

Protecting Quantum Circuits Through Compiler-Resistant Obfuscation

Pradyun Parayil^{1*}, Amal Raj^{2*}, and Vivek Balachandran²

¹ Department of Electronics and Communication Engineering, Amrita School of Engineering, Coimbatore Amrita Vishwa Vidyapeetham, India, 641112

`cb.en.u4cce23067@cb.students.amrita.edu`

² Singapore Institute of Technology, Singapore

`{amal.raj,vivek.b}@singaporetech.edu.sg`

Abstract. Quantum circuit obfuscation is becoming increasingly important to prevent theft and reverse engineering of quantum algorithms. As quantum computing advances, the need to protect the intellectual property contained in quantum circuits continues to grow. Existing methods often provide limited defense against structural and statistical analysis or introduce considerable overhead. In this paper, we propose a novel quantum obfuscation method that uses randomized U3 transformations to conceal circuit structure while preserving functionality. We implement and assess our approach on QASM circuits using Qiskit AER, achieving over 93% semantic accuracy with minimal runtime overhead. The method demonstrates strong resistance to reverse engineering and structural inference, making it a practical and effective approach for quantum software protection.

Keywords: Quantum Obfuscation, Basis Conjugation, Randomized Basis Transformation, QASM Compatibility, Compiler-Resistant Quantum Security, Quantum Circuit Security

1 Introduction

Quantum computing leverages the principles of quantum mechanics, such as superposition, entanglement, and quantum interference, to perform computations that are infeasible for classical computers. It is poised to revolutionize fields like cryptography, where it can break conventional encryption schemes, optimization, where it tackles complex combinatorial problems, and quantum simulation, where it models molecular and material properties for drug discovery and materials science [1]. Quantum computers are used in research labs, cloud-based quantum platforms, and increasingly in industrial applications to solve specialized problems that require massive parallelism and probabilistic computation [2]. As quantum software development accelerates, protecting intellectual property (IP) and securing quantum circuits against reverse engineering and structural

* Both authors contributed equally to the work.

analysis has become a critical concern. Quantum programs, typically expressed in gate-level formats like QASM, reveal their underlying structure, making it easier for adversaries to extract proprietary logic or algorithmic optimizations. Developing robust obfuscation techniques that preserve functionality while concealing the circuit’s logical structure is essential for ensuring the security and commercial viability of quantum software in practical applications.

1.1 Proposed Idea

This work introduces a quantum circuit protection technique that resists critical scrutiny and ensures operational functionality. The obfuscation of quantum circuit topology is accomplished by applying random $U3$ gate transformations. The main idea is to perform a full circuit transformation by enclosing it with a random basis transformation and its dual. Each internal gate is transformed as well with this same transformation, thereby masking their identity while maintaining functionality. All the gates in the circuit are transformed through single-qubit $U3$ matrices, which hide the circuit’s inner workings without affecting its primary function.

This method makes sure the circuit’s output stays the same because the changes cancel each other out. Traditional ways of doing obfuscation add more gates or make the circuit hard to understand, but our method adds only negligible extra overhead. This keeps the circuit quick and efficient on quantum hardware [2]. It fits right in with any quantum program, does not require modifications to the original algorithm, and is simple to deploy across different quantum setups. Due to these merits, our method is a flexible and practical approach for safeguarding quantum software in real-world scenarios. The overall workflow is illustrated in Fig. 1.



Fig. 1. Workflow of the proposed quantum circuit obfuscation using $U3$ -conjugation. The process begins with a QASM circuit, applies a randomized basis transformation, transforms each gate, and restores the original basis to produce an obfuscated circuit with preserved functionality.

Paper Organization

The rest of this paper is structured as follows. Section 2 provides background on quantum computing, including qubits, quantum gates, compilation, and measurement operations. It also reviews existing work on quantum circuit obfuscation, highlighting current techniques and their limitations. Section 3 outlines the

proposed methodology and presents our obfuscation technique. Section 4 illustrates the methodology through the obfuscation of a QAOA circuit. Section 5 presents evaluation metrics and results, demonstrating the effectiveness of the proposed technique. Section 6 details the security guarantees provided by our approach, formalizing the threat model and the protections achieved. Section 7 discusses the limitations of our approach and proposes directions for future research. Section 8 concludes the paper by summarizing the key contributions and their implications for secure quantum software development.

2 Background

2.1 Quantum Computing Fundamentals

Quantum computing diverges from classical computing by employing quantum bits, or qubits, which leverage quantum mechanics to process information. Unlike classical bits, restricted to a binary state of 0 or 1, qubits can exist in a superposition, embodying multiple states simultaneously. This capability enables quantum computers to perform parallel computations, vastly exceeding classical systems for certain tasks. Additionally, qubits can be entangled, meaning the state of one qubit is directly tied to another, allowing instant correlations across distances. Such properties enable quantum algorithms to tackle complex problems in cryptography, optimization, and beyond with remarkable efficiency [1]. Quantum gates, the core of quantum computation, manipulate qubit states through reversible operations, represented as unitary matrices. For instance, the Hadamard gate creates superposition, while the controlled-NOT gate fosters entanglement, enabling intricate quantum interactions [3].

Quantum circuits, composed of sequences of quantum gates, orchestrate computations by evolving qubit states toward a desired outcome. A single measurement yields one classical outcome, whereas repeated executions of the circuit generate a probability distribution over possible results. Quantum Assembly Language (QASM), particularly its evolution from OPENQASM 2.0 to OPENQASM 3.0, standardizes circuit descriptions, supporting classical control flow and dynamic operations [4]. This facilitates circuit portability across quantum platforms but also raises concerns about protecting intellectual property embedded in these circuits. Compilation, a vital step for executing circuits on physical hardware, translates high-level designs into hardware-specific instructions. Tools like Qiskit’s transpiler optimize gate sequences and map qubits to physical hardware, minimizing circuit depth to enhance performance on Noisy Intermediate-Scale Quantum (NISQ) devices. Yet, this process exposes circuits to potential reverse-engineering, as adversaries could analyze gate structures to extract proprietary algorithms [5].

