Counting Bases from Number of Qubits: Inferring VRP from Quantum Circuits

Jessie Chen
Yale University
New Haven, CT, USA
zixin.chen@yale.edu

Jakub Szefer
Yale University
New Haven, CT, USA
jakub.szefer@yale.edu

***Abstract*—The security and confidentiality of sensitive information processed by quantum computers are of paramount importance, especially given quantum computers' potential to efficiently solve classically-hard optimization problems. At the heart of many transport optimization tasks lies the Vehicle Routing Problem (VRP), a complex combinatorial optimization problem classified as NP-hard. However, a promising avenue for approximating solutions to VRP is found in the Quantum Approximate Optimization Algorithm (QAOA). This paper demonstrates how leaking and learning simple parameters from QAOA quantum circuit structures, enables attackers to learn about the problem being solved. In routing optimization scenarios used by the military, for example, the attacker can learn location or connection of military bases. By exploiting information leakage during the QAOA execution, attackers can potentially breach security and retrieve sensitive VRP details, posing profound implications for civilian and national security.**

***Index Terms*—quantum computing, optimization algorithms, information leaks, security**

I. Introduction

Quantum computers promise to deliver exponential speedups over their classical counterparts for certain classes of computational problems [9], [19], [21], [32]. Among the most notable examples are eigenvalue and optimization problems, which have key applications in finance, machine learning, and simulations of quantum chemistry [4], [10], [24], [25], [31]. Moreover, certain NP-hard combinatorial problems like the Vehicle Routing Problem (VRP) can be encoded into Ising Hamiltonians, allowing for solutions via eigenvalue optimization [2], [12]. VRP generalizes the well-known Travelling Salesman Problem (TSP) as it allows for multiple vehicles, as opposed to just one salesman. The goal of VRP is to find optimal routes for multiple vehicles visiting a set of locations. The "vehicles" could be not just cars, or trucks, but the problem can be applied to airplanes or ships. The VRP can in particular help in solving logistics and routing problems for the army, air force, or other branches of the military. As one example, during a mission, army has need to optimize routing between bases or airports to deliver cargo, ammunition, etc. If the location of the bases or airports can be discovered from the computation, then the secrecy of the mission and national security are compromised – this is what this work aims to bring light to.

One means of leveraging quantum computers is to help solve the VRP. Most current physical quantum computers are still small-scale machines and largely prone to errors and decoherence beyond the fault-tolerance threshold. These so-called Noisy Intermediate-Scale Quantum (NISQ) devices operate with fewer than 100 qubits and shallow quantum circuit depths [5], limiting the size and scope of quantum algorithms that can be implemented on them [27]. Despite these hardware limitations, the advent of NISQ-era quantum computing has spurred research into short-depth quantum circuits and hybrid quantum-classical algorithms that make use of quantum computers in conjunction with classical optimization techniques. Such hybrid quantum-classical algorithms allow for the possibility of performing useful computational tasks, even with NISQ devices. A crucial milestone in this direction was the invention of the Quantum Approximate Optimization Algorithm (QAOA), proposed by Farhi et al. [11]. as a general quantum algorithm that provides approximate solutions for combinatorial optimization problems, including VRPs.

While QAOA efficiently approximates VRPs, its deployment introduces novel security threats, especially those exploiting quantum circuit structures. Previously, various attacks leveraging quantum-specific features, such as reset attacks, fingerprinting tomography, and side-channel attacks, have been identified, raising concerns about information leakage beyond classical computing paradigms, e.g., [8], [23].

In this research, we investigate how information leaks which can reveal quantum circuit structures can lead to learning secrets from VRPs executed on quantum computers using QAOA. On a quantum computer, the VRP routing problem is represented as a graph, such as graph of the army bases or airports. The edges can represent the capacity of the links between the nodes in the graph. The graph being used in VRP is typically a subgraph of some larger, so-called mother graph. For example, the army bases or airports used in a mission are a subset of all the army bases or airports. The objective of this work is to show that if attacker is able to learn some information about QAOA structure from side-channels on a quantum computer, then he or she can reverse-engineer parameters of VRP, and finally that can lead to leakage of mission-critical information.

