Evaluating Decomposition Strategies for Bosonic Qudit Compilation

Katie Harrison*, Quinn Langfitt*, Yuchen Zhu*, Mu-Te Lau*, Nikos Hardavellas*, Kate Smith*,
Leonardo Bove†, Tanay Roy†, Doga Kürkçüoglu†, José Cruz Serrallés‡
*Northwestern University, Evanston, IL, USA
†Fermi National Accelerator Laboratory, Batavia, IL, USA
‡New York University, New York, NY, USA
Emails: {kharrison, qlangfitt, yzhu7, mu-telau2030, nikos, kns}@northwestern.edu; {lbove, roytanay, dogak}@fnal.gov;
jose.cruzserralles@nyulangone.org

***Abstract*—Many QED quantum systems naturally support more than two levels, allowing them to function as qudits and potentially enable more efficient quantum computation. Realizing these advantages requires compilation methods that map target unitaries onto hardware-executable operations. In superconducting cavity systems, Selective Number-dependent Arbitrary Phase (SNAP) and displacement (Disp) gates form a useful control set. We compare three strategies for decomposing qudit unitaries into SNAP and Disp gate sequences: a numerical optimization method, a structured decomposition method, and an analytical Givens-rotation-based method. Using ideal fidelity and Lindblad simulations with photon loss and dephasing with systems capable of $d = 20+$, we find that the decomposition with the highest noiseless fidelity does not necessarily perform best on realistic hardware [8]. Shorter decompositions often outperform longer analytical circuits due to reduced noise accumulation, highlighting the need for hardware-aware compilation strategies.**

***Index Terms*—bosonic qudits, quantum compilation, SNAP gates, displacement gates, superconducting cavities, gate decomposition**

I. Motivation

In bosonic systems, information can be encoded into Fock states of a harmonic oscillator, allowing a single mode to represent a high-dimensional logical system. This expanded state space can enable compact encodings and more efficient circuit implementations. However, these advantages depend on whether target qudit operations can be compiled into experimentally realizable controls. For superconducting cavities, SNAP and Displacement (Disp) gates together make up a universal gate set where SNAP gates apply number-dependent phases to Fock levels, while Disp gates mix oscillator states in phase space [3], [6], [7]. Our goal is to decompose a target unitary into a sequence of SNAP and Disp gates that is accurate, short, and robust to hardware noise.

II. Compilation Methods

Our pipeline begins with a target unitary acting on a logical qudit of dimension, d. The operation may be embedded into a larger Hilbert space of total dimension $N = d + b$, where b is the number of buffer levels. Buffer states allow higher Fock-state leakage to be incorporated into the synthesis process, producing decompositions that more accurately reflect hardware behavior. However, they also enlarge the optimization space, increasing the difficulty of finding high-fidelity decompositions.

Method 1: Numerical optimization. Fixes a SNAP-displacement circuit structure and optimizes all gate parameters simultaneously to minimize infidelity. This approach is highly flexible and can often produce relatively short circuits, but it is also computationally expensive. As N (total dimension) and the gate count increase, the number of free parameters grows rapidly, leading to a high-dimensional, nonconvex optimization problem. In practice, multiple random restarts are often required to avoid poor local minima and obtain high-fidelity decompositions. As a result, our current dataset contains fewer Method 1 results than Method 2 or Method 3. This difference reflects the significantly longer runtimes and higher computational cost associated with Method 1.

Method 2: Structured decomposition. Reduces the size of the optimization problem by incorporating prior knowledge about the target unitary and constraining portions of the parameter space. Rather than optimizing every SNAP angle and Disp parameter simultaneously, several parameters are fixed or initialized using properties of the target operator. In particular, the phases of the central SNAP gate are initialized from the eigenvalue decomposition of the target unitary. Additionally, for all subsequent SNAP gates, the first phase angle is fixed to zero. Because an overall global phase does not affect the implemented operation, fixing one phase removes this redundancy and prevents the optimizer from exploring many equivalent solutions that differ only by a global phase. The decomposition is also constructed symmetrically about the central SNAP gate, reducing the number of independent parameters that must be optimized.

Method 3: Analytical Givens-based decomposition. Follows the direct compilation approach proposed by Job [4]. This method decomposes an arbitrary $SU(N)$ operation into Givens-like rotations and implements those rotations using SNAP and displacement gates. The construction avoids numerical optimization at compile time and produces deterministic circuits. Its main drawback is gate count: the analytical construction often requires substantially more SNAP and Disp gates than the other methods.

III. Evaluation Methodology

We evaluate each decomposition in both noiseless and noisy settings. Noiseless fidelity measures whether the synthesized unitary matches the target operation. Noisy fidelity is computed using a Lindblad master-equation simulation with photon loss and dephasing [5]. Photon loss changes the Fock-state population distribution, while dephasing destroys coherence between levels. These channels are important for bosonic encodings because the computation is stored directly in oscillator states.

This distinction is central to our study. A circuit with near-perfect ideal fidelity may still perform poorly on hardware if it requires many gates. Conversely, a shorter circuit with slightly worse ideal fidelity may be more robust because it leaves less time for decoherence to accumulate. Decompositions were evaluated using both ideal simulations and a Lindblad master equation framework [5].

