2023 IEEE International Conference on Quantum Computing and Engineering (QCE) | 979-8-3503-4323-6/23/$31.00 ©2023 IEEE | DOI: 10.1109/QCE57702.2023.10200

Quantum Computer Metrics and HPC Center Environmental Sensor Data Analysis towards Fidelity Prediction

Hossam Ahmed[1] Xiaolong Deng[1] Helmut Heller[1] Carla Guillen[1] Asim Zulfiqar[1], Martin Ruefnacht[1], Amit Jamadagni[1], Matthew Tovey[1], Martin Schulz[1,2] and Laura Schulz[1]

[1] Leibniz Supercomputing Centre, Boltzmannstraße 1, Garching, Germany
[2] Technical University of Munich, Boltzmannstraße 3, Garching, Germany

***Abstract*—A tight integration of quantum accelerators with high-performance computers (HPC) necessitates relocating the often highly sensitive quantum devices from controlled laboratory environments to noisy data centers. Various factors impacting fidelity must be tracked and mitigated to ensure the performance of quantum computers in such environments. This paper describes a monitoring system utilizing Internet of Things (IoT) technology incorporating multiple sensors to gather and feed data into a data center-wide database. The data collected is then analyzed to identify correlations between fidelity and environmental factors. Using this data, we propose a novel, adaptive calibration mechanism in which a machine learning model predicts the optimal calibration timing for the quantum computer based on fidelity prediction with varying environmental conditions. In the larger scope, this research aims to enhance our understanding of fidelity-altering factors and develop strategies for maintaining optimal quantum computer performance in real-world, noisy data center environments.**

***Index Terms*—High-Performance Computing, Quantum Computing, Monitoring, Operational Data Analytics, Sensor**

I. Introduction

Various quantum technologies and systems exist to build quantum computers or quantum processing units (QPUs). Some of the more prominent include superconducting qubits [1], trapped ions [2], neutral atoms [3], nitrogen vacancies (NV) in diamonds [4], silicon quantum dots [5], and photonic platforms [6]. Such quantum systems generally have (effective) two-level states with long coherence times capable of being initialized, coherently manipulated, and read out [7]. Moreover, the peculiar features of quantum systems, such as superposition and entanglement, make quantum computing more attractive than classical computing for some algorithms [8]. However, while in the noisy intermediate-scale quantum (NISQ) era [9], these QPUs are very restricted by their number of qubits and their relatively short decoherence time [10].

In an ideal situation, quantum systems should be isolated from the environment. In reality, however, they cannot avoid coupling to the environment when controlling signals route to the QPUs. Thus, they may lose their coherence, and the coupling may increase the error rate or make the machine unusable. The challenge here is to minimize such couplings [11]. Further, many surrounding environmental factors – including temperature, magnetic/electric fields, mechanical vibration/rotation, acoustical noise, and pressure – may also affect the performance of QPUs. To fully understand the operating conditions and the impact they can have, it is, therefore, necessary to track them. For this purpose, we set up a comprehensive multi-sensor system to monitor the surrounding environment of quantum computers (QCs) and their supporting systems to investigate how the environment alters the behavior of QPUs.

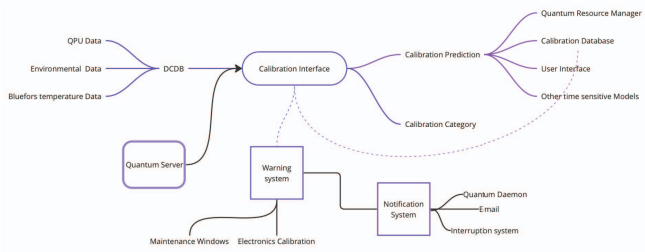

Fig. 1. Workflow description for the inputs from DCDB and quantum results and outputs of the calibration interface for calibration prediction.

The Data Center Data Base (DCDB) [12] provides a convenient and scalable monitoring solution in a uniform format for data retrieval and storage from various sensor sources. DCDB is already in large-scale production use for existing HPC systems, albeit with different sensors, often in the form of hardware performance counters. It uses modular plugins for data acquisition from arbitrary data sources and relies on the Message Queuing Telemetry Transport (MQTT) [13] protocol for data transmission. As part of this work, we have adopted DCDB for use on QPUs by developing and deploying new plugins capable of tracking quality metrics of QPUs as well as environmental sensors from the operation environment in the Leibniz Supercomputing Centre (LRZ) testbed known as BEAST (Bavarian Energy, Architecture, & Software Testbed). Access to DCDB's monitoring data is provided via a Grafana frontend that allows easy visualization with a well-established tool. Data can also be queried via command line tools and a

DOI 10.1109/QCE57702.2023.10200

shared library that enables integration with third-party software and analysis models.

