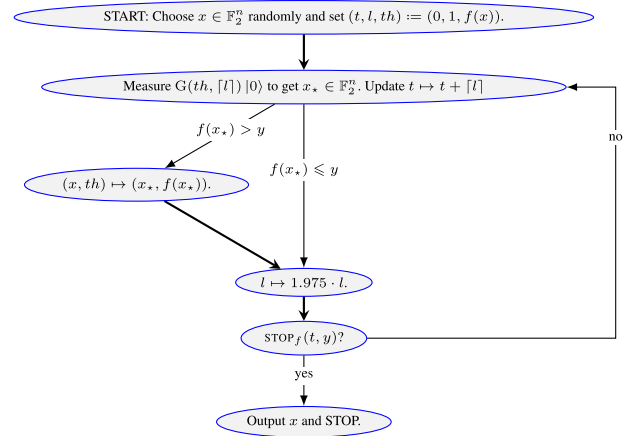


FIGURE 3. Fixed-point Grover adaptive search for QUBO.

the FPGS with our marker oracle design outperforms the classical method in [5]. Further investigation of this question is, however, beyond the scope of this article.

IV. FIXED-POINT GROVER ADAPTIVE SEARCH FOR QUBO

In this section, we define our classical-quantum method for QUBO problems, where the classical part is an adaptive method that assigns the parameters to the quantum search. This is done by building on ideas of [9] and [14], with further improvements coming from Section III. The full method is summarized in Fig. 3.

Let us assume that we are given an n -dimensional instance of a QUBO problem, f , so that its values are representable on $d = O(n)$ bit numbers. Assume also that we have a “stopping condition” described by a function

$$\mathbb{R}_+ \times \mathbb{Z} \ni (t, th) \mapsto \text{STOP}_f(t, th) \in \{\text{yes}, \text{no}\}$$

where t is the elapsed time and th is the highest value found at time t . Our goal is to find the highest value parameter before the stopping condition triggers.

Let $G(th, l)$ be the FPGS for QUBO circuit with $\delta = 0.4038$, threshold th , and l queries. With the aforementioned arguments in mind, we propose an adaptive quantum–classical hybrid method, which we call the *fixed-point Grover adaptive search for QUBO*, as follows.

- **Start** Choose $x \in \mathbb{F}_2^n$ at random and set the threshold to $th = f(x)$. Set $l = 1$.
- **Search** While the stopping condition is not satisfied, search for large values of f using FPGS for QUBO with $\delta = 0.4038$. After each search, update l to $1.975 \cdot l$ and, if a large value is found, update x to the corresponding configuration.
- **Stop** When the stopping condition is satisfied, output the configuration corresponding to the largest observed value.

The flowchart of the fixed-point Grover adaptive search for QUBO is shown in Fig. 3.

The aforementioned algorithm has a further benefit that is not present in the original one-time search case: as the threshold value increases, the set of marked states decreases, and hence, the query complexity with which the new highest value configuration is found is a suitable candidate for the next trial-and-error round. We summarize this point in Lemma 2. Before that, let us make the following definition: Fix a QUBO problem $f : \mathbb{F}_2^n \rightarrow \mathbb{Z}$. Recall that of $y \in \mathbb{Z}$

$$\mathcal{M}_{f,y} = \{x \in \mathbb{F}_2^n | y > f(x)\}.$$

Now, for each $\epsilon \in (0, 1)$, let

$$\mathcal{F}_{f,\epsilon} := \{x | (1 - \epsilon) 2^n \leq |\mathcal{M}_{f,f(x)}|\}$$

that is, $x \in \mathcal{F}_{f,\epsilon}$ exactly when x is in the top $\epsilon 2^n$ configurations (with respect to the values of f). Now, we are ready to state and prove our lemma.

Lemma 2: For each f and ϵ as above, the fixed-point Grover adaptive search for QUBO finds a configuration in $\mathcal{F}_{f,\epsilon}$ in $O(\frac{\log_2(n)}{\sqrt{\epsilon}})$ time.

Heuristic Proof of Lemma 2: Let $g_0 = 1.433$.

