

Received December 3, 2021, accepted December 15, 2021, date of publication December 28, 2021, date of current version January 7, 2022.

Digital Object Identifier 10.1109/ACCESS.2021.3139085

Shuttle-Exploiting Attacks and Their Defenses in Trapped-Ion Quantum Computers

ABDULLAH ASH SAKI¹, (Graduate Student Member, IEEE),
RASIT ONUR TOPALOGLU², (Senior Member, IEEE),
AND SWAROOP GHOSH¹, (Senior Member, IEEE)

¹Department of Electrical Engineering, The Pennsylvania State University, University Park, PA 16802, USA

²IBM, Hopewell Junction, NY 12533, USA

Corresponding author: Abdullah Ash Saki (axs1251@psu.edu)

This work was supported in part by the National Science Foundation (NSF) under Grant OIA-2040667, Grant DGE-2113839, Grant DGE-1723687, and Grant DGE-1821766 and in part by the Seed Grants from the Penn State Institute for Computational and Data Sciences (ICDS) and the Penn State Huck Institute of the Life Sciences.

ABSTRACT Trapped-ion (TI) quantum bits are a front-runner technology for quantum computing. TI systems with multiple interconnected traps can overcome the hardware connectivity issue inherent in superconducting qubits and can solve practical problems at scale. With a sufficient number of qubits on the horizon, the multi-programming model for Quantum Computers (QC) has been proposed where multiple users share the same QC for their computing. Multi-programming is enticing for quantum cloud providers as it can maximize device utilization, throughput, and profit for clouds. Users can also benefit from the short wait queue. However, shared access to quantum computers can create new security issues. This paper presents one such vulnerability in shared TI systems that require *shuttle* operations for communication among traps. Repeated shuttle operations increase quantum bit energy and degrade the reliability of computations (fidelity). We show adversarial program design approaches requiring numerous shuttles. We propose a random and systematic methodology for adversary program generation. Our analysis shows shuttle-exploiting attacks can substantially degrade the fidelities of victim programs by $\approx 2\times$ to $\approx 63\times$. Finally, we present several countermeasures such as adopting a hybrid initial mapping policy, padding victim programs with dummy qubits, and capping maximum shuttles.

INDEX TERMS Trapped-ion, qubit, quantum computing, shuttle, security, fidelity.

I. INTRODUCTION

Quantum computing has garnered immense attention from government, industry, and academia alike in recent years. It can be advantageous in domains like machine learning [1], drug discovery [2], molecule simulation [3], [4], and optimization [5]. With an active interest in the field, quantum computing is progressing at a rapid pace. On one end, researchers are proposing new quantum algorithms that leverage unique properties like superposition, entanglement, and interference to speed up computation. On the other end, researchers are pursuing various technologies like superconducting, trapped-ion (TI), and photonics to design quantum bits or qubits. Besides scientific breakthroughs, quantum computing is getting policy-level patronage from

governments as they understand the immense promise of the field. For example, the US government has established a *National Quantum Initiative Act* [6] and established the *National Quantum Coordination Office* [7] to spearhead the efforts. Quantum computing is poised to shine in the imminent future with all these concerted efforts.

The TI qubit is one of the most promising technologies for building a quantum computer. It offers several advantages such as identical qubits, long coherence times, and all-to-all connectivity among qubits [8]. Several companies such as IonQ and Honeywell are developing TI systems. Recently, Honeywell reported a trapped-ion system with quantum volume (QV) [9] of 1024 [10], highest thus far. Some of these TI qubit-based systems are also commercially available via IBM Qiskit [11], [12], Amazon Braket [13], and Microsoft Azure Quantum [14]. People are using TI devices to run their proof-of-concept applications that include a wide range

The associate editor coordinating the review of this manuscript and approving it for publication was Siddhartha Bhattacharyya¹.

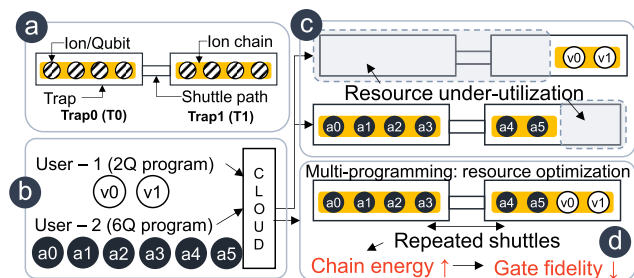


FIGURE 1. (a) Overview of a TI system. (b) The basic concept of multi-programming. (c) Issues with single programming. (d) Attack overview: the adversary program shares a trap with the victim. The adversary program tries forcing repeated shuttles between traps which increase chain-energy and degrade gate fidelity in the shared trap.

of practical topics like machine learning [15]–[17], generative adversarial networks [18]–[20], finance [21], and chemistry [22]–[26]. The user-base of TI systems not only includes academic institutions but also prominent industry names like Nippon Steel [27], BMW [28], JP Morgan Chase [21], and Goldman Sachs [29]. They are using the trapped-ion quantum computers to test their prototype applications so that when sufficiently large devices are available, they can immediately use the devices to solve the practical version of their problem to gain commercial advantages. Towards the goal of enabling real-world applications, the hardware makers of ion-trap computers are also making promising advances. The quality of qubits is improving [30], [31] and number of ions are increasing. Recently, IonQ announced a 64-ion trapped-ion device [32]. Similarly, Honeywell has a well-established roadmap for scaling their ion-trap hardware [33]. Furthermore, recent comparative studies [34], [35] between TI and other qubit technologies show promising results for TI systems. With efforts on both application and hardware fronts, the future of the TI system looks promising.

Confining many ions in a single trap becomes problematic from a control and gate implementation perspective. Therefore, the pathway to scalability in TI systems involves multiple interconnected traps. A technology named quantum charge-coupled device (QCCD) is proposed in [36] for scalable and modular trapped-ion systems. In this light, Murali et al. [37] performed extensive architectural studies for multi-trap trapped ion systems. They developed a compiler and a simulator [38] for such systems with experimentally calibrated values.

Figure 1a shows a TI quantum computer diagrammatically. In a TI system, qubits are realized using ions. Data is encoded as ion’s internal states. Ions are confined inside traps using direct current and oscillatory potentials (thus, we use the terms ions and qubits interchangeably for TI systems in this paper). Figure 1a shows a 2-trap system interconnected by a shuttle path that allows movement (shuttle) of ions between traps. Here, we are assuming that each trap can accommodate a maximum of 4 ions i.e., trap capacity = 4 per trap for illustration purpose only. Note, the work in [37] proposed a

trap capacity between 15-25 qubits for practical systems and we use this range in our analysis as well. Ions are first cooled and initialized. Then, laser pulses are applied in sequence on the ions to manipulate ions’ states to perform computation (quantum gates). Sometimes computation is required on data from ions in different traps. In such cases, one ion is shuttled (moved) from one trap to another so that the ions are co-located, and the gate can be performed. Finally, light is shined on ions and the data is measured as either 0 or 1 based on presence or absence of fluorescence.

Besides the computational aspects, security is an equally pivotal aspect for any computing paradigm. Several academic studies identifying security issues and fixes are surfacing in the quantum domain [39]–[44]. In [40], authors present an attack model where a rogue element in the cloud can report inaccurate device calibration data. Based on the incorrect data, a user may run his/her program on inferior qubits leading to poor results. They propose to include checkpoint circuits in the program to indicate unwanted variation in device calibration data. The attack model in [42] assumes an untrusted element in the cloud could schedule a program on inferior hardware instead of allocating the requested hardware. They propose a quantum physical unclonable function (QuPUF) to authenticate the requested hardware. Depending on the noise characteristics, each quantum computer demonstrates a unique signature. The authors in [42] leverage such signature to design the QuPUF.

A new type of attack vector can emerge in the multi-programming [45] setting (Figure 1d) for quantum computers. Suppose, two users are submitting their programs to the quantum cloud. User-1 program has 2 qubits, and user-2 program has 6 qubits. The cloud can schedule the programs individually on hardware. In such cases, device resources will be underutilized. For the 2-qubit program 6 qubits will be unused, and vice-versa (i.e., 2 unused qubits for the 6-qubit program). Multi-programming taps into this gap and proposes scheduling multiple programs together in the same hardware to maximize resource utilization (Figure 1d). The concept of multi-programming finds its application in commercial quantum clouds such as Rigetti’s Quantum Cloud Service (QCS) [46] where a user can reserve a lattice [47]. A lattice can be a partial set of qubits from a larger device. Thus, multiple users can run their programs on a different set of qubits from a larger device. Although running multiple programs can optimize resource usage, throughput, and profit for the cloud, it can create security issues. One such security vulnerability in a multi-programming environment is reported in [43] for superconducting qubits. The authors demonstrate a crosstalk-induced fault injection attack where crosstalk from the adversary program affects a victim program. However, this attack is not applicable for TI systems due to negligible crosstalk [30].

In this paper, we present an attack in the multi-programming setting for TI systems by exploiting a new vulnerability in terms of shuttle operations. Figure 1b-d provides an overview of the proposed attack model. We assume

qubits from the adversary program span over two traps, and they share a trap with qubits with the victim program. For example, adversary and victim qubits share Trap-1 (T1). The adversary can design his/her program such that it requires computation (gate) between ions from different traps that will need frequent shuttles between traps. Repeated shuttling adds energy to an ion and increases an ion-chain's energy. This elevated chain-energy degrades the reliability of computation (known as *gate fidelity*). As victim qubits share a chain (ion-chain in T1) with the adversary qubits, they also suffer from this shuttle-induced fidelity degradation. Although the premise seems simple, there are architectural policies that curbs shuttling and make the attack challenging. Thus, the attack culminates into designing a program that will trick the architectural policies and enforce repeated shuttles. The attack can be launched in a *white-box* setup where the attacker knows the policies and beats them to achieve repeated shuttles (Section V). Or, it can be a *black-box* type attack where no prior information is known (Section IV).