This work improves on existing work [7] where authors have made stronger assumptions on the ability of attackers to

This work was supported in part by National Science Foundation grant 2312754.

retrieve information from the quantum circuits. Here, we relax these assumptions, and show that even if they are afforded limited insights into the quantum circuits structure, such as only qubit counts learned through circuit connectivity, the attackers can still infer nontrivial information about the VRPs being optimized by the victim. As VRP optimization directly impacts civilian lives and military security, our work addresses crucial security concerns in the quantum computing landscape.

A. *Contributions*

The main contributions of this work are listed as follows:

1) Security assessment under few assumptions:
   - We evaluate the threat posed by attackers with very limited insights into the quantum circuit. In particular, we only grant attackers with knowledge of qubit counts from circuit connectivity, rendering the attack more realistic for near-future machines, compared to existing work [7].
2) Success probability analysis with two measures:
   - We investigate the success probability of attackers based on the size of the mother graph, evaluating their performance based on both the overall matching success rate and node recovery proportionality. In particular, we find that the attacker can stably recover $> 60\%$ of the nodes at relatively large graph sizes.
3) Impact of subgraph size on matching success:
   - We explore the relationship between subgraph size and matching success rate. We observe that, both the matching success rate and node recovery proportionality increase with subgraph sizes, topping at 1.00 for subgraphs the same size as the mother graph.

II. Background

A. *Vehicle Routing Problem (VRP)*

The Vehicle Routing Problem represents a fundamental challenge in logistics and optimization, seeking to efficiently plan the routes of a fleet of vehicles to service a set of geographically dispersed locations. This NP-hard combinatorial optimization problem has pervasive applications in transportation, distribution, and supply chain management. Traditionally formulated as a graph theory problem, where nodes represent delivery points and edges denote feasible routes, VRP aims to minimize overall transportation costs or time while satisfying constraints such as vehicle capacity and delivery time windows [3], [14], [34]. Moreover, the versatility of VRP extends beyond its graph-based representation; it can be expressed in terms of Ising Hamiltonians, a formalism originating from quantum mechanics. The VRP can be encoded into an Ising model, allowing for a quantum-inspired approach to optimization [2]. This unique representation provides a bridge between classical and quantum optimization methodologies, offering new perspectives and potential solutions for tackling the complexities inherent in the VRP.

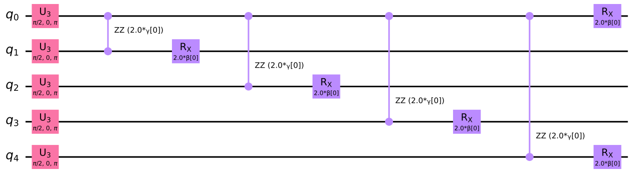

Fig. 1: The circuit representation of a sample variational ansatz for QAOA with $p = 1$. In the circuit, five qubits are used; two-qubit gates consist of RZZ rotational gates and single-qubit gates consist of Rx rotational gates.

B. *VRP Instance Representation*

In this paper, we will represent a problem instance of VRP using (n, k), where n is the number of locations and k is the number of vehicles. For simplicity, we consider the existence of a single depot D. We impose two minimum constraints: each location is visited exactly once, and all vehicles begin from and return to the depot D. Here D can represent mission headquarters, and the locations are the army bases or airports. The vehicles can be airplanes delivering mission-cortical cargo.

In Fig. 2, we present the graph representation of a VRP problem constructed using a real-world dataset from Kaggle [13]. For the efficient execution of sufficiently many experiments, we generate our VRP instances randomly obeying basic real-world principles. In particular, for some fully-connected graphs, we choose the depot randomly from all the nodes and sample the edge weights uniformly from $(0, \infty)$ using a pseudorandom number generator. Moreover, we note that the resulting graphs representing VRPs are not directed. Finally, since the edge weights can only be obtained from gates in the quantum circuits, we perform all arithmitics with reduction $\mod 2\pi$.