If, in the fixed-point Grover adaptive search for QUBO in Fig. 3, a new configuration is found using a FPGS search with l queries, then the expected number of queries needed for the next successful search is given by

$$\begin{aligned} \mathbb{E}[l_{\text{next}} | l_{\text{previous}} = l] &\approx \frac{g_0}{\sqrt{\lambda_{\text{next}}}} - l \\ &\approx g_0 \left(\frac{1}{\sqrt{\lambda_{\text{next}}}} - \frac{1}{\sqrt{\lambda_{\text{previous}}}} \right) \\ &\approx g_0 \frac{dF}{2(1 - F(y_\lambda))^{3/2}} \end{aligned}$$

where

- 1) $\lambda := \lambda_{\text{previous}}$;
- 2) $F : \mathbb{R} \rightarrow [0, 1]$ is the cumulative distribution for f , i.e.,

$$F(y) := \frac{|\mathcal{M}_{f,y}|}{2^n}$$

- 3) $y_\lambda := F^{-1}(1 - \lambda)$;
- 4) $dF := \lambda_{\text{previous}} - \lambda_{\text{next}}$.

We assume that F is well approximated by a smooth distribution.

The above yields that

$$\mathbb{E}[l_{\text{total},\epsilon}] \approx g_0 \frac{1}{\sqrt{1-F}} \Big|_0^{1-\epsilon} = g_0 \left(\frac{1}{\sqrt{\epsilon}} - 1 \right).$$

Since the shallowest implementation of a single query for our FPGS for QUBO is $O(\log_2(n))$, we get that the expected amount of time needed is $O(\frac{\log_2(n)}{\sqrt{\epsilon}})$. \square

Remark 9: As remarked earlier, our FPGS for QUBO is only expected to outperform random guessing when $\lambda = o(\frac{1}{\log_2(n)^2})$; thus, one can potentially further improve the aforementioned method by initially doing more rounds of random sampling (up to $\sim \log_2(n)^2$), using the largest




FIGURE 4. Test graph.

observed value as the first threshold and starting with $l = O(\log_2(n))$ queries.

Remark 10: The aforementioned results help cut the total resource cost and total time of the algorithm by a constant factor, but, of course, does not change the time complexity in the big- O sense. Note, however, that compared to the Grover adaptive search for QUBO method of Gilliam et al., we have improved the non-Clifford complexity of each query by (at least) a factor of 2 and improved the expected number of queries by a factor of $\approx \frac{2.25}{1.433} \approx 3.14$, yielding a total improvement of ≈ 6.28 . Similarly, both of these results yield reductions in time the overall time complexity of the method, especially when ancillas are used.

V. COMPARISON TO GROVER ADAPTIVE SEARCH

In this section, we give a detailed comparison of the Grover adaptive search of [9] and our fixed-point method. We first contrast the oracle designs and then construct a simulation to compare performances up to $n \leq n$ bit QUBOs.

Gilliam et al. [9] gave a design similar to our $U_{f,d}$ without ancillas. However, their construction uses simply controlled (for diagonal terms) and doubly controlled (for off-diagonal term) phase gates, which are more costly than ours using fanout gates and uncontrolled phase gates. Furthermore, in their design, each control qubit, $|x_j\rangle$, needs to access each target qubit, $|y_a\rangle$, sequentially. Thus, even assuming a charitable interpretation of their design (that is, assuming that instead of controlled phase gates, controlled R_Z gates are used) and a good compilation of the circuit, our designs yields an at least a twofold improvement in the number of non-Clifford gates (as a controlled R_Z gate costs two single qubit R_Z gates), an at least $\approx d$ -fold reduction in R_Z depth (as now d of them can be run parallel), and an at least $\approx \frac{d}{\log_2(d)}$ -fold reduction in total depth.

Finally, our design can make use of ancilla qubits to achieve further depth reduction without any extra non-Clifford cost, a feature not present in [9].

Next, we benchmark the performances of these methods on a concrete QUBO problem, by rewriting the evolution of the distributions as time-dependent discrete-time Markov chains. We describe the details of the simulation as follows.




FIGURE 5. Cut distributions.

