

Fast Fingerprinting of Cloud-based NISQ Quantum Computers

Kaitlin N. Smith*[§], Joshua Viszlai[†], Lennart Maximilian Seifert[†], Jonathan M. Baker[‡]
Jakub Szefer^{||} and Frederic T. Chong*[†]

**Super.tech, a division of Infleqtion*

[†]*Dept. of Computer Science, University of Chicago*

[‡]*Duke Quantum Center, Duke University*

^{||}*Dept. of Electrical Engineering, Yale University*

Abstract—Cloud-based quantum computers have become a reality with a number of companies allowing for cloud-based access to their machines with tens to more than 100 qubits. With easy access to quantum computers, quantum information processing will potentially revolutionize computation, and superconducting transmon-based quantum computers are among some of the more promising devices available. Cloud service providers today host a variety of these and other prototype quantum computers with highly diverse device properties, sizes, and performances. The variation that exists in today’s quantum computers, even among those of the same underlying hardware, motivate the study of how one device can be clearly differentiated and identified from the next. As a case study, this work focuses on the properties of 25 IBM superconducting, fixed-frequency transmon-based quantum computers that range in age from a few months to approximately 2.5 years. Through the analysis of current and historical quantum computer calibration data, this work uncovers key features within the machines, primarily frequency characteristics of transmon qubits, that can serve as a basis for a unique hardware fingerprint of each quantum computer.

I. INTRODUCTION

We are at a transformative moment in quantum computing. After significant investment from government, industry, and academia, the widespread speculation surrounding quantum information processing (QIP) is evolving into overwhelming optimism – much of the scientific community believes that quantum computers (QCs) will revolutionize computation in a matter of decades. This excitement surrounding quantum computing can be credited to the fact that machines once thought of as an elusive theoretical concept [1] are beginning to emerge in labs across the world [2], [3]. These prototypes consist of a variety of technologies from devices that compute using superconducting (SC) circuits to implementations that encode information within atomic particles.

QCs and their quantum bits, or qubits, offer a novel means to solve some of today’s most challenging problems in computation through the careful application of quantum superposition, interference, and entanglement. Quantum advantage over classical computers is on the horizon as larger QCs emerge with exponentially larger computational spaces than their classical counterparts. Anticipated applications for robust QCs include cryptography [4], big data [5], chemistry [6], optimization [7], and machine learning [8].

[§]Email correspondence: kaitlin.smith@infleqtion.com

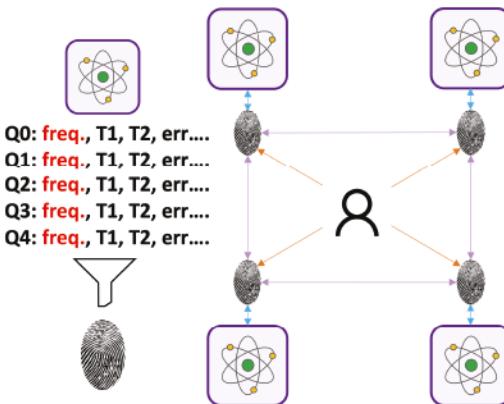


Fig. 1. Quantum hardware fingerprints are constructed from robust and reproducible property signatures. Potential candidates for fingerprint features include frequency of qubits (freq.), decoherence characteristics ($T1 / T2$), and various error rates (err.).

A. Today’s Quantum Computers

Recent advances have allowed the development of QCs with tens to over a hundred physical qubits [9]. The debut of QCs with a modest number of qubits marks a breakthrough in QIP, but unfortunately these devices, sometimes referred to as noisy intermediate scale quantum (NISQ) machines [10], are still plagued by high gate infidelity and short coherence windows. Regardless, the promise of quantum computing, and breakthroughs it can achieve, has led to the development of many prototype devices that compute with atoms, ions, superconducting circuits, and more. Several of these aforementioned implementations have been co-developed with quantum software development kits (SDKs) that allow for these QCs to be programmed remotely from a classical machine.

B. Cloud-based Quantum Computing Today

Quantum computing is under rapid development, and today there are several cloud-based QCs. Although these QCs, along with their highly-specialized, associated classical infrastructure, are in limited supply, they are readily available for remote access. These QCs and corresponding software stacks are primarily used by researchers in academia and industry

but are programmable by anyone with internet access and runtime credit. Current cloud vendors with their own QCs include industry giants like IBM, Google, and Quantinuum, as well as startups such as Xanadu, Rigetti, IonQ. Further, Amazon Braket and Microsoft Azure Quantum provide quantum computing as a service via multiple other quantum hardware vendors.

Quantum computing as a cloud service is expected to grow considerably over the next decade, continuing to be the main gateway to quantum computation targeted toward sensitive applications such as financial modeling, cryptography, and genome data analysis. The wide availability of quantum cloud devices has many benefits but has also introduced security vulnerabilities. For example, multi-tenant [11] and temporal [12] covert communication has been showcased in academia.

C. Towards Hardware Fingerprinting of Quantum Computers

A hardware fingerprint is a security primitive that includes information collected about a remote computing device for the purpose of identification. A robust fingerprint includes features that are unique, stable, and collision resistant. In classical computing, recent work fingerprints cloud FPGAs through Physical Unclonable Functions (PUFs) based on the decay of DRAM modules attached to the FPGAs [13]. Separately, DRAM PUFs have been widely used to identify and authenticate DRAM chips [14], [15], or generate keys [16], [17]. Unfortunately, hardware fingerprinting can also allow a malicious party to navigate a cloud architecture, enabling attacks that survey computation scheduling, track resource allocation, or leak information [18].

Not unlike the fingerprinting of classical computers in the cloud, identifying device-level features that are distinct across different QCs and are reasonably stable is a challenge. QCs for the foreseeable future will be cloud-based resources, and QC fingerprinting could serve both offensive and defensive purposes. For example, the ability to uniquely identify cloud QC resources could enable malicious parties to monitor or spoof quantum devices and learn or disrupt sensitive quantum applications. Conversely, if a unique identifier were used within a cryptographic protocol, QC authentication would be enabled. Fig. 1 demonstrates how different properties of QCs, such as frequency of qubits (*freq.*), decoherence characteristics (*T1 / T2*), or various error rates (*err.*) could be candidates for a fingerprint basis.

In this work, we are specifically interested in *in-situ* quantum hardware fingerprinting – the quantum device itself provides a means of identification as opposed to authentication provided by a co-located classical peripheral. We are motivated to use the quantum chip itself for identification as the potential exists for the classical peripheral to be disconnected from the QC of interest and reattached to another. Similar to the classical fingerprinting discussed earlier, a fingerprint of critical importance must comprise of unique properties of a QC which have sufficient consistency.

To better understand what QC features can be applied toward a hardware fingerprint, we must understand the evolving

quantum infrastructure on the cloud and, most importantly, their device characteristics. Here, we focus on IBM quantum systems that consist of fixed-frequency transmon qubits. There are currently over 20 machines of various generations and sizes that are available for public experimentation via the IBM Quantum Cloud [19], [20], and the list of available QCs is constantly evolving. Many similarities among QCs such as their underlying topology and technology might make the devices seem interchangeable. However, the QCs differ in their operational characteristics and performance as a result of their hardware, control, and programming infrastructure. This motivates the application of QC properties toward QC fingerprinting. A successful fingerprinting scheme should generalize across the spectrum of available quantum devices and device generations.

D. Paper Contributions

To evaluate potential device characteristics that generalize across device size and generation, demonstrate reliability over long periods of time, and are unique to each device, we study historical data of real IBM QCs and learn the property signatures unique to individual qubits. This study assimilates the properties of 25 IBM QCs that range in age from a few months to approximately 2.5 years. This required the collection and cleaning of 11,544 records of calibration data, each containing information about QC properties.

