

Chuanqi Xu Yale University New Haven, CT, USA chuanqi.xu@yale.edu

Allen Mi Yale University New Haven, CT, USA allen.mi@yale.edu

ABSTRACT

Enabling the sharing of quantum computers among different users requires a secure reset operation that can reset the state of a qubit to ground state $|0\rangle$ and prevent leakage of the state to a post-reset circuit. This work highlights that the existing reset operations available in superconducting qubit NISQ quantum computers are not fully secure. In particular, this work demonstrates for the first time a new type of higher-energy state attack. Although NISQ quantum computers are typically abstracted as working with only energy states $|0\rangle$ and $|1\rangle$, this work shows that it is possible for unprivileged users to set the qubit state to $|2\rangle$ or $|3\rangle$. By breaking the abstraction of a two-level system, the new higher-energy state attack can be deployed to affect the operation of circuits or for covert communication between circuits. This work shows that common reset protocols are ineffective in resetting a qubit from a higher-energy state. To provide a defense, this work proposes a new Cascading Secure Reset (CSR) operation. CSR, without hardware modifications, is able to efficiently and reliably reset higher-energy states back to $|0\rangle$. CSR achieves a reduction in $|3\rangle$ -initialized state leakage channel capacity by between 1 and 2 orders of magnitude, and does so with a 25x speedup compared with the default decoherence reset.

CCS CONCEPTS

• Security and privacy \rightarrow Security in hardware; • Hardware \rightarrow Quantum technologies.

KEYWORDS

quantum computers, reset gate, higher-energy state

ACM Reference Format:

Chuanqi Xu, Jessie Chen, Allen Mi, and Jakub Szefer. 2023. Securing NISQ Quantum Computer Reset Operations Against Higher-Energy State Attacks. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), November 26–30, 2023, Copenhagen, Denmark. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3576915.3623104



This work is licensed under a Creative Commons Attribution International 4.0 License.

CCS '23, November 26–30, 2023, Copenhagen, Denmark © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0050-7/23/11. https://doi.org/10.1145/3576915.3623104 Jessie Chen Yale University New Haven, CT, USA zixin.chen@yale.edu

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

1 INTRODUCTION

Today's quantum computers are commonly called Noisy Intermediate-Scale Quantum (NISQ) quantum computers [30]. So far they are too small for supporting quantum error correction, but already have promising applications in optimization, chemistry, and other important areas [16, 18, 23]. Further, NISQ quantum computers are being rapidly developed, with 433 qubit machines available today, and the industry projecting 4000 qubit or larger devices before the end of this decade.

Many different types of quantum computers exist, with superconducting qubit quantum computers being one of the types. The superconducting qubit machines are developed by numerous companies, such as IBM, Rigetti, and Quantum Circuits, Inc. These machines implement quantum computing with superconducting electronic circuits which are operated at about 20mK temperatures. They are typically considered a two-level system where qubits can exist in any quantum superposition of two independent and physically distinguishable quantum states. These states are typically denoted as $|0\rangle$ and $|1\rangle$. There is, however, nothing that physically prevents the hardware from being excited into higher-energy states such as $|2\rangle$ or $|3\rangle$.

The higher-energy states have potentially useful applications, and can actually be used in improving quantum computation. For example, ternary quantum systems are being studied recently because these provide more computational state space per unit of information, known as qutrit. A qutrit has three basis states, $|0\rangle$, $|1\rangle$, and $|2\rangle$. Using qutrits can result in circuit cost reductions for important algorithms like quantum neurons and Grover search [12]. Researchers have also developed a quantum computer design based on quantum digits or "qudits" that have even more energy states [19]. All of these require the use of higher-energy states beyond $|0\rangle$ and $|1\rangle$.

In parallel to the growing research on using higher-energy states, there is parallel research on, and practical deployment of, quantum computers as cloud-based accelerators. Cloud-based services such as IBM Quantum, Amazon Bracket, and Microsoft Azure already provide access to superconducting qubit quantum computers remotely for users. Although today only one user or program can run at a time, one of the next possible advances in quantum computers enabled by the rapidly increasing qubit numbers is the sharing of the devices among different users, or among different quantum programs of the same user, e.g., [5]. However, to realize this, there needs to be a secure way to reset the state of individual qubits to prevent any prior state of the qubit from being leaked or affecting the post-reset state. A full system reset deployed between different jobs in IBM's terminology (i.e. different programs of the same or different users) takes today on the order of $1000\mu s$ and further resets all the qubits, preventing sharing of the devices among different users, or among different quantum programs of the same user where the different programs execute on different qubits and not all start and stop at the same time. Fortunately, there is already a reset operation available in IBM machines, which resets the state of a qubit back to $|0\rangle$.

However, this reset, as we demonstrate in this work, is fully ineffective in resetting higher-energy states. A new higher-energy state attack is possible because the current reset gate, and other gates, do not work properly on higher-energy states $|2\rangle$ and above. To counter the new attack, a new Cascading Secure Reset operation, abbreviated CSR(n), where *n* is the highest energy level reset by the CSR, is introduced, prototyped, and evaluated in this work.

1.1 Insecurity due to Wrong Abstraction

The core idea explored in this paper can be summarized as one that a system can become insecure because the wrong abstraction of the hardware is used or assumed. Especially, the idea that the NISQ quantum computer is a two-level system, with energy states $|0\rangle$ and $|1\rangle$, is only an abstraction, but underlying hardware is not limited to only these states. We show that by pushing the qubits of the quantum computers into higher-energy states such as $|2\rangle$ and $|3\rangle$, undesired or possibly malicious behavior can be achieved. On one hand, it is possible to set qubits into higher-energy states, on the other, existing gates or operations of the quantum computer are not designed for, and are mostly ineffective when they are applied to the higher-energy states. A possible analogy to classical computer terms is that the architecture assumes certain hardware behavior (i.e. $|0\rangle$ and $|1\rangle$ states) while the micro-architecture implements additional, hidden behavior (i.e. extra $|2\rangle$ and $|3\rangle$ states).

1.2 Contributions

The contributions of this work are:

- Demonstration of the ineffectiveness of default reset gate in resetting higher-energy states, and use of this finding as the basis of a new security attack.
- Establishing higher-energy state attack as a new type of attack in superconducting qubit quantum computers, this attack is demonstrated against algorithms such as VQE, Deutsch-Jozsa, inverse quantum Fourier transformation, and Grover's search.
- Design of the new Cascading Secure Reset (CSR) operation which can defend against higher-energy state attacks and be a possible building block of future shared quantum computers.
- Evaluation of both the attack and defense on real, publicly available superconducting qubit quantum computers from IBM.

2 BACKGROUND

This work focuses on superconducting qubit quantum computers, with specific evaluation and analysis done on publicly-accessible IBM quantum computers. There are also other vendors developing superconducting qubit machines, such as Rigetti or Quantum Circuits, Inc.. There are also other types of quantum computers such as ones using trapped ion qubits [3]. We believe that secure reset operations for all types of machines will need to be developed, and will be examined in our future work, while this work focuses on IBM machines as being representative of superconducting qubit quantum computers.

2.1 Principles of Quantum Computing

The most basic unit in quantum computing is the quantum bit, or qubit for short, which is an analogous concept of the bit in modern classical computing. Similar to a bit, a qubit has two basis states, which are represented as $|0\rangle$ and $|1\rangle$ using the bra-ket notation. However, a bit can only be either 0 or 1, while a qubit can be any linear combination of $|0\rangle$ and $|1\rangle$ with norm 1. To be more specific, a qubit $|\psi\rangle$ can be represented as:

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle$$

where α and β are complex numbers and $|\alpha|^2 + |\beta|^2 = 1$.

Qubits can also be formalized with vector representation. For one qubit, the basis states can be represented as two-dimensional vectors, e.g., $|0\rangle = [1, 0]^T$ and $|1\rangle = [0, 1]^T$, and thus one state $|\psi\rangle$ above can be represented as $|\psi\rangle = \alpha |0\rangle + \beta |1\rangle = [\alpha, \beta]^T$. More generally, the space of *n*-qubit states has 2^n basis states starting from $|0...0\rangle$ to $|1...1\rangle$, and a *n*-qubit state $|\phi\rangle$ is given by:

$$|\phi\rangle = \sum_{i=0}^{2^n - 1} a_i \,|i\rangle$$

where $\sum_{i=0}^{2^{n}-1} |a_i|^2 = 1$.

In quantum computing, the input qubits are changed by unitary operations known as quantum gates. The effect of a quantum gate U on a quantum state $|\psi\rangle$ is $|\psi\rangle \rightarrow U |\psi\rangle$, where U is a unitary operation satisfying $UU^{\dagger} = U^{\dagger}U = I$. With the previous vector representation, *n*-qubit quantum gates can be represented as $2^n \times 2^n$ matrices. Examples of Qiskit qubit order are shown below.

$$ID = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, CX = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$
$$RZ(\theta) = \begin{bmatrix} e^{-i\frac{\theta}{2}} & 0 \\ 0 & e^{i\frac{\theta}{2}} \end{bmatrix}, SX = \frac{1}{2} \begin{bmatrix} 1+i & 1-i \\ 1-i & 1+i \end{bmatrix}$$

It is proved that any unitary quantum gate can be built from a few quantum gates [8]. Accordingly, currently accessible quantum computers can only have several basis gates instead of all unitary operations. For example, ID, RZ, SX, X, CX are the basis gates for most IBM Quantum quantum computers. Their matrix representations are shown above. Other quantum gates, such as the Hadamard gate, need to be decomposed into these basis gates before being executed on the real quantum computer hardware.