2.2 Related Works

Several approaches have been proposed to secure quantum circuits through obfuscation, each with distinct trade-offs. Suresh et al. introduced a method that

inserts a dummy gate at a specific circuit location to obscure its functionality, intentionally disrupting results during execution [6]. Post-compilation, this gate is removed to restore the original computation. To prevent optimization, barriers are placed around the dummy gate, and additional barriers are inserted around other gates to mask its identity from adversaries. However, this reliance on barriers compromises circuit optimization, increasing depth and degrading performance on NISQ hardware [6]. Similarly, Rehman, Langford and Liu proposed an obfuscation technique that modifies gate parameters using secret keys to conceal circuit logic [7]. While effective in hiding structure, this method requires significant classical processing for key management, complicating deployment and scalability, particularly for resource-constrained quantum systems. Another approach by Bartake et al. uses redundant gate insertion to obscure the intent of the circuit. This improves security, but the added gates significantly increase the depth of the circuit, making it more prone to errors on NISQ devices and limiting practical applicability [5]. These methods highlight the challenge of balancing security with performance, as increased complexity often undermines execution fidelity on current quantum hardware [2,6].

3 Methodology

Our quantum circuit obfuscation technique leverages basis transformations via U3 conjugation to obscure gate identities while preserving functional correctness. For clarity, the formulation presented in this paper adopts the weakest variant of the obfuscation framework, where a single random basis transformation is applied uniformly across the circuit. This setting intentionally models the worst case for security analysis and served as the baseline design during development. A strictly stronger variant assigns an independently sampled U3 basis to a randomly selected subset of gates, which removes global structural correlations and drives an adversary’s inference probability to negligible values. A detailed discussion of the resulting security guarantees is presented in Section 6.

The technique proceeds in four phases: input parsing, basis generation, gate transformation, and circuit reconstruction. Together, these stages preserve QASM compatibility, maintain semantic equivalence, and enhance resistance to reverse engineering.

3.1 QASM Input Parsing

The first step in our quantum circuit obfuscation process is to accurately interpret the user-provided quantum circuit, which is specified in Quantum Assembly Language (QASM). QASM serves as a standardized, platform-independent format for describing quantum circuits, including gate operations, qubit allocations, and measurement instructions [4]. To ensure compatibility with a wide range of quantum circuits, our system is designed to parse input strings written in either `OPENQASM 2.0` or the more advanced `OPENQASM 3.0` specification, accommodating both legacy and modern quantum programming workflows.

The input QASM string is first trimmed to remove any extra whitespace. Based on its prefix, the QASM version is identified: Qiskit uses `qiskit.qasm2.loads()` for OPENQASM 2.0 and `qiskit.qasm3.loads()` for OPENQASM 3.0, which supports advanced features like control flow and custom gate definitions [4,8]. If the version is unrecognized, an error is raised. This process produces a Quantum-Circuit object with the circuit’s gates, qubits, and measurements.

3.2 Basis Generation

The core of our obfuscation relies on a randomized basis transformation applied across all qubits. We generate random parameters ($\theta \in [0, \pi]$, $\phi \in [0, 2\pi]$, $\lambda \in [0, 2\pi]$) for a U3 gate using a uniform random distribution. The U3 gate is a universal single-qubit gate that can represent any unitary operation by tuning its parameters, making it well-suited for implementing a wide range of basis transformations. Randomization ensures that each circuit instantiation produces a unique basis, enhancing obfuscation by preventing predictable patterns that could be exploited in reverse engineering. As shown in Equation 1, the gate’s matrix representation is defined by these parameters.

$$U3(\theta, \phi, \lambda) = \begin{bmatrix} \cos\left(\frac{\theta}{2}\right) & -e^{i\lambda} \sin\left(\frac{\theta}{2}\right) \\ e^{i\phi} \sin\left(\frac{\theta}{2}\right) & e^{i(\phi+\lambda)} \cos\left(\frac{\theta}{2}\right) \end{bmatrix} \quad (1)$$

3.3 Gate Transformation

To obfuscate a quantum circuit, each gate G in the original circuit is transformed using unitary conjugation, expressed as $U^\dagger \cdot G \cdot U$, where U is the unitary matrix of a U3 gate with randomly chosen parameters (θ, ϕ, λ). Measurements and resets are left unchanged because the obfuscation method applies only to unitary operations. Since measurement is non-unitary, conjugating it with U^\dagger and U would rotate the measurement basis and alter the circuit’s behavior. For circuits containing mid-circuit measurements, we simply treat the computation as a sequence of unitary segments separated by measurement boundaries and obfuscate each unitary segment independently.

Using Qiskit’s `Operator` class, the gate G is first converted into its matrix representation to facilitate this transformation. For multi-qubit gates, the basis matrix U is constructed by computing the Kronecker product of single-qubit U3 matrices for each qubit. The Kronecker product, a tensor operation that combines matrices to form a larger matrix, ensures that the transformation scales appropriately for systems with multiple qubits, preserving the gate’s functionality across the circuit’s qubit register [9]. The resulting matrix $U^\dagger \cdot G \cdot U$ is then encapsulated as a `UnitaryGate` with a unique label, such as `Obf_H_0` for a Hadamard gate, effectively concealing the original gate’s identity. This process is applied to all gates, as detailed in Algorithm 1.

Algorithm 1 Gate Transformation via U3-Conjugation**Require:** Gate G , qubits, U3 parameters (θ, ϕ, λ)

- 1: Convert G to matrix using `Operator(G)`
- 2: Compute $U = U3(\theta, \phi, \lambda)^{\otimes n}$ for n qubits
- 3: Compute $U^\dagger = U3(-\theta, -\lambda, -\phi)^{\otimes n}$
- 4: Compute transformed matrix $U^\dagger \cdot G \cdot U$
- 5: Encapsulate as `UnitaryGate` with label `Obf_G_id`
- 6: Replace original gate G with the transformed `UnitaryGate` in the circuit

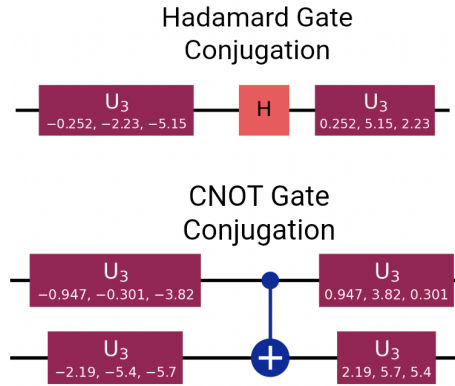


Fig. 2. Diagram of the conjugation process $U^\dagger \cdot G \cdot U$ for a single-qubit gate (Hadamard) and a multi-qubit gate (CNOT). This figure illustrates the underlying matrix operations and does not imply additional quantum gates or circuit depth. The result of this matrix multiplication ($U^\dagger \cdot G \cdot U$) is encapsulated as a single, opaque `UnitaryGate` object, thereby incurring no space overhead in the quantum circuit.