C. *Quantum Approximate Optimization Algorithm*

In recent years, the rise of quantum computation has introduced a quantum solution to VRPs, exemplified by the Quantum Approximate Optimization Algorithm. First proposed by Farhi et al. [11], QAOA stands as a general quantum algorithm offering optimal or near-optimal solutions for various combinatorial optimization problems [6], [17], [18], [20], [37], including VRPs [2].

Notably, QAOA leverages variational ansatzes, which represent the combinatorial problem under consideration. More specifically, the problem is encoded into a cost Hamiltonian H_c in QAOA. Besides H_c, the Ising Hamiltonian for QAOA also contains a heuristic mixer H_m for each optimization layer. By convention, the mixer Hamiltonian H_m is chosen as:

$$H_m = -\sum_i \sigma_i^x, \tag{1}$$

where σ^x (σ^z) is the Pauli X (Z) operator. Generally, a complete variational ansatz for a VRP instance of (n, k) takes the form:

$$\left|\vec{\beta}, \vec{\gamma}\right\rangle = \prod_{i \leq p} e^{-iH_m\beta_i} e^{-iH_c\gamma_i} \left|+\right\rangle^{\otimes n(n-1)}, \tag{2}$$

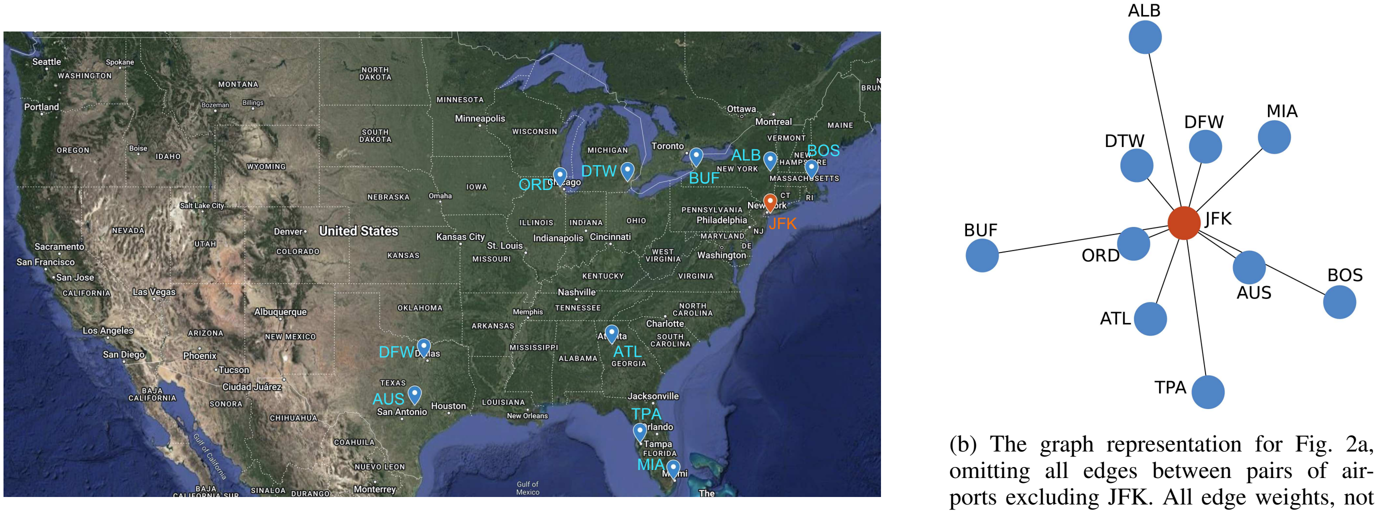

(a) The satellite map for the U.S. mainland, pinned with real locations of the airports used in the study. Made with Google Maps [15]. The pin for JFK (depot) is orange, and blue pins are used for the remaining airports.

(b) The graph representation for Fig. 2a, omitting all edges between pairs of airports excluding JFK. All edge weights, not shown, correspond to the cost of transportation and manifest as distances between nodes.