Using the integrated quality and sensor metrics via the DCDB database, we propose a novel calibration scheme (shown in Fig. 1), which applies the collected data to predict the needed recalibration times of the QPU. This information can then be fed to several sources, such as a quantum resource manager, for use in scheduling decisions or a calibration database to enhance the performance of calibration prediction models further.

Our prediction model identifies fidelity changes influenced by environmental or QPU-related factors. If needed, it raises a calibration flag to indicate the need for adjustments in environmental factors or QPU calibration. The system notifies the responsible individuals and sends the calibration prediction to the resource manager. This calibration prediction plays a crucial role in HPC environments, as it helps identify the optimal balance between quantum computer up-time and circuit fidelity.

The structure of the paper is as follows: Section II discusses a superconducting circuit quantum computer and its quality metrics, including coherence time, the fidelity of gates and calibration schemes. Section III covers the sensor project at LRZ, discussing the various kinds of sensors installed or planned. As one example, the section focuses on temperature monitoring. Section IV introduces our novel prediction model for calibrating a quantum system. Section V concludes the paper.

II. Quantum Computer in LRZ

At LRZ, we have installed a superconducting circuit quantum computer from the company IQM, referred to as System 1, which currently hosts a 5-qubit (Q5) QPU. An upgrade to 20-qubit QPU will occur in late 2023. The QPU comprises superconducting qubits, couplers, control, and readout lines. It can properly work when cooled to 10-15mK by a dilution refrigerator (cryostat) with a mixture of Helium 3 and Helium 4. The QPU is manipulated by controlling microwave electronics (in the rack), which is operated at room temperature and connected to the QPU through the co-axial wires and amplifiers, attenuators, filters, and mixers (see Fig. 2).

The qubits are based on transmons [11] with typical microwave frequencies ($5-10$GHz). In such superconducting qubits one can get fast single- and two-qubit gates on nanosecond (ns) time scales with high fidelity of gate operations (e.g. $\sim$ 99.5% for two-qubit gates). The gate operations on the qubits are rather stable and the number of qubits can be easily scaled in the fabrication.

A. Quality Metrics of QPU

To calibrate the QPU, the resonator spectroscopy of readout lines has firstly to be determined. Then we characterize single qubits by sending a sequence of pulses and measuring their outputs. In the Q5 chip there are flux lines that can tune the qubit frequencies. The drive frequency of an individual qubit is determined by the Rabi and Ramsey sequences. The

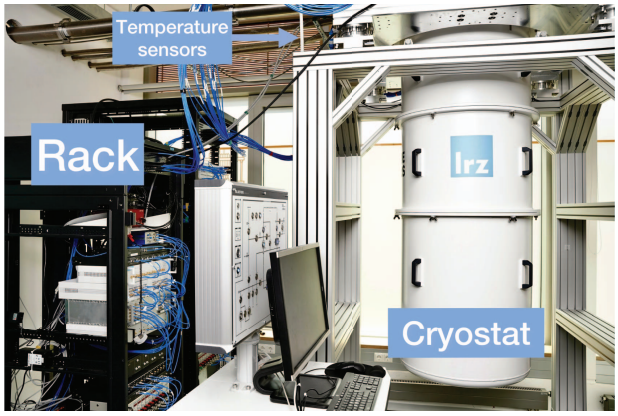

Fig. 2. The Quantum computer System 1 (Q5: 5 qubits) in LRZ. There is an installation of 35 temperature sensors on the cryostat frame.

corresponding coherence time T1 (energy relaxation) and T2 (dephasing) are further analysized by the means of X_π pulse and $X_{\pi/2}$ pulse, respectively [11]. Within T1, the qubit in a state will remain without decaying to another state. Within T2, the qubit in a superposition state will keep the phase relationship between the two terms without dephasing to an orthogonal state. Therefore, good qubits with long coherence time are needed to achieve useful calculation times, and calibration is essential to achieve optimal T1 and T2. Fig. 3 shows T1 and T2 after one complete calibration.