3.4 Basis Transformation and Obfuscation

For a circuit with n qubits, a global basis transformation $U = U3(\theta, \phi, \lambda)^{\otimes n}$ is applied to all qubits at the circuit's start, with its inverse $U^\dagger = U3(-\theta, -\lambda, -\phi)^{\otimes n}$ applied at the end. The parameters θ , ϕ , and λ are randomly sampled from a uniform distribution to maximize obfuscation. For each gate G , a local basis transformation is applied, replacing G with its conjugated form $G' = U3(-\theta, -\lambda, -\phi) \cdot G \cdot U3(\theta, \phi, \lambda)$, which encapsulates the original gate in a new basis while preserving its functionality. This transformation is implemented using Qiskit's circuit manipulation tools, where each gate is directly replaced with its transformed version $U^\dagger G U$. To ensure the circuit's original functionality, the inverse basis transformation U^\dagger is applied across all qubits as a single instruction labeled `InvBasis` after all gates are processed. Barriers from the original circuit are retained to maintain structural hints for hardware compatibility. The reconstructed circuit is then transpiled and executed using the Qiskit AER simulator, producing results statistically equivalent to the original circuit. Fig. 2 illustrates

this process with examples of Hadamard and CNOT gate transformations. The implementation code is available in the GitHub repository in [10].

3.5 Preservation of Functionality: Hadamard and CNOT Gate Examples

To illustrate how our obfuscation technique preserves the functionality of quantum circuits, we examine the transformation of two fundamental quantum gates: the single-qubit Hadamard (H) gate and the two-qubit controlled-NOT (CNOT) gate. These gates are critical building blocks in quantum algorithms, with the H gate creating superposition and the CNOT gate introducing entanglement [1]. By applying our basis conjugation approach to each gate individually, we demonstrate that their computational behavior remains unchanged despite structural obfuscation.

For the Hadamard gate, represented by the 2×2 unitary matrix $H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$, we apply a basis transformation using a U3 gate with randomly chosen parameters (θ, ϕ, λ) . The transformation is defined as $U = U3(\theta, \phi, \lambda)$, and the obfuscated gate is computed as:

$$H_{\text{obf}} = U^\dagger \cdot H \cdot U \quad (2)$$

To verify preservation of functionality, we apply the inverse transformation:

$$U \cdot H_{\text{obf}} \cdot U^\dagger = U \cdot (U^\dagger \cdot H \cdot U) \cdot U^\dagger = (UU^\dagger) \cdot H \cdot (UU^\dagger) = H \quad (3)$$

This confirms that the obfuscated Hadamard gate, when wrapped by the basis transformation and its inverse, behaves identically to the original H gate, preserving its ability to create superposition. The matrix-level transformation of the Hadamard gate is shown in Fig. 3.

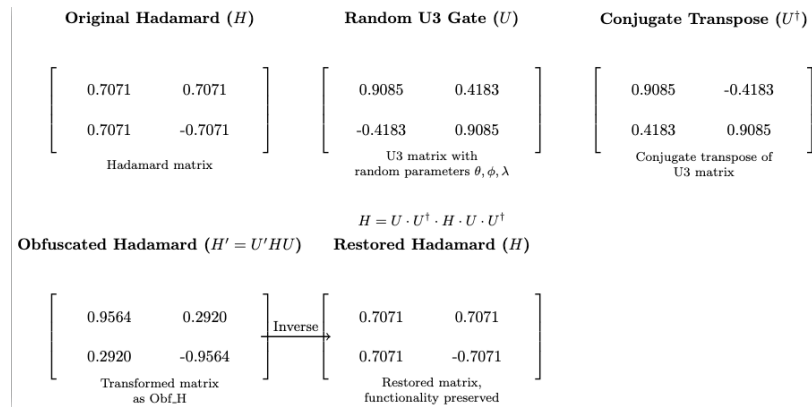


Fig. 3. Matrix-level transformation of a Hadamard gate using U3 conjugation. The functionality remains unchanged due to the inverse transformation restoring the original gate.

Pradyun et al.

For the two-qubit CNOT gate, represented by a 4×4 unitary matrix, we use a basis transformation $U = U3(\theta, \phi, \lambda)^{\otimes 2}$, formed by the Kronecker product of two single-qubit U3 matrices to scale the transformation across both qubits [9]. The obfuscated CNOT gate is:

$$\text{CNOT}_{\text{obf}} = U^\dagger \cdot \text{CNOT} \cdot U \quad (4)$$

Applying the inverse transformation yields:

$$\begin{aligned} U \cdot \text{CNOT}_{\text{obf}} \cdot U^\dagger &= U \cdot (U^\dagger \cdot \text{CNOT} \cdot U) \cdot U^\dagger \\ &= (UU^\dagger) \cdot \text{CNOT} \cdot (UU^\dagger) \\ &= \text{CNOT} \end{aligned} \quad (5)$$

This demonstrates that the CNOT gate’s entanglement-generating functionality is fully preserved, as the basis transformations cancel out, leaving the original gate’s behavior intact. This process, applied to individual gates, ensures that the circuit’s overall semantics remain unchanged, as the global basis shift and its inverse cancel out during execution. The corresponding matrix-level transformation for the CNOT gate is shown in Fig. 4.

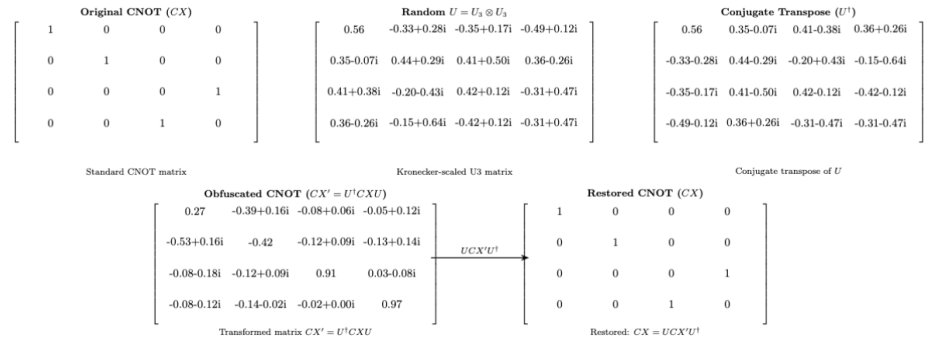


Fig. 4. Matrix-level transformation of a CNOT gate using U3 conjugation, with the Kronecker product scaling the basis for two qubits.

