Privacy-Preserving Quantum Annealing for Quadratic Unconstrained Binary Optimization (QUBO) Problems

Moyang Xie¹, Yuan Zhang¹, Sheng Zhong¹, Qun Li²

¹State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China ²Department of Computer Science, William & Mary, Williamsburg, VA, USA xie_moyang@foxmail.com, zhangyuan@nju.edu.cn, zhongsheng@nju.edu.cn, liqun@cs.wm.edu

Abstract-Quantum annealers offer a promising approach to solve Quadratic Unconstrained Binary Optimization (QUBO) problems, which have a wide range of applications. However, when a user submits its OUBO problem to a third-party quantum annealer, the problem itself may disclose the user's private information to the quantum annealing service provider. To mitigate this risk, we introduce a privacy-preserving QUBO framework and propose a novel solution method. Our approach employs a combination of digit-wise splitting and matrix permutation to obfuscate the QUBO problem's model matrix Q, effectively concealing the matrix elements. In addition, based on the solution to the obfuscated version of the OUBO problem, we can reconstruct the solution to the original problem with high accuracy. Theoretical analysis and empirical tests confirm the efficacy and efficiency of our proposed technique, demonstrating its potential for preserving user privacy in quantum annealing services.

Index Terms—Privacy-Preserving, Quadratic Unconstrained Binary Optimization, Quantum Annealing.

I. INTRODUCTION

Leveraging the unique quantum mechanics, such as superposition, entanglement, and quantum interference, quantum computers provide a new computing paradigm for many new applications. For example, quantum annealers, based on a process of simulated annealing, can search the solution space efficiently and find the optimal solution to optimization problems [1] that are difficult for classical computers to solve.

In particular, quantum annealers are known for effectively solving the Quadratic Unconstrained Binary Optimization (QUBO) [2], [3], which has broad-ranging applications in fields such as finance, economics, etc. [4]. As an NP-hard problem, QUBO poses a significant computational challenge to classical computers. It has been applied to various classical problems in theoretical computer science, including the maximum cut, graph coloring, and the partition problem, all of which can be formulated as QUBO embeddings [5]. Additionally, QUBO has extended its utility to machine learning, with embeddings designed for support-vector machines, clustering algorithms, and probabilistic graphical models [6], [7]. The close relationship between QUBO and Ising models positions QUBO as a central problem class in adiabatic quantum computation, where solutions are sought through the physical process known as quantum annealing. A diverse array of industry leaders, including D-Wave, Google, IBM, Microsoft, and Amazon, alongside esteemed public institutions such as Oak Ridge National Laboratory, Los Alamos National Laboratory, and NASA's Ames Research Center, have dedicated substantial resources to exploring the utilization of quantum computers for resolving QUBO problems [5]. This extensive body of research underscores the immense value and untapped potential of leveraging quantum computational resources for QUBO problem-solving.

Unfortunately, to date, it is hard for ordinary individuals or even companies to have their own quantum computer. Noticing this unmet demand, companies such as D-Wave have started to provide cloud services of quantum annealing. Although this cloud-based service is convenient for users to solve their QUBO problems remotely, it also brings immediate privacy concerns. Specifically, the QUBO instances stemming from users' real-world challenges may encapsulate sensitive information, ranging from financial data and corporate secrets to personal health records. Consequently, when users upload their OUBO problems, there exists a risk of privacy infringement by the service provider, who could potentially scrutinize these problems. Hence, the development of a robust method to safeguard user privacy while enabling utilization of third-party quantum annealing services for solving QUBO problems is paramount.

This paper addresses the privacy-preserving QUBO solving problem by introducing a novel approach. Our method generates an obfuscated version of the original QUBO problem, which, when solved, yields a solution that closely approximates the original problem's solution with high accuracy. Instead of submitting their original problem to the quantum annealing service provider, users can send this obfuscated version. This strategy allows users to leverage the computational power of quantum annealing for obtaining high-quality solutions to their QUBO problems while significantly reducing the risk of privacy breaches. Our approach strikes a balance between utilizing advanced quantum annealing capabilities and safeguarding sensitive information, potentially paving the way for more widespread and secure adoption of cloud-based quantum annealing services.