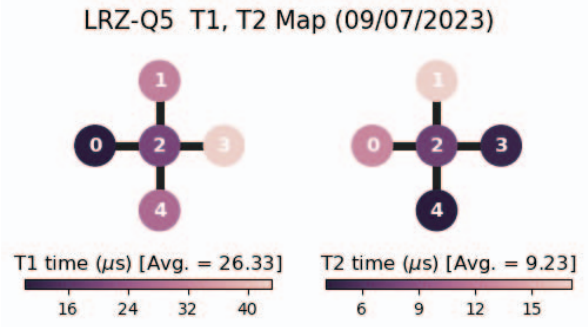

Fig. 3. The decay time (T1) and the dephasing time (T2) of the Q5 chip with a star topology. The time unit is microseconds.

Given the proper qubit characteristics the readout may be further optimized by tuning its frequency and pulse amplitude and phase through a couple of sweeps. In our Q5 chip the native gates are rotation gates (R) and the Controlled-Z gate (CZ). The CZ gate fidelity is determined by the interleaved randomized benchmarking experiment, in which a sequence of Clifford gates $C_1, CZ, C_2, CZ, \ldots, C_{N-1}, CZ$ followed by its inverse is measured. The procedure is repeated for a number of sequences with different lengths. Then the depolarizing parameter extracted from the survival probability of the initial state as a function of the sequence lengths can give us the gate fidelity. The detailed quality of the gates and readout in Q5 is shown graphically in its error map, see Fig. 4.

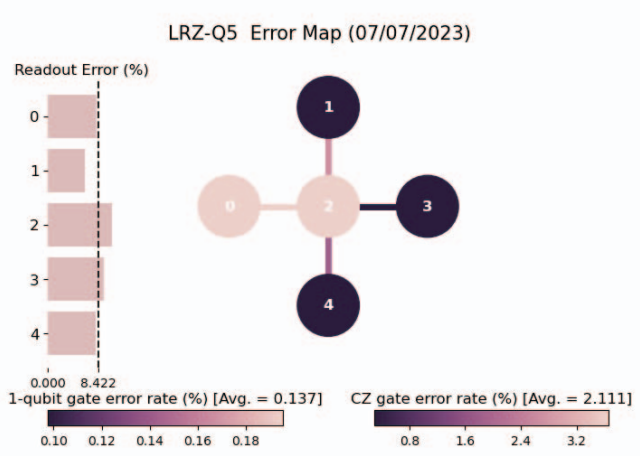

Fig. 4. The error map of the Q5 chip after one calibration: readout errors, single-qubit gate and two-qubit gate (CZ) errors. The error map is generated using Qiskit [14].

B. Calibration Scheme: Static and Dynamic

In LRZ a standard routine of calibrating QPUs is being set up. Our goal is to build up scalable calibration matrices and integrate them into our HPC workflow, database and job performance management.

A set of measurements have been added to our calibration setup as a quality metric set, including routines to extract T1, T2, Rabi, Ramsey, single-shot readout fidelity, single-qubit gate fidelity, two-qubit gate fidelity, and state tomography [11]. The number of measurement shots is fixed. The systematic re-calibration procedure is currently designed as the following. When the performance is still within bounds since the last calibration, we re-characterize the qubit readout, single-qubit gates, and CZ gates with a small number of runs. If the system worsens, we execute an extended process, adding more detailed runs. When the calibration is complete, a new quality metric data set is created, and the quality of QPU is updated with these fresh values. These static data are collected, ported into the DCDB, and traced for further executions, e.g., error mitigation.

Besides the static quality metric set, we also build a dynamic data set, which monitors dynamically the quality of qubits and gates. As soon as the quality is lower than some thresholds, a re-calibration will be triggered.

III. Environmental Factors Monitored by Sensors

The Q5 quantum computer resides in LRZ's Quantum Integration Centre (QIC). Two centrally placed air conditioning units (HVAC: Heating, Ventilation, Air Conditioning) maintain temperatures of approximately 20°C. Also located in the same room are the related microwave control electronics and controllers for fridge temperature and vacuum. On the other side of the room, two compute racks house the conventional access computers for the QPU and a Raspberry Pi 4B (RPI) for gathering environmental data. The RPI is connected via a 100Mb LAN cable to a private VLAN for security, while we can still ssh to the RPI for maintenance and updates.