Our study reveals that qubit frequency is unique to each device and is most suitable for use as the basis for a QC fingerprint. Qubit frequency is dependent on intrinsic device fabrication properties, making it very hard to duplicate exact frequencies even by the QC manufacturer. Thus, we show that qubit frequency within a QC's property signature can be applied in a fingerprint due to its reliability, whereas other characteristics, such as gate error rates, cannot.

II. THREAT MODEL

This paper assumes quantum computing will primarily be a cloud service that operates on the principle that users target a specific machine for execution of their quantum program or circuit. The QC is usually selected due to specific hardware performance requirements (i.e. coherence times, error thresholds, number of qubits, etc.) required for a certain computational task. This is analogous to the infrastructure-as-a-service model (IaaS) used by classical cloud service providers [21]. We assume that a QC is effectively “rented” for a job. As a direct result of requesting specific hardware, users must often wait in a queue until their job is serviced by their selected QC. In the future, users may have less control on what quantum hardware their job runs if multiple QCs satisfy the minimum QC performance requirements. In this scheme, known as server-less computing [22], the end user has minimal control over the underlying infrastructure used. In a server-less computing environment, workloads ideally run homogeneously from one partition to the next as providers attempt to hide intimate details of the underlying hardware from the user. This prevents users from gleaning details

of the cloud resources or tracking hardware scheduling and utilization.

In our work, we assume remote, cloud-based QCs based on fixed-frequency transmon qubits, basing our device modeling on the IBM QCs featured in Table I. The goal of the attacker is to infer the identity of a QC based on intrinsic characteristics directly tied to hardware. In particular, the attacker creates a fingerprint with carefully selected features from a QC's unique property signature that includes information about coherence time, gate fidelity, and error rates, among others. While we focus our study on fixed-frequency transmon QCs, we hypothesize that the presented techniques will be applicable to alternative QC platforms. We assume that the attacker has privileges to submit quantum circuits written at either the gate or pulse abstraction level in order to develop their QC fingerprint. Once a malicious party can reliably locate compute resources within the cloud, offensive behavior such as covert and side channel attacks, information leakage, and interference with data processing are facilitated. The victim is the benevolent QC user that accesses expensive cloud resources for the quantum acceleration of sensitive applications proposed for quantum such as cybersecurity, finance, drug design, and weather forecasting [23].

III. QUANTUM COMPUTING FUNDAMENTALS

A. Quantum Information

Classical bits hold a static value of either 0 or 1, but quantum bits (qubits) can demonstrate states of superposition in the form of $\alpha|0\rangle + \beta|1\rangle$ where $|\alpha|^2 + |\beta|^2 = 1$. Superposition enables n qubits to represent up to 2^n states simultaneously. This phenomenon, along with quantum interference and entanglement, allow select problems to be solved with significant reductions in complexity. Some common single-qubit transformations include $R_x(\pi) = X$, $R_z(\pi) = Z$, and $R_y(\pi) = Y$, that cause qubit basis-flip, phase-flip, and combination basis-flip and phase-flip, respectively. These are typically the single qubit “gates” that users can implement in quantum circuits. Examples of multi-qubit operations include the logical *SWAP* operation that exchanges quantum state between qubits. Controlled gates, such as *CX* or *CZ*, execute an operation on a target qubit depending on the state of one (or more) control qubit(s). The set of basis gates available to a particular QC is technology-dependent since quantum hardware has certain operators that are more favorable for physical implementation. The basis gates usually consist of a set of single-qubit operations that are capable of implementing arbitrary rotations, R_x , R_y , and R_z , along with an additional two-qubit operation such as *CX* or *CZ*.

B. Today's Physical Qubits

Some of the biggest challenges that limit QC scaling include decoherence, gate errors, readout noise, and crosstalk. Additionally, most NISQ devices suffer from limited qubit-qubit communication since architecturally constrained hardware only supports nearest-neighbor connectivity.

Coherence time represents how long qubits maintain state for useful computation. T_1 and T_2 times correspond to amplitude and phase information, respectively. Decoherence can present itself in the form of a qubit in a high-energy state decaying to a low-energy state or a qubit losing superposition by coupling to its surrounding environment. For SC QCs, coherence times have improved from one nanosecond to 100s of microseconds in the last decade.

Imprecise control of QCs can lead to gate and readout errors. From public IBM information, single-qubit instruction error rates are of the order of 10^{-3} , whereas for two-qubit instructions, such as *CX*, are 10^{-2} [24]. Crosstalk noise arises from unwanted interactions between neighboring qubits and from leakage of the control signals [25].

Fixed-frequency transmons, a type of SC circuit qubit realization, are of focus in this work. Fixed-frequency means that each transmon's operating frequency is fixed at fabrication and does not change, as opposed to tunable-frequency transmon qubits. Transmon qubits and their fabrication are further described in Section IV.

IV. TRANSMON-BASED QCS: FABRICATION & VARIATION

A. Device Fabrication

Transmons provide a viable platform for realizing physical qubits [26]. Transmon qubits employ Josephson Junctions (JJs) along with other hardware at ultra-cold temperatures to act as mesoscopic-scale artificial atoms with an anharmonic energy spectrum. Compared to alternative technologies, transmons are extremely promising due to recent improvements in device coherence, operation fidelity, and addressability [27], [28]. Progress in transmon technology has enabled early demonstrations of quantum algorithms [6], [29], quantum error correction [30]–[32], and quantum primacy [33].

Transmon QCs are promising for qubit scaling because their fabrication takes advantage of well-established techniques that are streamlined by decades of classical device manufacturing. During processing, qubits are constructed in layers [34]–[36], and just like in classical fabrication, some degree of device variation is injected. Inevitable QC variation is due to process imprecision or stochastic defects related to the manufacturing environment or materials.

While every SC quantum device lab may have a unique transmon qubit recipe, at a high level, qubit fabrication involves: ① Substrate cleaning and preparation, ② Base metal placement, ③ Lithography-based feature definition, ④ JJ definition, creation and placement, and ⑤ Packaging and final test [34]–[36]. Each of these fundamental steps presents an opportunity to introduce inter-chip and intra-chip QC feature variation.

Many SC qubit manufacturers implement custom procedures that attempt to reduce heterogeneity and boost qubit quality. For instance, prior work [35] demonstrates improvements in qubit relaxation and coherence times by replacing the often-used niobium with tantalum in the initial metal layer of the chip. Other research [34] found that removing a thin layer of the substrate as a pre-fabrication step helps decrease qubit

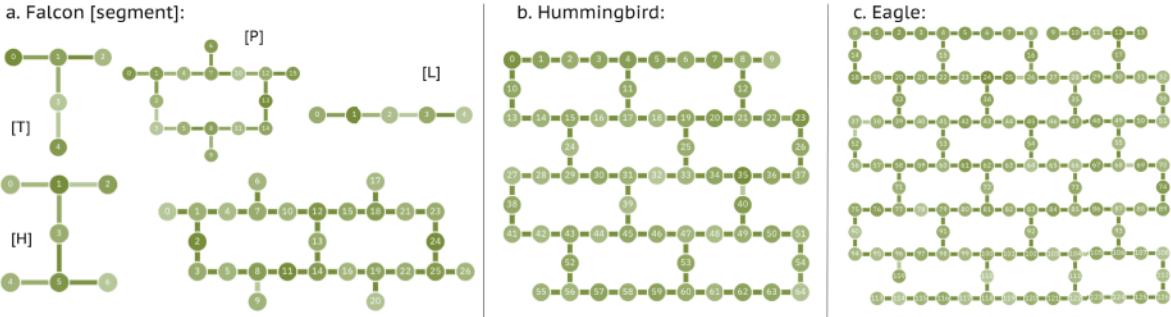


Fig. 2. IBM QC topology of the (a) Falcon, (b) Hummingbird, and (c) Eagle processors. Device [segment] is indicated, where applicable. Canary processor topology has not been included since it contains a single qubit.

loss. These examples show that the route to fault-tolerant (FT) quantum computing with SC qubits necessitates refined materials and methodologies for producing QCs. Unfortunately, improved fabrication that allows QCs to scale are still likely to fail at creating perfectly homogeneous systems. Unavoidable variability in QC fabrication, even if seemingly small, elicits unique differences between devices, opening doors for robust fingerprinting of QCs based on intrinsic device properties.