Fig. 2: The real-world map with airport locations and its graph representation constructed from the real-world dataset sourced from Kaggle, specifically the "USA Airport Dataset" [13].

where p represents the total optimization layers of the quantum circuit, and β_i, γ_i are variational parameters for the layer i. They act as overall scaling factors for all the rotational angles in H_c, H_m in the layer i. A circuit representation of a sample variational ansatz with $p = 1$ is given in Fig. 1.

As can be seen from the variational ansatz, QAOA involves multiple steps, each comprising the application of a layer of parametrized gates, followed by the evolution of the quantum state. The optimization task is embedded in the minimization of the eigenvalue for the cost Hamiltonian H_c associated with the specific combinatorial problem at hand.

As a hybrid quantum-classical variational algorithm, QAOA's time complexity is intricately tied to both its quantum optimization circuit and classical components, especially the initialization of variational parameters in the quantum circuit. Despite this intricacy, studies have demonstrated the acquisition of quasi-optimal solutions in $O(\text{poly}(p))$ time, where p represents the optimization level of the quantum circuit [37]. Additionally, adaptive variants of QAOA exhibit noteworthy performance improvements [16], [33], [38]. Continual advancements highlight QAOA's potential in efficiently solving real-world optimization problems compared to classical algorithms, positioning it as a leading candidate for achieving quantum advantage in practical applications.

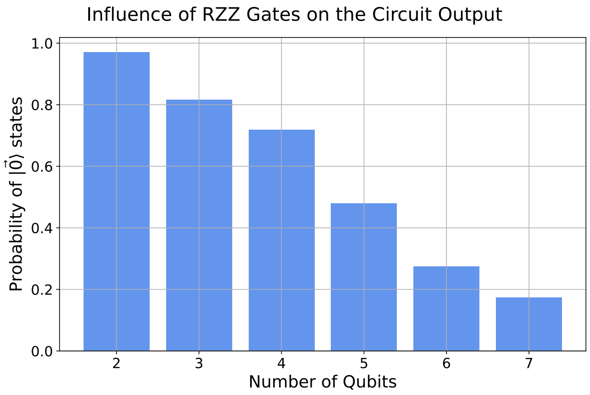

Fig. 3: Relationship between the circuit output, the RZZ gate count, and number of qubits in variational ansatzes for VRP. For simplicity, we have set the rotational angles of all Rx gates to be 0 and those of all RZZ gates to be π. The output fidelity decreases with the number of qubits, while the number of RZZ gates increases.

III. Threat Model

In this paper, we address the threat posed by a potent attacker possessing sufficient knowledge of a mother graph, denoted as G, and targeting victims engaged in optimizing a VRP defined on some subgraph G_s, where $G_s \subseteq G$. Assuming the attacker possesses ample computing resources, they can precompute and obtain all relevant parameters pertaining to any $G_s \subseteq G$. Armed with these precomputed parameters, the attacker endeavors to deduce the specific subgraph G_s the victim is optimizing, leveraging information extracted from the variational quantum circuits under. We assume the attacker uses some form of side-channel in the quantum computer to learn the information, recent works have demonstrated numerous side-channels on quantum computer controllers which can leak information about types of gates being executed [36], for example.

Under our assumption, the attacker is afforded very limited insights into the victim's circuit structures, namely only the qubit count. The determination of the qubit count is facilitated through the examination of circuit connectivity, or the number of two-qubit RZZ gates to be more precise.

As shown in Fig. 1, the only two-qubit gate used in the

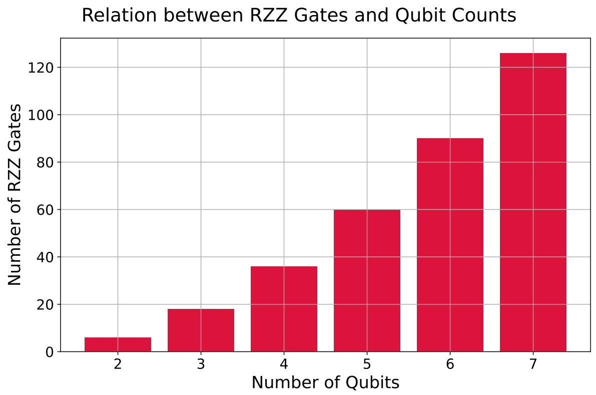

Fig. 4: Relationship between the circuit output, the RZZ gate count, and number of qubits in variational ansatzes for VRP. For simplicity, we have set the rotational angles of all Rx gates to be 0 and those of all RZZ gates to be π. The output fidelity decreases with the number of qubits, while the number of RZZ gates increases.