Next to the main room is its accessory machine room, which houses the cryostat's compressors, coolers, and gas handling system (GHS). Since this equipment produces mechanical vibrations, electromagnetic fields, and acoustic noise, they reside in this separate space. With various physical sensors installed, we seek to understand which physical qualities influence the QPU.

A. Environmental Sensors: Installed

As the monitoring project at LRZ is currently in the build-up phase, we currently rely on four sensor groups, which we expect to be among the most critical.

1) Temperature: On the frame of the cryostat, we provide 35 temperature sensors, one every ten centimeters, covering the room height from 0 cm (floor) to 340 cm (almost ceiling height). Each sensor is of the type DS18B20. They are all connected to a 1-wire bus routed to the RPI, where a DS2480B connects it to the serial port of the RPI. This way, the long 1-wire lines having to traverse the whole room are met with a controlled impedance, which minimizes ringing and transmission errors. All temperatures are measured every minute. Fig. 2 shows the sensors on the frame.

2) Movement: Humans near the fridge can introduce an unwanted coupling to the QPU, e.g., through phonons (vibrations) that might transmit through the frame and the cryostat vessel to the QPU. They can also be the source of increased temperature, especially in larger groups. In order to detect human presence in the room, we use a Passive Infrared (PIR) sensor, type HC-SR501, mounted to a cable duct close to the ceiling and centrally in the room. This sensor has a wide field of view and detects the movement of heat sources (e.g., humans) anywhere in the room. Its on/off output is connected to a General-purpose Input/output (GPIO) input pin on the RPI. Every change in the logical state of this line triggers an interrupt, which in turn writes a data point to our database.

3) Fridge Temperatures: Several temperatures and pressures inside the cryostat are being measured and recorded using vendor interfaces. Variations in the QPU temperature may indicate problems with the cryostat system.

4) Liquid Nitrogen: In the machine room, liquid nitrogen (LN2) is used by the GHS to freeze out and remove impurities from the cooling system. Therefore, LN2 must continuously pump into the GHS, and it is necessary to monitor the volume of the supply vessel as the LN2 is consumed and evaporates. To this end, we weigh the supply vessel automatically with an industrial scale of type IFB 60K-3L-2023e from KERN every hour and compute the amount of LN2 still in the vessel. This amount is then stored in the DCDB.

B. Environmental Sensors: Planned

While the above sensors give us a first set of critical metrics for monitoring, optimizing, and calibrating QCs in data center environments, many more potential sources for coupling exist. We are, therefore, in the process of adding additional sensors or sensor groups.

1) Air Humidity: Humidity is also controlled by the HVAC units; however, keeping a record of it is important, as electronic circuits may fail if the humidity is too low or too high.

2) Air Pressure: Variations in air pressure, which are not controlled, may indicate the opening or closing of doors or windows or a change in weather, like impending thunderstorms, which may influence the QPU.

3) Light: While light might not directly influence the QPU due to its opaque shielding by the vacuum vessel, light sensor is triggered to indicate the fluorescent ceiling lights, which may create electromagnetic stray fields.

4) Vibration: A very sensitive industrial seismometer/accelerometer, mounted directly to the cryostat frame, can detect minute vibrations that might couple to the QPU.

5) Acoustic Noise: Acoustic noise, be it from people chatting or surrounding systems and their fans operated on the same machine floor, might also influence the QPU and will be measured by a sensitive microphone. Only its amplitude will be registered, not speech contents (to fulfill data privacy laws). We will also perform an audio spectrum analysis on the noise to isolate predominant frequencies that might interfere, especially with trapped ion QPUs.

6) Electromagnetic Fields: Despite the heavy metallic shielding, electromagnetic fields might leak to the QPU. We will measure electric and magnetic fields in the low-frequency range of 20 Hz to 10kHz.

7) Radioactivity: Cosmic rays are very powerful and might penetrate the QPU shielding. We will install a Geiger counter above the cryostat to register abnormal events.

8) Water Temperatures in the Cooling System: At several critical points, we will measure the temperatures of our cooling systems via MODBUS industrial sensors to detect malfunctions or errors early on.

9) AC Power Condition: We use several 3-phase AC power circuits to supply our equipment. We will measure not only power consumption but also power noise on these lines.