B. Qubit Variation

QCs are fabricated in layers, and there are many steps in processing where a slight deviation from the target specification impacts physical qubit features. Physical qubit variation causes performance variation that is observable across and between QCs. QC variation often results from fabrication imprecision, and physical abnormalities in a QC are called defects. QC defects can be challenging to pinpoint and characterize since their impact can vary. For instance, a defect can have a subtle effect on qubit performance by slightly diminishing fidelity or coherence. At the other extreme, a defect can cause severe reliability issues or, in the worst case, can prevent qubits from participating in meaningful computation. Variation can broadly be categorized into two forms – static and dynamic.

Static variation can occur across devices (inter-) as well as within devices (intra-), but the corresponding device characteristics are fairly stable over time. Fabrication imprecision [37], [38] is a prime example of static variation. JJs have incredibly small feature sizes that are hundreds of nanometers in scale [37]. Thus, small imperfections that appear in JJ positioning, component dimension, or surrounding layers influence operational characteristics of the transmon [38]. In turn, variation in individual qubit characteristics influences how well qubits can work together during computation.

Current QC fabrication techniques lack the precision required to produce devices that exactly match targeted design specifications. A major effect of static variation is that the post-fabrication frequencies of transmon QCs are unique to each qubit and across each device. By the nature of static variation, the uniqueness of these frequencies are fairly stable over the entire life cycle of the Qc.

Dynamic variation, as its name alludes, fluctuates over time inter- and intra-QC. An important example of dynamic variation is a two-level system (TLS) [39]. While TLS stems from fabrication imperfection, its effect transiently and stochastically appears across a QC. TLS is caused by impurities inside materials or irregularities within atomic crystalline lattice structures that can appear unexpectedly in oxide layers or on the QC's surface. TLS parasitically couples to qubits in an unpredictable manner if it appears within close proximity to the critical components of transmons, or the JJs. When TLS appears close to, and is near-resonant with, active transmon elements, the TLS interferes with qubit activity. Qubit energy is absorbed, and the qubit experiences significant reductions in coherence that is observable in truncated T_1 and T_2 times. Coupling strength to TLS varies over time, causing time-dependent qubit property fluctuation [40], [41]. Apart from TLS, other dynamic variations are observed from thermal fluctuations, magnetic flux, and quasi-particles [40], [42].

Variation effects on fingerprinting: Differentiating variation as static and dynamic is particularly important to fingerprinting. While all variations produce unique qubit and device signatures, robust fingerprints require the use of device characteristics stemming from static variation so that fingerprints are reproducible over time. It should be noted that even features thought to be predominantly affected by static variation, like qubit frequency, are not entirely free from dynamic effects. Even the most reliable device features tend to show fine-grained fluctuations of fairly small magnitudes as well as occasional, coarse-grained spikes of larger magnitudes. Thus, effective fingerprinting schemes must be robust to constant, fine-grained feature fluctuations and tolerant to occasional, coarse-grained spikes.

V. PROPERTY ANALYSIS OF THE IBM QCs

A. The IBM Quantum Ecosystem

IBM has developed many transmon-based SC QCs, and since 2016, the company has allowed their prototype devices to be used by the public via the IBM Quantum Cloud services [19]. There are currently over 20 machines of various generations and sizes available for experimentation [20], and from 2016 to 2022, IBM QCs have increased from 5 to

TABLE I

IBM QCS OF THIS STUDY. * INDICATES UNAVAILABLE FEATURE, † INDICATES RETIRED QC, ^ INDICATES PREMIUM QC (7/12/2022).

Machine Name	Processor Type	Qubits	Quantum Volume	CLOPS	Date Online	Total Records
Armonk†	Canary r1.2	1	1	*	Oct. 16, 2019	941
Bogota†	Falcon r4L	5	*	*	June 3, 2020	713
Santiago†	Falcon r4L	5	*	*	June 3, 2020	718
Manila	Falcon r5.11L	5	32	2.8K	April 28, 2021	426
Lima	Falcon r4T	5	8	2.7K	Jan. 8, 2021	518
Belem	Falcon r4T	5	16	2.5K	Jan. 8, 2021	510
Quito	Falcon r4T	5	16	2.5K	Jan. 8, 2021	519
Casablanca†	Falcon r4H	7	*	*	Aug. 7, 2020	511
Jakarta	Falcon r5.11H	7	16	2.4K	April 28, 2021	440
Lagos^	Falcon r5.11H	7	32	2.7K	May 20, 2021	293
Perth^	Falcon r5.11H	7	32	2.9K	July 22, 2021	245
Nairobi	Falcon r5.11H	7	32	2.6K	May 20, 2021	399
Oslo	Falcon r5.11H	7	32	2.6K	March 25, 2022	55
Guadalupe	Falcon r4P	16	32	2.4K	Jan. 8 2021	490
Montreal	Falcon r4	27	128	2K	June 3, 2020	748
Toronto	Falcon r4	27	32	1.8K	June 3, 2020	743
Sydney†	Falcon r4	27	*	*	Sept. 2, 2020	446
Mumbai	Falcon r5.1	27	128	1.8K	Nov. 13 2020	572
Hanoi	Falcon r5.11	27	64	2.3K	April 24, 2021	407
Cairo	Falcon r5.11	27	64	2.4K	May 5, 2021	357
Auckland	Falcon r5.11	27	64	2.4K	July 22, 2021	238
Kolkata	Falcon r5.11	27	128	2K	Nov. 13, 2020	520
Geneva	Falcon r5.8	27	32	1.9K	March 25, 2022	56
Brooklyn†	Hummingbird r2	65	*	*	March 4, 2021	404
Washington	Eagle r1	127	64	850	Sept. 16, 2021	275

127 qubits in size [43]. As milestones are reached on IBM's quantum scaling roadmap [44], new QCs that feature cutting-edge technology are released to the quantum cloud while legacy devices are either updated or retired. Active devices can be found on the services page of IBM Quantum website [20]. Select information about retired QCs is also available [45], [46]. As a note, at time of writing, all multi-qubit QCs demonstrate heavy-hex connectivity in anticipation for the eventual adaptation of the hybrid surface/Bacon-Shor error correcting code [43], [47]. Additionally, the IBM transmon hardware natively supports the basis gate set that includes the single-qubit \sqrt{X} , X , and $R_z(\theta)$ operations along with the two-qubit CX gate.

B. Investigating the IBM QCs

QC calibration data is valuable for optimizing the use of quantum hardware. Through the analysis of machine properties, especially over time, unique property signatures can be used during circuit compilation to tailor applications to each QC [48]. However, the study of QC properties could also reveal candidate features that have potential for hardware fingerprinting. With this motivation, we gathered available historical data that contained daily calibration records for the IBM QCs.

Historical data was gathered and analyzed for 25 IBM QCs released from 2019 to 2021. The QCs under investigation are described in Table I. Processor type, the feature that groups the QCs, is found in column 2 of Table I. Processor type describes the general hardware qualities that go into the quantum machines. Processor family, where Eagle is the newest and Canary is the most mature, refers to the chip architecture. Revisions indicated with an "r" are design variants within

a given family. Segment, such as "L" or "T," differentiates processors comprising different sub-sections of a larger chip. Table I features many revisions and segments of the Falcon family of processors. The multi-qubit device coupling maps are included in Fig. 2. As a note, at time of writing, six devices, Armonk, Bogota, Santiago, Casablanca, and Sydney were retired from IBM Quantum and were no longer accessible for data retrieval. Two devices, Lagos and Perth, became premium-access devices and were also unavailable for data retrieval and analysis.