Our major contributions can be summarized as follows.

- We formulate the privacy-preserving QUBO problem and propose an obfuscation solution.
- We provide theoretical analyses on the security and efficiency of our method.
- Additionally, we conduct experiments to evaluate the efficacy and efficiency of our proposed technique. Results demonstrate that our method achieves good accuracy.

The rest of this paper is organized as follows. In Sec. II, we formulate the privacy-preserving QUBO problem. In Sec. III, we explain our method in details. Theoretical analyses and experimental evaluations are presented in Sec. IV and Sec. V respectively. After discussing related work in Sec. VI, we conclude our paper in Sec. VII.

II. THE PRIVACY-PRESERVING QUBO SOLVING PROBLEM

The Quadratic unconstrained binary optimization (QUBO) is a model for solving combinatorial optimization problems that can be applied to a variety of important and difficult combinatorial problems such as Multiple Knapsack Problems, P-Median Problems, Maximum Independent Set Problems, Maximum Cut Problems, Graph Coloring Problems, SAT problems and so on. The problem has the following form:

minimize
$$y = x^t Q x_t$$

where $x \in \{0,1\}^n$ is an vector of binary elements, and $Q \in \mathbb{R}^{n \times n}$ is the "model matrix" which is a square matrix of constants that determines the objective function to be optimized or minimized over all possible values of x in our case.

For example, one might aim to minimize a function $y = f(x_1, x_2, x_3, x_4)$ as follows.

$$y = 6x_1^2 - 18x_3^2 - 2x_4^2 + 10x_1x_3 - 18x_1x_4 - 6x_2x_3 + 4x_2x_4 + 4x_3x_4 = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}^T \begin{pmatrix} 6 & 0 & 5 & -9 \\ 0 & 0 & -3 & 2 \\ 5 & -3 & -18 & 2 \\ -9 & 2 & 2 & -2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix}.$$

Accordingly, the above optimization problem can be determined by a square matrix:

$$Q = \begin{pmatrix} 6 & 0 & 5 & -9 \\ 0 & 0 & -3 & 2 \\ 5 & -3 & -18 & 2 \\ -9 & 2 & 2 & -2 \end{pmatrix},$$

and a brute-force search can find the solution to the problem: $x_1 = 0, x_2 = 1, x_3 = 1, x_4 = 0.$

In this paper, we consider the *privacy-preserving QUBO* solving problem as follows.

A quantum annealing *server* solves QUBO problems. The *client* submits a QUBO model matrix and receives the optimal solution. The client aims to solve its QUBO problem with matrix Q while keeping Q private from the server. To achieve

this, the client sends a series of matrices $\{M\}$ to the server and receives optimal solutions $\{x'\}$, which are used to derive the optimal solution x of Q. Our goals are: (1) prevent the server from inferring Q from $\{M\}$ for privacy preservation, and (2) accurately derive x from $\{x'\}$ to solve the client's problem.

III. OUR APPROACH

Our approach generates an obfuscated version of the model matrix. We split the QUBO model matrix Q into matrices whose weighted sum equals Q. We represent each element of Q using a radix r ($r \in \mathbb{N}$) and split it into a string of r-radix digits, decomposing Q into "digit matrices". Each digit matrix contains elements' digits from the same digit place (e.g., the 1st digit's place, the 2nd digit's place, ...).

We randomly permute rows and columns of these matrices and send them to the server in random order. The server cannot recover the original matrix without knowing the permutations or sending order.

Our approach preserves the relationship between the original matrix and sent matrices. The client can reconstruct the original matrix by summing the sent matrices with correct permutations and weights. After the server returns the optimization solutions, which are a series of binary vectors, corresponding to the matrices it received from the client, the client can recover the binary vectors in the correct order, and combine them to obtain a solution to its original QUBO problem.

In the following, we describe in details: 1) how the client constructs obfuscated matrices to send; and 2) how the client reconstructs the final optimal solution x from returned binary vectors.