10) Data from Building Sensors: We will measure and record the power consumption of nearby large power consumers, including building elements like elevators and large-scale HPC systems with their power fluctuations.

All of these sensors, which will be handled through additional DCDB plugins and fed into the same database to enable correlation analysis, will help us fully track the environmental conditions in the QIC setup and the QC systems in HPC production environments. This new operational monitoring will be crucial for successfully operating and optimizing hybrid HPCQC setups.

C. Environmental Data Monitoring and Alerting

In this subsection, we present the monitoring of temperature parameters, including room temperatures, cryostat data, and liquid nitrogen containers. All temperature data is stored in the DCDB database. Using the versatility of DCDB, we have developed a custom script running as a system service on the same machine, which uses the Prometheus client library [15] and collects real-time data from sensors. The script scrapes the latest values from the DCDB database every five seconds. These values are transformed into numerical metrics using the gauge identified by the label of sensors.

The scraped metrics are made accessible via a dedicated HTTP server. Open-source monitoring and alerting tools Prometheus and Alert Manager [15] are employed on another virtual machine to monitor various real-time metrics continuously. We define a temperature threshold of 25.5°C for the room and 15mK for cryostats. Beyond those points, potential risks to the QPU may arise. Whenever the temperature rises above this threshold, Prometheus triggers an alert managed by the Alert Manager configured to send notifications through email using the capabilities of webhook integration. Additionally, alert messages are pushed to an internal chat channel for the operations team. Through our internal platform for real-time collaboration and communication, our monitoring system allows recipients to promptly address any environmental anomalies that may impact the quantum computing infrastructure. Fig. 5 shows one recent event as an example.

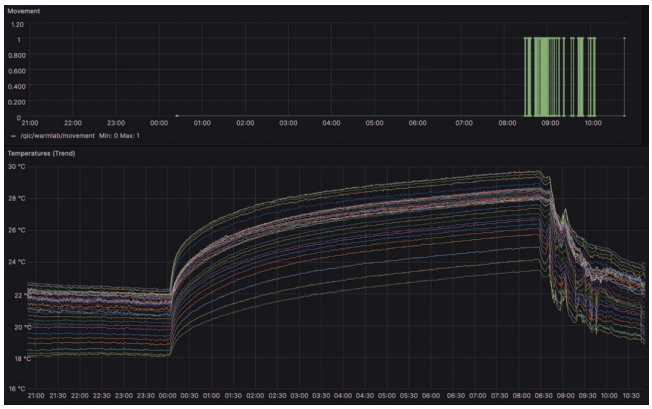

Fig. 5. Display of room temperatures and movement in the lab. On June 23rd, 2023, at midnight, both HVAC units were lost and the room temperature increased to almost 30°C.

IV. Data Prediction Model for Calibrating QC

The setup described above enables gathering a wide range of data from the quantum computer and its environment. The next step is to utilize this data to improve system operation. For this, we introduce a novel prediction scheme to estimate calibration times and enable an automatic calibration setup. Such capability is essential for operating QC systems in production data center environments, which typically have access restrictions, Service Level Agreement (SLA) uptime requirements, and environments geared towards HPC systems instead of QCs.

A. Assessing Gate/Circuit Fidelities

The fidelity of the quantum gates is one of the key features of benchmarking QPUs. There are multiple approaches to computing gate fidelity. The most common one, which we follow as well, is to apply a sequence of unitaries to retrieve the initial state, i.e., given a gate U the fidelity can be

captured by the expression $\langle\psi|UU^{\dagger}|\psi\rangle$. This is the so-called randomized benchamrking protocol when U includes a random sequence of gates.

For this approach, we implement several experimental circuits using random unitary gates to examine the fidelity decay with increasing depth. This intends to retrieve the original state using these random circuits. The fidelity of each circuit run is then measured using the Qiskit Ideal quantum circuit state vector simulator using the Hellinger fidelity technique.

The experimental circuit is implemented in Qiskit, utilizing the native gates of the QPU to avoid any transpilation processes that may affect the circuit's processing. We use the QuantumCircuit.r(θ, ϕ, qubit) function to apply the R gate to all five qubits of the system. For this experiment, we set θ to $-\pi/2$ and ϕ to π.

