The energetic cost of loading classical data in quantum computers

Joao F. Doriguello[1], Alessandro Luongo[2,3], and Harshit Verma[2,4],
[1] HUN-REN Alfréd Rényi Institute of Mathematics, Budapest, Hungary
[2] Centre for Quantum Technologies, National University of Singapore, Singapore
[3] Inveriant Pte. Ltd., Singapore
[4] MajuLab, CNRS-UCA-SU-NUS-NTU International Joint Research Laboratory, 117543 Singapore

Introduction. Loading large classical data sets into a quantum processor remains a central bottleneck for real-world applications of quantum computers: any speedup offered by a quantum routine can be undermined if the input data cannot be loaded in the quantum computer. Quantum random access memory (QRAM) [12, 13, 2] can be thought as a mapping allowing coherent access to a classical dataset, performing the mapping $|i\rangle |0\rangle \mapsto |i\rangle |m_i\rangle$, there m_i is the i-th element of the memory of size N (e.g., a bistring of 64 bits, like in classical computers). Physically, we can think if it as an external device, sharing registers with the quantum computer. It enables access to classical data in superposition, and is a central assumption in the vast majority of quantum algorithms in machine learning [23, 21, 20, 16, 15], quantum simulation [4], and cryptographic applications [8, 9]. To date, we are not aware of any study on the energetic cost of QRAM circuits. We provide the first application-level study of the energetic cost of QRAM-based data loading, focusing on anomaly detection through the quantum version of a foundational quantum machine learning algorithm (Principal Component Analysis). Using the metric-noise-resource (MNR) framework, for quantum computation primitive such as gate on a superconducting platform, we can define the error (fidelity), and the energy cost as a function of controllable physical control parameters [10]. We use this to derive QRAM energetics and fidelity (Fig. 1) which can be used to determine the optimal control parameters that enable feasible applications (based on fidelity), while also optimizing the energy consumption while keeping a fixed targeted fidelity.

QRAM architectures. A baseline implementation for the mapping discussed above is the multiplexer architecture (sometimes known as QROM): a linear sequence of controlled unitaries that checks whether the address register is in a prescribed state. In non-error corrected computers, this circuits has an infidelity that scales as $\mathcal{O}(\varepsilon N \log N)$ with memory size N and gate error rate ε. A different implementation, called *bucket-brigade* [12, 13] instead routes address and target qubits through a binary tree, achieving logarithmic depth at the cost of a substantial ancilla overhead. Crucially, this architecture also exhibits noise resilience properties: for a very general error channel, the query infidelity scales only as $\mathcal{O}(\varepsilon \log^2 N)$ [14, 19], a consequence of the limited entanglement generated among routing components. This makes bucket-brigade QRAM especially appealing for high-fidelity loading of large data sets without full quantum error correction.

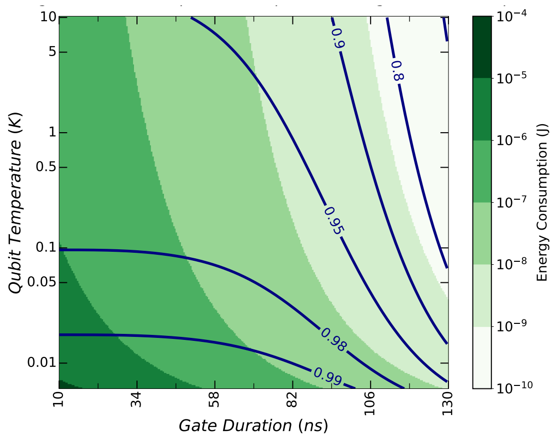

Figure 1: Heatmap of QRAM energy consumption with annotated QRAM fidelity on blue contours; as a function of controllable physical quantities such as temperature of the qubit and gate duration. QRAM energetics, fidelity shown here are based on gate fidelities, energetics for superconducting qubits available from literature. The query fidelity – a function of QRAM fidelity – determines the feasibility of an application, whereas the energy consumption and the physical parameters might impose physical constraints on the operating regime. Such heatmaps enable minimization of energy while keeping a targeted fidelity fixed, while also being compatible with any practical limits mentioned above. Each shade of green spans a single order of magnitude in energy consumption, thus, as an example the energy consumption of a QRAM can be reduced by 100× while keeping fidelity fixed at 0.99. In this example, Bucket Brigade QRAM is considered along with IBM native gate-set.

MNR framework. To connect these architectural guarantees with experimentally controllable quantities, we adopt the metric noise resource (MNR) framework [10] for transmon-based superconducting qubits [17, 5]. The framework links QRAM fidelity and energetic cost to physical parameters such as control-line attenuation A, qubit operating temperature T_{qb}, ambient temperature T_{ext}, and gate timing, while incorporating the dressed cryogenic power consumption of the refrigerator. Because faster gates reduce thermal-noise exposure but generally require more power, and larger attenuation suppresses thermal photon occupation while increasing the heat load that must be removed, the resulting model naturally supports constrained optimization for efficiency ($= metric/energy$) while keeping the metric fixed.

Results and applications. We use the MNR framework to study the energetic cost of performing anomaly detection on superconducting quantum computers using a novel bucket-brigade QRAM circuit, which we decompose into native superconducting gate-sets from Google, IBM, and Rigetti processors [1, 3, 11, 18].