A. Obtaining the Matrices to Be Sent

Following the operations below, the client computes k matrices M_1^{σ} , M_2^{σ} ,..., $M_k^{\sigma} \in \mathbb{R}^{n \times n}$ from the model matrix Q, and sends them to the server in a random order, which can be completed in one single message transmission.

1) Matrix normalization. To enhance both privacy and efficiency in our approach, we normalize the model matrix Qwhile preserving the relative ratio between its elements. This normalization process scales all elements to fall within the interval (-1, 1), with the element having the largest absolute value approaching either -1 or 1. Consequently, the integer parts of all elements become 0, simplifying our subsequent digit-wise splitting operation to only consider digits after the decimal point. By ensuring the element with the largest absolute value is near -1 or 1, we try to maximize the number of elements with a non-zero first decimal digit. This first digit is crucial as it contributes significantly to the overall value of each element. As a result, the matrix formed by these first digits holds the highest weight among all matrices derived from the radix decomposition. Its corresponding optimal binary vector also tends to carry the greatest weight in the subsequent weighted summation process. Maximizing non-zero first digits enhances the significance of both the corresponding matrix and

its optimal binary vector, thereby improving result accuracy. Moreover, this normalization step helps to obscure the actual element values within the model matrix Q, thereby enhancing the security of our approach. Specifically, the client computes the normalized matrix Q^* as:

$$Q^*[i,j] := \frac{Q[i,j]}{(1+\epsilon)max(Q)}, \ \forall i,j \in \{0,1,...,n\}, \quad (1)$$

where $\max(Q)$ denotes the maximum absolute value of all elements in Q, Q[i, j] and $Q^*[i, j]$ denote the *i*th-row-*j*th-column element of Q and Q^* resp., and $\epsilon \to 0^+$ is a small positive real number that is close to 0.

2) Matrix decomposition with base r. The client represents (the non-integer part of) normalized model matrix Q^* using radix r, takes all elements' most significant digits to form M_1 , the second most significant digits as M_2 , ..., and the kth digit as M_k :

$$M_m[i, j] := \text{the } m\text{th digit of } Q^*[i, j],$$
(2)
$$\forall m \in \{1, 2, ..., k\}, \forall i, j \in \{1, 2, ..., n\}$$

3) Random matrix permutation. The client generates k random permutations $\sigma_1, \sigma_2, ..., \sigma_k$ and uses them to permute the rows and columns of $M_1, M_2, ..., M_k$ respectively to obtain $M_1^{\sigma}, M_2^{\sigma}, ..., M_k^{\sigma}$ as:

$$M_m^{\sigma}[i,j] := M_m[\sigma_m(i), \sigma_m(j)]$$
(3)
$$\forall m \in \{1, 2, ..., k\}, \forall i, j \in \{1, 2, ..., n\}$$

4) Sending in a random order. To perform this, the client generates a random permutation σ' and sends $M_{\sigma'(1)}^{\sigma}$, $M_{\sigma'(2)}^{\sigma}$, ..., $M_{\sigma'(k)}^{\sigma}$ in order.

Note: To further obfuscate the matrix Q, the client can also generate and send additional random matrices to the server along with the decomposed and permuted matrices derived from Q. These random matrices serve as decoys, making it more difficult for the server to identify or extract sensitive information about the original matrix Q. Upon receiving the solutions from the server corresponding to all the matrices, including the random ones, the client can simply discard the solutions pertaining to the decomposed and permuted matrices of Q. This extra layer of obfuscation further enhances the privacy protection of the sensitive data within the matrix Q. For presentation simplicity, we omit this step in our approach description.

B. Reconstruction of the optimal solution

After receiving k obfuscated matrices from the client, the server sends back the corresponding k optimization solutions or binary vectors to the client. We denote these binary vectors as $v_1^{\sigma'}$, $v_2^{\sigma'}$, ..., $v_k^{\sigma'}$, which are the solutions to the binary optimization for $M_{\sigma'(1)}^{\sigma}$, $M_{\sigma'(2)}^{\sigma}$, ..., $M_{\sigma'(k)}^{\sigma}$. The client uses inverse permutations of $\sigma_1, \sigma_2, \ldots, \sigma_k$ and σ' to recover them

from the server's response, and then sum them with corresponding weights. Finally, the client computes a probability distribution from the weighted sum vector, and samples the final optimal x over the distribution. Specifically, the above operations are completed by executing the following steps:

1) Inversing permutations. The client first permutes $v_1^{\sigma'}, v_2^{\sigma'}, \dots, v_k^{\sigma'}$ using σ'^{-1} (i.e. the inverse of σ') to get the correctly ordered $v_1^{\sigma}, v_2^{\sigma}, \dots, v_k^{\sigma}$ which are binary optimization solutions to $M_1^{\sigma}, M_2^{\sigma}, \dots, M_k^{\sigma}$:

$$v_m^{\sigma} := v_{\sigma'^{-1}(m)}^{\sigma'}, \ \forall m \in \{1, 2, ..., k\}$$
(4)

Then the client permutes the elements of v_1^{σ} , v_2^{σ} , ..., v_k^{σ} using σ_1^{-1} , σ_2^{-1} , ..., σ_k^{-1} respectively to recover v_1 , v_2 , ..., v_k which are binary optimization solutions to M_1 , M_2 , ..., M_k :

$$v_m[i] := v_m^{\sigma}[\sigma_m^{-1}(i)], \ \forall m \in \{1, 2, ..., k\}, \ \forall i \in \{1, 2, ..., n\}$$
(5)

2) Weighted average calculation. The client sets a weight vector $w \in \mathbb{R}^k$ and computes the weighted average \tilde{x} of v_1 , v_2 , ..., v_k with weights of w[1], w[2], ..., w[k] as:

$$\tilde{x} := \frac{\sum_{m=1}^{k} w[m] \cdot v_m}{\sum_{m=1}^{k} w[m]}.$$
(6)

 \tilde{x} is used as an estimation on the probability that each entry in the optimal x equals 1.

Intuitively, the matrices with smaller subscripts in M_1 , M_2 , ..., M_k are split from higher digits in the model matrix Q and contribute more to x^tQx , and the corresponding optimization solutions should also have larger weights in x. A possible w is

$$w[m] := \frac{1}{r^m}, \ \forall m \in \{1, 2, ..., k\}.$$
(7)

We will discuss the choice of w again later in the accuracy analysis.

3) Sampling and selection. Finally, the client independently samples t vectors $X_1, X_2, ..., X_t$ from the distribution:

$$X \leftarrow [0,1]^n, \Pr[X[i]=1] = \tilde{x}[i],$$
 (8)

and selects the one that minimizes the objective function

$$\underset{x \in \{X_1, X_2, \dots, X_t\}}{\operatorname{argmin}} x^t Q x \tag{9}$$

as the final optimal x.

IV. PERFORMANCE ANALYSIS

Our approach has the following principal parameters:

- *n* the order of the matrix, determined by the model;
- *r* the radix or base in our numeral system, determined by the client;
- k the total number of "digit matrices", determined by the client;
- t the total number of random samplings on \tilde{x} , determined by the client.

In this section, we discuss the privacy protection, accuracy and efficiency of our approach and analyze the effects of the above parameters.

A. Privacy protection

First of all, due to the setup of the normalization step, Q^* hides the element-specific values of Q and only retains the proportionality between them, so neither the element-specific values of Q nor the optimal $x^T Q x$ values will be leaked. Therefore, there can only be leakage of the proportionality of the elements in Q, and we analyze the leakage of this below.

An optimistic, but without loss of generality, assumption is that the distribution of the digits of each element of Q^* in each *r*-base position is indistinguishable from a uniformly independent distribution, such that the absolute values of the elements in M_1, M_2, \ldots, M_k are also uniformly independently distributed. In this case, the server can only estimate a permutation by indices of the sign matrix of Q based on the sign matrices of any of the received matrices. It is discussed next whether the server can get more valuable information from the received matrices $M_{\sigma'(1)}^{\sigma}, M_{\sigma'(2)}^{\sigma}, \ldots, M_{\sigma'(k)}^{\sigma}$.