To facilitate the analysis, the tests are automated to run the circuit every five minutes, incrementing the depth by one for each qubit and saving the results in a database for further analysis. The decay of fidelity as the depth increases is depicted in Fig. 6.

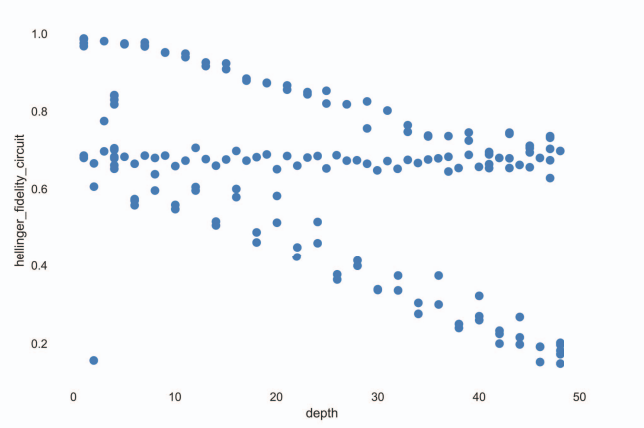

Fig. 6. The degradation of fidelity as the depth of the circuit increases, reaching a maximum depth of 48.

B. Correlations

We then use analysis models to establish correlations between various parameters. As anticipated, the primary correlation observed is between fidelity and depth. Investigating the correlation between fidelity and time in this context is less informative since depth increases with time. Before executing the main circuit, a secondary circuit comprising two unitary gates is executed on each qubit to assess the fidelity of minimal circuits on different qubits and detect any fidelity variations over time.

The correlations among various factors are presented in the table represented in Fig. 7. Analyzing these correlations facilitates the development of a prediction model for estimating fidelity as different factors vary across circuits and parameters.

To investigate the impact of time on fidelity, we devise an additional experiment that maintains a constant circuit depth

	timestamp	depth	hellinger_fidelity_circuit	hellinger_fidelity_Q0
timestamp	1.000	0.190	-0.080	-0.563
depth	0.190	1.000	-0.410	0.050
hellinger_fidelity_circuit	-0.080	-0.410	1.000	-0.007
hellinger_fidelity_Q0	-0.563	0.050	-0.007	1.000

Fig. 7. Illustration table of the correlation coefficients among different parameters in the experimental circuits. The "time_stamp" column represents the time, the "depth" column represents the circuit depth, the "hellinger_fidelity_circuit" column represents the fidelity of the circuit, and the "hellinger_fidelity_Q0" column represents the fidelity of the minimal circuit applied on each qubit.

while executing the circuit every five minutes under identical conditions as the previous experiment. In this iteration, we use a two-qubit native Controlled-Z gate, taking into account the star layout of the QPU to minimize transpilation and prevent any optimization applied to the designed circuit. The correlation between fidelity and time is shown in Fig. 8. The time is normalized between 0 to 1 by prepossessing for better analysis.

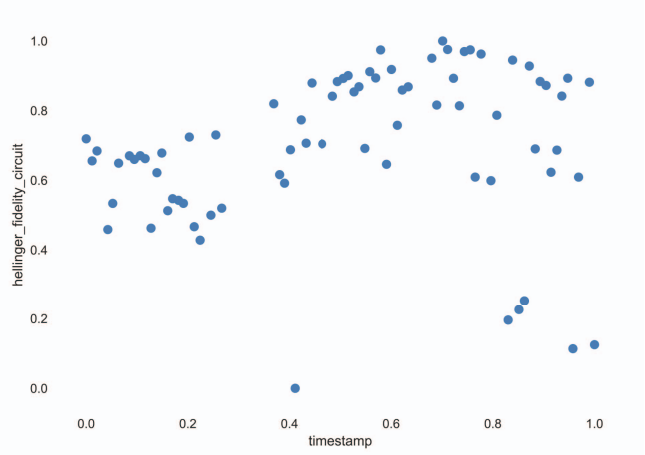

Fig. 8. Illustration of the correlation between the "time_stamp" representing the normalized time (horizontal-axis), and the "hellinger_fidelity_circuit" representing the fidelity of the circuit (vertical axis).

