



Long-Term Analysis of the Dependability of Cloud-based NISQ Quantum Computers

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

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

ABSTRACT

Numerous public cloud infrastructure providers today allow for access to Noisy Intermediate-Scale Quantum (NISQ) computers. Changes in the environment or the machine configuration may affect their dependability. Through analysis of real quantum computer calibration data, this work demonstrates that quantum computers available from IBM Quantum experience periods of fluctuation or abrupt qubit frequency changes. This work further analyzes the correlation between the frequency change events, decoherence times, gate errors, and machine maintenance or offline periods. The results highlight that the properties of NISQ computers change over time, affecting their dependability, but not all of the changes can be explained with publicly available data.

CCS CONCEPTS

• **Hardware** → **Quantum technologies.**

KEYWORDS

cloud-based quantum computing, dependability

ACM Reference Format:

Chuanqi Xu and Jakub Szefer. 2023. Long-Term Analysis of the Dependability of Cloud-based NISQ Quantum Computers. In *The 18th International Conference on Availability, Reliability and Security (ARES 2023), August 29–September 01, 2023, Benevento, Italy*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3600160.3600192>

1 INTRODUCTION

Noisy Intermediate-Scale Quantum (NISQ) quantum computers are being rapidly developed, with over 400-qubit machines available today [7], and the industry projecting 4000-qubit or larger devices before the end of the decade [14]. So far, they are too small for quantum error correction, but already have promising applications in optimization, chemistry, and other important areas [8, 10, 11].

Among the different types of technologies, NISQ computers based on superconducting qubit technology are developed by numerous companies, such as IBM, Rigetti, or Quantum Circuits, Inc. These machines implement quantum computing with superconducting electronic circuits which are operated at about 20mK temperatures. Even with extreme cooling, the devices are sensitive to

different types of environmental disturbance, as well as noise in the control within the quantum computer itself. Changes due to the environment, or changes within the devices themselves, can affect the dependability and operation of the NISQ computers.

In order to make these NISQ quantum computers accessible to users, a growing number of companies are deploying them as cloud-based accelerators. Public cloud infrastructures today already allow for easy, on-demand access to a range of quantum computers from a variety of manufacturers. Cloud-based services such as IBM Quantum [6], Amazon Bracket [1], and Microsoft Azure [12] are among the ones providing access to superconducting quantum computers remotely for users. While users can gain easy access to the devices thanks to cloud-based services, they have no control over the physical devices and their environment. Understanding the operation of the hardware, and any physical or environmental changes in the hardware is essential for users who want to use them.

Through analysis of real quantum computer historical calibration data from IBM Quantum [6], we demonstrate that computers available from IBM Quantum experience numerous events, where there are abrupt changes in the qubit frequencies, or where the frequencies fluctuate for periods of time, for example. Changes to the frequencies themselves are not critical. However, they can be used as indicators of environmental or physical changes to the machines. These changes can in turn correlate to changes in other properties such as gate errors. By tracking the changes in qubit frequencies, the users can use significant frequency change events as triggers for the re-optimization of their algorithms. In addition, it is known that superconducting qubit devices are sensitive to temperature changes. Frequency changes could then indicate the thermal cycling of the cryogenic refrigerators used to hold the superconducting qubit devices. As a result, tracking of the qubit frequency changes can be used to detect physical changes or tampering with the machines. Therefore, one of the key takeaways of our study is that, to fully characterize the behavior of quantum computers, users need to understand and track properties beyond the traditional metrics, such as qubit decoherence times and gate error rates.

Further, through historical device data, we analyze how device offline periods are correlated among different devices and how they are correlated to frequency changes. We discover some correlation where a number of devices tend to be offline at the same time, indicating they may share cooling, control, or other infrastructure. Shared infrastructure could be a potential failure point, and users looking for reliable execution of their programs on NISQ quantum computers may want to avoid using machines that may share all or some of the same infrastructure.

On the other hand, we also discover that many frequency change events are not correlated to periods when the devices may have been offline, indicating unknown and unpublished reasons for the



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

ARES 2023, August 29–September 01, 2023, Benevento, Italy
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0772-8/23/08.
<https://doi.org/10.1145/3600160.3600192>

frequency changes. This highlights that maintenance or firmware update periods may not be recovered from the historical data about the physical properties of the qubits. Yet, any changes to the control hardware or software will have an important impact on the dependability of NISQ quantum computers.