Theorem IV.1. If the absolute values of all digit matrices' elements follow independently uniform distribution over the digit value space, the probability of recovering a correct Q^* is $\frac{1}{\alpha^{k-1}k!n!}$, where α denotes the number of automorphisms of the sign matrix of Q under the row/column index permutation transformation.

Proof. Note that before the random permutations on rows and columns, all k digital matrices $\{M_m\}_{m=1,...,k}$ share the same sign matrix (i.e. the matrix formed by all elements' signs) with the Q or Q^{*}. Therefore the correct digital matrices can be recovered only by a combination of permutations $\{\sigma_m^*\}_{m=1,...,k}$ such that applying them on the permuted matrices $\{M_m^{\sigma}\}_{m=1,...,k}$ respectively would make their sign matrices become identical. Let Σ be the set of all such combinations.

Since the absolute values of the elements in digital matrices are uniformly and independently distributed, the server is not able to distinguish between the correct combination of permutations (i.e. $\{\sigma_m^{-1}\}_{m=1,...,k}$) that recovers the true digital matrices and other incorrect combination of permutations in Σ .

For any σ_1^* , there are α^{k-1} possible $(\sigma_2^*, ..., \sigma_k^*)$ that the sign matrices of the resulting matrices are the same after permutation.

Therefore, we know the size of Σ equals $\alpha^{k-1}n!$, and the probability that the server correctly guesses $\{\sigma_m^{-1}\}_{m=1,...,k}$ and recovers $\{M_{\delta'(m)}\}$ that have correct rows and columns equals $\frac{1}{\alpha^{k-1}n!}$.

Finally, there are k! possible orders that the digit matrices are sent with. Therefore, the probability that the server recovers the correct Q^* equals

$$\frac{1}{\alpha^{k-1}k!n!}.$$

In the worst-case scenario, where the absolute values of the elements in M_1 , M_2 , and M_k are not uniformly and

independently distributed, we can employ the optional privacyenhancing method mentioned earlier. By sending the server specific confusion matrices, we ensure that the matrices appear to match a uniformly independent distribution or another nonsensical distribution. This approach is effective but requires contextual analysis.

When a large k is not necessary, we can achieve the privacy preservation effect of a larger k by sending more confusion matrices in this manner.

From the analysis above, we find that larger values of n and k improve privacy performance. With a fixed n, a larger k means the server receives more matrices related to the model, which intuitively provides more information. However, paradoxically, it becomes more challenging to recover the private data. This suggests that increasing the amount of information increases the difficulty for the server to organize and analyze it, due to the larger probability space created by more matrices.

Although the parameter r does not appear in the equation, its impact on privacy is significant. As discussed, "recovering the first k'(k' < k) correct digits of Q^* " shows that a larger rprovides more space for the first k' digits, thus containing more information. In an extreme case, where r is very large, M_1 alone may contain almost all the information in Q. Conversely, a smaller r makes the distribution of each element look more uniform and independent, thereby enhancing privacy performance. Therefore, smaller values of r improve privacy protection.

B. Accuracy

First, we make several qualitative observations regarding the impact of various parameters. A larger t results in a smaller error because a greater number of samples increases the likelihood of obtaining a better binary vector. The parameter k influences the error up to a point and then has a marginal effect. A too-small number of matrices primarily contain higher digits, potentially lacking sufficient information, whereas a too-large number of matrices include many lower digits, which contribute minimally to the results and appear more random. Increasing r reduces the error, as a larger base encapsulates more information in each digit. In an extreme case, if r is sufficiently large, sending only the first matrix to the server can yield an adequately accurate result.

We have conducted numerous experiments demonstrating that our approach achieves excellent accuracy (see Section V for details). This can be intuitively understood: higher digits contribute more significantly to the model matrix Q, thus having a greater impact on x^tQx and the optimal x. The influence of a higher digit on the optimal x is reflected in the optimal solution of the matrix formed by that digit. The weight of this effect on the optimal x is captured in the final result through the weights we assign to that binary vector in the weighted average. Each lower digit reduces the weight in Q by a factor of r, leading us to propose that its weight in the final result should also be reduced by a factor of r. This is the basis for the generic weight vector $w[i] := \frac{1}{r^i}, \quad \forall i \in \{1, 2, ..., k\}.$ However, this weight vector w is not universally applicable. The selection of w should be related to the distribution of the elements in Q. For example, when the absolute values of the non-zero elements in Q are similar in magnitude, focusing only on the highest position may suffice, and considering lower positions might reduce accuracy, especially when r is small. When the magnitude distribution of elements in Q is uncertain, determining an optimal w is challenging. Conversely, if the magnitude distribution is known, techniques such as multiple linear regression or deep learning can be employed to learn a more appropriate w.