Let us assume that an instance of QUBO is given, called f , let $F : \mathbb{R} \rightarrow [0, 1]$ be the cumulative distribution of f as above, and let $y \in \mathbb{F}_2^n$. When one runs FPGS for a threshold $f(y)$, with parameter δ and l queries, then the probability of success is given by

$$\mathbb{P}_{\text{FPGS},y,\delta,l} = 1 - \left(T_{2l+1} \left(T_{1/2l+1} (1/\delta) \sqrt{1 - F(y)} \right) \right)^2 \delta^2.$$

When $x \in \mathbb{F}_2^n$ with $f(x) > f(y)$ is found, then the threshold changes in our fixed-point Grover adaptive search for QUBO and the role of y (as the “best configuration seen”) is now taken by x ; otherwise, y keeps its role. Thus, the probability of a given x replacing y is

$$\mathbb{P}_{\text{FPGS}}(y \rightarrow x, \delta, l) = \frac{\mathbb{P}_{\text{FPGS},y,\delta,l}}{(1 - F(y))2^n}$$

where the denominator $(1 - F(y))2^n$ is the number of configurations $x \in \mathbb{F}_2^n$ with $f(x) > f(y)$. The probability of y keeping its role is

$$\mathbb{P}_{\text{FPGS}}(y \rightarrow y, \delta, l) = 1 - \mathbb{P}_{\text{FPGS},y,\delta,l}.$$

All other transitions are impossible. Now, assume that y was randomly picked by a distribution $p : \mathbb{F}_2^n \rightarrow [0, 1]$. Then, the probability of changing having x after the above FPGS is given by

$$\mathbb{P}_{\text{FPGS}}(x|p, \delta, l) = \sum_{y \in \mathbb{F}_2^n} p(y) \mathbb{P}_{\text{FPGS}}(y \rightarrow x, \delta, l).$$

Similarly, using [2, Lemma 2], one sees that in the case of (non-fixed-point) Grover adaptive search, the analogous formulas are given by

$$\theta_y = \frac{\arcsin(\sqrt{1 - F(y)})}{\pi}$$

$$\mathbb{P}_{\text{GAS},y,m} = \frac{1}{2} \left(1 - \frac{\text{sinc}(4m\theta)}{\text{sinc}(2\theta)} \right)$$

TABLE 1. Test Results With Threshold 4

x	$f(x)$	$y_0 = 0$	$y_0 = 1$	Total
00000	0	3.44%	0.74%	4.18%
10000	2	1.68%	0.36%	2.04%
10000	1	3.18%	0.52%	3.70%
10000	1	4.08%	0.86%	4.94%
10000	2	2.16%	0.68%	2.84%
10100	4	2.02%	0.40%	2.42%
01100	3	1.34%	0.36%	1.70%
11100	3	2.90%	0.32%	3.22%
00010	2	1.34%	0.36%	1.70%
10010	2	5.26%	0.74%	6.00%
01010	3	2.52%	0.34%	2.86%
11010	1	2.84%	0.38%	3.22%
00110	4	1.58%	0.20%	1.78%
10110	4	2.94%	0.88%	3.82%
01110	5	1.06%	1.52%	2.58%
11110	3	2.20%	0.36%	2.56%
x	$f(x)$	$y_0 = 0$	$y_0 = 1$	Total
00001	2	1.68%	0.66%	2.34%
10001	4	1.06%	0.26%	1.32%
10001	3	1.04%	0.30%	1.34%
10001	3	2.02%	0.48%	2.50%
10001	2	3.90%	0.78%	4.68%
10101	4	2.30%	0.36%	2.66%
01101	3	3.42%	0.58%	4.00%
11101	3	4.54%	0.34%	4.88%
00011	4	1.14%	0.38%	1.52%
10011	4	2.82%	0.44%	3.26%
01011	5	0.52%	1.16%	1.68%
11011	3	1.60%	0.52%	2.12%
00111	4	1.88%	0.26%	2.14%
10111	4	5.94%	1.74%	7.68%
01111	5	0.94%	2.46%	3.40%
11111	3	4.18%	0.74%	4.92%

Values below or equal to 4 should not be marked ($y_0 = 0$) and above 4 should be marked ($y_0 = 1$). The green numbers are the percentages of the given state being measured with the correct marking, while the red ones are the percentages of incorrect markings. The total occurrence of a state is the sum of the corresponding green and red percentages which in the ideal case should all be $3/32 = 3.125\%$.