IV. Conclusions and Design Takeaways

We compared three strategies for decomposing qudit unitaries into SNAP and displacement gates for bosonic quantum systems. The method 3 analytical decomposition based on Job's direct $SU(N)$ compilation method is deterministic and highly accurate in noiseless simulations, but its larger gate count can make it vulnerable to photon loss and dephasing. The numerical method 1 can produce shorter circuits that may be more robust under realistic noise, but it's large search space significantly increases the time cost of convergence and introduces optimization cost and scaling challenges.

Our results reveal several important tradeoffs for bosonic qudit compilation. While a finite-dimensional quantum harmonic oscillator is mathematically well-defined, real bosonic systems possess an effectively infinite-dimensional Hilbert space. Consequently, performing decompositions using only the logical qudit subspace can lead to implementations that fail to accurately capture hardware behavior. Introducing buffer states allows leakage into higher Fock levels to be incorporated during synthesis, producing decompositions that more closely reflect physical devices. However, additional buffer states also increase optimization complexity and introduce new noise channels, leading to an optimal buffer size that balances physical realism, computational cost, and fidelity. We also find that fixing selected SNAP angles and removing redundant degrees of freedom substantially reduces the optimization space, enabling larger unitaries to be synthesized more efficiently.

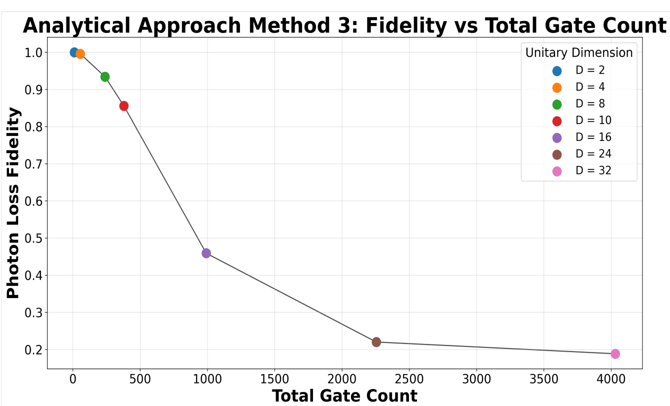

Fig. 1. Average logical fidelity as a function of total gate count for varying input Unitary dimensions

Acknowledgment

This work was supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Superconducting Quantum Materials and Systems Center (SQMS), under Contract No. 89243024CSC000002. Fermilab is managed by FermiForward Discovery Group, LLC, acting under Contract No. 89243024CSC000002.

References

[1] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, "Qudits and high-dimensional quantum computing," *Frontiers in Physics*, vol. 8, p. 589504, 2020.

[2] D. Kürkçüoğlu, J. E. Cruz Serrallés, T. Roy, L. Bove, and collaborators, "Compilation to bosonic qudit systems," Fermilab Superconducting Quantum Materials and Systems (SQMS) Center Poster, 2023.

[3] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, "Efficient cavity control with SNAP gates," arXiv:2004.14256, 2020.

[4] J. Job, "Efficient, direct compilation of SU(N) operations into SNAP & displacement gates," arXiv:2307.11900, 2023.

[5] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems*. Oxford, U.K.: Oxford University Press, 2002.

[6] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, "Universal control of an oscillator with dispersive coupling to a qubit," *Physical Review A*, vol. 92, no. 4, p. 040303, 2015.

[7] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, "Cavity state manipulation using photon-number selective phase gates," *Physical Review Letters*, vol. 115, no. 13, p. 137002, 2015.

[8] T. Kim, T. Roy, *et al.*, "Ultracoherent superconducting cavity-based multiqudit platform with error-resilient control," arXiv:2506.03286 [quant-ph], 2025.

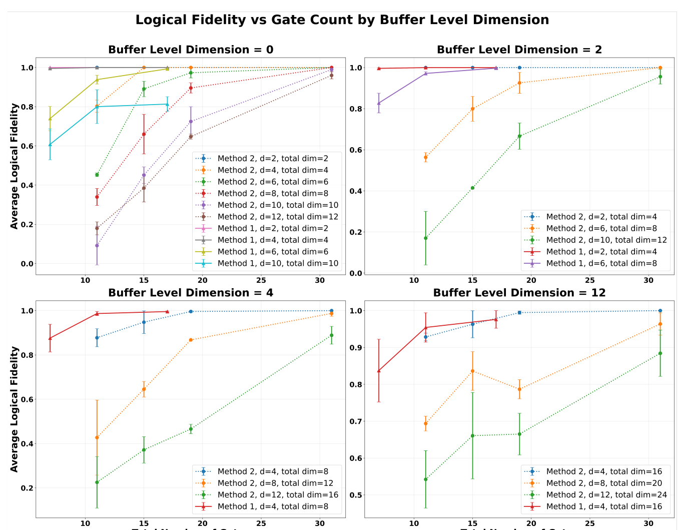

Fig. 2. Average logical fidelity as a function of total gate count for varying numbers of buffer states