Given the complexity and uncertainty of this problem, providing a rigorous theoretical analysis of the accuracy is difficult. A more in-depth analysis of the choice of w and its impact on accuracy will be the focus of our future research.

C. Efficiency

Our method demonstrates strong performance in terms of efficiency. For communication, only one round is required: the client sends k *n*-order matrices to the server, and the server returns k n-dimensional binary vectors, resulting in a communication overhead of $\mathcal{O}(kn^2)$. For *computation*, the client needs to generate and permute k n-order matrices, which is $\mathcal{O}(kn^2)$. After receiving the k n-dimensional binary vectors, the client computes their weighted average, which is $\mathcal{O}(kn)$. Following this, the client samples t candidate vectors and computes the values of the corresponding quadratic forms, amounting to $\mathcal{O}(tn^2)$. The sorting process to obtain the final result is $\mathcal{O}(t \log t)$, which is typically subsumed under $\mathcal{O}(tn^2)$. Thus, the total computational overhead for the client is $\mathcal{O}(kn^2 + tn^2)$. The server, on the other hand, needs to solve OUBO for k *n*-order model matrices, with a total overhead of $\mathcal{O}(k)$ assuming each solution takes $\mathcal{O}(1)$ time.

For *storage*, the client must store $\mathcal{O}(k)$ *n*-order matrices and $\mathcal{O}(t)$ *n*-dimensional vectors along with their corresponding quadratic values, leading to a total storage overhead of $\mathcal{O}(kn^2 + tn)$. The server needs to store *k n*-order matrices and the corresponding *k n*-dimensional vectors, resulting in a total storage overhead of $\mathcal{O}(kn^2)$.

As the number of transmitted matrices increases, the communication overhead increases linearly. Conversely, the difficulty for the server to learn information about the matrix Qincreases exponentially.

V. EVALUATION ON ACCURACY

To evaluate whether our approach achieves desirable accuracy, we conducted several experiments.

We randomly generated model matrices for testing under a normal distribution with a mean of 0 and a standard deviation of 4. Our observations indicate that optimizing these matrices is quite challenging, making them a rigorous test case. Given the current limitations in accessing a quantum annealer, we utilized the well-known commercial optimization solver IBM ILOG CPLEX to solve QUBO problems in our experiments. CPLEX is capable of providing optimal results at reasonably large problem scales [8]. We used the values of the optimization function x^tQx as the observed values and defined the accuracy of our solution as:

$$\operatorname{acc} = \frac{\operatorname{obtained value of } x^t Q x}{\operatorname{true optimal value of } x^t Q x},$$

and the error as:

$$err = \left| \frac{\text{obtained value of } x^t Q x - \text{true optimal value of } x^t Q x}{\text{true optimal value of } x^t Q x} \right|$$

We focus on the difference between the obtained and true optimal values of the optimization function, as this difference is more critical than the difference between the obtained and true optimal solutions.



Figure 1. Error with different matrix orders and total numbers of matrices when sampling 200 times with a base of 2.

In Fig. 1, we illustrate how the error varies as the matrix order increases. The results indicate that the error grows slowly with increasing order, following a trend that appears to be no more than linear, suggesting O(n)-like behavior. These encouraging results imply that we can achieve sufficiently accurate solutions by adjusting other parameters, with an acceptable overhead. Fig. 1 also demonstrates the effect of increasing k on accuracy improvement, consistent with theoretical predictions. However, beyond a certain point, further increases in k offer diminishing returns in accuracy. Therefore, we can select a modest k and employ the transmission of confusion matrices to enhance security without significantly compromising accuracy.