$$\mathbb{P}_{\text{GAS}}(y \rightarrow x, m) = \begin{cases} \frac{\mathbb{P}_{\text{GAS},y,m}}{(1 - F(y))2^n}, & \text{if } f(x) > f(y) \\ 1 - \mathbb{P}_{\text{GAS},y,m}, & \text{if } x = y \\ 0, & \text{otherwise} \end{cases}$$

$$\mathbb{P}_{\text{GAS}}(x|p, m) = \sum_{y \in \mathbb{F}_2^n} p(y) \mathbb{P}_{\text{GAS}}(y \rightarrow x, m).$$

In the first step of both our fixed-point Grover adaptive search for QUBO and the Grover adaptive search for QUBO of Gilliam et al., one chooses $y \in \mathbb{F}_2^n$ at random, that is, $p_0(y) = \frac{1}{2^n}$ and then runs the appropriate version with $l = m = 1$, updates y if a larger value was found, and finally updates l and m . Given f , the aforementioned formulas allow us to compute the probabilities of getting configurations in either of these methods.

Let now f be the cut function of the randomly generated (connected) Erdős-Rényi graph with $n = 32$ vertices shown in Fig. 4. The graph has a maximum cut of 146, and the expected value of a random cut is 114.5, which is 78.42% of the maximum cut. The variance is 5.18%.

We ran both our fixed-point Grover adaptive search for QUBO and the Grover adaptive search for QUBO of Gilliam et al. for four iterations. Straightforward computations, using the results of the previous sections, yield that the resource requirements of the two methods are similar in this case, so that the comparison is meaningful. We found that the expected value of cut that our method produces is 88.19% of the maximum cut, compared to the 84.86% of the method of Gilliam et al., and the variances are 2.6% and 3.33%, respectively. Our method amplified the chance of finding the maximum cut by a factor of more than 88, compared to random guesses, and by a factor of 8, compared to the Grover adaptive search for QUBO of Gilliam et al. The cut distributions of a randomly picked configuration (“before Grover”), the distributions after the Grover adaptive search for QUBO of Gilliam et al., and distributions after our fixed-point Grover adaptive search for QUBO are given in Fig. 5.

Remark 11: While, in the simulations, we used the aforementioned α parameter for our fixed-point Grover adaptive search for QUBO and $\Lambda = \frac{6}{5}$ for the Grover adaptive search for QUBO of Gilliam et al., which is recommended in [2, Sect. 4], we show that for small bit numbers ($n \lesssim 32$), different values might work better, albeit the change is marginal. The code to reproduce our testing is available in [17].

Remark 12: Word of caution: For the aforementioned small bit numbers, repeated random guess $O(n^2)$ times (so at the same time complexity of our method) performs comparably.

VI. CONCLUSION

In this article, we constructed a marker oracle for QUBO problems with improved complexities to the previously known designs. These oracles are expected to be useful for both Grover- and QAOA-type quantum circuits. We also studied adaptive optimization methods, using our constructions in conjunction with FPGs. By the results of [14] and our improved oracle design, it is immediate that our fixed-point Grover adaptive search for QUBO has better performance guarantees than that of the original Grover adaptive search of [9]. However, more extensive studies of

the theoretical foundations (such as a better understanding of the distribution of values in QUBO problems) as well as further circuit simplification and resource estimates are needed.

APPENDIX EXPERIMENTAL TESTING OF THE ENCODER ORACLE TESTING ON IONQ'S QUANTUM COMPUTERS

The aforementioned marker oracle was tested on IonQ's Aria 2 QPU. We used nine qubits (five-bit QUBO with four ancillas) with 5000 shots. The matrix of the quadratic form was

$$Q = \begin{pmatrix} 2 & -1 & 0 & -1 & 0 \\ -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & -1 \\ -1 & 0 & 0 & 2 & 0 \\ 0 & 0 & -1 & 0 & 2 \end{pmatrix}$$

with threshold $y = 4$. Since the maximum of the integer-valued quadratic polynomial $f(x) = x^T Q x$ is 5, only maximal configurations should be marked, that is, $|01110\rangle$, $|01011\rangle$, and $|01111\rangle$. The oracle was decomposed into one-qubit and CNOT gates, yielding a circuit depth of 47 and a CNOT count of 66. The experimental results are given in Table I.