4 Case Study: Quantum Approximate Optimization Algorithm (QAOA)

4.1 Overview of QAOA

The Quantum Approximate Optimization Algorithm (QAOA) is a hybrid quantum-classical algorithm designed for solving combinatorial optimization problems such as MaxCut and Max3SAT [11]. It combines parameterized quantum circuits with classical optimization to approximate the solution to a given cost function. QAOA is particularly relevant due to its practical applicability on near-term

quantum devices and its structured circuit layout, making it a suitable candidate for testing circuit obfuscation techniques.

In this case study, we focus on a QAOA circuit targeting the MaxCut problem on a 5-node graph. The circuit involves alternating applications of the problem (cost) unitary and a mixing unitary, parameterized by angles γ and β , respectively. We implement this circuit in `OPENQASM 2.0`, simulate its behavior before and after obfuscation, and evaluate functional correctness and performance overhead.

4.2 Circuit Description

The QAOA circuit is constructed with five qubits representing the five graph nodes in the MaxCut problem, where the edges are $(0, 1)$, $(1, 2)$, $(1, 3)$, $(2, 4)$ and $(3, 4)$, as shown in Fig. 5. The initial state is prepared using Hadamard gates to create a uniform superposition over all possible bitstrings. The cost unitary is implemented using two-qubit R_{ZZ} rotations (decomposed into `CX-RZ-CX`) corresponding to the edges in the graph. Each edge contributes a term in the cost Hamiltonian, and the circuit applies these terms with an optimized parameter γ . This is followed by a layer of R_X gates applied to each qubit, representing the mixer Hamiltonian and controlled by the parameter β .

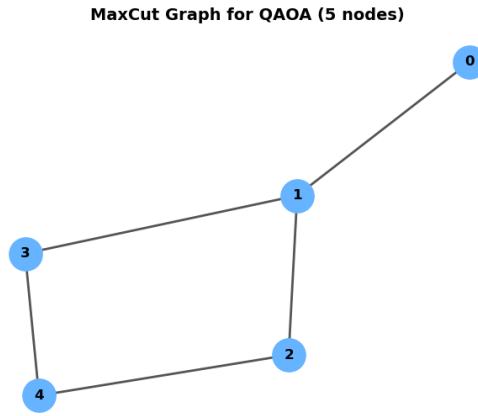


Fig. 5. MaxCut graph corresponding to the 5-node QAOA circuit.

In this implementation, the QAOA depth is $p = 1$, meaning one alternating layer of cost and mixer unitaries is applied. The parameters were optimized via classical simulation to $\gamma \approx 0.865$ and $\beta \approx 0.457$, producing a strong probability concentration on the optimal MaxCut solutions. The corresponding QAOA circuit structure for this problem is illustrated in Fig. 6. The two optimal MaxCut bitstrings for this graph are 01001 and 10110.




Fig. 6. QAOA circuit diagram for MaxCut on a 5-node graph, depth $p = 1$ with optimized parameters.

4.3 Circuit Obfuscation

To apply our obfuscation methodology (Section 3) to the QAOA circuit, we generated a random single-qubit basis transformation using $U3(\theta, \phi, \lambda)$ with parameters $\theta = 2.86$, $\phi = 2.33$, and $\lambda = 0.762$. This unitary was then extended to the five-qubit system as a global basis $U = U3(2.86, 2.33, 0.762)^{\otimes 5}$ applied across all qubits at the beginning of the circuit, with its inverse $U^\dagger = U3(-2.86, -0.762, -2.33)^{\otimes 5}$ appended at the end to preserve the circuit's overall functionality. Each gate G in the QAOA layers was obfuscated using the conjugation $G' = U^\dagger G U$ and encapsulated as an opaque `UnitaryGate` with a unique label (e.g., `Obf_H_0`). The transformed structure of the obfuscated QAOA circuit is illustrated in Fig. 7.

The resulting obfuscated circuit is functionally equivalent to the original but structurally unreadable, providing resistance against reverse engineering. Both the original and obfuscated circuits were then tested under identical simulation settings. The measurement outcomes for both the original and obfuscated QAOA circuits are shown in Fig. 8.

Protecting Quantum Circuits Through Compiler-Resistant Obfuscation



Fig. 7. Obfuscated QAOA circuit diagram, showing the transformed gate structure.




Fig. 8. Histogram of measurement outcomes for the original and obfuscated QAOA circuits. Optimal solutions are highlighted.

4.4 Testing Process

To evaluate the obfuscation technique, both the original and obfuscated QAOA circuits were executed using Qiskit’s `AerSimulator` (version 0.36.0). Each configuration was run with 1,024 shots to ensure sufficient sampling. The output distributions were recorded and plotted using histograms. Execution times were also measured to detect any latency introduced by obfuscation.

The comparison focused on:

- i. **Semantic accuracy** — probability assigned to each bitstring in the obfuscated circuit compared to the original.
- ii. **Total Variation Distance (TVD)** — quantifying statistical difference between distributions.
- iii. **Functional equivalence** — whether optimal solutions remain present with comparable probability.
- iv. **Execution time** — runtime overhead introduced by the obfuscation process

4.5 Results

Both circuits produce highly similar output distributions, with only minor statistical variation due to sampling noise. The semantic accuracy of the obfuscated circuit reached **93.30%**, and the TVD was calculated to be **0.0840**, indicating strong fidelity to the original behavior. The two optimal MaxCut bitstrings (01001 and 10110) emerged as the two highest-probability peaks in both runs, dominating the output distributions with comparable counts. Execution time measurements showed negligible overhead introduced by the obfuscation process.

These results confirm that the obfuscation strategy preserves the computational intent and functional output of the original circuit, with the chosen optimized QAOA parameters successfully amplifying the probability of the optimal solutions.