2.2 Achieving Higher-Energy States

Usually, there are only two energy levels, $|0\rangle$ and $|1\rangle$, for each qubit in the logic level of quantum computing. However, this is usually an artificial restriction because energy levels other than $|0\rangle$ and



Figure 1: Quantum computing workflow on IBM Quantum using Qiskit. The first step is to decide the math expression for the task at the logic-level. Then build the gate-level circuit from the math expression using Qiskit. After the circuit is fully implemented, it needs to be tranpiled to the circuit that can be executed on a given quantum device. In the end, the gate-level circuit is required to be translated to the pulse-level circuit which controls the equipment.

 $|1\rangle$ are inherent in the physical realizations of quantum computers. Generally, the π gate is used to excite the state from the lowerenergy state to the higher-energy state. If we denote the lower energy level as g and the higher-energy level as e, then the $\pi_{g,e}$ gate is the π gate that exchanges the populations on these energy levels. Suppose the state is $|\psi\rangle = a_g |g\rangle + a_e |e\rangle$, and then the $\pi_{g,e}$ gate transforms it as:

$$\pi_{q,e} |\psi\rangle = a_e |g\rangle + a_q |e\rangle$$

In particular, the X gate is also the $\pi_{0,1}$ gate that can excite $|0\rangle$ to $|1\rangle$ or in the reverse order. Higher-energy states can be obtained if relevant physical properties and equipment are given. For example, superconducting quantum computers can excite states to higher-energy states if the physical quantities of the microwave pulses of π gates are known and the existing experiment equipment can apply such π gates. In Section 3.1, we will introduce how to access higher-energy states on IBM quantum computers by calibrating and applying $\pi_{1,2}$ and $\pi_{2,3}$ gates.

2.3 Running Circuits on Quantum Computers

As introduced in Section 2.1, quantum computing at the logic-level can be described with vectors and matrices in the complex space. In addition, the process of how a quantum state evolves under quantum gates can be described by a quantum circuit. In a quantum circuit plot, lines that start from left to right represent qubits, and symbols on the lines represent operations. At the beginning of the quantum circuit, qubits are usually assumed to be in the $|0\rangle$ state. Then qubits evolve under operations in sequence from left to right. At the end of the quantum circuit, there are usually measurement operations to measure, acquire, and store qubit information in classical memory.

Similar to compilation in classical computers, quantum circuits need to be transpiled into circuits that can be performed on real quantum computers and have identical results at the logic-level. The transpiling process includes many stages, such as decomposing





Figure 2: Reset operation is composed of a measurement operation, followed by a conditional X gate, which flips the post-measurement state from $|1\rangle$ to $|0\rangle$ if the measurement result is 1. Here q is the target qubit and c is the respective classical register.

quantum gates to a combination of basis gates and mapping the logical qubits in the original circuits to the physical qubits.

All circuits discussed so far are gate-level circuits. Not surprisingly, to ultimately execute quantum circuits, gate-level circuits need to be translated into the quantum computer-specific language. This step depends on the physical realization and appliances of the quantum computer. In this paper, our experiment platform is IBM Quantum, which provides superconducting quantum computers and uses microwave pulses to control the evolution of qubits. The process that translates gate-level circuits to hardware-specific operations is called scheduling in Qiskit. Scheduling generates microwave pulse sequences based on previous calibrated data. Pulse sequences can also be called pulse-level circuits, and will be used to control the equipment of quantum computers to execute the original quantum circuits. The whole process of quantum computing on IBM Quantum is shown in Figure 1.

2.4 Operation of Jobs, Circuits, and Shots

Jobs are the basic unit to execute quantum circuits on IBM. A job is specified by a circuit or a list of circuits, the quantum device on which the circuits will execute, as well as a set of parameters specifying the details of how to run the job. Among these parameters, shots, rep_delay, init_qubits are related to this paper. shots specifies the number of repetitions of each circuit to execute on the quantum computer. rep_delay determines the delay time in seconds between the shots. init_qubits indicates whether to reset the qubits to the ground state for each shot, which implies whether to append a Reset gate and a small delay after each shot.

After jobs are submitted, a series of processes will be performed before the equipment starts to control quantum computers, such as checking the input arguments and circuits. Circuits will be executed in an interleaving scheme when they are running on quantum computers. More specifically, if the job contains only one circuit, this circuit will execute shots times. The delay between each shot is rep_delay, together with or without the initialization operations depending on init_qubits. On the other hand, if the job consists of a list of circuits, then the list of circuits will run one shot by one shot sequentially, and then wrap around to run for the next iteration.

2.5 Existing Reset Protocols

2.5.1 The Reset Gate. The existing Reset gate available on IBM quantum computers is shown in Figure 2. It consists of a measurement operation that yields the classical bit *c* from the qubit *q*. Following the measurement, there is a conditional *X* gate which will set the qubit to the $|0\rangle$ state if it is not already in that state.



Figure 3: Visualization of different delays and Reset (*R*) gates between different jobs, and between circuits within the same job, in IBM quantum computers' default operation. The delays are not drawn to scale.

When conditioned on the measurement outcome, the *X* gate will not be invoked if the qubit returns a measurement result of \emptyset and its post-measurement state is already in $|0\rangle$. On the contrary, if the qubit returns a result of 1 and is collapsed to $|1\rangle$, the *X* gate will flip the state back to $|0\rangle$. In the ideal scenario, this effect ensures that the qubit is always in the $|0\rangle$ after the reset. However, the reset is not perfect in real-world scenarios. The duration of the reset gate is about $1\mu s$, but the exact value depends on the specific machine.

2.5.2 Full System Reset. Existing IBM quantum computers perform a full system reset or full system wipe (FSW) between the execution of different circuits or different shots of the same circuits. Based on our analysis, the full system reset includes a Reset gate, followed by a $250\mu s$ default repetition delay (rep_delay). In addition, between different jobs, there is a load delay (load_delay). Figure 3 shows a visualization of the different delays between different jobs, and between circuits within the same job.

The users are able to control whether the Reset gate is used or not (setting init_qubits=False disables the use of the Reset gate). The users can also control the duration of the repetition delay by setting rep_delay=n, where n, is the delay in seconds. The default settings are init_qubits=True and rep_delay=250e-6 (units of seconds are used in Qiskit).

The inter-job load delay (load_delay) is at least on the order of seconds in our observations, but this is probably largely due to classical data loading in the controller, or other processing or some other operations that do not necessarily have to do with the quantum machines. If IBM has sufficient improvements to their processing pipelines, they may also cut down on the inter-job delay to the timescale of inter-circuit delay.

2.5.3 Depopulation Sequence. The depopulation sequence is a common set of operations in superconducting transmon engineering. The operation Depop consists of a sequence of π -pulses, arranged in descending order in energy level, such that:

$$\mathsf{Depop}(n) = \pi_{n-1,n} \circ \pi_{n-2,n-1} \circ \cdots \circ \pi_{0,1} \tag{1}$$

where \circ denotes composition. The sequence is designed to act on a pure $|n\rangle$ state. In the ideal case, each $\pi_{i-1,i}$ lowers the state from $|i\rangle$ to $|i-1\rangle$, with a final outcome of $|0\rangle$. This operation relies heavily on the purity of $|n\rangle$ and the fidelity of the π -pulses, since for each $0 \le i < n$, the amplitude of $|i\rangle$ is assigned as the amplitude of $|i+1\rangle$ via $\pi_{i,i+1}$ through the transformation. Consequently, the depopulation sequence is not stable under repetition: If the first execution successfully depopulates all higher-energy states, the second execution would switch the amplitude of $|0\rangle$ and $|1\rangle$ and produce $|1\rangle$ as the outcome. Each additional execution would further elevate the state by one energy level.

3 HIGHER-ENERGY STATES PROPERTIES

To set qubits into a higher-energy state, the calibration data is needed to configure the π gate for each pair of the energy levels, which swaps the amplitudes of two basis states. The calibration data can be obtained from the metadata provided by IBM. Using the calibration data, custom pulse gates can be set up to excite qubits into higher-energy states, after which we can test their properties, which are very different compared with $|0\rangle$ and $|1\rangle$.

3.1 Setting Qubits into Higher-Energy States

IBM quantum computers allow for the use of pulse gates, and thus it is possible to excite qubits into higher-energy states other than $|0\rangle$ and $|1\rangle$. The method to achieve higher states such as $|2\rangle$ and $|3\rangle$ is to calibrate $\pi_{1,2}$ and $\pi_{2,3}$. This is done by calibrating frequencies needed for $\pi_{1,2}$ and $\pi_{2,3}$ by using the frequency sweep circuit, and then calibrating amplitudes of the pulses using the Rabi experiment.

The frequency sweep circuit is employed to sweep a range of frequencies of a qubit and search for signs of absorption. Since the qubit can be excited to the next energy level with the resonant frequency between the original and final states, and this excitation takes in energy from the pulse, the frequency between energy levels can be approximated to be the frequency in which the minimum or the maximum of the signal intensity is located.