Although TI systems may not be used for secure real-world applications yet, it is paramount to *proactively* identify security flaws and mitigation measures. While *reactive* measures are an option in security research where vulnerabilities and solutions are researched post-adaptation of a technology, it may be expensive to fix the issues re-actively and may even be impossible. In this regard, proactive measures are beneficial as they unearth forthcoming issues, and our work is inspired by this mindset. With the rapid advancement of quantum hardware, quantum algorithm, and quantum architectures, quantum cloud services will become more practical and popular, and evidence [34], [35] suggests that TI technology will be at the forefront. Therefore, now is an opportune time to proactively identify vulnerabilities in and devise appropriate defenses for imminent multi-programming quantum clouds to prepare for practical scale deployments. This paper is one of the first efforts towards this goal.

We make the following contributions in this paper:

- Identify repeated shuttle operations as a mode of attack.
- Present two malicious program generation methodologies - systematic and random. Systematic attack uses prior knowledge about architectural policies to design a strong attack. The random attack does not require any prior knowledge, albeit losing some attack potency.
- Modify the QCCD-Compiler [37], [38] to accommodate multi-programming. We use QCCD-simulator and QCCD-compiler interchangeably in this paper since the compiler and the simulator are a part of the same software tool-chain.
- Analyze the impact of trap capacity and victim size¹ on shuttle number and fidelity reduction of the victim.
- Discuss three countermeasures to thwart attacks.

¹We use the following definitions for *program size* and *program length*: program size is the number of qubits in a program and program length is the number of 2-qubit gates in the program

The outline of the paper is as follows: Section II discusses the basics of quantum computing and TI systems. Section III describes the attack model and the simulation setup for analyses in this paper. Section V delineates the methodology of systematic malicious program generation. Section IV discusses the random attack program designing principles. Section VI reports the results and discussions. Section VII presents several countermeasures. Finally, Section VIII draws conclusion.

II. BASICS

In this section, we discuss the basics of trapped-ion quantum computers and terminologies used in the paper.

A. QUBIT AND QUANTUM GATE

1) QUBITS

Quantum bits or qubits are the building block of a quantum computer. Qubits store data (i.e., $|0\rangle$ and $|1\rangle$) as various internal states. A qubit can be in both $|0\rangle$ and $|1\rangle$ simultaneously due to quantum superposition property. A qubit state is represented as $|\psi\rangle = a|0\rangle + b|1\rangle$ where a and b are probabilities amplitudes. Measuring the qubit will collapse one state and return classical bits 0 or 1 with probabilities $|a|^2$ and $|b|^2$, respectively.

2) QUANTUM GATES

Quantum gates manipulate information stored in qubits to perform computation. Quantum gates are realized using pulses such as radio frequency (RF) and laser pulses. Gate pulses modify the probability amplitudes of a qubit. For example, a quantum NOT (X) gate pulse when applied to a qubit at state $|0\rangle$ will change amplitudes $a = 1$ and $b = 0$ to $a = 0$ and $b = 1$. Quantum gates are reversible in nature and represented by unitary matrices mathematically. At present, the physically realized gates are 1-qubit and 2-qubit. A quantum program is a sequence of quantum gates. Figure 3a shows a sample quantum program consisting of 2-qubit Mølmer-Sørensen (MS) gates [48].

3) GATE FIDELITY

Quantum gates in existing quantum computers are erroneous. They incur a finite and non-negligible error rate (ϵ) when executed. Suppose the aforementioned X gate is applied on state $|0\rangle$ for 10,000 times. One would, in theory, get output 1 all the time. However, due to gate errors, the user may end up with 9,900 1's (correct) and 100 0's (incorrect). The gate fidelity (F) is usually defined as the complement of the error rate i.e., $F = 1 - \epsilon$. It depends on variations of control pulses and environmental interference. Qubits are kept and operated in a controlled environment to shield them from various noises so that gate fidelities can be high. A lower gate fidelity will introduce more errors in the output and can completely decimate the result.

$$GPI2(\phi) = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & -ie^{-i\phi} \\ -ie^{i\phi} & 0 \end{bmatrix} \quad MS = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & 0 & -i \\ 0 & 1 & -i & 0 \\ 0 & -i & 1 & 0 \\ -i & 0 & 0 & 1 \end{bmatrix}$$

$$GPI(\phi) = \begin{bmatrix} 0 & e^{-i\phi} \\ e^{i\phi} & 0 \end{bmatrix}$$

$$GZ(\theta) = \begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}$$

Gate fidelity eq.

$$F = 1 - \Gamma\tau - A(2\bar{n} + 1)$$

$\Gamma = \text{Trap heating rate}$
 $\tau = \text{Gate time}$
 $\bar{n} = \text{vibrational energy}$

FIGURE 2. Matrices of gates used in IonQ trapped-ion quantum computer, and the gate fidelity equation presented in [37], [49].




FIGURE 3. (a) A sample quantum program consisting of 2-qubit MS gates. We primarily focus on the 2-qubit gates as it has a lower fidelity and requires shuttle. (b) Ions in same trap: gate can be directly executed. (c) Ions in separate traps: one ion needs shuttling before execution.

B. TRAPPED-ION QC

1) TRAP DETAILS

In a trapped-ion system, atoms like Yb or Ca are ionized and trapped between electrodes using electromagnetic fields. Hence, the name trapped-ion quantum computer. Data $|0\rangle$ and $|1\rangle$ are encoded as internal states such as hyper-fine or Zeeman states of the ions. Figure 1a shows the schematic of a trapped-ion system. It has 2 traps: Trap 0 or T0 and Trap 1 or T1. Inside the traps, ions form chains. The traps are connected by a shuttle path which allows movement (shuttle) of an ion from one trap to another if needed. Traps can accommodate a certain number of ions known as *trap capacity*. For example, traps in Figure 1a have a trap capacity of 4 ions per trap. Besides, some capacity is reserved for incoming ions from other traps known as *communication capacity*. The communication capacity is not shown explicitly in Figure 1a, however, an ion can be moved from T0 to T1, and in that case, T1 will hold 5 ions. Thus, the *trap capacity + communication capacity* defines the absolute maximum number of ions a trap can hold. *communication capacity* is much smaller than *trap capacity* in general [37]. Finally, the excess cap (EC) of a trap is defined as *trap capacity + communication capacity - ions in the trap*.

2) GATE DETAILS

Laser pulses are used to perform quantum gate operations on the qubits/ions. Mølmer-Sørensen (MS) gate is the typical native 2-qubit gate of trapped-ion systems [48]. It is accompanied by several 1-qubit gates which are mainly rotation gates to form a universal gate-set. For example, the 1-qubit gates are GPI, GPI2, and GZ in the IonQ system [50]. The matrices for these gates are shown in Figure 2. Inside a trap, all the

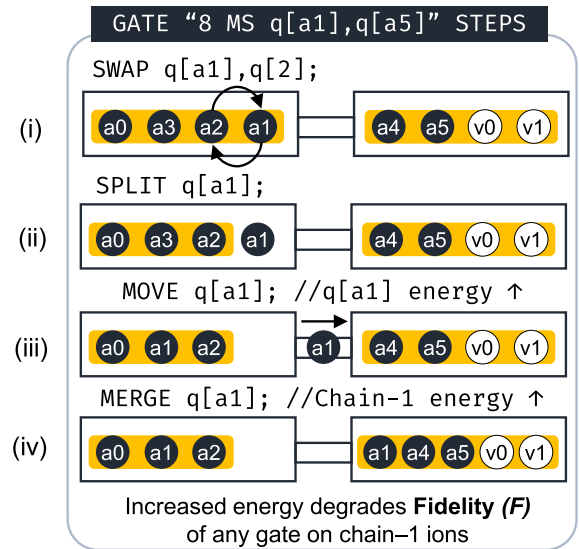


FIGURE 4. Shuttle steps to bring ions a1 and a5 in the same trap.

qubits are connected meaning a 2-qubit gate can be performed between any two qubits in that trap.

On one hand, single-qubit gates have a higher fidelity (error rate, ϵ in range of 10^{-3} to 10^{-4}) [8], [30]. Besides, they can be performed *in-place*. On the other hand, 2-qubit gates typically have an order of magnitude lower fidelity than 1 qubit gates ($\epsilon \sim 10^{-2}$) [8], [30]. Figure 2 shows the gate fidelity equation [37] for TI systems. This is an experimentally validated gate fidelity model. Here, Γ is the trap heating rate, τ is the gate time, and \bar{n} is the vibrational energy or motional mode of a chain. A is a scaling factor that depends on the number of ions in the chain as $N/\ln(N)$. The gate fidelity will degrade if gate time and/or motional mode of the chain increases. A 2-qubit gate cannot be applied to ions from different traps. It requires a shuttle.