C. Structural Analysis for Environmental Factors

Further, we add several environmental measurements from our sensor setup to enable further analysis and explore their correlation with fidelity. Due to the limited duration of our initial experiments (approximately 17 hours) and the frequency of experimental circuit runs (every five minutes), certain environmental factors exhibited minimal variance and are not included in this analysis. The analysis, then, focuses on the following environmental factors: room temperature, mixing chamber temperature in the cryostat, and cryostat pressure (P1). Other environmental data captured by sensors, such as laboratory movement and various pressures and temperatures at different stages of the cryostat, are excluded due to their

limited variance within the experiment duration but will play a crucial role in our planned longtime studies.

The analysis reveals a notable correlation between vacuum can pressure in the cryostat measured in (pBars) and fidelity, with a correlation coefficient of +0.495, as shown in Fig. 9.

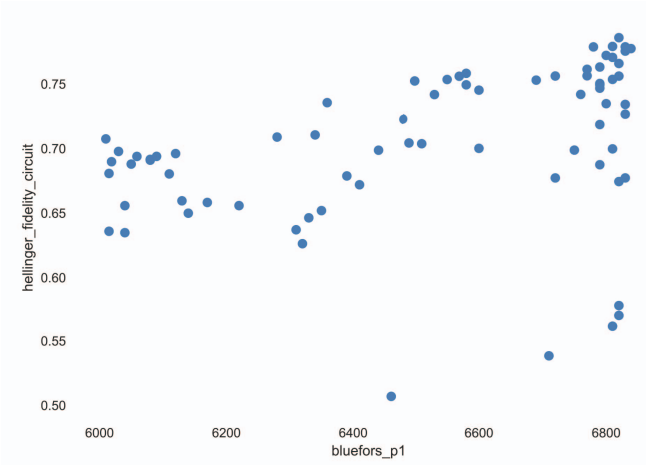

Fig. 9. Correlation between Pressure vacuum can pressure in cryostat (pBar) and fidelity.

D. Prediction

To determine how fidelity decays over time we run an experiment using a CZ gate configuration that matches the topology of our QPU on all 5 qubits at a depth of three. We include the environmental and QPU-related factors with time-stamps as indices to predict the fidelity over 24-hour window. We continuously ran the circuit every five minutes without conducting any changes to the circuit to observe how fidelity decays over time. Given the nature of our time-series data, several considerations influenced our choice of the prediction model. First, the presence of impulse signals, such as movement sensors data within the laboratory environment (see Fig. 5), necessitates a model capable of detecting data shifts. Second, while our fidelity prediction strive for utmost accuracy, a margin of approximately 2%-5% in fidelity estimation is acceptable, though this range may be subject to adjustment based on specific application requirements. Consequently, a fundamental requirement is the model's ability to generate uncertainty intervals encompassing the projected values, accommodating this acceptable variability. A third consideration is the ease of parameter tuning, ensuring that adjustments do not compromise the model's accuracy. Equally vital is the fourth aspect: scalability. The model must seamlessly accommodate the integration of new sensors and the escalated frequency of quantum computer job submissions. This necessitates the model's proficiency in handling substantial datasets, while remaining resilient to both missing data instances and outliers that may arise. Consequently, our chosen approach revolves around a general forecasting model, aligning with our overarching utilization – forecasting prediction. To this end, we opt for the implementation of the "Prophet" forecasting model [16] to predict fidelity 24 hours after the experiment run-time. We divide the data into 80for training and 20% for testing. The model shows an accuracy of 73.5510335%. The prediction graph is shown in Fig. 10

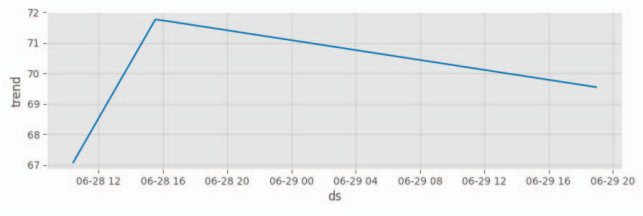

Fig. 10. Prediction of the fidelity over a 24-hour period using forecasting prediction model.

V. Conclusion

In summary, we have investigated the influence of the environment on a quantum computer by installing various sensors. We have developed workflows that seamlessly integrate the data collection from various sensors into a database, DCDB. Further, a machine learning-based prediction model has been built based on the data collected into the DCDB, resulting in optimal triggering of the re-calibration of the quantum system. In the future, as more data is collected over a longer time, we aim to establish stronger correlations between various environmental factors and the performance of the quantum computer.