The analysis in this paper highlights frequent changes in real, cloud-based NISQ quantum computer properties, while at the same time, there is a lack of information from the vendors about the sources of the changes possibly raising questions about the accuracy of the computation and dependability of the machines if there are frequent changes in their properties.

2 BACKGROUND

In this section, we provide background on cloud-based NISQ quantum computers available from IBM Quantum. We also discuss calibration data and how physical changes can affect the properties of the cloud-based NISQ devices.

2.1 Cloud-based NISQ Quantum Computers

NISQ computers nowadays available from IBM Quantum vary in size from 5 qubits to 127 qubits. IBM has also announced a 433 qubit machine in November 2022, as well as a projection for 4000-qubit or larger devices by the end of the decade. These machines are available for researchers and also the general public. In addition to free access for research, pay-as-you-go access has been made available, with users paying 1.60 USD per runtime second with a credit card or IBM Cloud credits. As a reference, a 2-qubit Grover’s search algorithm [4] requires about $2\mu\text{s}$ per shot, with a usual 4096 shots per experiment the cost would be 0.013 USD, not considering any time for initialization, thermalization, and finalization of the experiment.

Most algorithms require thousands of shots to execute. Each shot is one execution of the algorithm, and the outputs of all the shots are collected to compute the final output probabilities. The output probabilities can be the final answer of the algorithm, as is the case for Grover’s search. Or the output probabilities can be used to drive optimization of the algorithm and its future iterations, as is the case for VQE [13] and quantum machine learning [2]. Any changes in the environment or operational errors of the NISQ computer will impact the output probabilities of the algorithms.

2.2 Historical Calibration Data

IBM Quantum continuously calibrates the properties of their devices to offer insight into current processor states. The properties of a quantum device are dynamical quantities typically measured and calibrated once in a 24-hour period and are updated once the calibration sequence is complete. These properties include T1 (relaxation time), T2 (dephasing time), qubit frequency, anharmonicity, readout error, single-qubit gate error, and so on. There are typically two ways to retrieve historical data. The first one is to manually look up the IBM Quantum Lab website, and the second is to use Qiskit APIs. In general, the data is the same by either means, however, we would like to mention that some properties, such as T1 or T2, may be updated more frequently on the IBM Quantum Lab website (around 2 hours on many devices) than Qiskit APIs (usually at least one day), while other properties keep the same update frequency.

Though the IBM Quantum Lab website may have more calibration data, users can only obtain the most recent calibration data from it. Therefore, in this paper, historical data is collected automatically through Qiskit APIs.

An interesting pattern of the historical calibration data is that there may be periods where no updated calibration data is provided and duplicated data is used instead. While the documentation does not provide details of why this happens, we assume that the calibration sequence did not execute during the days containing identical calibration data. Given that these calibrations are automatically performed, we project that the machine, or at least the control logic used for calibration, is offline or disabled during these periods.

2.3 Qubit Frequency Changes

One important feature of the superconducting qubits is their frequencies and changes in the frequencies. Because of its lengthy coherence and noise tolerance, fixed-frequency transmons are appealing in the current architecture of superconducting quantum computers. Most of the machines have qubits that have a rather unique and well-proposed set of frequencies. Nevertheless, in fixed-frequency superconducting machines, such as from IBM Quantum, the resonant frequencies are determined by the manufacture of the qubits, and they cannot be set at run time. In addition, even though the transmon size can be adjusted for a target frequency, manufacturing variation means that the final frequencies are not identical to the target frequency. Besides, frequency crowding needs to be tackled as superconducting quantum processors scale up [3].

The frequencies could be changed with physical alteration to the qubits. To overcome these problems, laser-annealing [5, 9, 15], a recently developed method, can effectively adjust the frequency of superconducting qubits post-fabrication. This method has already been implemented on recently publicized IBM’s Falcon processors [15]. Laser annealing could be used to lower the resonant frequency, but not raise it. This is an example of intentional frequency change, e.g., done after manufacturing to adjust the frequency. An unintentional frequency change can occur when the qubit device is brought to cryogenic temperature and then warmed up again. This can happen, for example, during any maintenance that requires the cryogenic refrigerator to be warmed up and opened up. The qubit frequencies can go up or down, depending on changes to the wiring or the attenuators in the cryogenic refrigerator.