3) NEED OF SHUTTLE OPERATION

Consider the sample program in Figure 3a. The 4th gate in the program MS q[a0], q[a1] involves ions from the same trap (Trap 0) and can be executed in-place or directly (Figure 3b). However, the 8th gate MS q[a1], q[a5] involves ions from different traps (Figure 3c). Therefore, a shuttle is needed to bring the ions in the same trap.

4) SHUTTLE STEPS

The shuttle operation involves several steps as depicted in Figure 4. First, a1 and a2 are swapped so that a1 is transferred near the shuttle path. Then, a1 is split from the chain-0 and shuttled/moved from T0 to T1. The shuttle operation adds energy to the ion. Then, a1 is merged to the chain-1. This merge operation increases the vibrational energy (\bar{n}) of chain-1. Finally, MS q[a1], q[a5] can be executed as the ions are in the same trap (T1) now. As chain-1's \bar{n} is now higher, the subsequent gate operations in this chain (either on aX ions or vX ions) will experience lower fidelity (F).

Increasing a chain's motional mode by repeated shuttles is the basis of the attack proposed in this paper.

C. INITIAL MAPPING POLICY

The initial mapping - in this context - entails the assignment of program qubits in traps (i.e., mapping of logical qubits to physical traps) and the relative position of qubits inside a trap. For example, the program qubits (a0 to a5) from the sample program in Figure 3a are initially mapped as $T_0: [a_0, a_3, a_1, a_2]$, $T_1: [a_4, a_5]$ (explained in Example 1). The program execution will start with this allocation, and the mapping will be updated based on shuttles.

The initial mapping policy in [37] is the *greedy* policy where qubits are allocated in the descending order of *edge weights* (wt). A quantum program can be treated as a graph where each node represents a qubit and an edge between two qubits represents a 2-qubit gate. Thus, the edge weight represents the frequency of a 2-qubit gate between a pair of qubits. In the Greedy policy, the qubits of the most frequent gates are allocated first.

Example 1: The mapping policy and the result can be explained with the sample program in Figure 3. The edge weights of the program are as follows: $wt(a_0, a_3) = 3$ (as the $MS\ q[a_0], q[a_3]$ gate appears 3 times throughout the program), $wt(a_1, a_2) = 2$, and $wt(a_0, a_1) = wt(a_4, a_5) = wt(a_1, a_5) = 1$. Therefore, ions a0 and a3 are allocated first, then a1 and a2, and finally, a4 and a5. In this example, we assume a trap capacity of 4. Thus, ions a4 and a5 are in T1.

D. SHUTTLE DIRECTION POLICY

Shuttle direction policy dictates which ion will be moved to execute a 2-qubit gate. The shuttle direction policy used in the QCCD compiler [37] is illustrated in Listing 1. In this paper, we follow the same shuttle direction policy. In the example from Figure 4, both traps have an equal number of ions. Hence, they have the same excess capacity (=trap capacity - # ions in the trap). Thus, the first ion a1 in the $MS\ q[a_1], q[a_5]$; is shuttled from T0 to T1.

```

1 if excess_cap0 < excess_cap1:
2     Move Trap0 --> Trap1
3 elif excess_cap0 == excess_cap1:
4     Move 1st ion of the gate
5 else:
6     Move Trap1 --> Trap0

```

Listing 1. Shuttle direction policy [37], [38].

E. MULTI-PROGRAMMING

The proposed attack model exploits the multi-programming setup. We modify the initial mapping policy [37], [38] to allow for multi-programming as follows: Suppose, we have two traps (T0 and T1) and two programs (prog-0 and prog-1). We allocate prog-0 from one end of the T0 to T1 (if needed) and prog-1 from the opposite end of the T1 to T0. Qubits in a single program are allocated per the greedy mapping

policy [37]. Allocating multiple programs from opposite ends and different traps ensures that qubits of one program are not mixed with the other program.

III. ATTACK MODEL AND SIMULATION SETUP

A. ATTACK MODEL

The following assumptions are made in the attack model:

- More than one program are running on the QC.
- The adversary program spans more than one trap, and it shares one trap with another (victim) program.
- The adversary knows device specification such as trap capacities and communication capacities. This information (especially, trap capacities) are usually public information. For example, the Honeywell H0 TI system has a capacity of 6 [51].
- Adversary knows architectural policies: initial mapping policy and shuttle direction policy (can be relaxed. Rationale behind the assumption is discussed in Section V-E3).
- The adversary can access the *compiled* program. This is a reasonable assumption because present quantum clouds provide such access. It allows a user to identify bottlenecks and optimize their programs. For example, in AWS Braket the compiled program is available to the user as the `MetaData` [52].

The adversary designs his/her program so that it requires repeated shuttles increasing the vibrational energy of the shared ion-chain and degrading gate fidelities. Note that, the adversary program fidelity takes a hit as a byproduct. However, the objective of the attack is to affect the victim program.

We present two techniques of devising the attack programs: systematic (Section V) and random (Section IV). On one hand, systematic program generation requires several prior information such as initial mapping and shuttle direction policies. However, it guarantees a linear increase in the number of shuttles (desired) with increased program length. On the other hand, random generation does not guarantee a linear increase of shuttles but requires no prior knowledge. Nevertheless, both methods can degrade victim performance (Section VI).

B. SIMULATION SETUP

In this paper, we use the QCCD compiler-simulator [38] accompanying the paper [37] to perform simulations. The QCCD compiler takes care of the initial mapping, shuttle insertion, gate scheduling, and fidelity computation of a program. We add our modification on top of the QCCD compiler to allow for multi-programming. Our tweaks include: (i) modifying initial mapping to map multiple programs (as in Section II-E) to the device and (ii) reporting individual program fidelities.

For all simulations, we assume a device with 2 traps connected in a linear fashion as in Figure 1a. The trap capacity is 15 ions per trap with 2 additional spaces per trap for incoming (shuttled) ions (communication capacity). We also

show analysis for trap capacities 20 and 25 (with the same communication capacity of 2). The trap capacity values of 15 to 25 are selected as per [37] as they observed better performance in this range.

We set the adversary program size to 18 qubits (trap capacity $15 + 3$) so that it spans two traps. We vary victim program sizes from 2 to maximum remaining space in shared trap T1 i.e., 12 qubits. Note that the adversary and the victim program sizes will change accordingly for trap capacities 20 (23 and 2 to 17 respectively) and 25 (28 and 2 to 22 respectively).

2-qubit gates mostly affect the program fidelity as they have an order of magnitude lower fidelity than 1-qubit gates. Therefore, we consider only the 2-qubit gate fidelities without any loss of generality in our analysis. The QCCD-simulator [37], [38] includes experimentally calibrated parameters [30], [49], [53], [54] for the gate fidelity equation in Figure 2. The program fidelity is computed from individual gate fidelities (F). For a program with k 2-qubit gates (g_i), i being an enumeration parameter across the 2-qubit gates, the program fidelity is $F(g_0) \times F(g_1) \times \dots \times F(g_{k-1})$.

IV. RANDOMIZED MALICIOUS PROGRAM

In this section, we discuss randomized malicious program generation. The randomized attack programs are advantageous as they treat the compiler as a black-box and do not need information about compiler policies. We present a methodology to find and refine an effective random program to launch attacks. The only requirement is that the adversary can submit many programs to the cloud and can access the final compiled program. The final *shuttle-inserted* compiled-program will tell the adversary which program resulted in maximum shuttles. The adversary will pick that program to launch future attacks.

A. GENERAL METHODOLOGIES

We populate the program with $\binom{18}{2}$ gates (all 2 qubit combinations from 18-qubits adversary size with trap cap 15). Then, we randomly shuffle the gate orders to generate the random program.

We generate 1000 random circuits and compile them with pseudo-victim programs of sizes from 2 to 12 qubits. The idea of a pseudo-victim program is that the adversary will send two programs to the cloud to mimic an adversary-victim pair. After analyzing the collected results, the adversary can select the random circuit that gives the highest average number of shuttles across victim programs of all sizes.

B. PRUNING THE RANDOM CIRCUIT

The best random circuit can be pruned further as not all the gates contribute to shuttling. The intuition is that we can remove some gates from the random circuit without lowering the number of shuttles. The pruning logic is as follows: we remove one gate from the original random circuit starting from the first gate, compile it, and check the number of shut-

tles. If the number of shuttles does not drop from the original case, we permanently remove the gate from the circuit and move on to the next gate. If removing the gate lowers the number of shuttles, we reinstate the gate and move on to checking the next gate. Following this step-by-step check, we can remove some redundant gates without affecting the number of moves. The pruning on average removed ≈ 48 gates from the program of 153 gates.

V. SYSTEMATIC MALICIOUS PROGRAM

A. BASIC IDEA

The systematic method of malicious program generation uses the following 3 ingredients to craft a strong attack program: (i) initial mapping policy, (ii) shuttle direction policy, and (iii) information on the victim size. As mentioned earlier, a gate will require a shuttle when the ions belong to two different traps. This principle is leveraged in the systematic method, and gates are added in the malicious program with ions from different traps. However, this approach requires knowledge about ion locations, and the above 3 ingredients facilitate the tracking of ion locations.