The Rabi experiment is based on the idea of Rabi oscillation [13]. The step is to drive the qubit with different amplitudes. If the pulse frequency is set to be the frequency of two energy levels, then the amplitudes of the final state on two energy levels will be trigonometric functions of the amplitude of the pulse. The π pulse is used to excite qubit from a lower-energy state to a higher-energy state, or the opposite way, and thus half of the period of the trigonometric function is the amplitude for the π pulse.

Higher-energy states can be achieved one by one sequentially. Suppose we can prepare a qubit into state $|i\rangle$, the next steps to prepare the qubit into state $|i + 1\rangle$ is to first use frequency sweep to find the frequency between these two energy levels (*i* and *i* + 1), and then do the Rabi experiment to find the amplitude of π pulse based on the frequency acquired in the previous step. With errors accumulating, it is harder and harder to prepare the higher-energy state with high fidelity, and the equipment may also limit the range of the frequency we can set. Therefore, in our experiments, we choose $|2\rangle$ and $|3\rangle$, which can already show the idea of this paper. If not explicitly stated, the data is collected from ibm_lagos, which is a seven-qubit quantum computer provided for public access by IBM. The calibrated data is shown in Figure 4.

3.2 Higher-Energy States and Measurement

Even though we can obtain higher-energy states, it is hard to distinguish them with current measurement equipment, because in

CCS '23, November 26-30, 2023, Copenhagen, Denmark



Figure 4: Calibration data of frequency sweep and Rabi experiments performed on the ibm_lagos backend to obtain $\pi_{1,2}$ and $\pi_{2,3}$ pulses needed to drive $|1\rangle$ to $|3\rangle$ for Qubit 0 to 6 on ibm_lagos respectively. The frequency sweep experiments provide the frequency for the π pulses. The Rabi experiments provide the amplitudes for the π pulses. The exact experiments displayed are specified in the title for each line of the graphs.



Figure 5: Higher-energy states decoherence patterns on the IQ plane. The circles show measured signal amplitudes with different delay times. The coordinates of the centers of circles represent the average I and Q parts of the measurement results with given delays, and the radii indicate the standard deviation of 1024 measurement data points. The dashed line represents the default discriminator. The interval between circles is $\approx 11\mu s$. The qubit decays $0\mu s$ at the "Start" point, and decays about $555\mu s$ at the "End" point.

quantum computing it is usually assumed to only have $|0\rangle$ and $|1\rangle$ as the basis states. Figure 5 shows the measured signal amplitudes on the IQ plane of preparing each state. Further, it shows the decay of the higher-energy states. Initially, at points labeled "Start", $|1\rangle$ data points are separated from $|2\rangle$ and $|3\rangle$ data points. However, $|2\rangle$ and $|3\rangle$ data points are mixed with each other. As decay progresses, the states become mixed until the "End" state of $|0\rangle$ is reached.

Figure 5 also shows the default linear discriminator used by ibm_lagos as the dashed vertical line. This is used for the normal measurement which only separates $|0\rangle$ state (same as the 'End'

state) from $|1\rangle$, but it is not able to properly distinguish among $|1\rangle$, $|2\rangle$ and $|3\rangle$. Thus when a higher-energy state is measured, it may be reported as $|1\rangle$ even though the component of $|1\rangle$ is very small.

3.3 Effect of Higher-Energy States on Basic Gates used by Quantum Computers

The basic gates on IBM quantum computers are usually I, RZ, SX, X, CX (see Section 2.1 for their definitions), as well as the Reset gate. In Table 1, we show the effects of basic gates on state $|1\rangle$, $|2\rangle$, and 3). The experiment performed is to first prepare the excited states, then immediately apply one of the basic gates after the prepared states, and finally measure the results with the default measurement operation (i.e. with the linear discriminator used by ibm_lagos). Given that CX gate is a two-qubit gate, each value in the rows of CX gate in Table 1 is the measurement result where the control qubit is in state $|1\rangle$ and the target qubit is in the specified state. In this case, the CX gate effect is to maintain the control qubit on $|1\rangle$ and have the effect similar to X gate on the target qubit. The 'No gate' rows show the probabilities of measuring $|1\rangle$ directly after preparation with no gates between preparation and measurement. The other rows show the probability differences of measuring $|1\rangle$ from 'No gate' when there is one of the RZ, SX, X, Reset, or CX gate between preparation and measurement. Because existing quantum computers assume there are only $|0\rangle$ and $|1\rangle$ states in the system, the measurement only discriminates between them as discussed above. This means that even though we can prepare higher-energy states, the measurement results attributed to them will be either $|0\rangle$ or $|1\rangle$ based on the data points on the IQ plane.

As can be seen from the evaluation, these basic gates are designed based on the assumption of only $|0\rangle$ and $|1\rangle$ states in the system, and thus do not have much effect on qubits when the qubits are

CCS '23, November 26-30, 2023, Copenhagen, Denmark

Chuanqi Xu, Jessie Chen, Allen Mi, and Jakub Szefer

Init. State	Gate	Qubit 0	Qubit 1	Qubit 2	Qubit 3	Qubit 4	Qubit 5	Qubit 6	Average
	No gate	0.9924	0.9903	0.9932	0.9899	0.9874	0.986	0.9899	0.9899
	RZ	0.0012	-0.0018	-0.0028	-0.0012	-0.0008	0.0005	-0.0007	-0.0008
11)	SX	-0.4907	-0.4839	-0.497	-0.4706	-0.4777	-0.4743	-0.5025	-0.4852
1/	Х	-0.9798	-0.9791	-0.9884	-0.9503	-0.9738	-0.9678	-0.9824	-0.9745
	Reset	-0.9717	-0.9686	-0.9712	-0.9309	-0.9404	-0.9633	-0.9772	-0.9605
	СХ	-0.9678	-0.9671	-0.9761	-0.9434	-0.9522	-0.9516	-0.9703	-0.9612
	No gate	0.9617	0.99	0.9821	0.9311	0.9758	0.9663	0.9816	0.9698
	RZ	0.0003	0.0011	-0.001	0.0027	0.0036	-0.0049	0.0009	0.0004
12	SX	-0.0071	-0.0005	-0.0041	-0.0026	-0.0019	-0.0096	-0.0053	-0.0044
127	Х	-0.0188	-0.0004	-0.0122	-0.0113	-0.0099	-0.0221	-0.0124	-0.0124
	Reset	-0.0539	-0.0333	-0.1143	-0.0618	-0.0432	-0.1351	-0.0582	-0.0714
	СХ	-0.022	-0.0053	-0.0298	-0.02	-0.0022	-0.0604	-0.0053	-0.0207
3>	No gate	0.8693	0.7625	0.8596	0.7877	0.9118	0.9602	0.9628	0.8734
	RZ	0.0034	0.0057	0.0027	-0.0063	-0.0121	-0.0012	-0.0003	-0.0012
	SX	-0.0057	-0.0044	-0.0012	-0.0089	-0.0166	-0.007	-0.001	-0.0064
	Х	-0.0064	-0.0026	-0.0064	-0.0198	-0.0145	-0.0177	-0.0091	-0.0109
	Reset	-0.0236	0.0149	-0.0319	-0.019	-0.0246	-0.0366	-0.017	-0.0197
	СХ	-0.0215	0.0581	-0.035	-0.0303	-0.0189	-0.0183	0.002	-0.0091

Table 1: Probabilities of reading $|1\rangle$ state, i.e. $P(|1\rangle)$, as it is affected by different initial energy states of: $|1\rangle$, $|2\rangle$, $|3\rangle$. 'No gate' shows $P(|1\rangle)$ immediately measured after the state preparation. 'RZ, SX, X, Reset, CX' show $P(|1\rangle)$ differences from 'No gate' after applying the specified gate. 'CX' is the result where the target qubit is the denoted qubit and the control qubit is the adjacent qubit, and the control qubit is set to $|1\rangle$.



Figure 6: Probabilities of measuring $|0\rangle$ for different numbers of: (a) the Reset gates and (b) Cascading Secure Reset (CSR) sequences introduced later in this paper.



Figure 7: $|1\rangle$, $|2\rangle$, and $|3\rangle$ decoherence patterns.

in higher-energy states. For $|2\rangle$ and $|3\rangle$ states, all basic gates have very little effect on them. In Table 1 we can see that for $|2\rangle$ and $|3\rangle$

the gates behave almost as 'No gate'. For $|2\rangle$ the largest average difference from 'No gate' is only 7.14% when applying the Reset

CCS '23, November 26-30, 2023, Copenhagen, Denmark

gate. For $|3\rangle$, the interference of basic gates is even smaller, with the largest average difference to be less than 2%.

Among the basic gates, the Reset gate is important for qubit initialization. The Reset gate can be used inside the circuit to actively initialize the qubit, or it can be inserted between shots to quickly initialize the qubit and get rid of the delay time of waiting for the qubit to decohere to the ground state $|0\rangle$. This becomes more and more important as the decoherence time becomes longer and longer with the improvement of qubit manufacturing. Figure 6 (a) shows the probabilities of measuring $|0\rangle$ with different numbers of Reset gates. Similar to other basic gate results, the Reset gate can successfully reset $|0\rangle$ and $|1\rangle$ with small errors. However, it cannot effectively reset qubits in states $|2\rangle$ or $|3\rangle$.