The IBM data set was created by accessing the historical data for each IBM QC. We sampled each machine's properties at 24-hour intervals, collecting all available data up to the QC's debut, or date online, on the IBM Quantum cloud services. The choice for this sampling interval is that IBM systems are typically calibrated once in a 24 hour period [49]. The total number of records gathered for each QC can be found in the rightmost column of Table I. The complete data set sums to 11,544 records of QC calibration data. Each record contains a wealth of information that provides insight to the QC's operational characteristics and performance. This data combined forms a QC's property signature as much of this calibration data is different for every QC and changes each cycle, showcasing the machine variation described in Section IV-B. QC properties include information about qubit frequency, T_1 , T_2 , measurement success, operation duration, and gate error, among others.

Raw data contains flaws such as incomplete, incorrect, or corrupted records that can skew statistical analysis if not properly managed. This was observed in the IBM data in the form of 1) records that were missing data, such as a

coherence time or gate error for a single qubit, and 2) records that had incorrect data, such as a *CX* error reported for two qubits that are not connected according to the device coupling map. Thus, we developed a methodology for cleaning the QC property data set to draw conclusions about properties eligible for quantum hardware fingerprinting:

- Remove all duplicate QC cycles. This was observed when a calibration cycle surpassed the 24-hour sampling period. A possible explanation for this would be periods where the QC was offline or under maintenance.
- Remove invalid data or incorrect formatting. An example of this case was when a calibration cycle returned an error value of 1 for all QC *CX* operations. It was assumed that this was a period where the machine was inoperable.
- Remove all incomplete or incorrect records. Examples of these inconsistencies included qubits that were missing attributes or when *CX* error was reported between unconnected qubits.

C. Observations Between Processor Revisions

IBM QC processors experience periodic upgrades that reflect recent advances in quantum hardware. Updates to a family of processors, such as the Falcon devices, result in a new revision [50]. Of the QCs included in this study, we defined two Falcon categories, r4 and r5, and compared the properties. Devices were grouped according to size and segment. We are motivated to discover identifying features for fingerprinting that transcend hardware revisions. As an initial investigation, the QC metrics chosen for a case study were qubit coherence time, T_1 and T_2 , and readout error. It should be noted that the r5 category of devices does not contain either a five-qubit T-segment device or a 16-qubit P-segment device.

Statistics about T_1 and T_2 in the form of box-and-whisker plots for each type of Falcon processor in the r4 and r5 category are described in Fig. 3(a). Plots contain three bars in each group, early-life distribution, latest distribution, and distribution over all records, to provide insight to performance at different periods of time as well as over all time. Higher is better when considering qubit coherence. When examining the plots in Fig. 3(a), we once again see a large amount of variation in coherence times, not only within processor families but also over time. This finding shows the transitive nature of coherence time and suggests that this feature is not ideal for device fingerprinting. Thus, we search for alternative QC properties.

Next, we consider readout error. Qubit readout error is an important QC property to consider because of its ability to distort the final outcome of quantum computation. Lower readout error leads to more reliable QC results. Details about Falcon processor readout error are found in Fig. 3(b). There are clear differences between the r4 and r5 processors when focused on readout error – the improvements are more dramatic than those for T_1 and T_2 . With both types of processors, however, we still see a large amount of overlap with error values and variation over time. As a result, we continue to search for more desirable QC fingerprint features.

D. Observations Between Processor Generations

Fig. 4 provides statistics for *CX* infidelity over time for three IBM processor generations. In general, the median values are similar for all processors with the Falcon device, 27-qubit Auckland, having the lowest *CX* error. A possible explanation for this is the fact that Auckland is a more mature revision than 65-qubit Brooklyn and 127-qubit Washington, allowing more time for the technology to be refined. Another important observation of Fig. 4 is that *CX* error correlates with chip size – the larger devices demonstrate higher *CX* error upper bounds over the 12 calibration cycles. Additionally, the *CX* error of Washington generally shows a larger spread over the 12 samples.

Two important takeaways stem from Fig. 4. First, although a machine might seem preferable because of one feature, such as qubit capacity, another feature critical to performance, such as high *CX* infidelity, might make that machine less desirable. This motivates the importance of carefully pairing the right QC for a QC workload because the success of an algorithm could depend on whether or not it was run on the appropriate hardware. Second, *CX* error, while unique for each featured processor, demonstrates a significant amount of day-to-day variation, indicating that the feature might not provide the most reliable means of device identification.

E. Preliminary Fingerprinting Insights

In the context of fingerprinting, one would naively assume that the error characteristics of NISQ devices might be suitable for fingerprinting. This is not the case – the results of the two prior sections clearly indicate that device signature components such as coherence times and error rates are unsuited for application within a hardware fingerprint, which conflicts with prior work [51]. T_1 , T_2 , measurement error, and gate error demonstrate temporal variance, an unfavorable characteristic for a reproducible, unique fingerprint. This motivates the study of more intrinsic device characteristics for fingerprinting, such as qubit frequency – this is discussed in Section VI.

VI. TOWARDS NOVEL AND FAST FINGERPRINTING BASED ON QUBIT FREQUENCIES

As described in Section V, some QC features related to gate infidelity, coherence, and readout error might not be ideal for fingerprinting purposes due to issues associated with overlap among different devices and inconsistency over time. Here, we explore the viability of qubit frequency as a basis for constructing a unique QC fingerprint.

A. Feature Variation Over Time

Here, we compare the transient nature of single-qubit error and frequency on IBM's single-qubit Canary processor, Armonk. Fig. 5 includes two plots. Single-qubit gate error, i.e. *X* or *SX* gate error, is indicated on the left axis in blue while qubit frequency is indicated on the right axis in orange. The curves describe the error and frequency data normalized by their means, respectively, for calibration cycles collected over

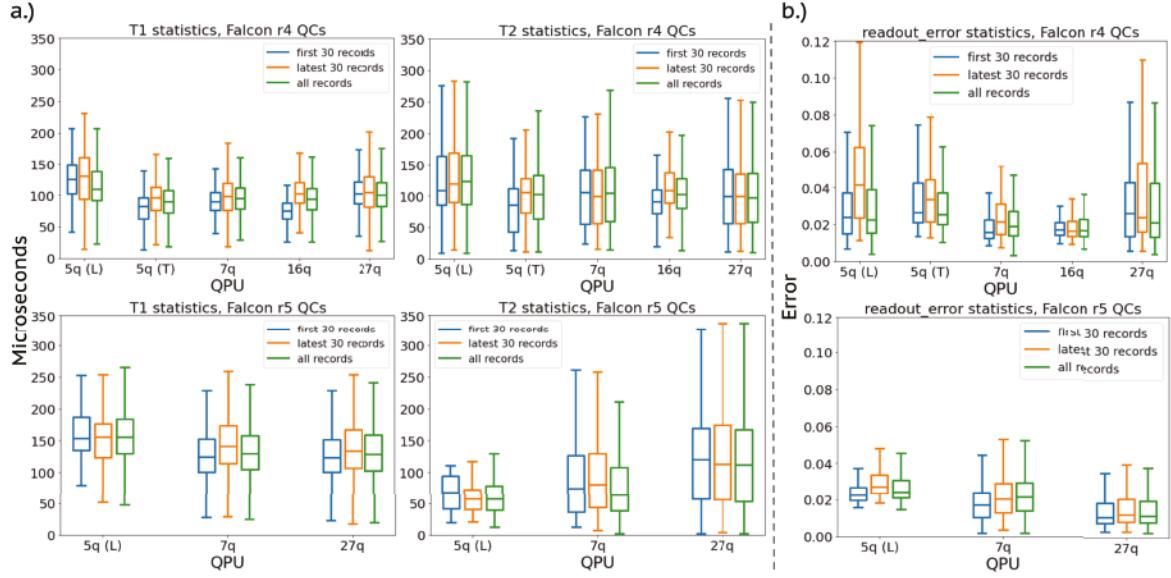


Fig. 3. Comparison of (a) coherence and (b) readout error for the Falcon r4 and r5 processors. Box-and-whisker plots show distributions for the first 30 records, the latest 30 records, and all records combined. Bars grouped by processor size (qubits) and segment.