Algorithm 1: Create Initial Mapping Controller Block

Input: trap capacity

Output: initial mapping controller

```

1 ion_list = [0 to (trap capacity - 1)];
2 ion_a, ion_b ← 2 arbitrary ions from ion_list;
3 add gate (ion_a & ion_b) twice in the initial mapping
  controller block;
4 remove ion_a & ion_b from ion_list;
5 while ion_list is not empty do
6   ion_a ← ion_b from last gate;
7   ion_b ← next ion from ion_list;
8   add gate (ion_a & ion_b) twice in the initial
  mapping controller block;
9   remove ion_b from ion_list;
10 end
```

Our proposed systematic malicious program consists of three blocks: (i) shuttle controller (SC), (ii) a bridging gate, and (iii) initial mapping controller (IMC). Each block is generated using specific logic as explained later in this section. After all blocks are generated, they are stitched to create the complete malicious program (i.e., malicious program = shuttle controller + a bridging gate + initial mapping controller).

The IMC block is generated first, then the SC block, and finally the bridging gate, although they appear in a different order in the program. This ensures no gate from the SC block and the bridging gate have a higher edge weight than gates from the IMC block (explained more in Section V-E1).

B. INITIAL MAPPING CONTROLLER

With knowledge about the initial mapping policy, the adversary can intelligently add gates in the program to force a known initial mapping. As described in Section II-C, the

#	GATE	#	GATE
imc1	MS $q[0], q[1]$;
imc2	MS $q[0], q[1]$;	imc23	MS $q[11], q[12]$;
imc3	MS $q[1], q[2]$;	imc24	MS $q[11], q[12]$;
imc4	MS $q[1], q[2]$;	imc25	MS $q[12], q[13]$;
imc5	MS $q[2], q[3]$;	imc26	MS $q[12], q[13]$;
imc6	MS $q[2], q[3]$;	imc27	MS $q[13], q[14]$;
...	...	imc28	MS $q[13], q[14]$;

FIGURE 5. IMC block of the malicious program. Arrows showing dependency between gates.

initial mapping policy in the QCCD-compiler is a greedy one that allocates ions based on gate frequencies (edge weights). Therefore, the adversary can judiciously increase edge weights between certain nodes (ions/qubits) which he/she wants to be allocated first. Algorithm 1 illustrates the gate selection methodology. We explain the algorithm with Example 2.

Example 2: Suppose, the trap capacity is 15. Thus, the ion_list will be $[0, 1, 2, \dots, 13, 14]$. Next, we arbitrarily select two ions 0 (ion_a) and 1 (ion_b) from the list. Note that any two ions can be selected. We add gate MS $q[0], q[1]$ twice in the program with these ions. This will make edge weight of $(0, 1) = 2$. Next, ions 0 and 1 are removed from the ion_list which now becomes $[2, 3, 4, \dots, 13, 14]$.

Next, ion_a and ion_b values are updated. ion_a 's value becomes the previous ion_b value (i.e., 1), and ion_b 's value becomes the next value from the ion_list (i.e., $ion_b = 2$). Then, the gate with these two ions - MS $q[1], q[2]$ - are added twice in the block. Finally, $ion_b = 2$ is removed from the ion_list for this iteration, and ion_list becomes $[3, 4, 5, \dots, 13, 14]$. The above routine is repeated unless the ion_list becomes empty i.e., all the ions are added in the block.

The final IMC block will be similar to Figure 5. All the gates in the block have an edge weight of 2. The logic in the other blocks (SC block and bridging gate) ensures that no other edge weight exceeds 1 (i.e., all other gates will appear once). Due to the higher edge weights, gates in the IMC block will be allocated first according to the greedy policy. Thus, ions 0 to 14 (15 qubits) will be allocated first to T0 (the remaining 3 qubits of the 18 qubit adversary will be allocated to T1 by default).

Assuming a victim size of 12, the trap states after initial mapping will be $\{T0 (EC = 2): [0, 1, 2, \dots, 13, 14], T1 (EC = 2): [15, 16, 17] + [12Q\ victim]\}$. Here, EC = excess capacity. T0 has 15 ions from adversary program. T1 also has 15 ions, 3 from the adversary and 12 from the victim. Thus, each trap has an excess capacity of 2 (from communication capacity).

C. SHUTTLE CONTROLLER

After the IMC block ensures a known initial mapping, we use Algorithm 2 to add gates in the malicious program that

require shuttles. The flow consists of 5 steps. We explain each step in Example 3. The core idea is to track ion locations (using shuttle direction policy) after each gate and select ions from different traps for the next gate.

Algorithm 2: Create Shuttle Controller Block

Input: trap states, shuttle direction policy, node weights, edge weights, block length, prog. size

Output: shuttle controller block (sc_block)

```

1  $sc\_block \leftarrow$  empty; # of added gates  $\leftarrow$  0; flag  $\leftarrow$  0;
2 while # of added gates < block length do
3   // STEP - 1;
4   if # of added gates == 0 then
5     |  $ion\_a \leftarrow$  random ion  $\in \{0$  to prog. size-1};
6   else if flag == 0 then
7     |  $ion\_a \leftarrow$  moved ion from last gate;
8   else
9     |  $ion\_a \leftarrow$  non-moved ion from last gate;
10  end
11  // STEP - 2;
12  get  $ion\_a$ 's location; get opposite_trap;
13  // STEP - 3;
14  for  $ion\_b \in \{ions\ in\ the\ opposite\ trap\}$  do
15    | if edge weight ( $ion\_a, ion\_b$ ) == 0 then
16      | flag  $\leftarrow$  0; break;
17    | else
18      | flag  $\leftarrow$  1; continue;
19    | end
20  end
21  if flag == 1 then
22    | continue;
23  end
24  // STEP - 4;
25   $sc\_block \leftarrow sc\_block +$  gate ( $q[ion\_a], q[ion\_b]$ );
26  last_gate  $\leftarrow$  ( $ion\_a, ion\_b$ );
27  // STEP - 5;
28  identify moved ion, update node and edge weights,
   update trap states using shuttle policy;
29  # of added gates  $\leftarrow$  # of added gates+1
30 end

```

Example 3: Step-1: For the very first gate, we randomly select 1 ion from adversary's ion list (i.e., from $[0$ to $17]$ for the 18-qubit adversary). Suppose, the selected ion is 14 (i.e., $ion_a = 14$).

Step-2: From the initial mapping, we know ion 14' is in T0. Therefore, the opposite trap is T1..

Step-3: Select ion 15 from T1. Check if gate MS $q[14], q[15]$ does not exist in the program (i.e., edge weight of $(14, 15)$ is 0 in the program graph). As the gate does not exist in the program, we do not check more ions from T1 and can break from the loop with $ion_b = 15$.

Step-4: Add gate MS $q[14], q[15]$ in the program. This gate will require a shuttle when executed. Therefore, trap states need updating after the gate.

(a) Shuttler controller block

#	GATE	#	GATE
sc1	MS q[14],q[15];
sc2	MS q[0],q[14];	sc77	MS q[4],q[17];
sc3	MS q[14],q[16];	sc78	MS q[0],q[4];
sc4	MS q[1],q[14];	sc79	MS q[0],q[15];
...	...	sc80	MS q[0],q[10];

(b) Bridging gate

#	GATE
bg1	MS q[0],q[10];
imc1	MS q[0],q[1];
...	...

Annotations:
 - Blue arrow from sc80 to bg1: One ion from SC block's last gate
 - Red arrow from imc1 to bg1: One ion from IMC block's first gate

FIGURE 6. (a) The partial SC block, showing first and last few gates. (b) The bridging gate between the shuttle controller and the initial mapping controller block. Only one bridging gate is necessary between these blocks.

Step-5: Both T0 and T1 have equal excess capacity of 2 each. Therefore, according to the shuttle direction policy in Listing 1, the first ion in the gate MS q[14], q[15] i.e., ion 14 will move from T0 to T1. Thus, the updated trap states are {T0 (EC = 3): [0, 1, 2, ...,13], T1 (EC = 1): [14, 15, 16, 17] + ([12Q victim])}. (Note, the victim size information is required to compute the excess capacity of the shared trap T1 and to find the shuttle direction accurately.) Edge weights list is updated with the new gate.

Following the same routine, we keep adding gates in the malicious program until target number of gates are reached. For the next iteration, we pick ion 14 (moved ion from the last gate) as the *ion_a* in Step-1. As ion 14 is in T1 now, we pick the other ion from T0 (say, ion 2). The next gate is MS q[2], q[14]. As T0 has more EC (=3) than T1 (EC = 1), ion 14 will again move but this time from T1 to T0. Finally, the updated trap states after this gate will be {T0 (EC = 2): [0, 1, 2, ...,13, 14], T1 (EC = 2): [15, 16, 17] + [12Q victim]}. A partial shuttle controller block of block length 80 is illustrated in Figure 6a. Note the last gate in the block. It is required for the bridging gate.

D. THE BRIDGING GATE AND THE COMPLETE MALICIOUS PROGRAM

1) THE BRIDGING GATE

The bridging gate is formed by taking one ion from the *last gate* of the shuttle controller block and one ion from the *first gate* of the initial mapping controller block (Figure 6b). As the name suggests, this gate bridges two blocks and maintains the gate dependency. Only one bridging gate is necessary.

2) THE COMPLETE MALICIOUS PROGRAM

The complete malicious program is created by combining individual parts in the following order: gates from the SC block + bridging gate + gates from the IMC block.




FIGURE 7. Variation of shuttle numbers across different (actual) victim sizes. The adversary program assumes victim size = 2.