3.4 Decoherence of Higher-Energy States

Over time, higher-energy states will decohere to lower states and eventually approach $|0\rangle$. Researchers have previously presented measurements of coherence and successive decay dynamics of higher-energy levels of a superconducting transmon qubit [28]. In this work, we have analyzed and observed similar behavior on IBM quantum computers.

 T_1 is defined to describe the decoherence behavior, which is the half-life of the $|1\rangle$ state. Suppose a qubit is at state $|1\rangle$, and then after t, the probability of measuring it to be at state $|1\rangle$ is $P(|1\rangle) = e^{-t/T_1}$. Figure 7 shows the decoherence patterns of the different higher-energy states, when sufficient time has passed $P(|0\rangle)$ approaches 1. IBM commonly reports T_1 , and it can be confirmed from our figures that the T_1 for initial state $|1\rangle$ is generally between $100\mu s$ and $200\mu s$, similar to values reported by IBM for the *ibm_lagos* backend. However, we find that for higher-energy states, the decoherence time may increase a lot, which means they may propagate further through circuits.

4 THREAT MODEL

We assume a strong attacker in order to later make strong guarantees about the security of our defense. We assume that quantum computers can be shared where circuits of different users can be executed in an alternating fashion on a set of qubits of a quantum computer. We assume there is a victim user (circuit) and an attacker user (circuit). Typically, the attacker circuit is assumed to execute before the victim circuit so that the attacker can set the qubits into higher-energy states before the same qubits are later used by the victim.

Following the operation of existing IBM quantum computers, we assume there is a reset and delay between attacker and victim circuits. In particular, we assume two circuits are separated by a Reset gate, and rep_delay as has been discussed in Section 2.5.2. Further, if the circuits are part of different jobs, there is additional load_delay. As the load_delay adds significant delay time between jobs, since it is on the order of seconds, it results in a waste of computation time. Thus, we assume a future optimized scenario where through pre-loading of circuits, or other means, load_delay is reduced or even eliminated. However, some form of reset gate cannot be eliminated, otherwise, the qubits would not be properly reinitialized between users.



Figure 8: Diagrams of Reset, $\pi_{i,j}$ and CSR(3). Each line corresponds to an energy level. A red dot on a line indicates the non-zero occupation of the corresponding energy level at a point in time. Observe that starting from a mixed occupation of $|1\rangle$, $|2\rangle$, $|3\rangle$, CSR(3) is designed to clear all high-energy occupation up to $|3\rangle$ at the end.

5 CASCADING SECURE RESET DESIGN

In this section, we introduce the protocol of *Cascading Secure Reset* (CSR) and compare it to the existing reset protocols performed on IBM Quantum Computers.

Given a qubit q with an assumed highest energy level of n, where $n \ge 1$, CSR is performed on q via n reverse-ordered depopulation sequences, described before in Section 2.5.3. Schematically, a CSR sequence performed on a qubit with energy levels up to 3 is provided in Fig (8b). And more concise pseudocode for CSR can be found in Algorithm (1).

As defined in Section 2.2, a $\pi_{i,i-1}$ pulse drives $|i\rangle$ to $|i-1\rangle$ and is its own inverse, i.e. $\pi_{i,i-1} = \pi_{i-1,i}$ and can also drive $|i-1\rangle$ to $|i\rangle$. When performed on a general qubit state $|q\rangle = \alpha_n |n\rangle + \alpha_{n-1} |n-1\rangle + \ldots \alpha_0 |0\rangle$ that does not boast high-purity in one energy-level, the depopulation sequence Depop(n) will excite energy states lower than $|n\rangle$ instead of driving all states to $|0\rangle$, and hence results in a non-ideal reset outcome. This issue is addressed by the CSR sequence, which drives each excited state in $|q\rangle$ down to one state below it per reverse-ordered depopulation sequence and eliminates all excited states after *n* reverse-ordered depopulation sequences on an ideal, noiseless machine. Although no fault-tolerant, noiseless quantum computer exists today, the CSR can defend with reasonable accuracy and performance in resetting a general qubit state with high energy-level components. We evaluate the CSR protocol on an IBM machine in the next section.

As shown in Section 2.5.2, the default reset protocol used by IBM consists of an active reset and a delay afterward. Since reset gates from $|1\rangle$ to $|0\rangle$ do not account for higher-energy states, IBM's reset protocol relies on a sufficiently long delay for higher-energy states to decohere and fall back to $|0\rangle$ or $|1\rangle$. This constitutes a bottleneck for the time it requires to perform a full-system reset. Compared to a typical full-system reset of $251\mu s$, a CSR sequence takes much less time. For CSR(*n*), i.e. a CSR sequence up to energy-level *n*, *n* reset gates, and n(n - 1)/2 number of π pulses are required to perform the sequence. In IBM Qiskit, each custom pulse gate takes 160 dt (or 35.56 ns). This is the time that is allowed for any single-qubit gate, such as the X-gate, and determines the time for a single π pulse. Since CSR(1) is just the usual $|1\rangle \rightarrow |0\rangle$ reset, CSR(1) takes about 1 μs . CSR(2) consists of 2 $|1\rangle \rightarrow |0\rangle$ resets and 1 π

CCS '23, November 26-30, 2023, Copenhagen, Denmark

Algorithm 1	Cascading	Secure	Reset for	Qubit q
-------------	-----------	--------	-----------	---------

Require: Highest energy level of <i>q</i> is $n \ge 1$
Ensure: $ q\rangle \rightarrow 0\rangle$
$N \leftarrow n$
for $i = 1$ to N do
Reset on $ q angle$
for $j = 1$ to $N - i$ do
$\pi_{j,j+1}$ pulse on $ q\rangle$
end for
end for

pulse, and thus takes about 2.07 μ s. CSR(3) consists of 3 $|1\rangle \rightarrow |0\rangle$ resets and 3 π pulse, and thus takes about 3.11 μ s. For transmonbased superconducting qubits, a reasonable high-energy state of a qubit that can be maintained without immediately decohering typically does not exceed $|5\rangle$ (in our experiments we test up to $|3\rangle$ due to limitations of user-level access to IBM machines). Thus, with modest expectation, a CSR sequence requires significantly less time than the current IBM reset protocols, e.g., which use a Reset gate and delay that is on the order of 250 μ s. Moreover, it also achieves excellent performance, which is further analyzed in the next section.

6 CASCADING SECURE RESET EVALUATION

In this section, we evaluate the Cascading Secure Reset against existing reset protocols introduced in Section 2.5. In particular, we focus on four distinct types: thermalization (i.e., idling), Reset, the depopulation sequence Depop, and CSR.

Evaluation is done from three angles. In Section 6.1, we benchmark different reset protocols under a varying amount of end-toend delay between the end of the attacker circuit and the start of the victim circuit, and quantify the state leakage or retention through each reset protocol. Additionally, we examine the covert channel capacity through each reset protocol by modeling the leakage with a binary asymmetric channel. In Section 6.2, we evaluate the impact of higher-energy state attacks on common quantum algorithms and examine the effectiveness of each reset protocol in mitigating this impact.

6.1 State Leakage Across Reset Protocols

The time cost of reset operations between circuit executions typically constitutes a significant portion of the overall device utilization. Therefore, the end-to-end delay between two consecutive circuits is an important metric for device efficiency. Under the setting of high-energy state attacks, we evaluate the tradeoff between end-to-end delay and the extent of state leakage under various reset protocols.

6.1.1 Leakage Experiment Setup. All protocols under test are summarized in Table 2. The protocols are both labeled and described by their IDs: All are evaluated under |3⟩-initialization (p3) and measured at the end (-m). The operator strings -r, -depop3 and -csr3 represent Reset, Depop(3) and CSR(3) respectively. The number prepended to an operator string denotes back-to-back repetitions of the operation. The delay string -d denotes a variable amount of delay, which can be inserted either in the first (e.g., -d-r) or the last

Chuanqi Xu, Jessie Chen, Allen Mi, and Jakub Szefer

Туре	Protocol ID	Top	$N_{\rm P}$	N _R	Delay
Therm.	p3-d-m	0.00	0	0	N/A
Posot	p3-r-d-m (FSW)	1.00	0	1	Last
Neset	p3-d-r-m	1.00	0	1	First
Donon(2)	p3-depop3-d-m	1.07	2	1	Last
Depop(3)	p3-d-depop3-m	1.07	2	1	First
	p3-1csr3-d-m	3.11	3	3	Last
	p3-d-1csr3-m	3.11	3	3	First
CSR(3)	p3-2csr3-d-m	6.21	6	6	Last
	p3-d-2csr3-m	6.21	6	6	First
	p3-3csr3-d-m	9.32	9	9	Last
	p3-d-3csr3-m	9.32	9	9	First

Table 2: Protocols considered in state leakage evaluation. $T_{\rm op}$ denotes its operator length in seconds. $N_{\rm P}$ and $N_{\rm R}$ denote the number of π -pulses and Reset gates respectively. IBM's default full-system wipe is highlighted in bold font.

position (e.g., -r-d) of the protocol. For each protocol, the end-toend delay it induces is the sum of its operator length T_{op} and the variable delay.