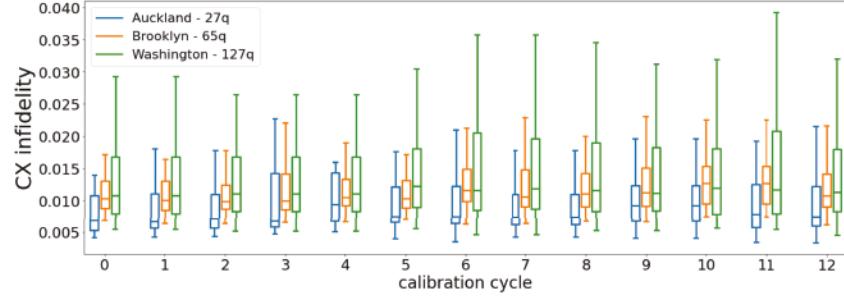


Fig. 4. Comparison of CX error statistics for three different types of IBM QC processors: Auckland (27-qubit Falcon), Brooklyn (65-qubit Hummingbird), and Washington (127-qubit Eagle). Plot includes the CX error distributions for the 12 latest calibration cycle records.

approximately 2.5 years. When analyzing Fig. 5, we observe that the recorded values for gate error are much more noisy and have a larger degree of variation as compared to frequency.

Fig. 5 demonstrates that qubit frequency does not fluctuate at the same extent as gate error. However, slight variation in frequency exists. As an example we analyze qubit frequency over device lifetime for the five-qubit IBM Bogota QC. Fig. 6 shows frequency vs. qubits for all of Bogota's qubits – it provides a better understanding of intra-chip frequency distribution over time. We see occasional spikes and shifts in qubit frequency, but overall, frequency for each qubit is stable. Variation for each qubit's frequency stays relatively close to the mean frequency value and significant changes are infrequent. It is important to note that all qubits also have adequate spacing between frequency values.

B. Inter-QC variation

Fig. 6 revealed the distinguishability and consistency of an individual QC's qubit frequencies over time. This moti-

vates the exploration of the extent that QC qubit frequencies vary among devices of similar size and performance. Fig. 7 features a box-and-whisker plot showing frequency vs. qubit for three, five-qubit IBM QCs: Bogota, Santiago, and Manila. These devices are all Falcon processors with an L-segment connectivity graph, illustrated in Fig. 2. One might conclude that all of these devices have the same operational characteristics because of their physical similarities, however, despite a common processor family and layout, the frequencies associated with each of the three Falcon processors varies drastically. Fig. 7 shows that frequency values continue to be constrained within a limited range and have definite values that are adequately spaced on each device. Further, we observe that the distributions of frequencies are different device-to-device with minimal overlap between QCs. This observation motivates us to target qubit frequency as a means to uniquely and consistently identify QCs.

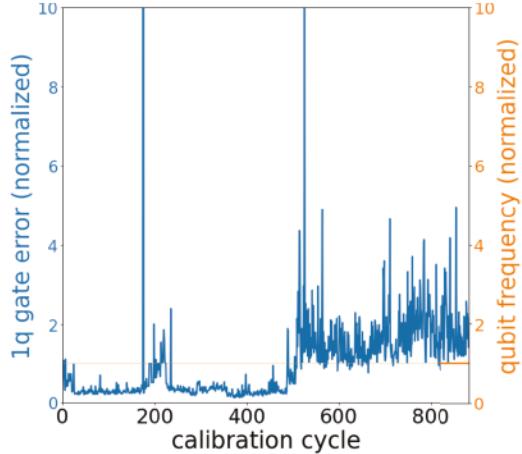


Fig. 5. Example of features unsuitable and suitable for fingerprinting QCs: 1-qubit gate error (left axis, unsuitable) and qubit frequency (right axis, suitable) across all records (normalized by their respective mean) for IBMQ Armonk QC, a 1-qubit Canary processor.

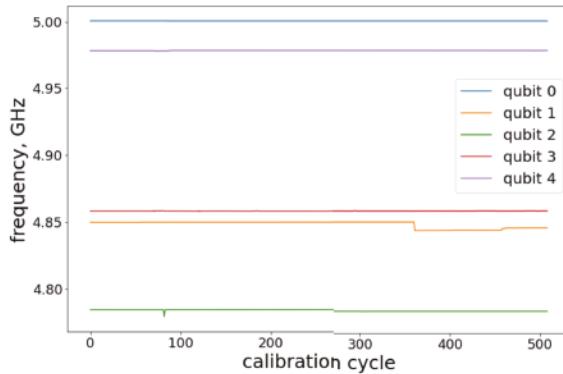


Fig. 6. Frequency vs. calibration cycles for the five-qubit Bogota IBM QC, showing the consistency of qubit frequency over many cycles.

C. Identifying Key Features for Fingerprinting

We continue our analysis of QC properties that contribute to a device's unique property signature to quantitatively discover features that are viable for QC fingerprinting. To this end, we compare frequencies, T_1 / T_2 coherence times, and readout errors between the qubits of the five-qubit devices Bogota, Manila and Santiago. Using the 100 most recent calibration records for a specific feature (i.e. frequency, T_1 , etc.), $\vec{x}_k = [x_{k,1} \ x_{k,2} \ \dots \ x_{k,100}]^\top$ of qubit k , we compute the scaled Euclidean distance between the records of two qubits i and j

$$d(\vec{x}_i, \vec{x}_j) = \frac{\|\vec{x}_i - \vec{x}_j\|_2}{\sqrt{100\Delta_{\max}}}. \quad (1)$$

Here, Δ_{\max} denotes the maximum variation for this feature over all qubits' 100 calibration cycles. Distances < 1 means that the qubits i and j are similar with respect to this feature, while distances $\gg 1$ indicate substantial dissimilarity.

Fig. 8 visualizes the feature distances for all qubits of all three machines, where the tick labels consist of the initial

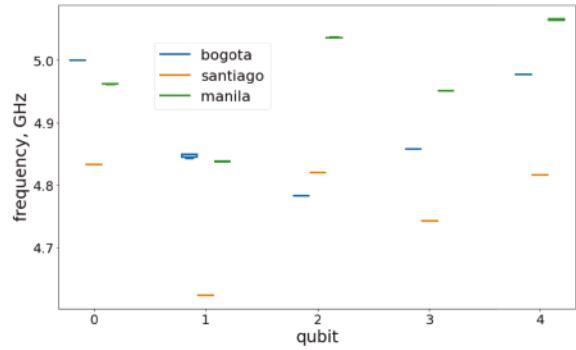


Fig. 7. Comparison of frequency distribution vs. qubit for the five-qubit IBM QCs Bogota, Santiago and Manila.

letters of the processors (B, M, S) and the qubit indices. In accordance with Fig. 7, we observe differences in qubit frequency intra- as well as inter-device as shown in Fig. 8(a). Red fields represent few instances of similarities, or collisions, between frequencies. Overall, this makes qubit frequency a strong candidate for a QC fingerprint. On the contrary, T_1 , T_2 , and readout errors overlap across different qubits, leading to low distances as can be seen in Fig. 8(b),(c). That means that these features do not vary more across qubits than they fluctuate over time for the same qubit, therefore they are less favorable properties for uniquely identifying QCs. This is in agreement with the distributions displayed in Fig. 3, where we note overlap of these features across different architectures. From this observation we conclude that the use of coherence times and readout error within a fingerprint would not be favorable as many instances of collision between identifiers would occur.