4.6 Implementation Setup

The simulations were conducted on a standard Macbook Air M1 with 16 GB RAM and MacOS 18.6. To provide a comprehensive evaluation of our obfuscation technique, Quantum Approximate Optimization Algorithm (QAOA) for combinatorial optimization was tested alongside nine other quantum circuits, representing a diverse set of quantum algorithms: Shor’s algorithm for integer factorization [12], the Bernstein-Vazirani (BV) algorithm for learning hidden linear functions [13], Grover’s algorithm for unstructured search [14], the Variational Quantum Eigensolver (VQE) for quantum chemistry simulations [15], the Phase Kickback Circuit [16], the Deutsch-Jozsa (DJ) algorithm for classifying balanced or constant functions [17], the Quantum Fourier Transform (QFT) for periodicity analysis [18], and Simon’s algorithm for finding periods in functions [19]. These algorithms, spanning applications from cryptography to quantum

simulations, provide a robust testbed for assessing the obfuscation method’s versatility.

Each circuit, including QAOA, was executed 100 times with 1,024 shots per execution under noise-free conditions using Qiskit’s AerSimulator. The Toffoli circuit used three qubits, while the other algorithms ranged from 4 to 12 qubits, depending on their complexity.

5 Evaluation Metrics

These circuits were executed under noise-free conditions using Qiskit’s AerSimulator to ensure consistent and reproducible results. The evaluation focuses on three primary metrics: *semantic accuracy*, *Total Variation Distance (TVD)*, and *execution time*, which together assess functional fidelity, statistical similarity, and computational overhead.

Semantic Accuracy

Semantic accuracy counts the overlapping probability mass between the two distributions, expressed as a percentage of the total probability in the original circuit’s output. It quantifies whether the obfuscated circuit preserves the functional behavior of the original circuit under realistic, finite-shot quantum execution. This metric was used by Bartake et al. to quantify preservation of circuit behavior after their circuit-level transformations [5]. Given two measurement distributions `original` and `obfuscated`, the semantic accuracy is calculated as:

$$\text{Semantic Accuracy} = \frac{\sum_{x \in X} \min(\text{original}[x], \text{obfuscated}[x])}{\sum_{x \in X} \text{original}[x]} \times 100\% \quad (6)$$

where X is the union of all bitstrings observed in either distribution.

Since quantum measurements are inherently probabilistic, two runs of the same circuit will exhibit slight differences in their empirical output distributions unless an infinite number of shots is used. Semantic accuracy captures this agreement by computing the fraction of output bitstrings whose measured probabilities in the obfuscated circuit fall within the statistical variation of those from the original circuit.

Values slightly below 100% are therefore expected even for functionally identical circuits and do not indicate any semantic deviation introduced by the obfuscation. Instead, they reflect natural run-to-run statistical fluctuations. In the idealized infinite-shot limit, the empirical distributions of equivalent circuits would converge exactly, and semantic accuracy would approach 100%. High semantic accuracy in our experiments (typically above 90%) demonstrates that the obfuscated circuits remain consistent with the intended computation within the limits of finite-shot quantum sampling.

Total Variation Distance (TVD)

Total Variation Distance (TVD) is a statistical measure of the difference between two probability distributions [20]. In this work, TVD is used to quantify the effect of gate additions introduced during obfuscation. TVD is also a standard and widely used metric for comparing output distributions of quantum circuits, and it has been adopted in several prior works on quantum obfuscation [21,22]. It is computed as the sum of the absolute differences between the measurement counts of the original and obfuscated circuits, normalized by the total number of shots, as given by:

$$\text{TVD} = \frac{\sum_i |x_{i,\text{obfus}} - x_{i,\text{orig}}|}{2 \times \text{Number of shots}} \quad (7)$$

where $x_{i,\text{orig}}$ and $x_{i,\text{obfus}}$ denote the counts of the i^{th} measurement outcome for the original and obfuscated circuits, respectively. A TVD of 0 means the two output distributions are exactly the same, showing no statistical difference at all, while a TVD of 1 means the distributions are completely distinct with no overlap, representing the maximum possible statistical difference.

Execution Time

Execution time refers to the duration required by the simulator to execute the circuit for a fixed number of shots (1024 in our case), excluding compilation time. It serves as a practical measure of the computational overhead introduced by circuit-level transformations, and its use is consistent with prior work in both quantum and classical settings. Studies on quantum program performance explicitly analyze execution time as a key metric to evaluate runtime and performance overhead in simulators and hardware [23], while recent work in binary obfuscation likewise evaluates execution time overhead to quantify the performance cost of structural code transformations [24]. Our results show that the overhead is minimal (typically < 1 ms), making the approach practical for real-world quantum workloads.

5.1 Obfuscation’s Compiler Resistance: A Key Advantage

A core strength of our U3-conjugation technique lies in its ability to effectively conceal the original gate identities from quantum compilers and reverse-engineering tools. As demonstrated in our QAOA case study, when the obfuscated circuit is passed to a transpiler or compiler (such as Qiskit’s transpiler), the individual original gates (e.g., Hadamard, CNOT) are no longer recognizable as their standard forms. Instead, they appear as opaque `UnitaryGate` objects, each representing the complex, randomized matrix resulting from the $U^\dagger \cdot G \cdot U$ transformation. Our method’s compiler resistance stems from this deliberate lack of identifiable internal structure. Adversaries are forced to treat each transformed gate as an unknown unitary operation, greatly hindering reverse engineering of the proprietary quantum logic. Fig. 9 visually exemplifies this transformation from a recognizable gate to an opaque block.

Protecting Quantum Circuits Through Compiler-Resistant Obfuscation

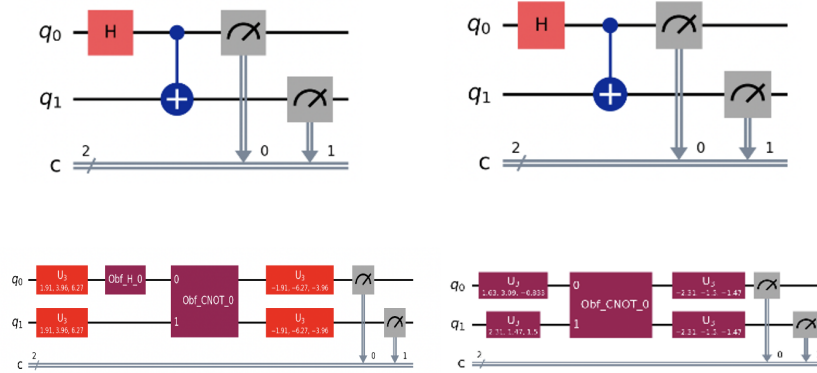


Fig. 9. Illustration of compiler resistance: On the left, the original Bell state circuit and its obfuscated counterpart. On the right, their respective compiled forms. While the original circuit structure is preserved after compilation, the obfuscated circuit appears as a single opaque UnitaryGate, concealing its functionality from quantum compilers and analysis tools

5.2 Overhead Analysis

To make the overhead of the obfuscation explicit, we summarize both the empirical execution-time overhead and the formal structural overhead introduced by the transformation. The execution-time results in Table 1 and Fig. 12 show that the runtime increase introduced by obfuscation is extremely small, typically well under one millisecond.