As discussed in Section 3.4, repeated execution of the Reset gate is ineffective in eliminating higher-energy states. Additionally, Section 2.5.3 shows the instability of the depopulation sequence under repetition. Therefore, we only evaluate the single-execution case for Reset and Depop(3). This includes IBM's default full-system wipe (p3-r-d-m). For CSR(3), we evaluate up to three repetitions to analyze its stability under repetition and potential performance gains.

We perform the evaluation on qubits 0 and 3 on Lagos. As shown in Figure 7, qubit 3 has the longest decoherence time among all qubits on Lagos, which may likely lead to increased state retention. Therefore, we select it for worst-case performance evaluation. Additionally, we select qubit 1 as a typical-case scenario.

6.1.2 Leakage Experiment Results. Figure 9 describes the amount of state leakage in terms of 1-output frequency measured after each reset protocol. In particular, Figure 9a compares protocols that contain a single CSR(3) with Depop(3) and Reset-type protocols, and Figure 9b extends the evaluation to repeated CSR(3) operations.

For each qubit, \lim_T and \lim_A respectively denote the minimum 1-output frequency at full thermalization and immediately after a Reset. Both values are obtained after fully thermalizing the system to $|0\rangle$. \lim_T establishes a lower bound on state leakage for protocols with variable delay after the operation. Similarly, \lim_A lower-bounds the leakage for protocols with a variable delay before the operation. These values are dependent on device architecture, calibration, and ambient factors. $\lim_T > \lim_A$ for qubit 0 and $\lim_T < \lim_A$ for qubit 3.

As expected, all protocols perform worse on qubit 3 due to its longer decoherence time. For both qubits, IBM's default FSW is not effective under higher-energy state attacks, as p3-r-d-m converges towards the thermalization protocol p3-d-m in each case. For Resettype protocols, placing variable delay before reset is a more effective strategy, evidenced by the faster delay in the pattern of p3-d-r-m.

The performance of Depop(3) varies significantly between qubits. As delay increases, p3-d-depop3-m fails to depopulate higher-energy



(a) Comparison of CSR against full-system wipe and depopulation.

(b) Effect of back-to-back execution of multiple CSRs.

Figure 9: State leakage of various reset protocols as a function of end-to-end delay. The dotted vertical line in each panel represents the end-to-end delay of $251\mu s$ in IBM's default full-system wipe. The colored horizontal regions represent the interval between \lim_{T} and \lim_{T} .

states due to its reliance on $|3\rangle$ purity. As time progresses during the variable delay in the protocol, some population of $|3\rangle$ decays to $|2\rangle$ or $|1\rangle$, which is then elevated to $|3\rangle$ or $|2\rangle$ by Depop(3). On the other hand, p3-depop3-m exhibits competitive performance on both qubits.

Limited to single-execution, the CSR(3) protocol p3-d-1csr3-m is the best performing protocol for the majority of the range of delays from $10\mu s$ to the typical $251\mu s$. At $10\mu s$, it achieves on qubits 0 and 3 state leakage factors of 0.39x and 0.27x of the default FSW at $251\mu s$. The repeated CSR(3) protocols lead to further performance improvements, achieving minimum leakage factors of 0.13x and 0.25x in the same setting. Compared to the default FSW, CSR(3) achieves large-factor decreases in state leakage at less than 1/25 of the time cost. Full details of leakage factors at delay values of $10\mu s$, $50\mu s$, $100\mu s$ and $251\mu s$ are presented in Table 3.

At large delay values, p3-3csr-d-m is the best-performing protocol on qubit 0, while p3-d-3csr-m performs the best on qubit 3. This distinction is likely due to the difference in the ordering of \lim_T and \lim_A by magnitude between qubits 0 and 3, as the two strategies converge to \lim_T and \lim_A respectively at full thermalization.

6.1.3 Capacity of Covert Channel via State Leakage. The capacity through a channel describes the amount of information (i.e., number of bits) one can recover through one use of the channel. In the context of reset protocols, lower channel capacity through the protocol leads to reduced information leakage through the channel, and consequently reduced security impacts from covert-channel and side-channel attacks. We evaluate the capacity of each protocol as a covert channel.

Denote the preparation and measurement of a reset protocol by X and Y respectively. Let X = 1 denote the event that the protocol is initialized to $|3\rangle$, and let X = 0 denote a $|0\rangle$ -initialization. Similarly, define Y = 1 and Y = 0 as the events that the measurement returns an output of 1 and 0 respectively. Then a reset protocol can be modeled as a binary asymmetric channel by specifying the

conditional error probabilities P(Y = 1|X = 0) and P(Y = 0|X = 1). In the context of leakage, P(Y = 0|X = 1) is approximated by $1 - P(|1\rangle)$, where $P(|1\rangle)$ denotes the 1-output frequency. To place an upper-bound on the capacity of the leakage channel, we lower-bound the other error probability by specifying P(Y = 1|X = 0) to be \lim_{T} for each qubit. Under this setup, the capacity of each reset protocol describes the maximum average number of bits that can be transferred across the protocol each time by modulating the initialization between $|0\rangle$ and $|3\rangle$.

Figure 10 compares the covert channel capacity of each reset protocol and shows results qualitatively similar to state leakage. Again, the CSR(3) protocol p3-d-1csr3-m is the best performing single-execution protocol for most of $10\mu s$ to $251\mu s$. At $10\mu s$, it achieves on qubits 0 and 3 capacity factors of 0.23x and 0.10x of the default FSW at $251\mu s$. With repetition, p3-d-3csr3-m achieves capacity factors of 0.01x and 0.09x in the same setting. Overall, CSR(3) achieves a channel capacity reduction of between 1 and 2 orders of magnitude compared to default FSW at less than 1/25 the time cost. Full details of capacity factors at delay values of $10\mu s$, $50\mu s$, $100\mu s$ and $251\mu s$ are presented in Table 3.

6.2 Attacks on Quantum Algorithms

We experiment with quantum algorithms to show that higherenergy states can be used by attackers. Because basic gates are only calibrated based on $|0\rangle$ and $|1\rangle$, they are ineffective on higherenergy states, especially since the Reset gate cannot successfully initialize some states to $|0\rangle$, higher-energy states will propagate through shots if the between-shot mechanism uses default reset gates, and thus may result in a wrong probability distribution in the subsequent circuits.

6.2.1 Attack Experiment Setup. All experiments consist of two circuits in sequence. The first circuit is the state preparation circuit to simulate the attacker, which prepares one of $|0\rangle$, $|1\rangle$, $|2\rangle$, $|3\rangle$. The second circuit is the circuit that implements one of the quantum

CCS '23, November 26-30, 2023, Copenhagen, Denmark

Chuanqi Xu, Jessie Chen, Allen Mi, and Jakub Szefer



(a) Comparison of CSR against default reset and depopulation.

(b) Effect of back-to-back execution of multiple CSRs.

Figure 10: Worst-case Shannon capacity across the state leakage covert channel of various reset protocols as a function of end-to-end delay. The dotted lines represent the default reset delay. Capacity values are calculated 1-output frequency and \lim_T via an asymmetric binary channel.

Туре	Drata col ID	10µs (Minimum)		50 <i>µs</i>		100 <i>µs</i>		251µs (Default FSW)	
	FIOLOCOLID	Leakage	Capacity	Leakage	Capacity	Leakage	Capacity	Leakage	Capacity
Therm.	p3-d-m	2.71	4.67	2.08	2.95	1.63	2.01	8.93 e-1	7.91 e-1
Reset	p3-r-d-m (FSW)	2.56	4.01	2.06	2.90	1.71	2.18	1.00	1.00
	p3-d-r-m	2.44	3.88	1.54	1.89	9.07 e-1	9.03 e-1	3.03 e-1	1.47 e-1
Depop(3)	p3-depop3-d-m	1.59	2.05	1.17	1.28	9.17 e-1	9.02 e-1	5.23 e-1	3.80 e-1
	p3-d-depop3-m	1.31	1.49	1.37	1.58	1.31	1.50	8.59 e-1	8.08 e-1
	p3-1csr3-d-m	5.20 e-1	3.90 e-1	4.39 e-1	2.94 e-1	3.63 e-1	2.13 e-1	2.08 e-1	6.96 e-2
	p3-d-1csr3-m	3.85 e-1	2.28 e-1	2.29 e-1	9.75 e-2	1.35 e-1	2.19 e-2	1.07 e-1	8.60 e-3
CSR(3)	p3-2csr3-d-m	2.26 e-1	7.31 e-2	1.75 e-1	4.52 e-2	1.54 e-1	3.18 e-2	1.19 e-1	1.36 e-2
	p3-d-2csr3-m	2.01 e-1	6.24 e-2	1.41 e-1	2.58 e-2	1.16 e-1	1.14 e-2	1.09 e-1	9.04 e-3
	p3-3csr3-d-m	1.54 e-1	3.07 e-2	1.32 e-1	1.90 e-2	1.21 e-1	1.33 e-2	9.95 e-2	6.14 e-3
	p3-d-3csr3-m	1.29 e-1	1.34 e-2	1.15 e-1	1.01 e-2	1.06 e-1	7.80 e-3	9.99 e-2	5.92 e-3
(a) Lagas Oubit 0									

(a) Lagos, Qubit 0.