3 ANALYSIS OF IBM CLOUD-BASED QUANTUM COMPUTERS

We analyze nine quantum computers available from IBM Quantum. If not specifically stated, all dates and times in this paper are in Eastern Standard Time (EST). The machines have been accessible for varying periods of time. The online dates, i.e. dates when the machines were first made available to the public, come in five sets:

- (1) 2022/03/25: `ibm_oslo`.
- (2) 2021/07/22: `ibm_perth`.
- (3) 2021/05/20: `ibm_lagos` and `ibm_nairobi`.
- (4) 2021/04/28: `ibmq_jakarta` and `ibmq_manila`.
- (5) 2021/01/08: `ibmq_quito`, `ibmq_belem`, and `ibmq_lima`.

Using historical data from these machines, we analyze periods of fluctuation or abrupt qubit frequency changes, as well as, examine

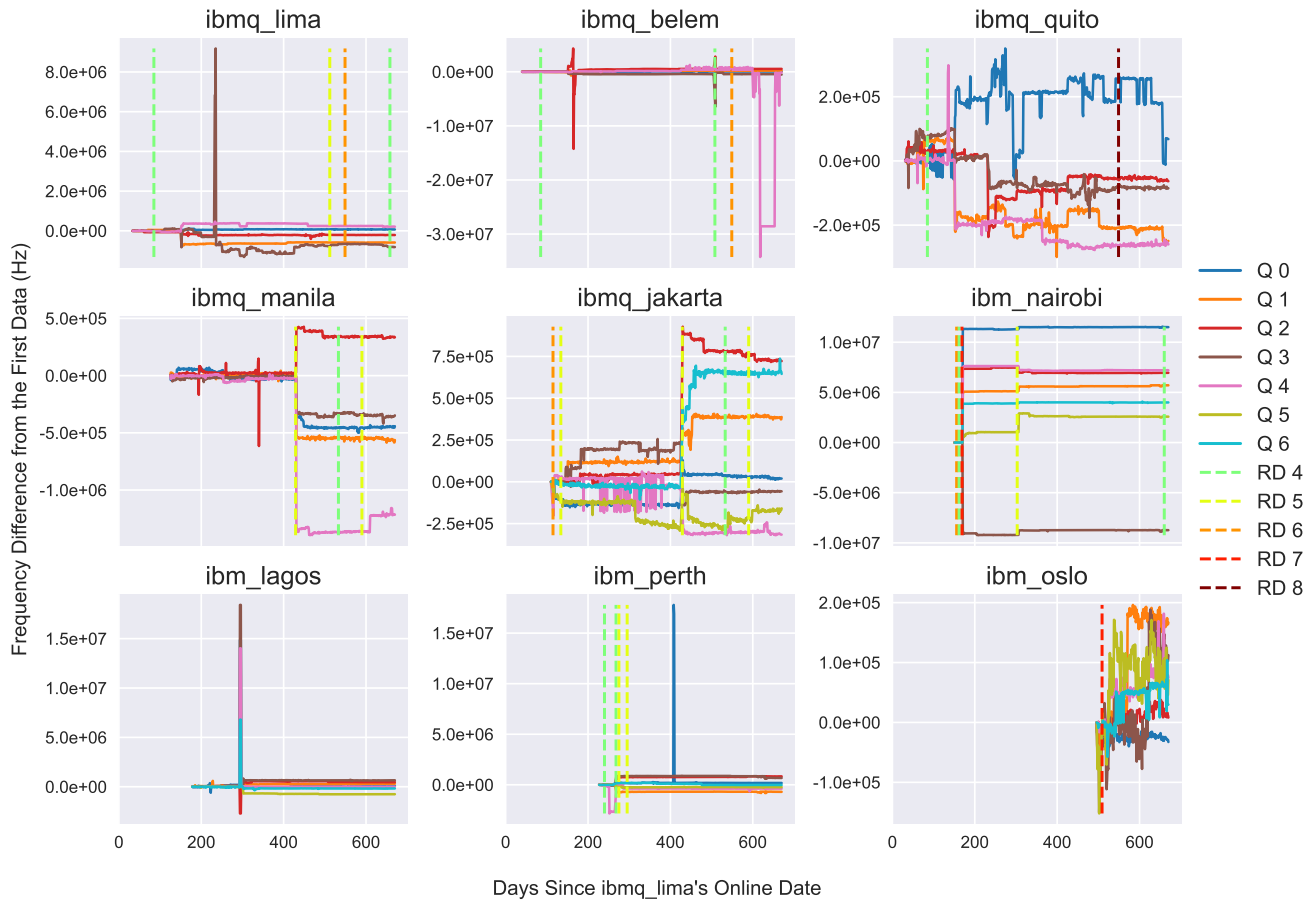


Figure 1: Frequency change over time from the first frequency calibration data on nine IBM quantum devices since `ibmq_lima`'s online date. Solid lines represent the frequency difference from the value on the first calibration date on each device and qubit (qubit i is shown as Q_i in the legend). Dashed lines are placed on the last date for the repeated calibration data, and their colors represent the number of repeated dates (RD).

any correlation between the frequency change events, decoherence times, gate errors, and machine maintenance or offline periods. The results highlight that the properties of NISQ computers change over time affecting their dependability, but not all of them can be explained with publicly available data.

3.1 Data Collection Setup