standard variational ansatz of QAOA is the two-qubit RZZ gate. From the experimental findings of Fig. 3, it is obvious that variations in the number of RZZ gates exert a discernible impact on the circuit output. Furthermore, from Fig. 4, an established correlation exists between the number of RZZ gates and the qubit count in a QAOA ansatz. In general, an ansatz of m qubits can have at most $m(m-1)$ RZZ gates in a single optimization layer. Therefore, we are afforded at least two ways to infer the number of qubits in the VRP quantum ansatz.

First, through crosstalk and other methods of circuit tomography [8], an attacker can try to learn number of RZZ gates directly, leading to the discovery of the structure of VRP and enabling recovery of information about the VRP problem, such as which airports or bases are being used. These could be learned through recently demonstrated quantum computer controller side channels [36], or through crosstalk among qubits [8].

Secondly, one may take advantage of the reset leakage [23, 35] to probe the output state fidelity, which in turns informs the attacker of the number of RZZ gates in the circuit, along with the qubit counts.

IV. Attacker Routine

A. Cost Hamiltonian and Weight Function

Following the convention used in [2], the quantum cost Hamiltonian for a VRP can be encoded as

$$H_c = -\sum_{i,j<i} J_{ij}\sigma_i^z\sigma_j^z + \sum_i h_i\sigma_i^z + d, \tag{3}$$

where the parameters J_{ij}, h_i, d are determined uniquely from each VRP. In the quantum circuit, the parameter d is proportional to an identity gate and drops out. The parameter of interest to us is h_i, which contains the weight for each edge in the VRP graph. By the standard mapping,

$$h_i = \frac{w_i}{2} + C, \tag{4}$$

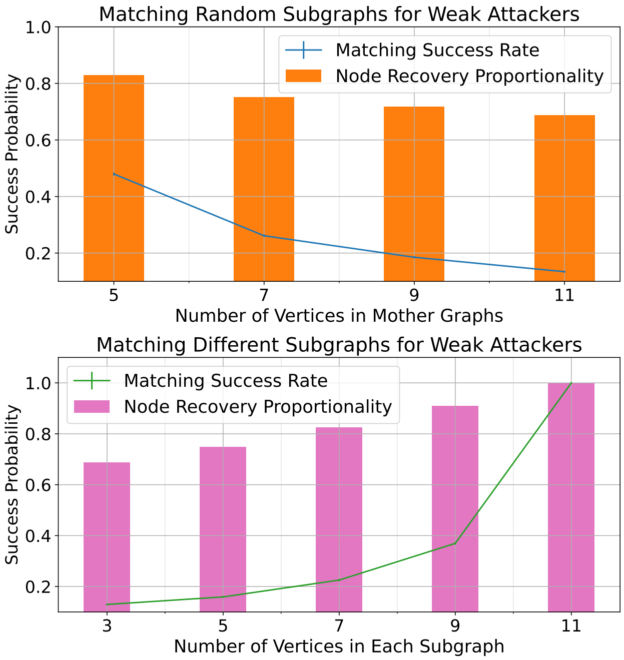

Fig. 5: Experimental findings of the matching success rate and node recovery proportionality for a weak attacker. Both their performance on random subgraphs and subgraphs of fixed sizes is evaluated. In the experiment of random subgraphs, we evaluate the performance of the attacker for all subgraphs with size ≥ 2 and $\leq$ the size of the mother graph. In the experiment of fixed subgraphs, we consider a mother graph $|G| = 11$.

where w_i is the weight function for the i-th edge in the graph for a VRP and C is a constant independent of the weight function of the VRP instance (n, k) under consideration. More specifically, C depends on only n and k. Thus, the attacker can treat C as a constant offset with the knowledge of n and k.