Consider an original circuit with m gates acting on n qubits. For each gate, we insert two single-qubit U_3 operations, resulting in an additional

$$G_m = 2m$$

single-qubit gates surrounding the original operations.

We further apply n single-qubit U_3 gates at the beginning and n at the end of the circuit, contributing:

$$G_n = 2n$$

additional gates.

Before fusion, the total number of gates is therefore

$$G_{\text{total, pre-fuse}} = m + 2m + 2n = 3m + 2n$$

The $3m$ gates surrounding the original m positions are fused into m effective gates by multiplying the three matrices acting at each position:

$$3m \longrightarrow m,$$

Pradyun et al.

yielding a final structural gate count of

$$G_{\text{final}} = m + 2n$$

The two global boundary layers contribute a fixed depth increase:

$$\Delta D = 2$$

For example, in the QAOA case study, the computational overhead was negligible because each transformed operation is represented as a single opaque unitary during simulation, contributing no additional internal cost. Because the structural additions are constant and independent of circuit size, the empirical overhead remains negligible even for large workloads.

5.3 Results Summary

Table 1 presents the evaluation metrics for all tested circuits, including QAOA. The results show high semantic accuracy ($> 93\%$) and low TVD (≈ 0) across all circuits, confirming that the obfuscation preserves functionality. Execution times remain comparable, with minimal overhead, highlighting the method’s efficiency. Visualizations in Fig. 10, 11, and 12 further illustrate the consistency of TVD, accuracy, and execution times, reinforcing the robustness of the obfuscation technique across diverse quantum algorithms.

Table 1. Evaluation metrics across quantum circuits.

Circuit	Original Time (s)	Obfuscated Time (s)	Semantic Accuracy (%)	TVD
BV (1011)	0.00309	0.00390	100.00	0.0000
DJ	0.00141	0.00147	100.00	0.0000
Grover (101)	0.00148	0.00148	97.17	0.0302
Phase Kickback	0.00165	0.00151	100.00	0.0000
QAOA	0.00156	0.00166	93.30	0.0840
QFT	0.00147	0.00178	94.73	0.0770
Shor	0.00179	0.00201	96.68	0.0146
Simon	0.00266	0.00183	94.24	0.0380
Toffoli	0.00167	0.00168	100.00	0.0000
VQE	0.00125	0.00132	99.51	0.0048

Protecting Quantum Circuits Through Compiler-Resistant Obfuscation




Fig. 10. Distribution of Total Variation Distance (TVD) across circuits, with low values indicating strong statistical similarity between original and obfuscated outputs.

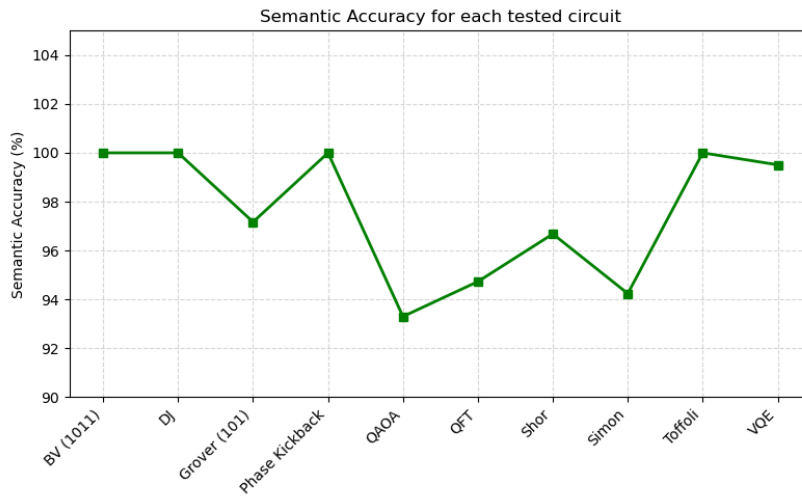


Fig. 11. Semantic accuracy for each tested circuit, with high values demonstrating functional correctness of obfuscated versions.

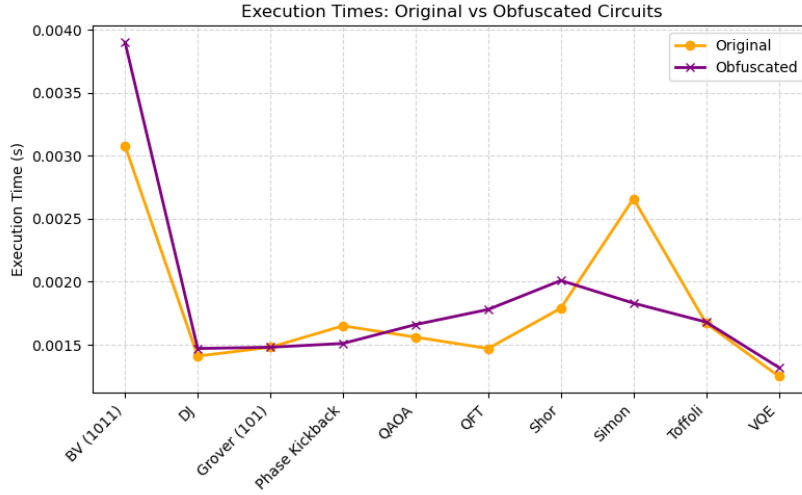


Fig. 12. Comparison of execution times between original and obfuscated quantum circuits, showing minimal overhead.

6 Security Guarantee

We quantify the difficulty of deobfuscation under two natural cases:

6.1 Black-box adversary (no circuit access)

In this model, the adversary does not see the obfuscated circuit description. Instead, they have only black-box access to its input-output behavior (e.g., they can prepare input states and observe measurement statistics). Our obfuscation replaces each protected single-qubit gate G by a basis-conjugated form $U^\dagger \cdot G' \cdot U$, where $G' = U \cdot G \cdot U^\dagger$ and the basis rotation U is sampled from the continuous $U3$ parameter domain $(\theta, \phi, \lambda) \in [0, \pi] \times [0, 2\pi) \times [0, 2\pi)$. An adversary who does not see the circuit would need to determine these continuous parameters solely from observed input-output statistics. Because the parameter space is uncountable, the probability of exactly guessing the correct triplet (θ, ϕ, λ) approaches 0.