In Ref. [17], the authors have established a method to estimate the fidelity of the quantum circuits by introducing the notion of Probability of Successful Trails (PST). The above quantity being monotonically related to the fidelity, we aim to leverage the above measure to estimate the fidelity of quantum computers better. There have also been recent proposals, for instance, Ref. [18], where a relationship between a random circuit and feature vectors of a neural network has been explored. In the current context, the above ideas can be inverted, wherein the neural network trained on the sensor data can be mapped to a random circuit, therefore resulting in a quantum circuit capturing the environmental effects.

Acknowledgments

The authors wish to thank fellow LRZ colleagues: Mahmoud Abuzayed and Robert Bakyayita (Quantum Computing and Technologies team); Dr. Josef Weidendorfer and Dr. Michael Ott (Future Computing team) for assistance setting up and implementing the DCDB infrastructure; and Pascal Weibel, Üseyin Sarikaya, Hiren Ghandi and Albert Kirnberger (Facilities team) for their assistance identifying various building and operations metrics.

Funded by the German Federal Ministry for Education and Research under grants 13N15689 (DAQC), 13N16063 (Q-Exa), 13N16188 (MUNIQC-SC), and 13N16078 (MUNIQC-ATOMS), and the Bavarian State Ministry of Science and the Arts as part of Munich Quantum Valley (MQV).

REFERENCES

[1] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, "State preservation by repetitive error detection in a superconducting quantum circuit," *Nature*, vol. 519, pp. 66–69, mar 2015.

[2] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Krieglsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler, and T. Monz, "Compact ion-trap quantum computing demonstrator," *PRX Quantum*, vol. 2, jun 2021.

[3] L. Henriet, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, "Quantum computing with neutral atoms," *Quantum*, vol. 4, p. 327, sep 2020.

[4] L. Childress and R. Hanson, "Diamond nv centers for quantum computing and quantum networks," *MRS Bulletin*, vol. 38, no. 2, p. 134–138, 2013.

[5] S. G. J. Philips, M. T. Mądzik, S. V. Amitonov, S. L. de Snoo, M. Russ, N. Kalhor, C. Volk, W. I. L. Lawrie, D. Brousse, L. Tryputen, B. P. Wuetz, A. Sammak, M. Veldhorst, G. Scappucci, and L. M. K. Vandersypen, "Universal control of a six-qubit quantum processor in silicon," *Nature*, vol. 609, pp. 919–924, sep 2022.

[6] S. Slussarenko and G. J. Pryde, "Photonic quantum information processing: A concise review," *Applied Physics Reviews*, vol. 6, p. 041303, dec 2019.

[7] D. P. DiVincenzo, "The physical implementation of quantum computation," *Fortschritte der Physik*, vol. 48, pp. 771–783, sep 2000.

[8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

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

[10] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, "Noisy intermediate-scale quantum algorithms," *Reviews of Modern Physics*, vol. 94, feb 2022.

[11] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "A quantum engineer's guide to superconducting qubits," *Applied Physics Reviews*, vol. 6, p. 021318, jun 2019.

[12] A. Netti, M. Müller, C. Guillen, M. Ott, D. Tafani, G. Ozer, and M. Schulz, "DCDB wintermute: Enabling online and holistic operational data analytics on HPC systems," in *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*, ACM, jun 2020.

[13] D. Locke, "Mq telemetry transport (mqtt) v3.1 protocol specification," tech. rep., IBM, August 2010.

[14] Qiskit contributors, "Qiskit: An open-source framework for quantum computing," 2023.

[15] "Prometheus Python Client." https://github.com/prometheus/client_python. Accessed: June 29,2023.

[16] E. Žunić, K. Korjenić, K. Hodžić, and D. Donko, "Application of facebook's prophet algorithm for successful sales forecasting based on real-world data," *International Journal of Computer Science and Information Technology*, vol. 12, pp. 23–36, apr 2020.

[17] S. S. Tannu and M. K. Qureshi, "A case for variability-aware policies for nisq-era quantum computers," 2018.

[18] H. Wang, P. Liu, J. Cheng, Z. Liang, J. Gu, Z. Li, Y. Ding, W. Jiang, Y. Shi, X. Qian, D. Z. Pan, F. T. Chong, and S. Han, "Quest: Graph transformer for quantum circuit reliability estimation," 2022.