

Swap Network Augmented Ansätze on Arbitrary Connectivity

Teodor Parella-Dilmé

Jakob S. Kottmann

Antonio Acín

Trainable and resource-efficient parametrizations of quantum states underpin hybrid quantum–classical algorithms. A central obstacle to their design is the limited qubit connectivity of real processors, which restricts the ability to build correlations between distant qubits in a way that is simultaneously economical in gates and trainable. The two standard responses fall short. On the one hand, hardware-efficient ansätze families [1] entangle only adjacent qubits and avoid SWAP overhead, but cannot spread long-range correlations and lose trainability (concentrating in local minima or barren plateaus) when the target state carries long-range entanglement. On the other hand, state-of-the-art dynamical routing compilers [2, 3, 4], insert many SWAP gates to compile circuits, significantly increasing circuit depth and noise.

In this work we co-design the ansatz and the routing simultaneously. First, we introduce a classical optimization algorithm that, for an arbitrary qubit connectivity graph, constructs an optimized swap network: a fixed, layered SWAP schedule after which every qubit pair has been nearest-neighbours at least once, enabling effective all-to-all connectivity. The algorithm is architecture-agnostic, excludes faulty qubits, and uses a simulated annealing optimization that tracks interaction between qubits. Second, we embed layered, connectivity-aware entangling blocks within the optimized swap network, so that routing is absorbed into the ansatz structure rather than dynamically appended afterwards (see Fig. 1).

We benchmark the optimized routing protocol against state-of-the-art compilers such as SABRE, reporting substantial reductions in circuit depth and gate count. We then demonstrate its benefits in simulations of strongly correlated ground states—random spin-glass Hamiltonians and the open-shell p-benzyne biradical, the latter in fermionic circuits for quantum chemistry. Across connectivities, the swap-augmented ansatz reaches lower energy errors with fewer entangling gates, shallower circuits, and fewer parameters than hardware-efficient baselines, translating directly into higher noise resilience at comparable accumulated noise. Because these gains stem from improved trainability and parameter efficiency, the method provides a general, hardware-aware tool for ansatz design in variational quantum optimization and quantum machine learning alike.




Figure 1: Diagram of the proposed algorithm for designing a layered ansatz embedded within a swap network with optimized qubit routing in arbitrary connectivity. **1)** The available qubit connectivity graph is taken as input. **2)** A classical optimization algorithm designs an optimized swap network for the given connectivity constraints. The optimized sequence of swaps ensures that by the end of the process each node has been adjacent to every other node at least once. **3)** The layered ansatz is embedded within the swap network, enabling direct information exchange between all qubits on the constrained connectivity hardware.

References

- [1] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.
- [2] Gushu Li, Yufei Ding, and Yuan Xie. Tackling the qubit mapping problem for nisc-era quantum devices, 2019.
- [3] Henry Zou, Matthew Treinish, Kevin Hartman, Alexander Ivrii, and Jake Lishman. Lightsabre: A lightweight and enhanced sabre algorithm, 2024.
- [4] Chenghong Zhu, Xian Wu, Zhaohui Yang, Jingbo Wang, Anbang Wu, Shenggen Zheng, and Xin Wang. Quantum compiler design for qubit mapping and routing: A cross-architectural survey of superconducting, trapped-ion, and neutral atom systems, 2025.