Туре	Protocol ID	10μs (Minimum)		50 <i>µs</i>		100 <i>µs</i>		251µs (Default FSW)	
	TOUCCOTID	Leakage	Capacity	Leakage	Capacity	Leakage	Capacity	Leakage	Capacity
Therm.	p3-d-m	1.61	2.46	1.56	2.28	1.45	1.94	9.96 e-1	9.92 e-1
Reset	p3-r-d-m (FSW)	1.27	1.51	1.33	1.62	1.35	1.67	1.00	1.00
	p3-d-r-m	1.27	1.48	1.13	1.23	9.15 e-1	8.75 e-1	4.48 e-1	2.66 e-1
Depop(3)	p3-depop3-d-m	3.82 e-1	1.97 e-1	3.80 e-1	1.94 e-1	3.58 e-1	1.75 e-1	2.56 e-1	9.11 e-2
	p3-d-depop3-m	5.40 e-1	3.58 e-1	7.42 e-1	6.21 e-1	9.31 e-1	8.80 e-1	8.48 e-1	7.61 e-1
	p3-1csr3-d-m	4.65 e-1	2.84 e-1	4.57 e-1	2.74 e-1	4.30 e-1	2.56 e-1	3.26 e-1	1.48 e-1
CSR(3)	p3-d-1csr3-m	2.69 e-1	1.03 e-1	2.53 e-1	8.93 e-2	2.37 e-1	6.65 e-2	2.52 e-1	8.82 e-2
	p3-2csr3-d-m	3.48 e-1	1.70 e-1	3.32 e-1	1.52 e-1	3.13 e-1	1.36 e-1	2.43 e-1	8.18 e-2
	p3-d-2csr3-m	2.93 e-1	1.21 e-1	2.84 e-1	1.13 e-1	2.76 e-1	1.07 e-1	2.69 e-1	1.01 e-1
	p3-3csr3-d-m	3.07 e-1	1.29 e-1	2.83 e-1	1.10 e-1	2.62 e-1	9.50 e-2	2.04 e-1	5.43 e-2
	p3-d-3csr3-m	2.49 e-1	8.55 e-2	2.43 e-1	8.14 e-2	2.42 e-1	7.73 e-2	2.28 e-1	7.02 e-2

(b) Lagos, Qubit 3.

Table 3: State leakage and covert channel capacity of various reset protocols at various end-to-end delays, scaled relative to IBM's default full-system wipe at 251µs. The minimum value of each column and its corresponding protocol are in bold.

algorithms, which serves as the victim. To make sure that states are

mostly initialized to $|0\rangle$ at the end of the second circuit, we append

a 2×10^6 dt delay (around 444 μ s) after the measurement of the second circuit to let states decohere to the ground state. Between these two circuits, three inter-shot mechanisms are evaluated:

- (1) Default: Include the Reset gate specified by init qubits = True together with a 250μ s delay specified by rep_delay = 250e - 6.
- (2) Reset: Only add the Reset gate at the end of the circuit and no delay.
- (3) CSR: Only add CSR at the end of the circuit and no delay.

To show that higher-energy states can ruin the results of victim circuits, we choose four well-known quantum algorithms as an example. As shown in Figure 11, these four quantum circuits are:

- (1) Deutsch-Jozsa (DJ) [7]: We used a 2-qubit Deutsch-Jozsa algorithm with a balanced oracle. The balanced oracle f(x)satisfies $f(|00\rangle) = f(|11\rangle) = 0, f(|01\rangle) = f(|10\rangle) = 1$. The preparation circuit prepares states on the non-measured qubit (q_1 in Figure 11a), while on the other two measured qubits (q_0, q_2) , there is no gate in the preparation circuit. The theoretical outcome for this circuit is $P(|11\rangle) = 1$.
- (2) Inverse Quantum Fourier Transformation (IQFT) [4]: We used a 3-qubit IQFT circuit. The initial state is chosen to be $\left|\tilde{0}\right\rangle = \frac{1}{\sqrt{8}} \sum_{i=0}^{7} |i\rangle$. The theoretical outcome for this circuit is $P(|000\rangle) = 1.$
- (3) Grover's Search (GS) [14]: We used a 2-qubit Grover's Search to search for $|00\rangle$. The theoretical outcome for this circuit is $P(|00\rangle) = 1.$
- (4) Variational Quantum Eigensolver (VQE) [27]: we used a 3qubit VQE circuit. The operator to be minimized is $I \otimes Z \otimes Z$. The ansatz, i.e., the parameterized circuit to be estimated and optimized, is the hardware efficient SU(2) 2-local circuit [17]. The optimizer is simultaneous perturbation stochastic approximation (SPSA) [37]. Random initial points are the same in all experiments with a random seed.

These four circuits can represent four scenarios of quantum computing tasks to which the higher energy state attack can be applied. For the Deutsch-Jozsa circuit, we only prepared state on the non-measured qubit (qubit q_1 in Figure 11), and in this case, we show that higher energy states on the qubits that control the circuit and are not directly measured can still attack the circuit. On the contrary, we prepared states on the measured qubits in Grover's search circuit, which presents how higher energy states can propagate through the circuits and influence the final results. As for the inverse quantum Fourier transformation circuit, we would like to show how higher energy states perform in the problem with continuous answers instead of the decision problems shown in Deutsch-Jozsa and Grover's search. In the end, VQE illustrates the variational algorithms that include both the classical and the quantum process, where the idea is also used in some optimization algorithms such as quantum machine learning [2].

6.2.2 Experiment Result. In Table 4 we present the probability distributions of DJ and GS with different inter-shot mechanisms and initialization states. For DJ, if we make decisions based on a threshold, i.e., f(x) is balanced if $P(|00\rangle) < t$, and is constant if $P(|00\rangle) > t$, then t determines whether the attack is successful. If we choose t = 0.5, then higher energy states only on the nonmeasured qubit may not change the decision since the decision

		Delaun	$ 2\rangle$	0.02	0.06	0.05
e	DJ		3>	0.03	0.07	0.06
		Reset	$ 0\rangle$	0.01	0.03	0.03
			$ 1\rangle$	0.01	0.02	0.04
			$ 2\rangle$	0.14	0.50	0.10
			3>	0.15	0.30	0.18
			$ 0\rangle$	0.01	0.03	0.04
		CSR	$ 1\rangle$	0.01	0.04	0.04
			$ 2\rangle$	0.02	0.04	0.04
			3>	0.04	0.07	0.06
			$ 0\rangle$	0.96	0.02	0.02
		Default	$ 1\rangle$	0.96	0.01	0.01
			$ 2\rangle$	0.84	0.03	0.02
			1 =)			

Mec.

Default

Init.

 $|0\rangle$

 $|1\rangle$

10)

 $|00\rangle$

0.01

0.01

0.02

Alg.

0.02 0.04 0.82 $|0\rangle$ 0.95 0.02 0.02 $|1\rangle$ 0.88 0.04 0.02 GS Reset $|2\rangle$ 0.02 0.04 0.03 $|3\rangle$ 0.05 0.06 0.14 $|0\rangle$ 0.94 0.02 0.02 $|1\rangle$ 0.90 0.03 0.02 CSR

 $|2\rangle$

 $|3\rangle$

Table 4: Probabilities of $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$ of Deutsch-Jozsa (DJ) and Grover's search (GS). The initial states are $|0\rangle$, $|1\rangle$, $|2\rangle$, and $|3\rangle$, and the between-shot mechanisms are the default mechanism on IBM Quantum, Reset gate with no repetition delay, and cascading secure reset (CSR) with no repetition delay. The bold values indicate the largest probabilities.

0.89

0.71

0.02

0.04

0.03

0.04

space consists of only two items, and the probability changes are small compared with the threshold. However, higher energy states prepared on measured qubits are more powerful. For GS whose decision space can be very large, higher energy states can change the decision. If we choose the state with the largest probability to be the good state, then the decisions are changed from $|00\rangle$ to $|11\rangle$ with the Reset mechanism if the prepared states are $|2\rangle$ or $|3\rangle$.

Figure 12 shows the probabilities of reading the theoretically 100% outcome states on Lagos. As shown in Figure 12, the results of the attacker setting the qubit states to $|0\rangle$ and $|1\rangle$ are very similar, which indicates all three inter-shot mechanisms can effectively initialize $|0\rangle$ and $|1\rangle$. On the other hand, the results of $|2\rangle$ and $|3\rangle$ vary a lot depending on the inter-shot mechanism. If only the Reset gate is used, the output probabilities for cases of the attacker setting $|2\rangle$ and $|3\rangle$ states will differ greatly from the correct values.

In contrast to IQFT where the outcomes are continuous, DJ and GS are decision problems and their decisions are discrete and need to be decoded from the results. Therefore, the decisions are depending on how the results are interpreted. DJ is to determine whether f(x) is a constant function, i.e., all outputs are the same, or a balanced function, i.e., half the outputs are 0 and half the outputs are 1. This is done by measuring $|0\rangle^{\otimes n}$. $P(|0\rangle^{\otimes n}) = 1$ if f(x) is constant and $P(|0\rangle^{\otimes n}) = 0$ otherwise. GS is to search one or several good states. The good state is the state with the highest probability when measuring at the end of the circuit.