B. General Scheme

A general routine adopted by the attacker can be summarized as follows:

1) Given a mother graph G, precompute the set of subgraphs G_s, where $G_s \subseteq G$.
2) Deduce (n, k) given the information available. In particular, n may be determined by the experiments described in Section III.
3) Select some subgraph G_s from the precomputed set with the same qubit number and deem it as the subgraph under consideration by the victim.

For clarity, we note that the knowledge of k is not relevant for the weak attacker under consideration, since the effect of k only manifests itself in the weight of the VRP graphs, which is not accessible for the weak attacker under consideration.

V. Results and Discussions

In this section, we present our experimental evaluations for the performance of the attacker. Mainly two types of experimennts are conducted: matching subgraphs of random and matching subgraphs of fixed sizes under uniform noise. Note that in all the experiments for matching random subgraphs, we have ignored subgraphs of sizes ≤ 2. For any VRP graph, there exists at least one node as the depot. Therefore, a subgraph of size 2 represents a VRP of one extra location other than the depot. This makes the optimization trivial and unnecessary to consider in realistic settings.

In all experiments, we consider both the matching success/error rate and the node recovery proportionality. Here, the matching success/error rate refers to the success/error rate for an attacker to find the complete subgraph that the victim is trying to optimize. However, although an attacker cannot retrieve the complete subgraph sometimes, it is meaningful to consider the portion of the subgraph that the attacker can recover. In this sense, it is useful to consider and evaluate the proportionality of node in a subgraph optimized by the victim recoverable by the attacker.

But this does not mean that the matching success/error rate is less useful alone in realistic settings. As we will show in Section V-A, even for a weak attacker, the node recovery proportionality is highly non-negligible for subgraphs of large sizes. This is because the node recovery proportionality is subject to combinatorial effects. In this sense, the matching success/error rate shows more intrinsically how good the performance can be for an attacker, independent of those combinatorial effects.

A. Weak Attacker Results

In Fig. 5, we present the experimental findings on the matching success rate and node recovery proportionality for the weak attacker. Both the case of matching subgraphs of random sizes > 2 and the case of matching subgraphs of fixed sizes are considered. For subgraphs of random sizes, it is obvious that both the success rate and the recovery proportionality drops rapidly with increasing size of the mother graphs, rendering the threat from the weak attacker negligible in face of large optimization problems.

However, if we consider subgraphs of fixed sizes, the weak attacker regains good performance at large subgraph sizes. This is a combinatorial effect of graph matching – the larger the subgraph size is, the fewer possible subgraphs of that size exist. Therefore, the error probability drops correspondingly.

On a side note, compared to the experiments conducted in [7], we have ignored imperfect resolution and other noise for the weak attacker considered in this paper. This is because the imperfect resolution and noise themselves on the matching results through the rotational angles and edge weights. However, a weak attacker is oblivious to those, since they are not provided with any information on the rotational angles. Therefore, it is unnecessary to test them for a weak attacker.

VI. Securing VRP

Based on our findings, attackers with access to information about QAOA executing on a quantum computer can recover secrets from the VRP instance. This can compromise secrecy and national security when VRP instances is used in context of routing between army bases or airports. Our work brings particular attention to the need for better understanding, and prevention, of side channels in quantum computers. As we demonstrate, even the weak attacker who can only obtain information about count of qubits used and the number of two-qubit gates in the QAOA, such attacker already can recover information about the VRP instance. As a result, future work needs to explore prevention of side channels, as well as protection of VRP itself.

VII. Related Work

There is recently, increasingly growing body of research on security of quantum computers. For superconducting quantum computers, recent work [1] shows that the crosstalk errors could be used in fault injection attacks. It also showed how an adversary can launch a denial of service attack on the victim circuit using crosstalk errors, similar to our evaluation. In addition, due to the difference of eigenstates, qubit-sensing employs malicious circuits to sense qubits of victim circuits based on already known statistical information [28]. Among side channel attacks, recent work proposed power side channel attacks that can help recover the control pulses of the quantum computers [36], leading to recovery gates being executed on the quantum computer. Researchers have also proposed different methods to fingerprint quantum computer hardware by characterizing error patterns unique to each device or qubit [22], [26]. In trapped-ion quantum computers, repeated shuttle operations can elevate the ion-chain's energy, which can damage the fidelity of victim circuits [29], [30].