IBM Quantum provides access to data about qubit frequency, anharmonicity, T1 time, T2 time, readout error, and CX gate error. As mentioned in Section 2.2, some of the properties may be calibrated as frequently as about every 2 hours on the IBM Quantum Lab website, however, through Qiskit APIs, the historical data is updated at least every day. The data was collected with Qiskit APIs in a Python script. The data is from the first date when calibration data was provided until 2022/11/11. The data was then locally collected and analyzed.

3.2 Repeated Entries in Calibration Data

An interesting pattern of the historical data is the repeated entries. During certain periods, the same calibration data is provided for

one or more consecutive days. While there is no official explanation, our surmise is that new calibration data is not collected these days, and old data is simply copied over. Thus days with repeated calibration data may represent times when the machines or some of their control equipment is offline or unavailable. In addition, the data of `cx_error` (CX gate error) is simply populated with a "1" for certain days, here again, the conjecture is that these are days when the machines or some of their control equipment is offline or unavailable so the system fills in an (unrealistic) default value.

3.3 Analysis of Qubit Frequency Changes

In this section, we used the historical calibration data from IBM to explore how the frequencies change over time. Figure 1 shows the qubit frequency changes for nine quantum computers. `ibmq_lima`, `ibmq_belem`, `ibmq_quito`, and `ibmq_manila` are 5-qubit machines, while others are 7-qubit machines. The vertical dashed lines represent the last date when repeated historical data or qubit frequencies are provided, and its color indicates how many days the data is repeated. The dashed lines are used to help visualize any correlations between repeated historical data and frequency changes.