CCS '23, November 26-30, 2023, Copenhagen, Denmark

 $|10\rangle$

0.04

0.05

 $|11\rangle$

0.91

0.90

0.87

0.84

0.93

0.92

0.26

0.37

0.91

0.91

0.91

0.83

0.00

0.02

0.11

0.13

0.01

0.07

0.91

0.75

0.02

0.05

0.07

0.20

 $|01\rangle$

0.03

0.04

0.06

CCS '23, November 26-30, 2023, Copenhagen, Denmark

$q_1: \longrightarrow X \longrightarrow H \longrightarrow \Phi$	+ + + +
<i>q</i> ₂ :H	
$C: /\frac{1}{2}$	



(a) 2-qubit Deutsch-Jozsa with a balanced oracle. (b) 3-qubit inverse quantum Fourier transformation with $|\tilde{0}\rangle$ as the initial state.



(c) 2-qubit Grover's search with $|00\rangle$ as the target state.

$q_0: - \mathbb{R}_{Y}(\theta[0]) - \mathbb{R}_{Z}(\theta[3])$	•	$R_{Y}(\theta[6])$	$-R_Z(\theta[9])$
$q_1: - \mathbb{R}_{Y}(\theta[1]) - \mathbb{R}_{Z}(\theta[4])$	⊢⊕_∙	$-R_{Y}(\theta[7])$	$R_Z(\theta[10])$
$q_2: - \mathbb{R}_{Y}(\theta[2]) - \mathbb{R}_{Z}(\theta[5]) -$	└─────	$R_{Y}(\theta[8])$	$R_Z(\theta[11])$

(d) 3-qubit VQE ansatz. The operator is chosen to be $I \otimes Z \otimes Z$.

Figure 11: Victim quantum circuits used in the attack experiments.



Figure 12: Probabilities of the theoretically largest state of Deutsch-Jozsa, inverse quantum Fourier transformation, and Grover's search used in our experiments. Different color bars indicate the qubit energy state set by the attacker prior to the execution of the victim circuit. Attackers setting high energy states $|2\rangle$ or $|3\rangle$ can in some cases significantly affect output probabilities.

As for VQE, we show how the optimization processes are different under 6 inter-shot mechanisms in Figure 13. These 6 inter-shot mechanisms differ in the type of the reset gate, either the Reset gate or CSR, and in rep_delay between shots, including $0\mu s$, $125\mu s$, $250\mu s$. In Figure 13, the three plots above show that with higherenergy states injected into the initial states, the optimizer cannot minimize the expected value if higher-energy states are not successfully reset. In this case, the gates in the ansatz cannot effectively transform the initial states to the desired ansatzes. In addition, with rep_delay increases, the expected values that cannot be minimized are decreased, which is due to the decoherence of higher-energy states over time. Note that because of the noise in the quantum device, we may not make any conclusion on whether the optimization process can be sped up. Intuitively, the degree that higher-energy states are reset to $|0\rangle$ is smaller for CSR with rep_delay=0 than CSR with rep_delay=125e-6 or rep_delay=250e-6. However, in Figure 13, the optimization lines in the first case are lower than those in the latter two cases. While we were reproducing our results, we also found that the optimization process may vary a lot even though using the same circuits and within a short period, which is also



Figure 13: The optimization process of VQE. Each subfigure shows the optimization process with one inter-shot mechanism which is indicated on the top of the subfigure as the title. Each inter-shot mechanism consists of one reset gate type from the normal Reset gate and CSR, and the number after the underscore symbol specifies rep_delay in μs . Reset_250 applies the default mechanism, where init_qubits=True, rep_delay=250e-6.

consistent with this statement. Nevertheless, whether the expected value can be optimized is consistent throughout all experiments.

6.3 Trojan Circuit Attack

As a further extended evaluation, we demonstrate a novel type of Trojan circuit. To show that higher-energy states can be used for Trojan attacks, we implement one scheme that can be used by the attack: at the end of the victim circuit, following the measurement operations, the Trojan can be secretly inserted which can add pulse gates to drive the qubit from $|1\rangle$ to $|n\rangle$. Since the frequency of the appended pulse gates is for higher-energy levels, if the measurement result is $|1\rangle$, it will be driven to $|n\rangle$, while it does almost nothing if the measurement result is $|0\rangle$. In this case, the information of the victim's measurement is encoded into the high-energy state. If the attacker can execute their circuit after the victim's circuit, they can



Figure 14: Measurement probabilities of the attacker when the victim circuit's individual qubits are driven from $|1\rangle$ to $|3\rangle$ (D) and when they are not driven (ND) following the measurement of the victim circuits. The figure shows the attackers' measurement probabilities for different inter-shot mechanisms. P(x|y) is the probability of xas the measurement result for the attack circuit under the condition that the measurement result of the victim circuit is y and then followed by the inter-shot mechanism.



Figure 15: Example the victim circuit when it is 2-qubit Grover search to search for $|11\rangle$. The figure shows the probability of attackers measuring the correct $|11\rangle$ with Trojan driving the qubits to higher-energy states (D) and without Trojan driving the qubits to higher-energy states (ND).

retrieve information about the measurement. The Trojan circuit could be inserted, for example, by a malicious compiler, or it can be hidden as part of code inadvertently downloaded by the user.

Figure 14 shows the probabilities of measuring $|0\rangle$ and $|1\rangle$ for qubits when they have been driven (D) and not driven (ND) into higher-energy states after measurement. From Figure 14 we can observe that the reset gate is not effective. Our CSR is much more effective than the reset gate. The default thermalization of 250 us is the best, but it is 25x slower than our CSR.

To evaluate the Trojan attack on a specific circuit, we inserted the Trojan gates at the end of the 2-qubit Grover search to search for $|11\rangle$. Figure 15 shows that the reset gate is again not effective, and the attacker can correctly recover the $|11\rangle$ state. With CSR, the probability of the $|11\rangle$ state being measured is very low. Default thermalization works better, but again takes 25x times as long.

7 RELATED WORK

7.1 Higher-Energy States

Support for higher-energy states has started to draw more attention recently. Lately, there is a new feature that Qiskit has introduced to enable work with qumodes, which are realized as higher-energy states in superconducting machines [36]. There is also software support for qutrits for Rigetti [32]. Amazon Braket API recently also provides support for pulse-level control [15] which may be used to generate high-energy states. The higher-energy state attacks can also be exploited in quantum computers with other qubit technology, such as trapped ions. Also, existing work has explored the use of qutrits [10, 12, 33], and simply disabling higher-energy states is not feasible as it limits research on novel algorithms.

7.2 Quantum Computer Security

Recent work [1, 24, 29, 34] has shown different attacks on quantum computers, from fault injection in multi-tenant settings, to fingerprinting devices. Motivated by the attacks, recent work has explored some preliminary defenses. One approach to crosstalk prevention is detecting crosstalk errors by analyzing the execution of a circuit [35], then mitigating them using connectivity reduction, qubit frequency tuning, and coupler tuning are proposed [9]. Other defenses include instruction scheduling modifications [26], or a proposal for quantum computer antivirus [6]. None of these attacks or defenses have explored higher-energy states.

The most closely related work to our research is work on secure reset gates [25]. The researchers have shown in their work that existing reset gates in IBM quantum computers cannot reliably reset the $|1\rangle$ state of the qubits to $|0\rangle$. As a result, there is an information leak across reset gates that could be abused by malicious users or programs. Applying multiple resets does not mitigate the problem they identify, and the authors proposed a new secure reset gate. Their work, however, did not consider higher-energy states, as we do. Applying multiple reset gates, or applying their "secure reset" gate, does not prevent the new problem uncovered in our work.

Considering higher-energy states, the higher-energy states have potentially useful applications. For example, researchers are exploring qutrit systems. A qutrit leverages three energy states, $|0\rangle$, $|1\rangle$, and $|2\rangle$. Using qutrits can reduce circuit cost for important algorithms like quantum neurons and Grover search [12]. Higherenergy states can also be used for new error correction code [21]. In parallel, quantum computer designers are working on new hardware that can effectively reset higher-energy states [11, 20, 22, 31]. However, these require either new hardware or more control of the systems, and thus cannot be deployed on current cloud-based quantum computers, such as from IBM. They have also not considered malicious users who set higher-energy states on purpose.

8 CONCLUSION

This work demonstrated for the first time a new type of *higher-energy state attack*, where a malicious user or circuit sets the qubit state to $|2\rangle$ or higher to interfere with other circuits or facilitate information transmission through a covert channel. We demonstrated that the common reset gate is ineffective in resetting the qubit state, resulting in state leakage from a previous circuit to the next. This effect opens a new vector where an attacker can set a higher-energy state that impacts computation afterward. To provide a defense, we proposed the *Cascading Secure Reset* (CSR) operation, which, without hardware modifications, is able to efficiently reset higher-energy states back to $|0\rangle$ with high fidelity. Both the new attack and the CSR operation were prototyped and evaluated on real IBM

quantum computers. Compared to the default full-system wipe of length $251\mu s$, CSR achieves a reduction in $|3\rangle$ -initialized state leakage channel capacity by between 1 and 2 orders of magnitude, and does so within $10\mu s$, a 25x speedup over the default operation. CSR can be an enabling technology for shared or multi-programmed quantum computers in light of higher-energy state attacks.