E. DISCUSSIONS ON THE SYSTEMATIC METHOD

1) OUT-OF-ORDER IMC BLOCK

In greedy policy, ions from a gate is primarily mapped based on their frequency of appearance and secondarily by their order of appearance in the program. Gates in the IMC block have the highest frequencies across the complete program because they each are deliberately added twice in the program, and gates from the other blocks each are added once. Generating the IMC block first ensures that other blocks can skip gates already in the IMC block. The compiler sorts gates in the descending order of their frequencies for mapping. Therefore, gates in the IMC block come first in the sorted list for mapping although they appear last in the malicious program.

2) NECESSITY OF VICTIM SIZE INFORMATION

The victim size is a parameter of the SC block generation algorithm. If the victim size is correct, each gate in the generated SC block will require one shuttle when the malicious program is executed. However, with an inaccurate victim size, some of the gates in the SC block will not force a shuttle.

Example 4: Consider the SC block in Figure 6a generated assuming a victim size of 12. Suppose, the actual victim size during run is 5. Then, the actual trap states and excess capacities will be as follows at the beginning: T0 (EC = 2): [0, 1, ...,14], T1 (EC = 7): [15, 16, 17] + [5Q victim]. Gate # sc1 will need a shuttle as ion 14 is in T0 and ion 15 is in T1. Updated trap states will be T0 (EC = 3): [0, 1, ...,13], T1 (EC = 6): [14, 15, 16, 17] + [5Q victim] (Note T1 EC > T0 EC, different than the assumption). Gate # sc2 will also require a shuttle as the ions are in different traps. New trap states: T0 (EC = 4): [1, ...,13], T1 (EC = 5): [0, 14, 15, 16, 17] + [5Q victim]. However, the next gate # sc3 will not require a shuttle as both ions 14 and 16 are in T1.

The above example illustrates that some gates in the SC block will skip shuttling when the actual victim size is different from the design assumption. As the adversary may not know the victim size, he/she needs to assume a value that gives the best number of shuttles across all possible victim sizes. To find the best assumption, we sweep the victim size




FIGURE 8. Effect of victim size assumption (trap capacity = 15). Assuming victim size 8 provides best result (highest ICV).

and propose selecting the value that gives the highest inverse coefficient of variation (ICV). The inverse of coefficient of variation is the ratio of the mean (μ) to standard deviation (σ). A higher ICV indicates a higher mean and/or lower standard deviation. For our purpose, we want a distribution that provides a higher mean number of moves with a tighter spread (a lower σ). Before computing ICV, we normalize both μ and σ with respect to the respective maximum value. The results are discussed in Section VI.

3) AVAILABILITY OF ARCHITECTURAL POLICIES

The systematic method relies on knowledge initial mapping and shuttle direction policies to ensure an effective attack. Details on such architectural policies are available in literature [45], [55]–[57] and/or compiler documentations [58], [59]. Architectural policies like speculative execution are available to the public even in the classical domain. Numerous papers on architectural policy flaws and their mitigation also exist [60]–[63]. *Spectre* [62] and *Meltdown* [63] are two such famous vulnerabilities that exploit the knowledge of an architectural policy. Thus, we envision such availability of architectural policies in the quantum domain as well. Finally, if the cloud provider does a good job in protecting their secrets i.e., restrict access to policies, the adversary can follow two options. *Option-1*: the adversary can adopt a trial-and-error route and design several attack programs each considering separate policies available in the literature. Then, he/she can launch an attack to check which one gives the best results (high number of shuttles). *Option-2*: the adversary can revert to the random attack described in Section IV which does not require any prior information.

VI. RESULTS AND DISCUSSIONS

A. VICTIM SIZE SWEEP

1) SYSTEMATIC ATTACK

Figure 7 shows the number of shuttles for various actual victim sizes (2 to 12, trap cap 15) where the attack program is designed assuming a 2-qubit victim. The plot shows the highest number of moves (80) is achieved when adversary assumption matches actual victim size (dark bar in the plot). We also observe that for other actual victim sizes shuttle numbers vary and typically drop from the highest value.

TABLE 1. Statistics from victim size-assumption sweep for trap caps 20 and 25.

Victim size asmp.	Trap cap = 20			Trap cap = 25		
	AVG	STD	ICV	AVG	STD	ICV
2	50.12	14.96	0.77	60.29	12.92	1.21
3	62.19	11.18	1.28	51.05	17.58	0.75
4	51.06	15.88	0.74	58.62	10.75	1.41
5	66.25	9.97	1.53	50.48	16.92	0.77
6	57.19	15.12	0.87	59.24	14.38	1.07
7	64.94	9.04	1.66	51.38	16.6	0.8
8	58	12.34	1.08	62.38	11.36	1.42
9	68.88	7.41	2.14	54.57	15	0.94
10	61.69	9.79	1.45	61.9	10.18	1.58
11	64.38	9.1	1.63	57.14	13.59	1.09
12	62	8.76	1.63	66.57	8.27	2.09
13	67.25	6.75	2.3	59.86	10.9	1.42
14	67.12	7.62	2.03	64.76	7.78	2.16
15	67.12	7.62	2.03	59.38	10.73	1.43
16	65.62	7.97	1.9	66.33	6.83	2.52
17	65.62	7.97	1.9	62.76	8.29	1.96
18	-	-	-	65.62	6.95	2.45
19	-	-	-	67.86	7.19	2.45
20	-	-	-	67.86	7.19	2.45
21	-	-	-	65.05	7.76	2.17
22	-	-	-	65.05	7.76	2.17

We report the mean (μ), standard deviation (σ), and inverse of coefficient of variation (ICV) of this distribution.

To find the best assumption, we sweep the victim size assumption from 3 to 12 (in addition to the assumed size of 2 in Figure 7) and record the statistics. The results are plotted in Figure 8. In all cases, we observe that the highest number of moves is achieved when the assumption matches the actual size. From the statistics, an assumed victim size of 8 qubits gives the best performance (highest ICV).

We perform the same analysis for trap capacities (cap) 20 and 25. For trap cap of 20, possible victim sizes vary from 2 to 17, and for trap cap 25 it varies from 2 to 22. We report the mean, standard deviation, and ICV for both trap caps in Table 1. For trap cap 20, an assumption of 13 qubits in the victim gives the highest ICV, and for trap cap 25 the best assumption is 16 qubits.

TABLE 2. Shuttle statistics for the random attack.

Victim size asmp.	Trap cap = 15 (random attack)			Victim size asmp.	Trap cap = 15 (random attack)		
	AVG	STD	ICV		AVG	STD	ICV
2	42	10.13	1.17	8	42.45	11.84	1.01
3	42.91	9.57	1.26	9	45.55	9.73	1.32
4	35.09	13.04	0.76	10	42.18	11.88	1
5	44.82	8.55	1.47	11	43.36	11.37	1.07
6	44.64	7.81	1.61	12	41.27	10.87	1.07
7	46.36	12.36	1.05				




FIGURE 9. Fidelity reduction for various victim lengths. The adversary program assumes victim size of 8. The trend exhibits a positive correlation between reduction and length.

2) RANDOM ATTACK

Table 2 shows the statistics for victim size sweep for random attack programs. The assumption of victim size 6 provides the best results in terms of ICV. Therefore, we select this program to launch future attacks.

B. FIDELITY REDUCTION

1) SYSTEMATIC ATTACK

Figure 9 shows fidelity reduction of the victim program under attack for 3 victim lengths - 60, 80, and 100. The SC block length, trap cap, and assumed victim size are 80 gates, 15 ions, and 8 qubits respectively. The plot shows that the fidelity reduction increases with higher victim lengths. Average fidelity reductions are 2.63x, 5.40x, and 13.11x for victim lengths 60, 80, and 100 respectively. Figure 10 qualitatively explains the positive correlation between victim length and fidelity reduction. The shuttle operations are spread across the length of the SC block. If the victim program completes before the SC block, shuttles at the later part of the block (hatched pattern in Figure 10) does not affect the victim program. The opposite happens if the victim program is longer. Repeated shuttles from adversary programs increase the chain energy. Therefore, the later gates (cross pattern in Figure 10) in the victim program experience excessive fidelity drops.

Table. 3 shows fidelity reductions for trap caps 20 and 25 for 3 victim program lengths - 60, 80, and 100. For trap cap 20, the average fidelity reduction for these 3 program lengths are 2.71x, 5.66x, and 13.51x. For trap cap 25, the average fidelity reductions are 4.77x, 15.65x, and 68.13x. In case of trap cap 25, we observe aggravated fidelities, especially for

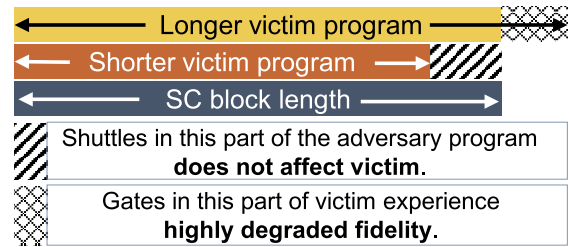


FIGURE 10. Explanation of the positive correlation between fidelity reductions and victim lengths.

TABLE 3. Fidelity reduction values for trap capacities 20 and 25.

Trap cap	60 gates	80 gates	100 gates
20	2.71X	5.66X	13.51X
25	4.77X	15.65X	68.13X




FIGURE 11. Fidelity reduction for random attack program.

larger victim sizes. Intuitively, this behavior can be attributed to the scaling factor A in the gate fidelity equation $F = 1 - \Gamma\tau - A(2\bar{n} + 1)$. Factor $A \propto N/\log(N)$ where $N =$ number of qubits in the chain. For larger victims in a larger capacity trap, A scales up making the motional mode (\bar{n}) more pronounced, and exacerbating gate fidelity (F).