While the oracle is two to five times more likely to correctly mark states, further experiments and simulations (not listed here) show that contemporary QPUs are too noisy to run our circuits without error correction. Furthermore, it is worth noting that the distribution of states is no longer uniform.

CODE AND DATA AVAILABILITY

Supplementary Qiskit code for this article is available in [17].

ACKNOWLEDGMENT

The authors are grateful to Amazon Web Services and QLab for providing credits to access IonQ's QPUs. The authors also thank Constantin Gonciulea, Tom Ginsberg, Jan Tułowicki, and Shahaf Asban for providing valuable feedback. Finally, the authors thank QuForce for bringing all the authors together for the project.

REFERENCES

- [1] N. Benchasattabuse, T. Satoh, M. Hajdušek, and R. Van Meter, “Amplitude amplification for optimization via subdivided phase oracle,” in *Proc. IEEE Int. Conf. Quantum Comput. Eng.*, 2022, pp. 22–30, doi: 10.1109/QCE53715.2022.00020.
- [2] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, “Tight bounds on quantum searching,” *Fortschritte Phys.*, vol. 46, nos. 4/5, pp. 493–505, 1998, doi: 10.1002/(SICI)1521-3978(199806)46:4/5
- [3] G. A. Bridi and F. de L. Marquezino, “Analytical results for the quantum alternating operator Ansatz with Grover mixer,” 2024, *arXiv:2401.11056*, doi: 10.48550/arXiv.2401.11056.
- [4] B. Claudon, J. Zylberman, C. Feniou, F. Debbasch, A. Peruzzo, and J.-P. Piquemal, “Polylogarithmic-depth controlled-NOT gates without Ancilla Qubits,” *Nat. Commun.*, vol. 15, no. 1, 2024, Art. no. 5886, doi: 10.1038/s41467-024-50065-x.