ACKNOWLEDGMENTS

The authors would like to thank IBM and Yale University for providing access to IBM's superconducting devices.

REFERENCES

- [1] Abdullah Ash-Saki, Mahabubul Alam, and Swaroop Ghosh. 2020. 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 (Boston, Massachusetts) (ISLPED '20). Association for Computing Machinery, New York, NY, USA, 25–30. https://doi.org/10.1145/ 3370748.3406570
- [2] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum machine learning. *Nature* 549, 7671 (2017), 195–202.
- [3] Juan I Cirac and Peter Zoller. 1995. Quantum computations with cold trapped ions. *Physical Review Letters* 74, 20 (1995), 4091.
- [4] D. Coppersmith. 2002. An approximate Fourier transform useful in quantum factoring. https://doi.org/10.48550/ARXIV.QUANT-PH/0201067
- [5] Poulami Das, Swamit S Tannu, Prashant J Nair, and Moinuddin Qureshi. 2019. A case for multi-programming quantum computers. In International Symposium on Microarchitecture (MICRO). 291–303.
- [6] Sanjay Deshpande, Chuanqi Xu, Theodoros Trochatos, Yongshan Ding, and Jakub Szefer. 2022. Towards an Antivirus for Quantum Computers. In Proceedings of the International Symposium on Hardware Oriented Security and Trust (HOST).
- [7] David Deutsch and Richard Jozsa. 1992. Rapid solution of problems by quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences 439, 1907 (1992), 553–558.
- [8] David Elieser Deutsch, Adriano Barenco, and Artur Ekert. 1995. Universality in quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences 449, 1937 (1995), 669–677.
- [9] Yongshan Ding and Frederic T. Chong. 2020. Quantum Computer Systems: Research for Noisy Intermediate-Scale Quantum Computers. Synthesis Lectures on Computer Architecture (2020).
- [10] Alexey Galda, Michael Cubeddu, Naoki Kanazawa, Prineha Narang, and Nathan Earnest-Noble. 2021. Implementing a ternary decomposition of the toffoli gate on fixed-frequencytransmon qutrits. arXiv preprint arXiv:2109.00558 (2021).
- [11] Kurtis Geerlings, Zaki Leghtas, Ioan M Pop, Shyam Shankar, Luigi Frunzio, Robert J Schoelkopf, Mazyar Mirrahimi, and Michel H Devoret. 2013. Demonstrating a driven reset protocol for a superconducting qubit. *Physical review letters* 110, 12 (2013), 120501.
- [12] Pranav Gokhale, Jonathan M Baker, Casey Duckering, Natalie C Brown, Kenneth R Brown, and Frederic T Chong. 2019. Asymptotic improvements to quantum circuits via qutrits. In Proceedings of the 46th International Symposium on Computer Architecture. 554–566.
- [13] David J Griffiths and Darrell F Schroeter. 2018. Introduction to quantum mechanics. Cambridge university press.
- [14] Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing (Philadelphia, Pennsylvania, USA) (STOC '96). Association for Computing Machinery, New York, NY, USA, 212–219. https://doi.org/10.1145/ 237814.237866
- [15] Jean-Christophe Jaskula, Kshitij Chhabra, Peter Karalekas, and Stefan Natu. [n.d.]. Amazon Braket launches Braket Pulse to develop quantum programs at the pulse level | Amazon Web Services – aws.amazon.com. https://aws.amazon.com/blogs/quantum-computing/amazon-braket-launchesbraket-pulse-to-develop-quantum-programs-at-the-pulse-level/. [Accessed 27-Feb-2023].
- [16] Jonathan A Jones, Michele Mosca, and Rasmus H Hansen. 1998. Implementation of a quantum search algorithm on a quantum computer. *Nature* 393, 6683 (1998), 344–346.
- [17] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. 2017. Hardware-efficient variational

quantum eigensolver for small molecules and quantum magnets. *Nature* 549, 7671 (2017), 242–246.

- [18] Benjamin P Lanyon, James D Whitfield, Geoff G Gillett, Michael E Goggin, Marcelo P Almeida, Ivan Kassal, Jacob D Biamonte, Masoud Mohseni, Ben J Powell, Marco Barbieri, et al. 2010. Towards quantum chemistry on a quantum computer. Nature Chemistry 2, 2 (2010), 106–111.
- [19] Yuan Liu, Jasmine Sinanan-Singh, Matthew T Kearney, Gabriel Mintzer, and Isaac L Chuang. 2021. Constructing qudits from infinite-dimensional oscillators by coupling to qubits. *Physical Review A* 104, 3 (2021), 032605.
- [20] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J.-C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff. 2018. Fast and Unconditional All-Microwave Reset of a Superconducting Qubit. *Phys. Rev. Lett.* 121 (Aug 2018), 060502. Issue 6. https://doi.org/10.1103/PhysRevLett.121.060502
- [21] Ritajit Majumdar, Saikat Basu, Shibashis Ghosh, and Susmita Sur-Kolay. 2018. Quantum error-correcting code for ternary logic. *Physical Review A* 97, 5 (may 2018). https://doi.org/10.1103/physreva.97.052302
- [22] Matt McEwen, Dvir Kafri, Z Chen, Juan Atalaya, KJ Satzinger, Chris Quintana, Paul Victor Klimov, Daniel Sank, C Gidney, AG Fowler, et al. 2021. Removing leakage-induced correlated errors in superconducting quantum error correction. *Nature communications* 12, 1 (2021), 1–7.
- [23] N David Mermin. 2007. Quantum computer science: an introduction. Cambridge University Press.
- [24] Allen Mi, Shuwen Deng, and Jakub Szefer. 2021. Short Paper: Device- and Locality-Specific Fingerprinting of Shared NISQ Quantum Computers. In Workshop on Hardware and Architectural Support for Security and Privacy (HASP).
- [25] Allen Mi, Shuwen Deng, and Jakub Szefer. 2022. Securing Reset Operations in NISQ Quantum Computers. In Proceedings of the Conference on Computer and Communications Security (CCS).
- [26] Prakash Murali, David C. Mckay, Margaret Martonosi, and Ali Javadi-Abhari. 2020. Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers. Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems (Mar 2020). https://doi.org/10.1145/3373376.3378477
- [27] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature communications* 5, 1 (2014), 1–7.
- [28] Michael J. Peterer, Samuel J. Bader, Xiaoyue Jin, Fei Yan, Archana Kamal, Ted Gudmundsen, Peter J. Leek, Terry P. Orlando, William D. Oliver, and Simon Gustavsson. 2014. Coherence and Decay of Higher Energy Levels of a Superconducting Transmon Qubit. Phys. Rev. Lett. 114, 010501 (2015). (2014). https://doi.org/10.1103/PhysRevLett.114.010501 arXiv:arXiv:1409.6031
- [29] Koustubh Phalak, Abdullah Ash Saki, Mahabubul Alam, Rasit Onur Topaloglu, and Swaroop Ghosh. 2021. Quantum PUF for Security and Trust in Quantum Computing. IEEE Journal on Emerging and Selected Topics in Circuits and Systems 11, 2 (2021), 333–342. https://doi.org/10.1109/JETCAS.2021.3077024
- [30] John Preskill. 2018. Quantum computing in the NISQ era and beyond. Quantum 2 (2018), 79.
- [31] Matthew D Reed, Blake R Johnson, Andrew A Houck, Leonardo DiCarlo, Jerry M Chow, David I Schuster, Luigi Frunzio, and Robert J Schoelkopf. 2010. Fast reset and suppressing spontaneous emission of a superconducting qubit. Applied Physics Letters 96, 20 (2010), 203110.
- [32] Rigetti. [n.d.]. manipulating_qutrits_in_quilt.ipynb gist.github.com. https://gist.github.com/rigetti-gist/df881f61bc2bf53c5c7821c94b2c8dcd#filemanipulating_qutrits_in_quilt-ipynb. [Accessed 27-Feb-2023].
- [33] Tanay Roy, Ziqian Li, Eliot Kapit, and David I Schuster. 2022. Realization of two-qutrit quantum algorithms on a programmable superconducting processor. arXiv preprint arXiv:2211.06523 (2022).
- [34] Abdullah Ash Saki and Swaroop Ghosh. 2021. Qubit Sensing: A New Attack Model for Multi-programming Quantum Computing. arXiv preprint arXiv:2104.05899 (2021).
- [35] Mohan Sarovar, Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout. 2020. Detecting crosstalk errors in quantum information processors. *Quantum* 4 (Sept. 2020), 321. https://doi.org/10.22331/q-2020-09-11-321
- [36] Kevin C. Smith, Eleanor Crane, Tim Stavenger, and S.M. Girvin. [n.d.]. Introducing Bosonic Qiskit: A package for simulating bosonic and hybrid qubit-bosonic circuits – medium.com. https://medium.com/qiskit/introducingbosonic-qiskit-a-package-for-simulating-bosonic-and-hybrid-qubit-bosoniccircuits-1e1e528287bb. [Accessed 27-Feb-2023].
- [37] James C Spall. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE transactions on automatic control* 37, 3 (1992), 332–341.