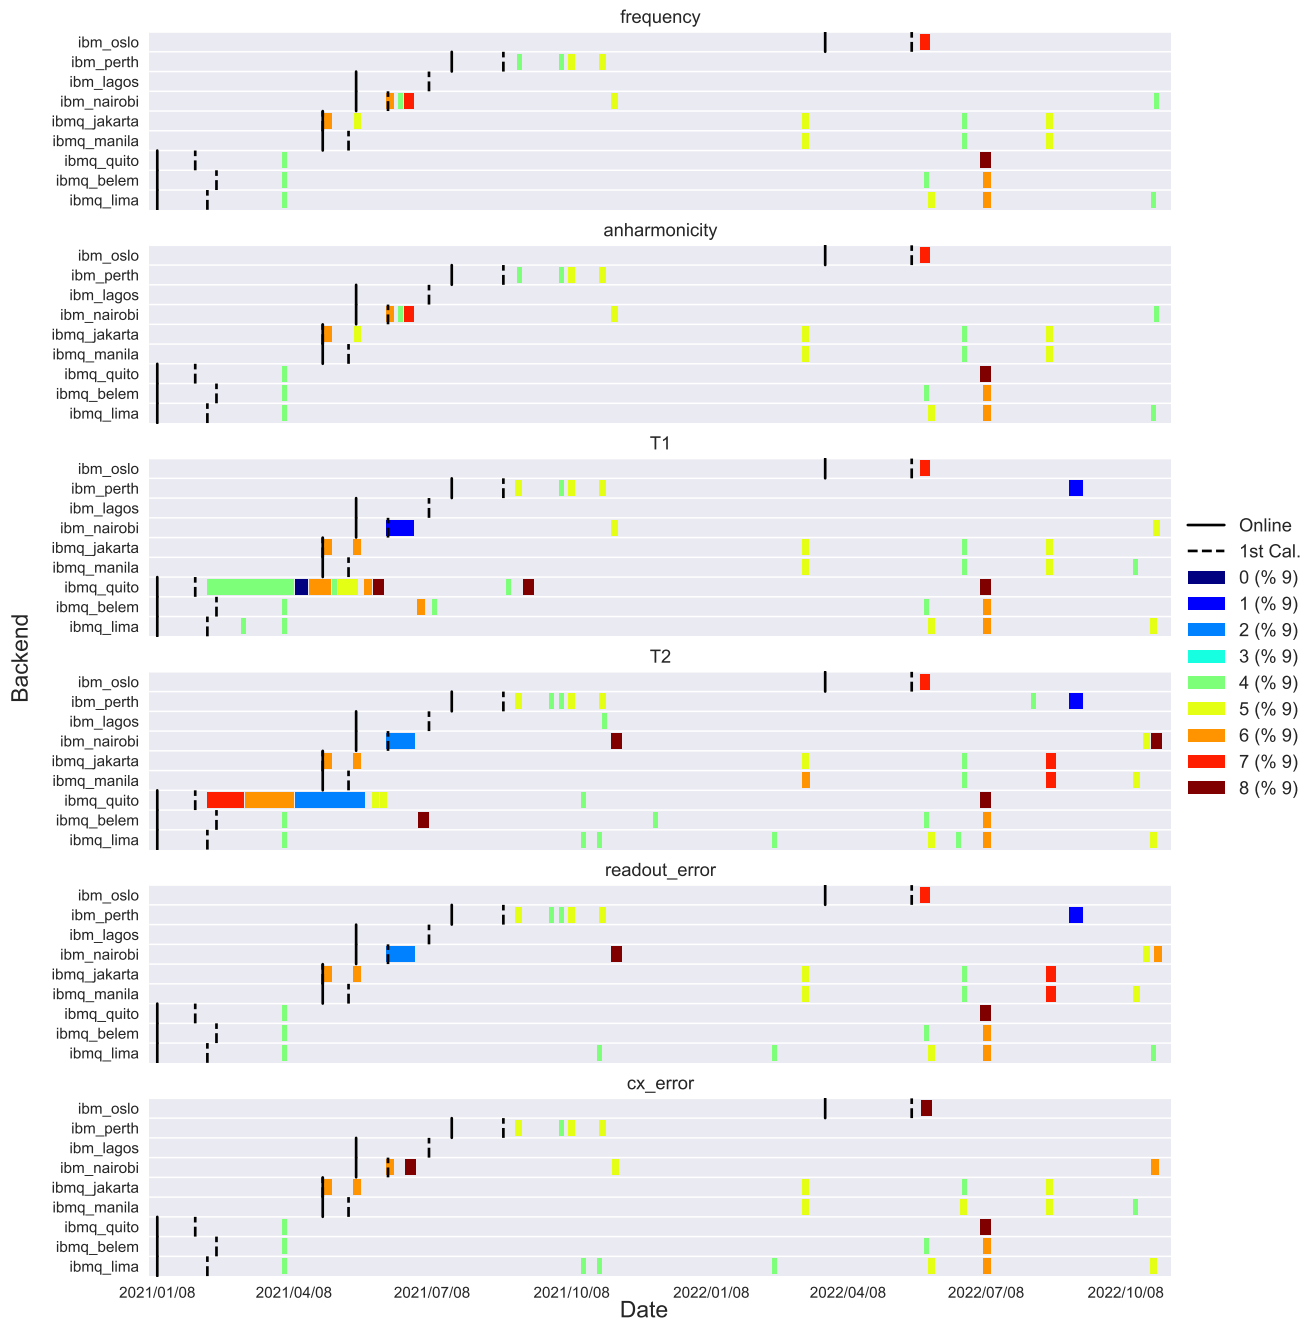


Figure 2: Repeated historical calibration data for different IBM machines. Labels show the number of days the same calibration data is used modulo 9, if the same data is repeated for n days, its label corresponds to $n \bmod 9$. Black solid lines represent machines’ online dates, and black dashed lines indicate the first dates that have the reported historical calibration data.