- First we obtained a superconducting-friendly BB circuit. The original circuit is based on controlled-SWAP gates, which are not native to superconducting platforms. Our implementation is based on controlled-iSWAP gates. Then, adapting ideas from [22], we derived a novel and efficient decomposition of the controlled-iSWAP gate into one and two qubit gates of the aforementioned platforms.

- Then, we selected the most efficient quantum algorithm for a simple (but non-trivial) quantum machine learning algorithm [6], for which we know the *threshold for quantum advantage* in terms of the number of queries to a QRAM-based oracle [7], on anomaly detection datasets.
- For a representative configuration of 2^{16} bits of memory, our bucket-brigade QRAM implementation requires $196,625$ total qubits with a circuit depth of $1,248$. Our model of superconducting qubits obtains a very generous gate error rate of $\sim 10^{-7}$ (assuming qubit frequency of 6Ghz Hz, 50 ns of gate duration, and 6 mK of qubit temperature), we achieve a query infidelity of 0.111, which meets the target threshold of 0.206 required for running 40 QRAM queries with a total failure probability below 0.9. The corresponding energetic footprint shows a maximum instantaneous power consumption of 625 W and a total energy cost of 0.0693 Joules per query.

Future work. Extending the analysis to other quantum machine learning algorithms, and to other QRAM architectures, and exploring the trade-off between fidelity, energy consumption, size and depth of different QRAM circuits, for different applications.

References

[1] Muhammad AbuGhanem. IBM quantum computers: evolution, performance, and future directions. *The Journal of Supercomputing*, 81(5):687, Apr 2025.

[2] Jonathan Allcock, Jinge Bao, Joao F. Doriguello, Alessandro Luongo, and Miklos Santha. Constant-depth circuits for Boolean functions and quantum memory devices using multi-qubit gates. *Quantum*, 8:1530, November 2024.

[3] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019.

[4] Ryan Babbush, Craig Gidney, Dominic W Berry, Nathan Wiebe, Jarrod McClean, Alexandru Paler, Austin Fowler, and Hartmut Neven. Encoding electronic spectra in quantum circuits with linear T complexity. *Physical Review X*, 8(4):041015, 2018.

[5] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O'Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.*, 111:080502, Aug 2013.

[6] Armando Bellante, Alessandro Luongo, and Stefano Zanero. Quantum algorithms for svd-based data representation and analysis. *Quantum Machine Intelligence*, 4(2):20, 2022.

[7] Armando Bellante, Alessandro Luongo, and Stefano Zanero. Quantum algorithms for SVD-based data representation and analysis. *Quantum Machine Intelligence*, 4(2):20, 2022.

[8] André Chailloux and Johanna Loyer. Lattice sieving via quantum random walks. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021*, pages 63–91, Cham, 2021. Springer International Publishing.

[9] Joao F Doriguello, George Giapitzakis, Alessandro Luongo, and Aditya Morolia. On the practicality of quantum sieving algorithms for the shortest vector problem. *arXiv preprint arXiv:2410.13759*, 2024.

[10] Marco Fellous-Asiani, Jing Hao Chai, Yvain Thonnart, Hui Khoon Ng, Robert S. Whitney, and Alexia Auffèves. Optimizing Resource Efficiencies for Scalable Full-Stack Quantum Computers. *PRX Quantum*, 4(4):040319, October 2023.

[11] M. Ganzhorn, G. Salis, D. J. Egger, A. Fuhrer, M. Mergenthaler, C. Müller, P. Müller, S. Paredes, M. Pechal, M. Werninghaus, and S. Filipp. Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform. *Phys. Rev. Res.*, 2:033447, Sep 2020.

[12] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Architectures for a quantum random access memory. *Phys. Rev. A*, 78:052310, Nov 2008.

[13] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Phys. Rev. Lett.*, 100:160501, Apr 2008.

[14] Connor T. Hann, Gideon Lee, S.M. Girvin, and Liang Jiang. Resilience of quantum random access memory to generic noise. *PRX Quantum*, 2:020311, Apr 2021.

[15] Iordanis Kerenidis, Alessandro Luongo, and Anupam Prakash. Quantum expectation-maximization for Gaussian mixture models. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5187–5197. PMLR, 13–18 Jul 2020.

[16] Iordanis Kerenidis and Anupam Prakash. Quantum Recommendation Systems. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:21, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[17] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Phys. Rev. A*, 76:042319, Oct 2007.

[18] Janusz Kusyk, Samah M. Saeed, and Muharrem Umit Uyar. Survey on quantum circuit compilation for noisy intermediate-scale quantum computers: Artificial intelligence to heuristics. *IEEE Transactions on Quantum Engineering*, 2:1–16, 2021.

[19] Rohan Mehta, Gideon Lee, and Liang Jiang. Analysis and suppression of errors in quantum random access memory errors under extended noise models. *arXiv e-prints*, pages arXiv–2412, 2024.

[20] Gernot Schaller and Ralf Schützhold. Quantum algorithm for optical-template recognition with noise filtering. *Phys. Rev. A*, 74:012303, Jul 2006.

[21] Ralf Schützhold. Pattern recognition on a quantum computer. *Phys. Rev. A*, 67:062311, Jun 2003.

[22] Peter Selinger. Quantum circuits of t-depth one. *Phys. Rev. A*, 87:042302, Apr 2013.

[23] C. A. Trugenberger. Phase transitions in quantum pattern recognition. *Phys. Rev. Lett.*, 89:277903, Dec 2002.