2) RANDOM ATTACK

Figure 11 shows the fidelity reduction values for random attack program (trap cap = 15). We omit the values for trap cap 20 and 25 for brevity (fidelity reductions will be even higher at these capacities). The results show an average fidelity reduction of 2.22x, 4.0x, and 8.94x. These values are lower than the systematic attack.

C. CHOICE BETWEEN SYSTEMATIC AND RANDOM ATTACK PROGRAMS

The choice between systematic and random attack is not an either-or proposition although the systematic approach provides a higher fidelity reduction. The adversary needs to submit many programs to the cloud to find a good attack program using the random approach. Using the systematic approach, a good attack program can be generated in one try. Thus, the choice between approaches will depend on the resources available to the adversary. If he/she has information about the architectural policies, adopting the systematic

approach is the fastest and the best choice. If an adversary has the resources to run many programs on the cloud (running programs will cost money) and/or does not have knowledge about necessary policies, then adopting the random approach will lead to stronger attacks.

VII. COUNTERMEASURES

A. RANDOM INITIAL MAPPING

The compiler can adopt a random initial mapping policy where each program at each iteration starts from a random allocation. Both malicious program creation methods rely on a consistent initial mapping i.e., the same program will be allocated in the same fashion for every run. In case of a random attack, if initial mapping changes randomly from one instance to another, the attack program generated at one instance (with one mapping) will not work effectively for another instance (with a different) mapping. A random mapping will invalidate the systematic method as traps will start from unknown states.

We validate this proposed technique with an 18-qubit adversary program designed for a 12-qubit victim. With the greedy policy, it forces 80 shuttles. Next, we switch the initial mapping policy to random and gather results for 1000 runs. We find that with a random initial mapping policy, the average number of moves drastically drops to ≈ 27 ($\sigma = 9.45$) ($2.96\times$ drop). Therefore, it proves the efficacy of random initial mapping in weakening the attack.

1) TRADE-OFF OF RANDOM INITIAL MAPPING

Although a random initial mapping policy can disarm an adversary, it may penalize a legitimate user. A good initial mapping policy tries to place ions with frequent gates together so that communication can be minimized in typical benchmarks (note that attack programs are not typical programs— they are artificially crafted to contain numerous shuttles by hacking an intelligent initial mapping policy). However, a random initial mapping policy does not exploit such intelligence and cannot always guarantee an optimal number of communications. To show the impact of random mapping, we simulate a suite of popular noisy intermediate-scale quantum (NISQ) benchmarks used in [37] with both greedy and random mapping policies. Benchmarks include quantum Fourier transform (QFT), quantum approximation optimization algorithm (QAOA) circuit, supremacy circuit from Google’s quantum supremacy experiment, and quantum adder circuit. The circuits are generated using [64]. The mean values for random policy are computed from 1000 random allocations. The results (Table 4) show that random mapping increases shuttles for NISQ benchmarks up to $6\times$. Thus, a random mapping policy can weaken attacks at a cost of penalizing legitimate users.

2) HYBRID APPROACH

To alleviate the issue, the cloud can adopt a hybrid approach. In the hybrid approach, the compiler compiles a program with

TABLE 4. Penalty of random mapping in NISQ benchmarks.

Benchmark name	# Moves (Greedy)	Avg # moves (Random)	Moves \uparrow
Supremacy_25Q (depth = 10)	56	65	1.16X
QAOA_20Q (depth = 10)	14	47	3.35X
QAOA_30Q (depth = 10)	19	114	6X
QFT_20Q	15	48	3.20X
QFT_30Q	30	162	5.4X
Adder_22Q	6	22	3.67X

both random and intelligent (e.g., greedy) initial mapping and discards the result with higher shuttles. It will execute the version with the lower shuttles.

B. DUMMY PAD QUBITS IN THE VICTIM PROGRAM

A user (victim) can protect his/her program by adding a sufficient number of dummy qubits to pad the unused qubits in a trap. Suppose, the actual user program needs 10 qubits, and he/she wants to execute the program on a system with a trap cap 15/trap. In such a case, the user can add 5 dummy pad qubits in his/her program to make the program size 15 which will fully occupy a trap. Thus, an adversary qubit cannot share a trap with the victim preventing shuttle-induced fidelity degradation. The user can apply virtual-Z gates [65] (e.g., GZ gate in the IonQ machine [50]) on the dummy qubits. It will ensure that the compiler considers the qubits during allocation. As the virtual-Z gate is a software gate, it has perfect fidelity, requires no physical, and will not affect the user program.

However, there can be a security vs. cost trade-off. The quantum cloud may charge the user more for using more qubits. Consider a linear cost model where requesting 1 qubit cost 1 unit. The cost of running the 10-qubit will be 10 units. However, with 5 dummy qubits for padding and security, now the user has to spend 15 units increasing the cost by $1.5\times$. Not to mention that the cost model could be based on an exponential relation with the qubits counts. Therefore, this defense will be more cost-effective for a low number of dummy qubits (i.e., where actual user program size is large and/or trap cap is low).

C. CAPPING MAXIMUM NUMBER OF ALLOWED SHUTTLES

The cloud can enforce a *max shuttle* to prevent shuttle-exploiting attacks. The cloud provider can check the required number of shuttles in a program, and if it exceeds the set maximum value, the cloud can schedule it separately (without any accompanying program). This means for certain programs the cloud will dynamically switch to a single-programming mode from the multi-programming mode. However, the cloud will lose some throughput due to this switching. It can cover the loss by charging extra for programs requiring a high number of shuttles. Suppose, the 18 qubit adversary program

exceeds the max set shuttles. The cloud needs to run it in a single-programming mode with 30 qubit resources (2 traps \times 15 qubits/trap). Thus, the cloud will now charge 30 units - considering the previous linear cost model - instead of 18 units before to cover the loss in device utilization. In this way, programs with high shuttles cannot affect other programs, and the cloud will not incur a loss.

VIII. CONCLUSION

In this paper, we present a vulnerability in multi-programming access to TI quantum computers and propose several defenses. With extensive numerical simulations, we establish the shuttle operation as a mode of fault injection attack. We present two methodologies, namely, systematic and random, to design malicious programs that will cause repeated shuttles between traps. The purpose is to increase the energy of an ion chain, which degrades the gate and program fidelity. We perform analysis for a wide range of design parameters such as trap capacity (15 to 25 ions per trap), program length (60 to 100 gates), and victim size (2 to 22 qubits). Our simulation results indicate that shuttle-exploiting attacks can degrade victim program fidelity by $\approx 2\times$ to $\approx 68\times$. The results prove that shuttle attacks are a significant problem unless appropriate countermeasures are deployed.

To protect against the shuttle-exploiting attacks, we propose several mitigation techniques, including adopting a hybrid initial mapping policy, adding dummy pad qubits in the victim program, and capping the maximum number of allowed shuttles. We show that using a random initial mapping policy can reduce the number of shuttles by $\approx 3\times$ in malicious programs and thus, can weaken the attack. However, this approach also penalizes legitimate programs, which is an undesired effect. To balance the opposite effects, we prescribe compilation using both a random and a systematic (e.g., greedy) initial mapping policy and discard the result with a higher number of shuttles. In this manner, malicious programs will be abated, whereas legitimate programs will not be penalized.

A. OPEN CHALLENGES AND FUTURE DIRECTION

For the countermeasures, there is a cost vs. effectiveness trade-off. For example, adding dummy padding qubits or switching to a single program mode when the number of shuttles exceeds a set maximum can isolate two programs and safeguard against nefarious effects of repeated shuttles. However, these reduce device utilization and throughput and increase the cost to run a program. Thus, there are scopes to explore low-cost countermeasures. In this regard, ion cooling (e.g., [66]) can be explored as a potential solution. It will require adding laser cooling of ions after a certain number of shuttles to negate the heating of the ion chain and restore the fidelity of gate operations. Added cooling steps may consume some time and can make program execution slower. Therefore, cooling operations need to be added as required. Exploration of the optimum number of cooling operations is a future research topic.

ACKNOWLEDGMENT

The authors would like to thank Prakash Murali (Princeton) for providing the codebase for the QCCD compiler and simulator.