Examining the historical data, we can observe that frequencies are usually fluctuating in the range of 0.5 MHz. Some of the devices have a peak in their historical data, such as the peaks in *ibmq_lima*, *ibmq_lagos*, and *ibmq_perth*. When we zoomed in and carefully checked these data, we found that these data may only last one or two days (also there are no repeated dates dashed line on these peaks), and thus we assume that these peaks may be due to the

calibration errors rather than a long-lasting changes in the environment or devices. Among the machines analyzed, *ibmq_manila*, *ibmq_jakarta*, and *ibmq_nairobi* experience the largest changes. Interestingly, *ibmq_manila* and *ibmq_jakarta* experience large frequency change at the same time, around day 400 since IBM Lima’s online date, as shown in Figure 1. It should be noted these two machines came online at the same time. Therefore, there was

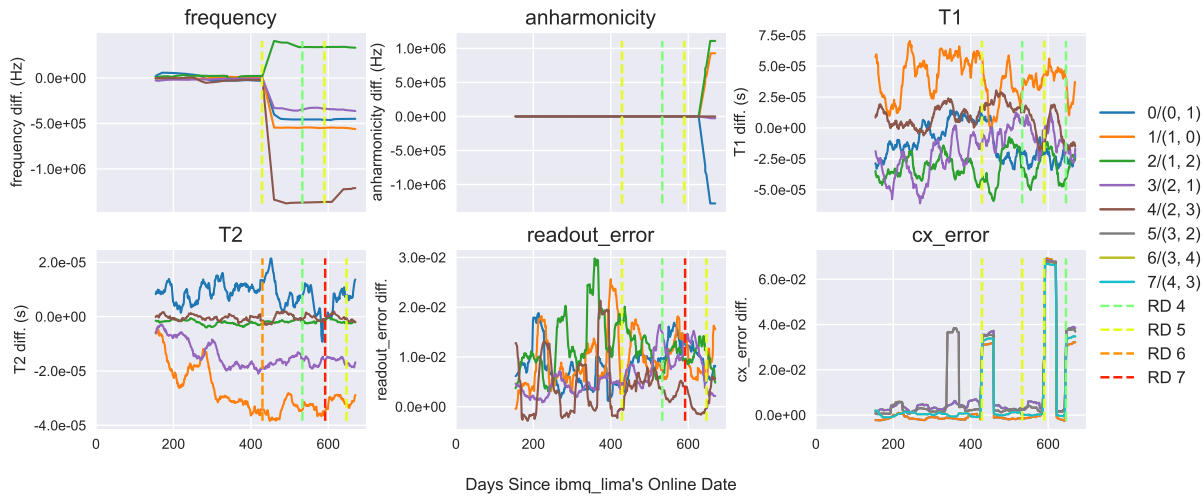


Figure 3: Properties of `ibmq_manila` over time from the online date of `ibmq_lima`. The data points show the value differences from the first calibration date and then smoothed with the 30-day moving average. Labels indicate qubit pairs for `cx_error` and qubit indices for other properties. Dashed lines represent repeated dates.

a correlated change in the frequencies. Separately, `ibmq_nairobi` experienced a large frequency change around day 180 since IBM Lima’s online date. Both `ibmq_nairobi` and `ibmq_lagos` came online at the same time, but here there is no correlation between them, as `ibmq_lagos` does not have a large frequency change except the peaks that may be due to calibration errors as discussed previously. The large changes in `ibmq_manila`, `ibmq_jakarta`, and `ibmq_nairobi` seem to correlate with periods when calibration data is repeated (see vertical dashed lines Figure 1).

While most of the smaller frequency changes may be stochastic, we assume that the larger frequency changes represent unspecified physical events. This necessitates NISQ quantum computer users to rely on frequent calibration data to ensure their algorithms run correctly. Assumptions that qubit frequency changes are only stochastic over time should not be made, considering the historical data we have analyzed.

It is not clear, nor specified, what the events causing the change may be. In some cases, as for `ibmq_manila`, `ibmq_jakarta`, and `ibmq_nairobi`, the frequency changes correlate slightly with the dates when there is repeated calibration data. This may or may not indicate that the machines were physically offline at this time, and that afterward, the qubit frequencies changed due to warm up and cool-down of the machines. To aid algorithm designers and users, the abrupt frequency changes should be better documented.