Under a practical finite-precision discretization of step size δ (that is, each angle quantized in increments of δ), the number of possible basis rotations is on the order of

$$\left(\frac{2\pi}{\delta}\right)^3$$

and therefore a random-guessing adversary succeeds with probability at most

$$\left(\frac{\delta}{2\pi}\right)^3$$

which is negligible for realistic values of δ . Thus, in the black-box setting, both full and partial obfuscation lead to a difficult continuous parameter-recovery problem.

6.2 White-box adversary (circuit access)

In this stronger model, the adversary has access to the full obfuscated circuit (e.g., QASM or a Qiskit circuit) and can directly observe the explicit gates U_i^\dagger , G'_i , and U_i inserted around protected gates. If all n gates in the circuit are obfuscated, the obfuscation pattern is trivial: the adversary may assume every gate is protected and straightforwardly recover the original gate as

$$G_i = U_i^\dagger \cdot G'_i \cdot U_i,$$

thus succeeding with probability 1.

The problem becomes non-trivial when only a selected subset of gates is obfuscated. Let the original circuit contain n gates, and let the obfuscator randomly choose a hidden subset $S \subseteq \{1, \dots, n\}$ of size $|S| = x$. Only positions in S are replaced by basis-conjugated blocks ($U_i^\dagger \cdot G'_i \cdot U_i$); the remaining gates stay unchanged. A white-box adversary must therefore first determine which gate positions belong to S before applying the algebraic simplification above.

If S is sampled uniformly from all $\binom{n}{x}$ choices of size x , then any adversary who attempts to guess the full obfuscation pattern in one shot succeeds with probability

$$p_{\text{succ}}(n, x) = \Pr[\widehat{S} = S] = \frac{1}{\binom{n}{x}}$$

The corresponding min-entropy of the obfuscation pattern is

$$H_\infty(S) = -\log_2 p_{\text{succ}}(n, x) = \log_2 \binom{n}{x}$$

This value is 0 in the trivial cases $x = 0$ (no obfuscation) and $x = n$ (all gates obfuscated), but strictly positive for $1 \leq x \leq n - 1$, and maximised when $x \approx n/2$, where $H_\infty(S) \approx n$ bits. Thus:

1. When all gates are obfuscated ($x = n$), a white-box adversary can de-obfuscate by simplifying every block.
2. When only a randomly selected subset of gates is obfuscated ($1 \leq x \leq n - 1$), the adversary must identify which gates are protected; this requires distinguishing among $\binom{n}{x}$ possible obfuscation patterns. The success probability $\frac{1}{\binom{n}{x}}$ is exponentially small when $x \approx n/2$.

7 Limitations and Future Work

While the proposed $U3$ -based basis conjugation technique effectively preserves functional correctness and provides strong resistance to compiler-level reverse engineering, it introduces structural overhead that impacts scalability on near-term quantum hardware.

7.1 Limitations

The obfuscation process can be computationally intensive, especially when applying basis transformations to multi-qubit gates. Constructing these transformations using Kronecker products, as outlined in [9], involves manipulating high-dimensional matrices (such as 8×8 for three-qubit gates), which is feasible for small circuits (e.g., Toffoli). However, for deeper circuits, the large number of gates leads to significant computational costs from repeated matrix operations, limiting scalability [25]. Generating new random $U3$ parameters for each gate further increases time and memory usage, posing challenges for algorithms with high depth or width.

As analyzed in Section 5.2, the obfuscation process replaces each protected gate with a conjugated form that introduces additional single-qubit basis rotations, together with two global basis transformation layers applied at the circuit boundaries. Under ideal, noise-free simulation, these additional operations are algebraically fused into opaque unitary blocks, resulting in negligible execution-time overhead and no loss of functional equivalence. Consequently, ideal simulation masks the practical impact of this added structure.

When deployed on NISQ hardware, the additional basis operations introduced by conjugation must be realized as physical gate executions with nonzero error probabilities. Each single-qubit rotation corresponds to a calibrated control pulse whose imperfections accumulate as circuit depth increases, imposing a practical scalability limit even when individual gate fidelities are relatively high. For example, a two-qubit CNOT gate in a depth- D circuit is transformed under basis conjugation into $U^\dagger \cdot \text{CNOT} \cdot U$, introducing single-qubit basis rotations on both qubits before and after the entangling operation. While these rotations are algebraically fused into a single opaque unitary during simulation, on physical hardware they must be executed as separate calibrated pulses. Consequently, a circuit containing m such entangling gates incurs $O(m)$ additional single-qubit operations, causing errors to accumulate with depth and limiting scalability for deep circuits such as the Quantum Fourier Transform (QFT) or Variational Quantum Eigensolvers (VQE), despite preserved functional correctness in ideal execution.

Finally, the current implementation does not exploit hardware-specific gate sets or connectivity constraints, which could mitigate some of the depth overhead through optimized decompositions. Incorporating hardware-aware synthesis and calibration-informed transformations remains an important direction for improving scalability on NISQ platforms.

7.2 Future Work

To overcome these challenges, we plan to pursue several targeted improvements to enhance the technique’s scalability, hardware compatibility, and security. First, to reduce the computational burden of matrix calculations for deep circuits like QFT, or VQE, we will cache precomputed U3-transformed gates (e.g., $U^\dagger \cdot G \cdot U$ for Hadamard or CNOT), minimizing repetitive Kronecker product operations. Second, we will investigate methods to reduce the fidelity losses introduced by the extra single-qubit rotations in the conjugated form $U \cdot G \cdot U^\dagger$, including hardware-calibrated decompositions and approximation-aware synthesis to limit error accumulation. In addition, to better support NISQ devices, we will tailor U3 transformations to match the native gate sets and connectivity of specific quantum hardware, reducing circuit depth and errors. Finally, we also plan to add special gates that trigger errors if tampered with, helping detect unauthorized analysis attempts. These enhancements will make the technique more efficient and secure, supporting applications like protecting intellectual property in cloud-based quantum computing to visualize secure quantum circuits.