As a conclusion of this analysis, we identify frequency in fixed-frequency transmon systems as a potential resource for developing a QC fingerprint. Frequency is just one of many properties in a QC property signature that characterises a QC as a result of device variation. We investigate this special feature in more detail in the following section.

D. Frequency Vector as a Fingerprint

We define our frequency-based fingerprint of a QC as a vector of its qubit frequencies, $\vec{f} = [f_0 \ f_1 \ \dots \ f_{N-1}]^\top$, where f_k is the frequency of qubit k . Given the frequency vectors \vec{f}_i and \vec{f}_j of two N -qubit processors i and j , we use a normalized Hamming distance

$$\frac{\# \text{ frequencies that differ}}{\# \text{ qubits}} = \frac{|\{k : |f_{i,k} - f_{j,k}| > \Delta_{\text{avg}}\}|}{N} \quad (2)$$

to estimate how similar the processors' frequency characteristics are. Here we consider two frequencies to be different if they are more than a threshold, Δ_{avg} , apart.

To demonstrate the efficacy of qubit frequencies as a unique and stable fingerprint, we evaluate frequency fingerprints from eight 27-qubit QCs over 100 calibration cycles. The QCs used represent IBM's Falcon r4 and r5 processors ranging in age

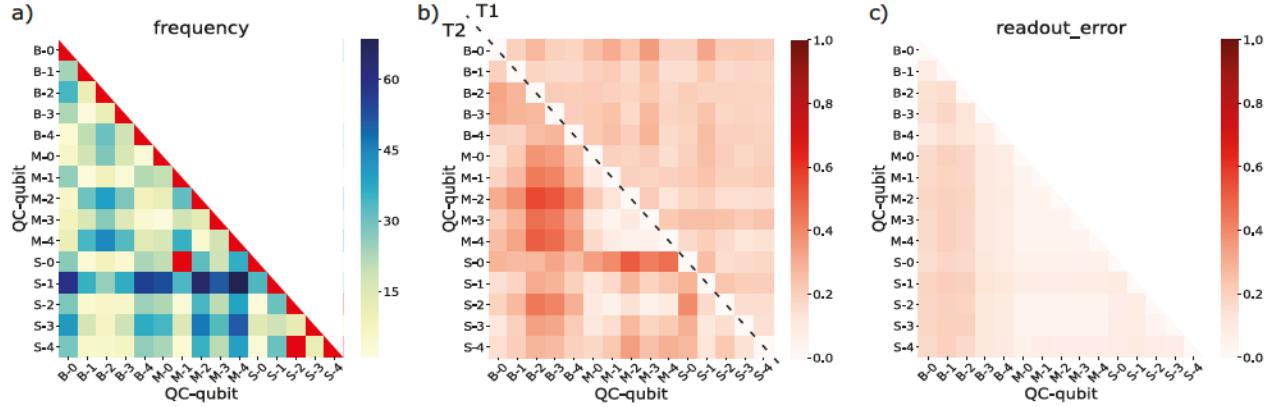


Fig. 8. Distances of qubit features across the five-qubit processors Bogota (B), Manila (M) and Santiago (S). Values < 1 correspond to similar attributes while values $\gg 1$ indicate dissimilarity. In all plots, darker color (higher value) means greater dissimilarity. (a) Qubits are fabricated from a wide range of frequencies, thus they generally differ significantly on the same device as well as across devices. Similar frequencies that dip below the lower bound of the legend are highlighted in red. (b) T_1 and T_2 times as well as (c) readout error fluctuate across qubits but not significantly more than they fluctuate for the same qubit over time. This leads to similarity (low distance) across devices, rendering those features less favorable for fingerprinting.

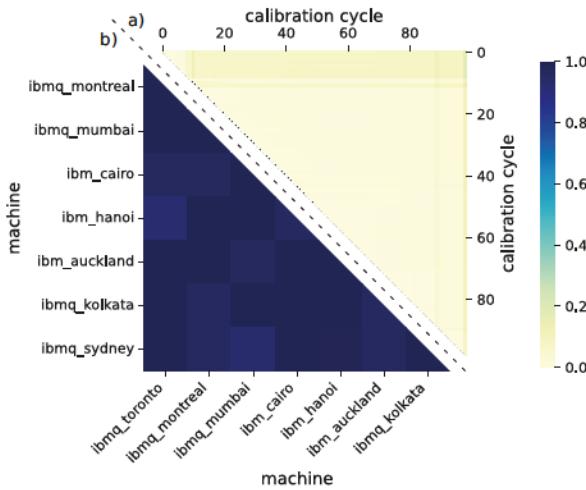


Fig. 9. Evaluation of frequency-based fingerprints on eight 27-qubit machines over 100 calibration cycles. (a) Normalized Hamming distances between calibration cycles for a given machine (values are averaged over all machines). The low observed distances indicate the fingerprint is stable for a given machine over time. (b) Distances between machines for a given calibration cycle (values are averaged over all cycles). The high observed distances indicate the fingerprints are unique between machines at a given time. All distances are calculated using Eq. (2).

from a couple months to over 2 years old. Fig. 9 shows the scaled Hamming distances between these fingerprints. The threshold, Δ_{avg} , was chosen as the maximum variation in a qubit's frequency across calibration cycles averaged over all machines and qubits.

The results in Fig. 9 support that a machine fingerprint can be constructed from qubit frequencies. The fingerprint is *stable*, there is low distance between fingerprints of the same machine, and *unique*, there is high distance between

fingerprints of different machines.

E. Considerations for a Frequency-based Fingerprint

Here, we discuss considerations for a frequency-based QC fingerprint.

Tuning of qubit frequency: Historical data analysis has shown the consistency of qubit frequency over time, however, there exist methods that purposefully adjust the frequency of fixed-frequency transmons. Laser annealing selectively tunes qubit frequency to bring it within a range that improves two-qubit interactions between nearest neighbors [37], [52]. Laser annealing has been applied to improve the yield of IBM QCs with average two-qubit gate performance of $\lesssim 1\%$. Since procedures such as laser annealing enable device improvements via changing critical fingerprint components, possibly causing authentication to fail, it must be assumed that fingerprint re-enrollment is necessary, especially after device maintenance. If qubit frequencies are altered after routine QC maintenance, a device will be considered an updated version of the QC with a “new” intrinsic fingerprint that incorporates the frequency changes. As a note, the laser annealing methods in [37] are imperfect, preventing annealing processes from achieving target frequencies with 100% accuracy. This means that creating two perfectly identical QCs via laser annealing is highly unlikely.

Implications of scaling: As QCs are manufactured with more qubits, it is likely that the majority of their frequency distributions will be unique. However, more populated distributions could result in qubit frequencies that nearly collide, causing confusion between qubits and errors in identification. In this case, it is essential to develop protocols that include enough information (i.e. enough qubit frequency permutations) to discern one device from another, even in the case of occasional qubit frequency overlap.

Access to low-level control: Qubit frequency is determined by sweeping a transmon qubit with discrete microwave control pulses that span a range of frequencies. The signal frequency that best resonates with the qubit is determined to be the qubit’s frequency.

Implementing frequency sweeps require low-level pulse control [53]. Currently, the IBM family of QCs allow pulse-level control. Further, qubit frequency is included in the QC properties that can be queried by an end-user that has access to the device. Thus, qubit frequency for a QC is easily determined, making the frequency fingerprint for IBM QCs in the quantum cloud open knowledge. Obscuring qubit frequency to discourage device fingerprinting is possible. For example, if low-level access to QCs was restricted, forcing users to write quantum programs at gate-level abstractions, qubit frequency would be challenging to pinpoint by attackers wishing to identify a QC. Alternatively, if the cloud provider wishes to allow low-level control for their machines while concealing frequency information, an QC antivirus [54] could be developed that is trained to detect and prevent frequency sweep operations that are initialized by an end-user.