REFERENCES

- [1] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, Feb. 2019, Art. no. 040504, doi: 10.1103/PhysRevLett.122.040504.
- [2] Y. Cao, J. Romero, and A. Aspuru-Guzik, "Potential of quantum computing for drug discovery," *IBM J. Res. Develop.*, vol. 62, no. 6, pp. 1–6, 2018.
- [3] B. Bauer, D. Wecker, A. J. Millis, M. B. Hastings, and M. Troyer, "Hybrid quantum-classical approach to correlated materials," *Phys. Rev. X*, vol. 6, no. 3, Sep. 2016, Art. no. 031045, doi: 10.1103/PhysRevX.6.031045.
- [4] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, 2017.
- [5] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*.
- [6] (Dec. 2018). *H.R.6227—National Quantum Initiative Act*. [Online]. Available: <https://www.congress.gov/bill/115th-congress/house-bill/6227/>
- [7] (Dec. 2018). *The National Quantum Coordination Office (NQCO)*. [Online]. Available: <https://www.quantum.gov/nqco/#THE-NATIONAL-QUANTUM-COORDINATION-OFFICE>
- [8] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N.-C. Piseni, M. Chmielewski, C. Collins, and K.-M. Hudek, "Benchmarking an 11-qubit quantum computer," *Nature Commun.*, vol. 10, no. 1, pp. 1–6, 2019.
- [9] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, "Validating quantum computers using randomized model circuits," *Phys. Rev. A, Gen. Phys.*, vol. 100, no. 3, Sep. 2019, Art. no. 032328, doi: 10.1103/PhysRevA.100.032328.
- [10] Honeywell. (Jul. 2021). *Honeywell Sets Another Record for Quantum Computing Performance*. [Online]. Available: <https://www.honeywell.com/us/en/news/2021/07/honeywell-sets-another-record-for-quantum-computing-performance>
- [11] IONQ. (Apr. 2021). *Getting Started Using Qiskit With IonQ Hardware*. [Online]. Available: <https://ionq.com/posts/getting-started-using-qiskit-with-ionq-hardware>
- [12] (Apr. 2021). *Quantum (Hardware) Service Providers*. [Online]. Available: <https://qiskit.org/documentation/partners/>
- [13] Amazon. (May 2021). *Amazon Braket*. [Online]. Available: <https://aws.amazon.com/braket/>
- [14] Microsoft. (May 2021). *Azure Quantum*. [Online]. Available: <https://azure.microsoft.com/en-us/services/quantum/>
- [15] T. Hubregtzen, D. Wierichs, E. Gil-Fuster, P.-J. H. S. Derks, P. K. Faehrmann, and J. J. Meyer, "Training quantum embedding kernels on near-term quantum computers," 2021, *arXiv:2105.02276*.
- [16] D.-B. Zhang, S.-L. Zhu, and Z. D. Wang, "Protocol for implementing quantum nonparametric learning with trapped ions," *Phys. Rev. Lett.*, vol. 124, no. 1, Jan. 2020, Art. no. 010506.
- [17] S. Johri, S. Debnath, A. Mocherla, A. Singk, A. Prakash, J. Kim, and I. Kerendis, "Nearest centroid classification on a trapped ion quantum computer," *npj Quantum Inf.*, vol. 7, no. 1, pp. 1–11, Dec. 2021.
- [18] C. Bravo-Prieto, J. Baglio, M. Cè, A. Francis, D. M. Grabowska, and S. Carrazza, "Style-based quantum generative adversarial networks for Monte Carlo events," 2021, *arXiv:2110.06933*.
- [19] M. S. Rudolph, N. Bashige Toussaint, A. Katarawa, S. Johri, B. Peropadre, and A. Perdomo-Ortiz, "Generation of high-resolution handwritten digits with an ion-trap quantum computer," 2020, *arXiv:2012.03924*.
- [20] E. Y. Zhu, S. Johri, D. Bacon, M. Esencan, J. Kim, M. Muir, N. Murgai, J. Nguyen, N. Piseni, A. Schouela, K. Sosnova, and K. Wright, "Generative quantum learning of joint probability distribution functions," 2021, *arXiv:2109.06315*.
- [21] R. Yalovetzky, P. Minssen, D. Herman, and M. Pistoia, "NISQ-HHL: Portfolio optimization for near-term quantum hardware," 2021, *arXiv:2110.15958*.

- [22] D. Zhu, S. Johri, N. M. Linke, K. A. Landsman, C. H. Alderete, N. H. Nguyen, A. Y. Matsuura, T. H. Hsieh, and C. Monroe, "Generation of thermofield double states and critical ground states with a quantum computer," *Proc. Nat. Acad. Sci. USA*, vol. 117, no. 41, pp. 25402–25406, Oct. 2020.
- [23] Z. Davoudi, M. Hafezi, C. Monroe, G. Pagano, A. Seif, and A. Shaw, "Towards analog quantum simulations of lattice gauge theories with trapped ions," *Phys. Rev. Res.*, vol. 2, no. 2, Apr. 2020, Art. no. 023015.
- [24] Y. Kawashima, E. Lloyd, M. P. Coons, Y. Nam, S. Matsuura, A. J. Garza, S. Johri, L. Huntington, V. Senicourt, A. O. Maksymov, J. H. V. Nguyen, J. Kim, N. Alidoust, A. Zaribafiyani, and T. Yamazaki, "Optimizing electronic structure simulations on a trapped-ion quantum computer using problem decomposition," *Commun. Phys.*, vol. 4, no. 1, pp. 1–9, Dec. 2021.
- [25] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, "Quantum chemistry calculations on a trapped-ion quantum simulator," *Phys. Rev. X*, vol. 8, no. 3, Jul. 2018, Art. no. 031022.
- [26] Y. Nam, J.-S. Chen, N. C. Pistenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, J. M. Amini, J. Apisdorf, and K. M. Beck, "Ground-state energy estimation of the water molecule on a trapped-ion quantum computer," *NPJ Quantum Inf.*, vol. 6, no. 1, pp. 1–6, Dec. 2020.
- [27] Honeywell. (Nov. 2021). *How Quantum Computing Can Help Nippon Steel Improve Scheduling at Plants*. [Online]. Available: <https://www.honeywell.com/us/en/news/2021/06/how-quantum-computing-can-help-nippon-steel-improve-scheduling-at-plants>
- [28] Honeywell. (Nov. 2021). *How BMW Can Maximize its Supply Chain Efficiency with Quantum*. [Online]. Available: <https://www.honeywell.com/us/en/news/2021/01/exploring-supply-chain-solutions-with-quantum-computing>
- [29] T. Giurgica-Tiron, S. Johri, I. Kerenidis, J. Nguyen, N. Pistenti, A. Prakash, K. Sosnova, K. Wright, and W. Zeng, "Low depth amplitude estimation on a trapped ion quantum computer," 2021, *arXiv:2109.09685*.
- [30] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, "Demonstration of the trapped-ion quantum CCD computer architecture," *Nature*, vol. 592, no. 7853, pp. 209–213, Apr. 2021.
- [31] R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter, "High-fidelity laser-free universal control of trapped ion qubits," *Nature*, vol. 597, no. 7875, pp. 209–213, Sep. 2021.
- [32] IonQ. (Nov. 2021). *IonQ Opens Door to Dramatically More Powerful Quantum Computers, Debuts Industry First: Reconfigurable Multicore Quantum Architecture*. [Online]. Available: <https://ionq.com/news/august-25-2021-reconfigurable-multicore-quantum-architecture>
- [33] Honeywell. (Oct. 2020). *Get to Know Honeywell's Latest Quantum Computer System Model H1*. [Online]. Available: <https://www.honeywell.com/us/en/news/2020/10/get-to-know-honeywell-s-latest-quantum-computer-system-model-h1>
- [34] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, "Experimental comparison of two quantum computing architectures," *Proc. Nat. Acad. Sci. USA*, vol. 114, no. 13, pp. 3305–3310, 2017. [Online]. Available: <https://www.pnas.org/content/114/13/3305>
- [35] T. Lubinski, S. Johri, P. Varosy, J. Coleman, L. Zhao, J. Necaie, C. H. Baldwin, K. Mayer, and T. Proctor, "Application-oriented performance benchmarks for quantum computing," 2021, *arXiv:2110.03137*.
- [36] D. Kielpinski, C. Monroe, and D. J. Wineland, "Architecture for a large-scale ion-trap quantum computer," *Nature*, vol. 417, no. 6890, pp. 709–711, 2002.
- [37] P. Murali, D. M. Debroy, K. R. Brown, and M. Martonosi, "Architecting noisy intermediate-scale trapped ion quantum computers," in *Proc. ACM/IEEE 47th Annu. Int. Symp. Comput. Archit. (ISCA)*, May 2020, pp. 529–542, doi: [10.1109/ISCA45697.2020.00051](https://doi.org/10.1109/ISCA45697.2020.00051).
- [38] P. Murali. (Jan. 2021). *QCCDSim*. [Online]. Available: <https://github.com/prakashmurali/QCCDSim>
- [39] S. M. Saeed, R. Wille, and R. Karri, "Locking the design of building blocks for quantum circuits," *ACM Trans. Embedded Comput. Syst.*, vol. 18, no. 5s, pp. 1–15, Oct. 2019.
- [40] N. Acharya and S. M. Saeed, "A lightweight approach to detect malicious/unexpected changes in the error rates of NISQ computers," in *Proc. 39th Int. Conf. Comput.-Aided Design*, Nov. 2020, pp. 1–9.
- [41] V. Saravanan and S. M. Saeed, "Decomposition-based watermarking of quantum circuits," in *Proc. 22nd Int. Symp. Quality Electron. Design (ISQED)*, Apr. 2021, pp. 73–78.
- [42] K. Phalak, A. Ash-Saki, M. Alam, R. O. Topaloglu, and S. Ghosh, "Quantum PUF for security and trust in quantum computing," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 11, no. 2, pp. 333–342, May 2021, doi: [10.1109/JETCAS.2021.3077024](https://doi.org/10.1109/JETCAS.2021.3077024).
- [43] A. Ash-Saki, M. Alam, and S. Ghosh, "Analysis of crosstalk in NISQ devices and security implications in multi-programming regime," in *Proc. ACM/IEEE Int. Symp. Low Power Electron. Design*, New York, NY, USA, Aug. 2020, pp. 25–30, doi: [10.1145/3370748.3406570](https://doi.org/10.1145/3370748.3406570).
- [44] A. Suresh, A. A. Saki, M. Alam, R. O. Topaloglu, and D. S. Ghosh, "A quantum circuit obfuscation methodology for security and privacy," 2021, *arXiv:2104.05943*.
- [45] P. Das, S. S. Tannu, P. J. Nair, and M. Qureshi, "A case for multi-programming quantum computers," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchitecture*, New York, NY, USA, Oct. 2019, pp. 291–303, doi: [10.1145/3352460.3358287](https://doi.org/10.1145/3352460.3358287).
- [46] Rigetti. (May 2021). *Quantum Cloud Services*. [Online]. Available: <https://qcs.rigetti.com/>
- [47] Rigetti. (May 2021). *Reservations*. [Online]. Available: <https://docs.rigetti.com/en/getting-started-with-your-ide/jupyterlab-ide>
- [48] K. Mølmer and A. Sørensen, "Multiparticle entanglement of hot trapped ions," *Phys. Rev. Lett.*, vol. 82, no. 9, pp. 1835–1838, Mar. 1999, doi: [10.1103/PhysRevLett.82.1835](https://doi.org/10.1103/PhysRevLett.82.1835).
- [49] Y. Wu, S.-T. Wang, and L.-M. Duan, "Noise analysis for high-fidelity quantum entangling gates in an anharmonic linear Paul trap," *Phys. Rev. A, Gen. Phys.*, vol. 97, no. 6, Jun. 2018, Art. no. 062325, doi: [10.1103/PhysRevA.97.062325](https://doi.org/10.1103/PhysRevA.97.062325).
- [50] IONQ. (May 2021). *Best Practices—Native Gates*. [Online]. Available: <https://ionq.com/best-practices>
- [51] Honeywell. (Jan. 2021). *Honeywell System Models H0 and H1*. [Online]. Available: <https://www.honeywell.com/us/en/company/quantum/quantum-computer>
- [52] AWS-Braket. (Mar. 2021). *Allocating Qubits on QPU Devices*. [Online]. Available: https://github.com/aws/amazon-braket-examples/blob/main/examples/braket_features/Allocating_Qubits_on_QPU_Devices.ipynb
- [53] C. J. Trout, M. Li, M. Gutiérrez, Y. Wu, S.-T. Wang, L. Duan, and K. R. Brown, "Simulating the performance of a distance-3 surface code in a linear ion trap," *New J. Phys.*, vol. 20, no. 4, Apr. 2018, Art. no. 043038.
- [54] M. Gutiérrez, M. Müller, and A. Bermúdez, "Transversality and lattice surgery: Exploring realistic routes toward coupled logical qubits with trapped-ion quantum processors," *Phys. Rev. A, Gen. Phys.*, vol. 99, no. 2, Feb. 2019, Art. no. 022330, doi: [10.1103/PhysRevA.99.022330](https://doi.org/10.1103/PhysRevA.99.022330).
- [55] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, "Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, 2019, pp. 1015–1029.
- [56] S. S. Tannu and M. K. Qureshi, "Not all qubits are created equal: A case for variability-aware policies for NISQ-era quantum computers," in *Proc. 24th Int. Conf. Architectural Support Program. Lang. Operating Syst.*, Apr. 2019, pp. 987–999.
- [57] M. Y. Siraichi, F. M. Q. Pereira, V. Dos Santos, and C. Collange, "Qubit allocation," in *Proc. Int. Symp. Code Gener. Optim.*, Jun. 2020, pp. 113–125.
- [58] Rigetti. (May 2021). *Compiler*. [Online]. Available: <https://pyquil-docs.rigetti.com/en/v2.22.0/compiler.html>
- [59] (May 2021). *Transpiler*. [Online]. Available: <https://qiskit.org/documentation/stubs/qiskit.compiler.transpile.html>
- [60] M. Yan, J. Choi, D. Skarlatos, A. Morrison, C. Fletcher, and J. Torrellas, "Invisispec: Making speculative execution invisible in the cache hierarchy," in *Proc. 51st Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, Oct. 2018, pp. 428–441.
- [61] O. Weisse, I. Neal, K. Loughlin, T. F. Wenisch, and B. Kasikci, "NDA: Preventing speculative execution attacks at their source," in *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchitecture*, New York, NY, USA, Oct. 2019, pp. 572–586, doi: [10.1145/3352460.3358306](https://doi.org/10.1145/3352460.3358306).
- [62] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2019, pp. 1–19.