8 Conclusion

Our quantum circuit obfuscation method, built on U3 gate conjugation, hides the structure of quantum circuits while preserving functionality. By applying randomized basis transformations, it ensures compatibility with OPENQASM 2.0 and OPENQASM 3.0. The evaluation presented in this work uses the conservative single-basis variant, while the per-gate randomized variant provides strictly stronger security without altering the functional behavior of the circuit.

We tested the method on a diverse set of circuits—Shor’s algorithm, QAOA, Bernstein-Vazirani, Grover’s algorithm, VQE, Phase Kickback, Deutsch-Jozsa, QFT, Toffoli, and Simon’s algorithm—covering applications from cryptography to quantum chemistry. The obfuscated circuits achieve an average semantic accuracy of at least 93% and a Total Variation Distance (TVD) below 0.035, indicating minimal changes in output distributions. Using Qiskit’s `UnitaryGate` to implement transformed gates ($U^\dagger \cdot G \cdot U$) enables seamless integration into quantum circuits, with runtime overhead typically under 1 ms.

These results demonstrate that the method is both practical and versatile, spanning circuits from small Toffoli gates to complex algorithms like QFT and Shor. Future enhancements will focus on improving scalability for larger algorithms, optimizing transformations for NISQ hardware. These advances will strengthen the method’s utility for protecting commercial quantum software, enabling secure cloud execution in quantum computing.

Declaration of Competing Interests

The authors have no competing interests to declare relevant to this article.

References

1. Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press, Cambridge (2010)
2. Preskill, J.: Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (aug 2018). <https://doi.org/10.22331/q-2018-08-06-79>
3. Rieffel, E.G., Polak, W.H.: Quantum Computing: A Gentle Introduction. MIT Press (2011)
4. Cross, A.W., Bishop, L.S., Smolin, J.A., Gambetta, J.M.: Open quantum assembly language (2017), <https://arxiv.org/abs/1707.03429>
5. Bartake, N., Jie, S.T.Z., Jiawen, C.W., Ren, D.F.Y., Kasper, M., Balachandran, V.: Obfusqate: Unveiling the first quantum program obfuscation framework. In: Proceedings of the 2025 Workshop on Research on Offensive and Defensive Techniques in the Context of Man At The End (MATE) Attacks. p. 10–19. CheckMATE '25, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3733817.3762700>
6. Suresh, A., Saki, A.A., Alam, M., Topalaglu, R., Ghosh, S.: A quantum circuit obfuscation methodology for security and privacy (2021), <https://arxiv.org/abs/2104.05943>
7. Rehman, A., Langford, V., Liu, Y.: Opaque: Obfuscating phase in quantum circuit compilation for efficient ip protection (2025), <https://arxiv.org/abs/2502.16605>
8. Cross, A., Javadi-Abhari, A., Alexander, T., De Beaudrap, N., Bishop, L.S., Heide, S., Ryan, C.A., Sivarajah, P., Smolin, J., Gambetta, J.M., Johnson, B.R.: Openqasm3: A broader and deeper quantum assembly language. *ACM Transactions on Quantum Computing* **3**(3), 1–50 (Sep 2022). <https://doi.org/10.1145/3505636>
9. Laub, A.J.: Matrix Analysis for Scientists and Engineers. SIAM, Philadelphia (2005)
10. Parayil, P., Raj, A., Balachandran, V.: Basis Transformation, <https://github.com/vivekianity/Basis-Transformation>
11. Farhi, E., Goldstone, J., Gutmann, S.: A quantum approximate optimization algorithm (2014), <https://arxiv.org/abs/1411.4028>
12. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* **26**(5), 1484–1509 (oct 1997). <https://doi.org/10.1137/S0097539795293172>
13. Bernstein, E., Vazirani, U.: Quantum complexity theory. *SIAM Journal on Computing* **26**(5), 1411–1473 (oct 1997). <https://doi.org/10.1137/S0097539796300921>
14. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th Annual ACM Symposium on Theory of Computing. pp. 212–219 (1996). <https://doi.org/10.1145/237814.237866>
15. Peruzzo, A., et al.: A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* **5**, 4213 (2014). <https://doi.org/10.1038/ncomms5213>
16. Cleve, R., Ekert, A., Macchiavello, C., Mosca, M.: Quantum algorithms revisited. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**(1969), 339–354 (Jan 1998). <https://doi.org/10.1098/rspa.1998.0164>
17. Deutsch, D., Jozsa, R.: Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A* **439**(1907), 553–558 (nov 1992). <https://doi.org/10.1098/rspa.1992.0167>

Protecting Quantum Circuits Through Compiler-Resistant Obfuscation

18. Coppersmith, D.: An approximate fourier transform useful in quantum factoring (2002), <https://arxiv.org/abs/quant-ph/0201067>
19. Simon, D.R.: On the power of quantum computation. *SIAM Journal on Computing* **26**(5), 1474–1483 (oct 1997). <https://doi.org/10.1137/S0097539796298637>
20. Verdú, S.: Total variation distance and the distribution of relative information. In: 2014 Information Theory and Applications Workshop (ITA). pp. 1–6. San Diego, CA, USA (2014). <https://doi.org/10.1109/ITA.2014.6804281>
21. Wang, Q., John, J., Dong, B., Liu, Y.: Tetrislock: Quantum circuit split compilation with interlocking patterns (2025), <https://arxiv.org/abs/2503.11982>
22. Liu, Y., John, J., Wang, Q.: E-loq: Enhanced locking for quantum circuit ip protection. In: 2025 IEEE International Symposium on Hardware Oriented Security and Trust (HOST). p. 67–77. IEEE (May 2025). <https://doi.org/10.1109/host64725.2025.11050014>
23. Ma, N., Li, H.: Understanding and estimating the execution time of quantum circuits (2025), <https://arxiv.org/abs/2411.15631>
24. Zhang, P., Wu, C., Hu, H., Jia, L., Peng, M., Xu, J., Xie, M., Lai, Y., Kang, Y., Wang, Z.: Shining light on the inter-procedural code obfuscation: Keep pace with progress in binary diffing. *ACM Trans. Archit. Code Optim.* **22**(1) (Mar 2025). <https://doi.org/10.1145/3701992>
25. Zhou, Y., Stoudenmire, E.M., Waintal, X.: What limits the simulation of quantum computers? *Phys. Rev. X* **10**, 041038 (Nov 2020). <https://doi.org/10.1103/PhysRevX.10.041038>