F. Alternative Applications of the QC Fingerprint

Here we propose a frequency-based fingerprint based on the intrinsic properties of fixed-frequency transmon machines. Our considered threat model, Section II, proposed QC fingerprinting for malicious practices. However, secure cryptographic protocols along with a hardware fingerprint would enable QC authentication. Such a scheme could facilitate the identification of expensive quantum cloud resources and the verification machine performance. Additionally, a quantum hardware fingerprint could be critical in distributed quantum applications that demand on-the-fly quantum authentication, without which any potential quantum advantage would be lost.

In this potential quantum authentication scheme, we assume that qubit frequencies are concealed by the cloud provider. Further, user access to low-level pulse control or frequency sweep experiments is restricted. When qubit frequencies are concealed, challenge-response based authentication schemes could be developed that combine permutations of qubit frequencies with a one-way function. In our analysis, frequency has shown to be unique, reliable, and collision-resistant, making it a good candidate for use within a physically unclonable function (PUF) [55] that identifies cloud-based quantum compute resources. We leave the complete design of this QC authentication scheme for future work, including the translation of QC qubit frequency to challenge-response pairs.

VII. RELATED WORK

The stability of quantum machines over time is described [56]. In [56], the authors quantify the similarity of several NISQ devices by comparing gate fidelities, duty cycles, and register addressability across temporal and spatial scales. Like our study, [56] agrees with the modest body of work that notes the variability of current quantum machines.

Other prior works that discuss variation in QCs, either among individual devices or over time, include [48], [57]–[59]. We also note that a body of work called ‘quantum fingerprinting’ exists that serves the purpose of checking the equivalence of distributed database files (rather than our goal here QC hardware identification) [60].

Schemes have been developed that apply quantum tunneling for classical electronic identification [61]. Further, prior work exists that outlines proposals for PUFs targeted for quantum computing applications. A classical PUF based on the quantum readout of unclonable quantum states is found in [62]. The work [51] presents quantum hardware PUFs that generate bit strings with functions that employ qubit superposition and decoherence. The work in [63] and [64] present theory for quantum PUFs that establish a secure communication channel using quantum properties and operations that are secure against quantum cryptographic attacks. Finally, there is prior art in the space of quantum device identification based on crosstalk [65], and proposals for quantum antivirus exist [54] that aim to flag suspicious programs that inject malicious crosstalk and degrade the quality of program outcomes. Our work differs from these past efforts by proposing a frequency-based fingerprint for QC identification on the basis of uniqueness, reliability, and collision resistance. Our results derive from the historical data analysis of today’s fixed-frequency transmon QCs.

VIII. CONCLUSION

With QCs evolving into a cloud-provided commodity, the ability to fingerprint quantum hardware within the quantum cloud has serious security implications. In this paper, we show how QCs based on fixed-frequency transmon qubits can be fingerprinted. We use historical data from IBM QCs over the past 2+ years and introduce a frequency-based fingerprint that is stable, unique, and generalizable across different device generations. We argue that these successful features emerge from the variability inherent in the fabrication of today’s fixed-frequency transmon QCs.

ACKNOWLEDGMENT

This work is funded in part by EPiQC, an NSF Expedition in Computing, under award CCF-1730449; in part by STAQ under award NSF Phy-1818914; in part by NSF award 2110860; in part by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; in part by the NSF Quantum Leap Challenge Institute for Hybrid Quantum Architectures and Networks (NSF Award 2016136); and in part based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. FTC is Chief Scientist for Quantum Software at Infleqtion and an advisor to Quantum Circuits, Inc.

REFERENCES

[1] D. P. DiVincenzo, "The physical implementation of quantum computation," *Fortschritte der Physik: Progress of Physics*, vol. 48, no. 9-11, pp. 771–783, 2000.

[2] J. M. Gambetta, J. M. Chow, and M. Steffen, "Building logical qubits in a superconducting quantum computing system," *npj Quantum Information*, vol. 3, no. 1, pp. 1–7, 2017.

[3] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina *et al.*, "Fault-tolerant control of an error-corrected qubit," *Nature*, vol. 598, no. 7880, pp. 281–286, 2021.

[4] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.

[5] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.

[6] 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.

[7] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, "Quantum optimization using variational algorithms on near-term quantum devices," *Quantum Science and Technology*, vol. 3, no. 3, p. 030503, 2018.

[8] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.

[9] "IBM quantum breaks the 100-qubit processor barrier," <https://research.ibm.com/blog/127-qubit-quantum-processor-eagle>.

[10] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[11] I. Giechaskiel, K. B. Rasmussen, and J. Szefer, "Measuring long wire leakage with ring oscillators in cloud fpgas," in *2019 29th International Conference on Field Programmable Logic and Applications (FPL)*, 2019, pp. 45–50.

[12] S. Tian and J. Szefer, "Temporal thermal covert channels in cloud fpgas," in *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, ser. FPGA '19, New York, NY, USA: Association for Computing Machinery, 2019, p. 298–303. [Online]. Available: <https://doi.org/10.1145/3289602.3293920>

[13] S. Tian, W. Xiong, I. Giechaskiel, K. Rasmussen, and J. Szefer, "Fingerprinting cloud fpga infrastructures," in *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, ser. FPGA '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 58–64. [Online]. Available: <https://doi.org/10.1145/3373087.3375322>

[14] A. Schaller, W. Xiong, N. A. Anagnostopoulos, M. U. Saleem, S. Gabmeyer, S. Katzenbeisser, and J. Szefer, "Intrinsic rowhammer pufs: Leveraging the rowhammer effect for improved security," in *2017 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, 2017, pp. 1–7.

[15] A. Rahmati, M. Hicks, D. E. Holcomb, and K. Fu, "Probable cause: The deanonymizing effects of approximate dram," in *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, 2015, pp. 604–615.

[16] A. Schaller, W. Xiong, N. A. Anagnostopoulos, M. U. Saleem, S. Gabmeyer, B. Škorić, S. Katzenbeisser, and J. Szefer, "Decay-based dram pufs in commodity devices," *IEEE Transactions on Dependable and Secure Computing*, vol. 16, no. 3, pp. 462–475, 2019.

[17] S. Rosenblatt, D. Fainstein, A. Cestero, J. Safran, N. Robson, T. Kirihata, and S. S. Iyer, "Field tolerant dynamic intrinsic chip id using 32 nm high-k/metal gate soi embedded dram," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 940–947, 2013.

[18] S. Tian, W. Xiong, I. Giechaskiel, K. Rasmussen, and J. Szefer, "Fingerprinting cloud fpga infrastructures," in *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, ser. FPGA, February 2020.

[19] "IBM Quantum," <https://quantum-computing.ibm.com>, 2022.

[20] "IBM quantum systems," <https://quantum-computing.ibm.com/services?systems=all>, 2022.

[21] A. Lenk, M. Menzel, J. Lipsky, S. Tai, and P. Offermann, "What are you paying for? performance benchmarking for infrastructure-as-a-service offerings," in *2011 IEEE 4th International Conference on Cloud Computing*. IEEE, 2011, pp. 484–491.

[22] I. Baldini, P. Castro, K. Chang, P. Cheng, S. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. Rabbah, A. Slominski *et al.*, "Serverless computing: Current trends and open problems," in *Research advances in cloud computing*. Springer, 2017, pp. 1–20.

[23] *Current Status and Next in Quantum Computing*. [Online]. Available: <https://www.eetimes.eu/current-status-and-next-in-quantum-computing/>

[24] 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 *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 987–999. [Online]. Available: <https://doi.org/10.1145/3297858.3304007>

[25] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari, "Software mitigation of crosstalk on noisy intermediate-scale quantum computers," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 1001–1016.