Fig. 2 demonstrates the relationship between accuracy and the number of samples. As depicted, accuracy improves with an increasing number of samples. For each order of test matrices utilized, with 5 matrices and a base of 4, we achieved accuracy exceeding 90% within 300 samples—a notably low overhead given the complexity of the problem. However, it's also evident that beyond a certain number of samples, the rate of accuracy improvement diminishes, indicating that achieving very high accuracy solely by increasing the number of random samples may still be challenging. The current method of random sampling is suboptimal, and we plan to explore more efficient techniques to derive the final x from the $[0, 1]^n$ vector \tilde{x} at a reduced cost in the future.

 Table I

 COMPUTATION, COMMUNICATION AND STORAGE COSTS

Communication	Rounds	Client Work	Server Work	Client Storage	Server Storage
kn^2	1	$kn^2 + tn^2$	k *	$kn^2 + tn$	kn^2
*Here we count the overhead of the server solving a OUBO as $\mathcal{O}(1)$.					



Figure 2. Accuracy with different total numbers of samples and orders of model matrix when the number of matrices is 5 with a base of 4.



Figure 3. Accuracy with different total numbers of samples and bases of numeral/rounding system when the number of matrices is 5 and the order of model matrix is 40.

Fig. 3 illustrates the increase in accuracy as the base, r, grows. The base is a flexible and influential parameter, allowing for significant accuracy improvements with its increase. However, this increase also comes at a cost to security, highlighting the trade-off between security, privacy, and overhead.

VI. RELATED WORK

Many quantum computing research projects focus on developing general quantum computing platforms for new applications, such as quantum machine learning [9]–[11]. This paper, however, examines quantum annealers, which are specialized quantum computers.

A considerable amount of research has been conducted on privacy-preserving methods for classical optimization. Notable works include [12]–[21]. However, these methods are not well-suited for our specific scenario, as quantum annealers differ fundamentally from general-purpose computation and continuous optimization, which these methods are designed for.

Specifically, the gradient descent method are considered in [14], [15], [17]–[19], while the simplex algorithm is studied in [20]. However, the methods proposed in these works are designed for continuous or linear optimization and are not applicable to discrete quadratic optimization. Additionally, the quantum annealer does not support general-purpose computation, making it impossible to adopt these methods.

The works [16], [21] achieve privacy preservation through matrix transformations; however, their methods are only applicable to continuous optimization problems, while we are dealing with a discrete optimization problem.

A few works [12], [13] adopt differential privacy mechanisms [22] to protect individual participant's model matrix privacy in a distributed joint optimization task where the objective function to be minimized corresponds to the sum of all participants' model matrices. Although differential privacy can protect individual privacy, it offers little protection to the overall model matrix. Therefore, it cannot solve our problem.

VII. CONCLUSIONS AND FUTURE WORK

In this work, we propose a novel approach that enables privacy-preserving access to a remote quantum annealer for solving QUBO optimization problems. Our approach balances privacy, accuracy, and efficiency, allowing for tradeoffs among these aspects by adjusting relevant parameters. Theoretical analysis demonstrates that our method effectively conceals the values of the elements in the model matrix, with the probability of an adversary accurately determining the ratio between them being negligible. Experimental evaluation shows that, with 5 matrices and a base of 4, our method achieves over 90% accuracy within 300 samples.

Future research will focus on refining the weight vector selection process, developing more sophisticated methodologies for deriving final results, establishing theoretical bounds on accuracy, and testing our approach on a broader range of existing QUBO benchmarks for large-scale evaluation.

ACKNOWLEDGMENT

The authors would like to thank all the reviewers for their helpful comments. Qun Li was supported in part by the Commonwealth Cyber Initiative (cyberinitiative.org).