3.4 Analysis of the Repeated Historical Calibration Data

Apart from qubit frequency, IBM Quantum provides access to data about anharmonicity, T1, T2, readout error, and CX gate error. This data is stored on a per-day basis in the historical records, which can be retrieved for analysis. Figure 2 shows the periods when the repeated historical data is reported for each type of property. In the figure, black solid lines represent machines’ online dates, and black dashed lines indicate the first dates that report historical calibration

data. The colored blocks represent the periods when repeated data is reported. Each color represents the duration of the period.

In Figure 2, we can see some correlations between the machines that came online at the same time. For example, `ibmq_lima`, `ibmq_belem`, and `ibmq_qui` to come online 2021/01/08, and for most of the properties, they have repeated data on the same days. This may indicate that the computers share control or other logic. Shared control or other logic may affect the dependability of the machines.

The dependability of the computers may be affected (and hopefully improved) by firmware updates. To analyze this, we examined data during the recent firmware update and offline periods announced by IBM. At the start of Nov. 2022, a number of machines were announced to undergo firmware updates:

- (1) 2022/11/01: `ibmq_perth`, `ibmq_canberra`, and `ibmq_auckland`.
- (2) 2022/11/02: `ibmq_oslo` and `ibmq_geneva`.
- (3) 2022/11/03: `ibmq_washington`.
- (4) 2022/11/07: `ibmq_algiers`, `ibmq_cairo`, `ibmq_nairobi`, and `ibmq_lagos`.

Analyzing the repeated calibration data did not reveal however any patterns. It can be then assumed that the machines did not fully go offline, but the firmware was updated during their online time. Without noticeable changes in the machines’ properties, the firmware update may not have affected the control logic that may impact the properties reported in the historical data.

In addition, we found in Figure 1 that except for `ibmq_nairobi`, repeated dates in frequency calibration data make the smallest repeated dates set, i.e. the repeated dates of the frequency are also the interception of repeated dates for all properties. Though one 4-day repeated data at the start of `ibmq_nairobi` calibration process is an exception and does not exist in `cx_error`, repeated dates of frequency data can represent nearly all of the commonly repeated dates for the other quantum device properties.

3.5 Analysis of Other Property Changes

Figure 3 shows the frequency, anharmonicity, T1, T2, readout error and CX gate errors for `ibmq_lima`. The frequency changes are the same as in Figure 1. Dashed vertical lines again show the repeated calibration data periods. Generally, there is a weak correlation between repeated calibration data and property changes. There is some correlation for the frequency, seeing abrupt frequency changes around day 400. Also, there is a strong correlation for CX gate errors with the repeated calibration data dates around day 420 and day 580 and 650. However, the large CX gate error change around day 350 does not have a corresponding repeated calibration data event. It is possible that with more information, repeated calibration data events could be used to forecast changes, such as in CX gate errors. However, similar to our conclusion about qubit frequency data, to aid algorithm designers in better understanding and anticipating hardware fluctuations, the abrupt changes in the other properties should be better documented.

4 DISCUSSION

In addition to dependability issues, we have also considered the security implications of the observations made from the data.

One motivation to analyze the security implications is to understand if attackers could learn any information about the physical setup or physical events from the data. For example, a frequency change can occur when the qubit device is warmed up and then brought to cryogenic temperature again. This can happen, for example, during any maintenance that requires the cryogenic refrigerator to be warmed up and opened up. Interestingly, we did not see any correlation between the offline data and very large frequency changes. Our intuition is that qubit frequency changes over a few MHz could indicate device warm-up. However, such frequency changes were not observed, nor did we observe long-term offline periods likely needed to warm up and cool back down the superconducting machines.

Another motivation to analyze the security implications is so that users could understand important infrastructure changes. As we have observed, users may be oblivious to important changes. For example, the firmware update periods we have attempted to analyze do not correlate to known changes in the properties of the quantum computers or repeated calibration data periods. As a result, unless IBM publicly announces firmware updates and stops the quantum computer job queues, it may be possible for IBM to update the firmware or make other low-level modifications without users being aware of them. This can be important, as users may be oblivious to equipment changes, which may cause unexpected jobs or results to the cloud platform, especially for noise-aware experiments. Also, users thus have no way to prove that the underlying hardware has not changed unless IBM publicly announces the changes. Undetectable changes to some parts of the quantum computer system could put the integrity and security of the computations performed on this system in jeopardy.