This work is a direct continuation on [7], where attackers are assumed to be more powerful in retrieving information from quantum circuits. The strong attackers described in [7], in addition to the qubit count, is provided information on gates executed in the quantum circuit and rotational angles associated with each gate, albeit with limited resolution. This additional information access empowers the attacker to infer about the weights between pairs of nodes in the VRP graph. Consequently, the attacker is equipped with better capacity to exploit vulnerabilities therein.

VIII. Conclusion

This paper focused on an unexplored domain: examining the potential compromise of transportation logistics, including civilian airports and military bases, through an analysis of quantum circuit structures. The examination of quantum circuit structures allows for the inference of underlying algorithms. The paper sheds light on vulnerabilities introduced by quantum circuits, exposing potential leaks of information that could lead to the recovery of VRP under optimization. The research underscores the urgency in developing robust techniques to safeguard sensitive information, especially in the face of rapid advancements in quantum computing. In particular, it highlights that attackers pose a nontrivial security threat with little insights into the quantum circuits, emphasizing the need for enhanced attention and the formulation of effective defense mechanisms to counteract potential risks.

References

[1] Abdullah Ash-Saki, Mahabubul Alam, and Swaroop Ghosh. Analysis of crosstalk in nisq devices and security implications in multi-programming regime. In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design*, ISLPED '20, page 25–30, New York, NY, USA, 2020. Association for Computing Machinery.

[2] Utkarsh Azad, Bikash K. Behera, Emad A. Ahmed, Prasanta K. Panigrahi, and Ahmed Farouk. Solving vehicle routing problem using quantum approximate optimization algorithm. *IEEE Transactions on Intelligent Transportation Systems*, 24(7):7564–7573, July 2023.

[3] Roberto Baldacci, Aristide Mingozzi, and Roberto Roberti. Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints. *European Journal of Operational Research*, 218(1):1–6, 2012.

[4] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.

[5] Sergey Bravyi, David Gosset, and Robert König. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, October 2018.

[6] Yagnik Chatterjee, Eric Bourreau, and Marko J. Rančić. Solving various np-hard problems using exponentially fewer qubits on a quantum computer, 2023.

[7] Jessie Chen and Jakub Szefer. Stealing vrp secrets from quantum circuit structures. In *International Symposium on Hardware Oriented Security and Trust*, HOST, May 2024.

[8] Sanjay Deshpande, Chuanqi Xu, Theodoros Trochatos, Hanrui Wang, Ferhat Erata, Song Han, Yongshan Ding, and Jakub Szefer. Design of quantum computer antivirus. In *Proceedings of the International Symposium on Hardware Oriented Security and Trust*, HOST, May 2023.

[9] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

[10] Nada Elsokkary, Faisal Shah Khan, Davide La Torre, Travis S Humble, and Joel Gottlieb. Financial portfolio management using D-wave quantum optimizer: The case of Abu Dhabi securities exchange, 2017.

[11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm, 2014.

[12] Sebastian Feld, Christoph Roch, Thomas Gabor, Christian Seidel, Florian Neukart, Isabella Galter, Wolfgang Mauerer, and Claudia Linnhoff-Popien. A hybrid solution method for the capacitated vehicle routing problem using a quantum annealer. *Frontiers in ICT*, 6, June 2019.

[13] FlashGordon. Usa airport dataset. https://www.kaggle.com/datasets/flashgordon/usa-airport-dataset, 2000-2009. Accessed: 2023-11-10.

[14] Bruce Golden, Saahitya Raghavan, and Edward Wasil. *The Vehicle Routing Problem: Latest Advances and New Challenges*, volume 43. 01 2008.

[15] Google Maps. Google map of airport locations, 2023. Accessed: 2023-12-22.

[16] Harper R. Grimsley, Sophia E. Economou, Edwin Barnes, and Nicholas J. Mayhall. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nature Communications*, 10(1), July 2019.

[17] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, February 2019.