REFERENCES

- A. Das and B. K. Chakrabarti, "Quantum annealing and analog quantum computation," *Reviews of Modern Physics*, vol. 80, no. 3, p. 1061, 2008.
- [2] D. A. Battaglia, G. E. Santoro, and E. Tosatti, "Optimization by quantum annealing: Lessons from hard satisfiability problems," *Physical Review E*, vol. 71, no. 6, p. 066707, 2005.
- [3] A. Rajak, S. Suzuki, A. Dutta, and B. K. Chakrabarti, "Quantum annealing: An overview," *Philosophical Transactions of the Royal Society A*, vol. 381, no. 2241, p. 20210417, 2023.
- [4] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, "The unconstrained binary quadratic programming problem: a survey," *Journal of combinatorial optimization*, vol. 28, pp. 58–81, 2014.
- [5] F. Glover, G. Kochenberger, and Y. Du, "Quantum bridge analytics i: a tutorial on formulating and using qubo models," *4or*, vol. 17, no. 4, pp. 335–371, 2019.
- [6] S. Mücke, N. Piatkowski, and K. Morik, "Learning bit by bit: Extracting the essence of machine learning." in *LWDA*, 2019, pp. 144–155.
- [7] P. Date, D. Arthur, and L. Pusey-Nazzaro, "Qubo formulations for training machine learning models," *Scientific reports*, vol. 11, no. 1, p. 10029, 2021.
- [8] IBM ILOG. ILOG CPLEX Optimization Studio Python tutorial. [Online]. Available: https://www.ibm.com/docs/en/icos/22.1.1?topic=tutorialspython-tutorial
- [9] J. Wu, T. Hu, and Q. Li, "Distributed quantum machine learning: Federated and model-parallel approaches," *IEEE Internet Computing*, vol. 28, no. 2, pp. 65–72, 2024.
- [10] J. Wu, Z. Tao, and Q. Li, "Scalable quantum neural networks for classification," in 2022 IEEE International Conference on Quantum Computing and Engineering (QCE). IEEE, 2022, pp. 38–48.
- [11] J. Wu, T. Hu, and Q. Li, "More: Measurement and correlation based variational quantum circuit for multi-classification," in 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), vol. 1. IEEE, 2023, pp. 208–218.
- [12] Z. Huang, S. Mitra, and N. Vaidya, "Differentially private distributed optimization," in *Proceedings of the 16th International Conference on Distributed Computing and Networking*, 2015, pp. 1–10.
- [13] S. Han, U. Topcu, and G. J. Pappas, "Differentially private distributed constrained optimization," *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 50–64, 2016.
- [14] Y. Shoukry, K. Gatsis, A. Alanwar, G. J. Pappas, S. A. Seshia, M. Srivastava, and P. Tabuada, "Privacy-aware quadratic optimization using partially homomorphic encryption," in 2016 IEEE 55th Conference on Decision and Control (CDC). IEEE, 2016, pp. 5053–5058.
- [15] A. B. Alexandru, K. Gatsis, and G. J. Pappas, "Privacy preserving cloudbased quadratic optimization," in 2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2017, pp. 1168– 1175.
- [16] L. Zhou and C. Li, "Outsourcing large-scale quadratic programming to a public cloud," *IEEE Access*, vol. 3, pp. 2581–2589, 2015.
- [17] S. Gade and N. H. Vaidya, "Private optimization on networks," in 2018 Annual American Control Conference (ACC), 2018, pp. 1402–1409.
- [18] Z. Zhang, X. Che, X. Jiao, W. Yu, and L. Wan, "Quadratic optimization using additive homomorphic encryption in cps," in 2022 13th Asian Control Conference (ASCC). IEEE, 2022, pp. 1995–2000.
- [19] S. Han, W. K. Ng, L. Wan, and V. C. Lee, "Privacy-preserving gradientdescent methods," *IEEE transactions on knowledge and data engineering*, vol. 22, no. 6, pp. 884–899, 2009.
- [20] T. Toft, "Solving linear programs using multiparty computation," in International conference on financial cryptography and data security. Springer, 2009, pp. 90–107.
- [21] S. Salinas, C. Luo, W. Liao, and P. Li, "Efficient secure outsourcing of large-scale quadratic programs," in *Proceedings of the 11th* ACM on Asia Conference on Computer and Communications Security, ser. ASIA CCS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 281–292. [Online]. Available: https://doi.org/10.1145/2897845.2897862
- [22] C. Dwork, "Differential privacy," in International colloquium on automata, languages, and programming. Springer, 2006, pp. 1–12.