5 CONCLUSION AND FUTURE WORK

This work demonstrated that NISQ quantum computers experience periods of fluctuation or abrupt qubit frequency changes. This work further analyzed the correlation between the frequency change

events, decoherence times, gate errors, and machine maintenance or offline periods. The results highlight that properties of NISQ computers change over time affecting their dependability, but not all of them can be explained with publicly available data. Based on the observations, future work should focus on further collection of data and analysis to explore patterns in the calibration and other data. There are also naturally other types of quantum computers beyond superconducting qubit quantum computers. Similar studies of other types of machines could help to paint a more complete picture of operation and property changes in all types of cloud-based quantum computers.

ACKNOWLEDGMENTS

This work was supported in part by NSF grant 1901901. The authors thank Sanjay Deshpande, Yongshan Ding, Yao Lu, and Theodoros Trochatos for helpful discussions. The authors thank IBM Quantum for the free access to the quantum computers and the data.

REFERENCES

- [1] Amazon. 2023. Amazon Braket. <https://aws.amazon.com/braket/>.
- [2] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum machine learning. *Nature* 549, 7671 (Sep. 2017), 195–202.
- [3] Yongshan Ding, Pranav Gokhale, Sophia Fuhui Lin, Richard Rines, Thomas Propson, and Frederic T. Chong. 2020. Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 201–214. <https://doi.org/10.1109/MICRO50266.2020.00028>
- [4] Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *Symposium on Theory of Computing (STOC)*. Association for Computing Machinery, 212–219.
- [5] Jared B. Hertzberg, Eric J. Zhang, Sami Rosenblatt, Easwar Magesan, John A. Smolin, Jeng-Bang Yau, Vivekananda P. Adiga, Martin Sandberg, Markus Brink, Jerry M. Chow, and Jason S. Orcutt. 2021. Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors. *npj Quantum Information* 7, 1 (Aug. 2021).
- [6] IBM. 2023. IBM Quantum. <https://quantum-computing.ibm.com/>.
- [7] IBM. 2023. IBM Unveils 400 Qubit-Plus Quantum Processor and Next-Generation IBM Quantum System Two. <https://newsroom.ibm.com/2022-11-09-IBM-Unveils-400-Qubit-Plus-Quantum-Processor-and-Next-Generation-IBM-Quantum-System-Two>.
- [8] Jonathan A Jones, Michele Mosca, and Rasmus H Hansen. 1998. Implementation of a quantum search algorithm on a quantum computer. *Nature* 393, 6683 (1998), 344–346.
- [9] Hyunseong Kim, Christian Jünger, Alexis Morvan, Edward S. Barnard, William P. Livingston, M. Virginia P. Altoé, Yosep Kim, Chengyu Song, Larry Chen, John Mark Kreikebaum, D. Frank Ogilvie, David I. Santiago, and Irfan Siddiqi. 2022. Effects of laser-annealing on fixed-frequency superconducting qubits. *Applied Physics Letters* 121, 14 (Oct. 2022), 142601.
- [10] Benjamin P Lanyon, James D Whitfield, Geoff G Gillett, Michael E Goggin, Marcelo P Almeida, Ivan Kassal, Jacob D Biamonte, Masoud Mohseni, Ben J Powell, and Marco Barbieri. 2010. Towards quantum chemistry on a quantum computer. *Nature chemistry* 2, 2 (2010), 106–111.
- [11] N David Mermin. 2007. *Quantum computer science: an introduction*. Cambridge University Press.
- [12] Microsoft. 2023. Azure Quantum. <https://azure.microsoft.com/en-us/products/quantum>.
- [13] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* 5, 1 (Jul. 2014).
- [14] H Riel. 2021. Quantum Computing Technology. In *International Electron Devices Meeting (IEDM)*. IEEE, 1–3.
- [15] Eric J. Zhang, Srikanth Srinivasan, Neereja Sundaresan, Daniela F. Bogorin, Yves Martin, Jared B. Hertzberg, John Timmerwilke, Emily J. Pritchett, Jeng-Bang Yau, Cindy Wang, William Landers, Eric P. Lewandowski, Adinath Narasgond, Sami Rosenblatt, George A. Keefe, Isaac Lauer, Mary Beth Rothwell, Douglas T. McClure, Oliver E. Dial, Jason S. Orcutt, Markus Brink, and Jerry M. Chow. 2022. High-performance superconducting quantum processors via laser annealing of transmon qubits. *Science Advances* 8, 19 (2022), eabi6690.