[26] J. Majer, J. Chow, J. Gambetta, J. Koch, B. Johnson, J. Schreier, L. Frunzio, D. Schuster, A. A. Houck, A. Wallraff *et al.*, "Coupling superconducting qubits via a cavity bus," *Nature*, vol. 449, no. 7161, pp. 443–447, 2007.

[27] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. Borgorin, M. Brink, L. Capelluto, O. Gunluk, T. Itoko, N. Kanazawa *et al.*, "Demonstration of quantum volume 64 on a superconducting quantum computing system," *Quantum Science and Technology*, 2021.

[28] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, "Superconducting qubits: Current state of play," *Annual Review of Condensed Matter Physics*, vol. 11, pp. 369–395, 2020.

[29] P. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding *et al.*, "Scalable quantum simulation of molecular energies," *Physical Review X*, vol. 6, no. 3, p. 031007, 2016.

[30] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Realization of three-qubit quantum error correction with superconducting circuits," *Nature*, vol. 482, no. 7385, pp. 382–385, 2012.

[31] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. Girvin, L. Jiang *et al.*, "Extending the lifetime of a quantum bit with error correction in superconducting circuits," *Nature*, vol. 536, no. 7617, pp. 441–445, 2016.

[32] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. Song, C.-L. Zou *et al.*, "Quantum error correction and universal gate set operation on a binomial bosonic logical qubit," *Nature Physics*, vol. 15, no. 5, pp. 503–508, 2019.

[33] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.

[34] A. Nersisyan, S. Poletto, N. Alidoust, R. Manenti, R. Renzas, C.-V. Bui, K. Vu, T. Whyland, Y. Mohan, E. A. Sete *et al.*, "Manufacturing low dissipation superconducting quantum processors," in *2019 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 2019, pp. 31–1.

[35] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman *et al.*, "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds," *Nature communications*, vol. 12, no. 1, pp. 1–6, 2021.

[36] M. A. Vrajitoarea, "Strongly correlated photonic materials: parametric interactions and ultrastrong coupling in circuit qed," Ph.D. dissertation, Princeton University, 2020.

[37] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, E. Magesan, J. A. Smolin, J.-B. Yau, V. P. Adiga, M. Sandberg, M. Brink, J. M. Chow *et al.*, "Laser-annealing josephson junctions for yielding scaled-up superconducting quantum processors," *npj Quantum Information*, vol. 7, no. 1, pp. 1–8, 2021.

[38] J. Kreikebaum, K. O'Brien, A. Morvan, and I. Siddiqi, "Improving wafer-scale josephson junction resistance variation in superconducting

quantum coherent circuits," *Superconductor Science and Technology*, vol. 33, no. 6, p. 06LT02, 2020.

[39] C. Müller, J. H. Cole, and J. Lisenfeld, "Towards understanding two-level-systems in amorphous solids: insights from quantum circuits," *Reports on Progress in Physics*, vol. 82, no. 12, p. 124501, 2019.

[40] J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niegce, M. Kudra, P. Delsing, and J. Bylander, "Decoherence benchmarking of superconducting qubits. npj quantum inf. 5," 2019.

[41] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, "Correlating decoherence in transmon qubits: Low frequency noise by single fluctuators," *Physical review letters*, vol. 123, no. 19, p. 190502, 2019.

[42] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and J. M. Martinis, "Fluctuations of energy-relaxation times in superconducting qubits," *Phys. Rev. Lett.*, vol. 121, p. 090502, Aug 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.121.090502>

[43] "The ibm quantum heavy hex lattice," <https://www.research.ibm.com/blog/heavy-hex-lattice>, 2021.

[44] J. Gambetta, "Ibm's roadmap for scaling quantum technology," <https://research.ibm.com/blog/ibm-quantum-roadmap>, 2021, accessed: 2021-09-29.

[45] "IBM Quantum systems - retired systems," <https://quantum-computing.ibm.com/lab/docs/iql/manage/systems/retired-systems>, 2022.

[46] "Qiskit/ibmq-device-information," <https://github.com/Qiskit/ibmq-device-information>, 2019.

[47] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, "Topological and subsystem codes on low-degree graphs with flag qubits," *Physical Review X*, vol. 10, no. 1, p. 011022, 2020.

[48] 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 *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 1015–1029.

[49] "System properties," <https://quantum-computing.ibm.com/lab/docs/iql/manage/systems/properties>, 2022.

[50] "Ibm quantum processor types," <https://quantum-computing.ibm.com/composer/docs/iqx/manage/systems/processors>, 2022.

[51] K. Phalak, A. Ash-Saki, M. Alam, R. O. Topaloglu, and S. Ghosh, 2019, pp. 987–999.

[52] 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 *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 987–999.

[53] "Quantum puf for security and trust in quantum computing," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 11, no. 2, pp. 333–342, 2019.

[54] E. J. Zhang, S. Srinivasan, N. Sundaresan, D. F. Bogorin, Y. Martin, J. B. Hertzberg, J. Timmerwilde, E. J. Pritchett, J.-B. Yau, C. Wang *et al.*, "High-fidelity superconducting quantum processors via laser-annealing of transmon qubits," *arXiv preprint arXiv:2012.08475*, 2020.

[55] T. Alexander, N. Kanazawa, D. J. Egger, L. Capelluto, C. J. Wood, A. Javadi-Abhari, and D. C. McKay, "Qiskit pulse: programming quantum computers through the cloud with pulses," *Quantum Science and Technology*, vol. 5, no. 4, p. 044006, Aug 2020. [Online]. Available: <http://dx.doi.org/10.1088/2058-9565/aba404>

[56] S. Deshpande, C. Xu, T. Trochatis, Y. Ding, and J. Szefer, "Towards an antivirus for quantum computers," *arXiv preprint arXiv:2203.02649*, 2022.

[57] R. Maes and I. Verbauwhede, "Physically unclonable functions: A study on the state of the art and future research directions," *Towards Hardware-Intrinsic Security*, pp. 3–37, 2010.

[58] S. Dasgupta and T. S. Humble, "Stability of noisy quantum computing devices," *arXiv preprint arXiv:2105.09472*, 2021.

[59] S. S. Tannu and M. Qureshi, "Ensemble of diverse mappings: Improving reliability of quantum computers by orchestrating dissimilar mistakes," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2019, pp. 253–265.

[60] G. S. Ravi, K. N. Smith, P. Murali, and F. T. Chong, "Adaptive job and resource management for the growing quantum cloud," in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2021, pp. 301–312.

[61] P. Gokhale, E. Anschuetz, C. Campbell, F. Chong, E. Dahl, P. Frederick, E. Jones, B. Hall, S. Issa, P. Goiporia *et al.*, "Supercheq: Quantum advantage for distributed databases," *arXiv preprint arXiv:2212.03850*, 2022.

[62] J. Roberts, I. E. Bagci, M. Zawawi, J. Sexton, N. Hulbert, Y. Noori, M. Young, C. Woodhead, M. Missous, M. Migliorato *et al.*, "Using quantum confinement to uniquely identify devices," *Scientific reports*, vol. 5, no. 1, pp. 1–8, 2015.

[63] B. Škorić, "Quantum readout of physical unclonable functions," *International Journal of Quantum Information*, vol. 10, no. 01, p. 1250001, 2012.

[64] M. Arapinis, M. Delavar, M. Doosti, and E. Kashefi, "Quantum physical unclonable functions: Possibilities and impossibilities," *Quantum*, vol. 5, p. 475, 2021.

[65] M. Doosti, N. Kumar, M. Delavar, and E. Kashefi, "Client-server identification protocols with quantum puf," *ACM Transactions on Quantum Computing*, vol. 2, no. 3, pp. 1–40, 2021.

[66] M. Allen, D. Shuwen, and S. Jakub, "Short paper: Device-and locality-specific fingerprinting of shared nisq quantum computers," in *Workshop on Hardware and Architectural Support for Security and Privacy*, 2021, pp. 1–6.