- [5] R. Crowston, M. Jones, and M. Mních, "Max-cut parameterized above the Edwards-Erdős bound," *Algorithmica*, vol. 72, no. 3, pp. 734–757, 2015, doi: [10.1007/s00453-014-9870-z](https://doi.org/10.1007/s00453-014-9870-z).
- [6] T. G. Draper, "Addition on a quantum computer," 2000, *arXiv:quant-ph/0008033*, doi: [10.48550/arXiv.quant-ph/0008033](https://doi.org/10.48550/arXiv.quant-ph/0008033).
- [7] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*, doi: [10.48550/arXiv.1411.4028](https://doi.org/10.48550/arXiv.1411.4028).
- [8] V. Gheorghiu, M. Mosca, and P. Mukhopadhyay, "T-count and T-depth of any multi-Qubit unitary," *npj Quantum Inf.*, vol. 8, no. 1, 2022, Art. no. 141, doi: [10.1038/s41534-022-00651-y](https://doi.org/10.1038/s41534-022-00651-y).
- [9] A. Gilliam, S. Woerner, and C. Goniculea, "Grover adaptive search for constrained polynomial binary optimization," *Quantum*, vol. 5, 2021, Art. no. 428, doi: [10.22331/q-2021-04-08-428](https://doi.org/10.22331/q-2021-04-08-428).
- [10] L. Giuffrida, "Solving the binary optimization problem for linear regression with Grover adaptive search," Ph.D. dissertation, Corso di laurea magistrale in Ingegneria Elettronica, Politecnico Torino, Turin, Italy, 2021.
- [11] J. Golden, A. Bärttschi, D. O'Malley, and S. Eidenbenz, "Threshold-based quantum optimization," in *Proc. IEEE Int. Conf. Quantum Comput. Eng.*, 2021, pp. 137–147, doi: [10.1109/QCE52317.2021.00030](https://doi.org/10.1109/QCE52317.2021.00030).
- [12] L. K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," *Phys. Rev. Lett.*, vol. 79, Jul. 1997, Art. no. 325328, doi: [10.1103/PhysRevLett.79.325](https://doi.org/10.1103/PhysRevLett.79.325).
- [13] L. K. Grover, "Fixed-point quantum search," *Phys. Rev. Lett.*, vol. 95, Oct. 2005, Art. no. 150501, doi: [10.1103/PhysRevLett.95.150501](https://doi.org/10.1103/PhysRevLett.95.150501).
- [14] T. Li et al., "Quantum search for unknown number of target items by hybridizing fixed-point method with trial-and-error method," *Chin. Phys. B*, vol. 28, no. 12, Nov. 2019, Art. no. 120301, doi: [10.1088/1674-1056/ab4e88](https://doi.org/10.1088/1674-1056/ab4e88).
- [15] G. H. Low, V. Kliuchnikov, and L. Schaeffer, "Trading T gates for dirty Qubits in state preparation and unitary synthesis," *Quantum*, vol. 8, 2024, Art. no. 1375, doi: [10.22331/q-2024-06-17-1375](https://doi.org/10.22331/q-2024-06-17-1375).
- [16] S. Marsh and J. B. Wang, "Combinatorial optimization via highly efficient quantum walks," *Phys. Rev. Res.*, vol. 2, no. 2, 2020, Art. no. 023302, doi: [10.1103/PhysRevResearch.2.023302](https://doi.org/10.1103/PhysRevResearch.2.023302).
- [17] Á. Nagy, "GitHub repository for the implementation of the fixed-point Grover adaptive search for binary optimization problems," 2024. [Online]. Available: https://github.com/akos-nagy/FPGAS_for_QUBO
- [18] Á. Nagy and C. Zhang, "Novel oracle constructions for quantum random access memory," 2024, *arXiv:2405.20225*, doi: [10.48550/arXiv.2405.20225](https://doi.org/10.48550/arXiv.2405.20225).
- [19] Y. Nam, Y. Su, and D. Maslov, "Approximate quantum Fourier transform with $O(n \log(n))$ T gates," *NPJ Quantum Inf.*, vol. 6, no. 1, 2020, Art. no. 26, doi: [10.1038/s41534-020-0257-5](https://doi.org/10.1038/s41534-020-0257-5).
- [20] J. Nie, W. Zi, and X. Sun, "Quantum circuit for multi-qubit Toffoli gate with optimal resource," 2024, *arXiv:2402.05053*, doi: [10.48550/arXiv.2402.05053](https://doi.org/10.48550/arXiv.2402.05053).
- [21] Y. Sano, K. Mitarai, N. Yamamoto, and N. Ishikawa, "Accelerating Grover adaptive search: Qubit and gate count reduction strategies with higher-order formulations," *IEEE Trans. Quantum Eng.*, vol. 5, 2023, Art. no. 3101712, doi: [10.1109/TQE.2024.3393437](https://doi.org/10.1109/TQE.2024.3393437).
- [22] P. Shyamsundar, "Non-Boolean quantum amplitude amplification and quantum mean estimation," *Quantum Inf. Process.*, vol. 22, no. 12, 2023, Art. no. 423, doi: [10.1007/s11128-023-04146-3](https://doi.org/10.1007/s11128-023-04146-3).
- [23] M. Szegedy, "What do QAOA energies reveal about graphs?," 2019, *arXiv:1912.12277*, doi: [10.48550/arXiv.1912.12277](https://doi.org/10.48550/arXiv.1912.12277).
- [24] F. M. Toyama, W. Van Dijk, and Y. Nogami, "Quantum search with certainty based on modified Grover algorithms: Optimum choice of parameters," *Quantum Inf. Process.*, vol. 12, 2013, Art. no. 18971914, doi: [10.1007/s11128-012-0498-0](https://doi.org/10.1007/s11128-012-0498-0).
- [25] T. J. Yoder, G. H. Low, and I. L. Chuang, "Fixed-point quantum search with an optimal number of queries," *Phys. Rev. Lett.*, vol. 113, 2014, Art. no. 210501, doi: [10.1103/PhysRevLett.113.210501](https://doi.org/10.1103/PhysRevLett.113.210501).
- [26] C. Zalka, "Grover's quantum searching algorithm is optimal," *Phys. Rev. A*, vol. 60, no. 4, 1999, Art. no. 2746, doi: [10.1103/PhysRevA.60.2746](https://doi.org/10.1103/PhysRevA.60.2746).