[18] Matthew P. Harrigan, Kevin J. Sung, Matthew Neeley, Kevin J. Satzinger, Frank Arute, Kunal Arya, Juan Atalaya, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Sean Demura, Andrew Dunsworth, Daniel Eppens, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Steve Habegger, Alan Ho, Sabrina Hong, Trent Huang, L. B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Martin Leib, Orion Martin, John M. Martinis, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Charles Neill, Florian Neukart, Murphy Yuezhen Niu, Thomas E. O'Brien, Bryan O'Gorman, Eric Ostby, Andre Petukhov, Harald Putterman, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Andrea Skolik, Vadim Smelyanskiy, Doug Strain, Michael Streif, Marco Szalay, Amit Vainsencher, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Leo Zhou, Hartmut Neven, Dave Bacon, Erik Lucero, Edward Farhi, and Ryan Babbush. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3):332–336, February 2021.

[19] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), October 2009.

[20] Itay Hen and Federico M. Spedalieri. Quantum annealing for constrained optimization. *Physical Review Applied*, 5(3), March 2016.

[21] A. Yu. Kitaev. Quantum measurements and the abelian stabilizer problem, 1995.

[22] Allen Mi, Shuwen Deng, and Jakub Szefer. Short paper: Device- and locality-specific fingerprinting of shared nisq quantum computers. In *Proceedings of the Workshop on Hardware and Architectural Support for Security and Privacy*, HASP, October 2021.

[23] Allen Mi, Shuwen Deng, and Jakub Szefer. Securing reset operations in nisq quantum computers. In *Proceedings of the Conference on Computer and Communications Security*, CCS, November 2022.

[24] Nikolaj Moll, Panagiotis Barkoutsos, Lev S Bishop, Jerry M Chow, Andrew Cross, Daniel J Egger, Stefan Filipp, Andreas Fuhrer, Jay M Gambetta, Marc Ganzhorn, et al. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology*, 3(3):030503, 2018.

[25] Roman Orus, Samuel Mugel, and Enrique Lizaso. Quantum computing for finance: overview and prospects. *Reviews in Physics*, 4:100028, 2019.

[26] Koustubh Phalak, Abdullah Ash Saki, Mahabubul Alam, Rasit Onur Topaloglu, and Swaroop Ghosh. Quantum puf for security and trust in quantum computing. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 11(2):333–342, 2021.

[27] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, August 2018.

[28] Abdullah Ash Saki and Swaroop Ghosh. Qubit sensing: A new attack model for multi-programming quantum computing, 2021.

[29] Abdullah Ash Saki, Rasit Onur Topaloglu, and Swaroop Ghosh. Muzzle the shuttle: Efficient compilation for multi-trap trapped-ion quantum computers, 2021.

[30] Abdullah Ash Saki, Rasit Onur Topaloglu, and Swaroop Ghosh. Shuttle-exploiting attacks and their defenses in trapped-ion quantum computers. *IEEE Access*, 10:2686–2699, 2022.

[31] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.

[32] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997.

[33] Ho Lun Tang, V.O. Shkolnikov, George S. Barron, Harper R. Grimsley, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. Qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *PRX Quantum*, 2(2), April 2021.

[34] Paolo Toth and Daniele Vigo. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, 2002.

[35] Chuanqi Xu, Jessie Chen, Allen Mi, and Jakub Szefer. Securing nisq quantum computer reset operations against higher energy state attacks. CCS '23, page 594–607, New York, NY, USA, 2023. Association for Computing Machinery.

[36] Chuanqi Xu, Ferhat Erata, and Jakub Szefer. Exploration of power side-channel vulnerabilities in quantum computer controllers. In *Proceedings of the Conference on Computer and Communications Security*, CCS, November 2023.

[37] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D. Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Phys. Rev. X*, 10:021067, Jun 2020.

[38] Linghua Zhu, Ho Lun Tang, George S. Barron, F. A. Calderon-Vargas, Nicholas J. Mayhall, Edwin Barnes, and Sophia E. Economou. Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer. *Phys. Rev. Res.*, 4:033029, Jul 2022.