- [63] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown: Reading kernel memory from user space," in *Proc. 27th USENIX Secur. Symp. Secur. (USENIX)*, Baltimore, MD, USA, Aug. 2018, pp. 973–990. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/lipp>
- [64] T. Tomesh. (May 2021). *Quantum-Circuit-Generator*. [Online]. Available: https://github.com/teaguetomesh/quantum_circuit_generator
- [65] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient Z gates for quantum computing," *Phys. Rev. A, Gen. Phys.*, vol. 96, Aug. 2017, Art. no. 022330, doi: [10.1103/PhysRevA.96.022330](https://doi.org/10.1103/PhysRevA.96.022330).
- [66] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, "Resolved-sideband Raman cooling of a bound atom to the 3D zero-point energy," *Phys. Rev. Lett.*, vol. 75, no. 22, pp. 4011–4014, Nov. 1995, doi: [10.1103/PhysRevLett.75.4011](https://doi.org/10.1103/PhysRevLett.75.4011).



ABDULLAH ASH SAKI (Graduate Student Member, IEEE) received the B.Sc. degree from the Bangladesh University of Engineering and Technology (BUET), in 2014, and the M.S. and Ph.D. degrees from The Pennsylvania State University, University Park, PA, USA, in 2020 and 2021, respectively. In his Ph.D., he worked on noise resilient and secure quantum computing. He is set to join Zapata Computing as a Quantum Research Scientist, since January 2022. He was a recipient

of Dr. Nirmal K. Bose Dissertation Excellence Award from the Department of Electrical Engineering, Penn State.



RASIT ONUR TOPALOGLU (Senior Member, IEEE) received the B.S. degree in EE from Bogazici University and the Ph.D. degree in computer science and engineering from the University of California at San Diego. He has worked for companies, such as Qualcomm, AMD, GLOBALFOUNDRIES and is currently with IBM. He works on next-generation computer technology and design currently as a Senior Hardware Developer. He was partially involved with qubit characterization laboratory work at IBM Research. He has over 60 peer-reviewed publications and over 60 issued U.S. patents, more than a 30 of which are on quantum technologies. He has chaired the IEEE/ACM DAC Workshop

on Design Automation for Quantum (DAQ). As of 2021, he is working on a Quantum Computing book. He serves on IEEE/ACM Design Automation Conference (DAC), IEEE/ACM International Conference on Computer-Aided Design (ICCAD), and IEEE International Symposium on Quality Electronic Design (ISQED) Technical Program Committees that cover quantum topics. He serves as the Chair for IEEE Mid-Hudson and the Secretary of ACM Poughkeepsie. He is an IEEE/ACM DAC Outstanding Innovator and an IBM Master Inventor.



SWAROOP GHOSH (Senior Member, IEEE) received the B.E. degree (Hons.) from IIT, Roorkee and the Ph.D. degree from Purdue University. He is currently an Associate Professor with Pennsylvania State University. Earlier, he was with the Faculty of University of South Florida (USF). Prior to that, he was a Senior Research and Development Engineer in Advanced Design, Intel Corporation. His research interests include emerging memory technologies, hardware security, quantum computing, and digital testing. He is a Senior member of the National Academy of Inventors (NAI) and an Associate Member of Sigma Xi and a Distinguished Speaker of the Association for Computing Machinery (ACM). He was a recipient of Intel Technology and Manufacturing Group Excellence Award in 2009, the Intel Divisional Award in 2011, the Intel Departmental Awards in 2011 and 2012, the USF Outstanding Research Achievement Award in 2015, the College of Engineering Outstanding Research Achievement Award in 2015, the DARPA Young Faculty Award (YFA) in 2015, the ACM SIGDA Outstanding New Faculty Award in 2016, the YFA Director's Fellowship in 2017, the Monkowsky Career Development Award in 2018, the Lutron Spira Teaching Excellence Award in 2018, the Dean's Certificate of Excellence in 2019 and 2021, and the Best Paper Award in American Society of Engineering Education (ASEE) Annual Conference in 2020. He has also served in the technical program committees of more than 25 ACM/IEEE conferences, including DAC, ICCAD, CICC, DATE, ISLPED, GLSVLSI, Nanoarch, and ISQED. He served as the Technical Program Co-Chair of ISLPED 2022, the General Chair of ISQED 2021, the Conference Chair of ISQED 2020, the Program Chair of DAC Ph.D. Forum 2016, and the Track (Co)-Chair of CICC 2017–2019, ISLPED 2017–2021, and ISQED 2016–2017. He served as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I: REGULAR PAPERS and IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN and as a Senior Editorial Board member for IEEE JOURNAL OF EMERGING TOPICS ON CIRCUITS AND SYSTEMS (JETCAS). He served as a Guest Editor for the IEEE JETCAS and IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS.

on Design Automation for Quantum (DAQ). As of 2021, he is working on a Quantum Computing book. He serves on IEEE/ACM Design Automation Conference (DAC), IEEE/ACM International Conference on Computer-Aided Design (ICCAD), and IEEE International Symposium on Quality Electronic Design (ISQED) Technical Program Committees that cover quantum topics. He serves as the Chair for IEEE Mid-Hudson and the Secretary of ACM Poughkeepsie. He is an IEEE/ACM DAC Outstanding Innovator and an IBM Master Inventor.

• • •