Energy Analysis of Classical and Quantum Algorithms for the Traveling Salesman Problem

Nathaniel James Pacey, Cherilyn Christen, Lindon Zymberi

The exponential growth in data center energy demand poses a challenge to the sustainability of future computation. According to the IEA, data centers may consume up to 8% of global electricity by 2030 [4]. Among the most energy-intensive workloads are NP-hard optimization problems like the Traveling Salesman Problem (TSP), which underlies logistics applications including routing for UPS and Amazon [5]. Certain quantum algorithms offer an alternative with potential reductions in both time and energy complexity.

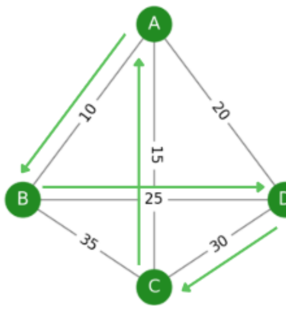

Figure 1: Traveling Salesman Graph (4 Cities)

The TSP seeks the shortest tour through n cities, visiting each once and returning to the start. Its cost is given by:

$$\text{Cost}(\pi) = \sum_{i=0}^{n-2} d(\pi(i), \pi(i+1)) + d(\pi(n-1), \pi(0))$$

Classically, this can be solved via brute-force search ($O(n!)$) or the Held-Karp algorithm ($O(n^2 \cdot 2^n)$) [6, 7], both used here as benchmarks.

To solve TSP on quantum hardware, the Quantum Approximate Optimization Algorithm (QAOA) was implemented using Qiskit and PennyLane. The cost and mixer Hamiltonians are derived from a MaxCut formulation, compiled into sequences of RZ, RX, and CNOT gates [1].

QAOA prepares the state:

$$|\psi(\vec{\gamma}, \vec{\beta})\rangle = \prod_{k=1}^{p} e^{-i\beta_k H_M} e^{-i\gamma_k H_C} |+\rangle^{\otimes n}$$

where H_C encodes the cost and H_M is the mixer Hamiltonian [1]. For a 4-city problem with depth $p = 2$, the circuit was constructed using real IBM Sherbrooke backend data [2].

Gate counts scale as follows for n cities: n Hadamards, n RX, $\frac{n(n-1)}{2}$ RZ, $n(n-1)$ CNOTs, and n measurements.

Energy usage was modeled using $E = P \cdot t$, with quantum systems assumed to draw 25 kW [8] and classical CPUs 200 W [3]. Using backend data, we extracted gate durations: 1Q = 57 ns, 2Q = 533 ns, and measurement = 1.216 μs [2].

For a 4-city TSP, one QAOA execution consumed 1.34×10^{-7} kWh (0.48 J), and the full cold-start training (200 steps) required 96.3 J. Pretrained (hot-start) QAOA eliminates this overhead. Brute-force and Held-Karp scale steeply, while QAOA becomes more efficient beyond 22 cities (hot start) and 30 cities (cold start), which matches expectations.

These findings highlight QAOA's potential as an energy-efficient alternative to classical optimization for large problem sizes. With realistic gate times and hardware models, quantum computing offers tangible energy advantages in domains such as logistics, where energy costs directly affect emissions and infrastructure demands.

This supports the UN Sustainable Development Goals: SDG 9 by promoting energy-aware computing innovation, SDG 11 by enabling more efficient transport systems, and SDG 13 by reducing the carbon footprint of computation-heavy operations [9].

References

[1] E. Farhi, J. Goldstone, and S. Gutmann, "A Quantum Approximate Optimization Algorithm," arXiv:1411.4028, 2014.

[2] IBM Quantum, "Qiskit Runtime Backend Properties," IBM Sherbrooke, Accessed via Qiskit API, 2025.

[3] J. Koomey, "Growth in Data Center Electricity Use 2005 to 2010," Analytics Press, 2011. https://www.analyticspress.com/datacenters.html

[4] International Energy Agency, "Data Centres and Data Transmission Networks," 2022. https://www.iea.org/reports/data-centres-and-data-transmission-networks

[5] UPS, "ORION: On-Road Integrated Optimization and Navigation," Technical Overview, 2016.

[6] T. H. Cormen et al., *Introduction to Algorithms*, 3rd ed., MIT Press, 2009.

[7] M. Held and R. M. Karp, "A Dynamic Programming Approach to Sequencing Problems," *SIAM Journal*, vol. 10, no. 1, pp. 196–210, 1962.

[8] IBM Quantum Sustainability Report, 2023. Internal estimate based on IBM Q System One design.

[9] United Nations, "Transforming our world: the 2030 Agenda for Sustainable Development," United Nations General Assembly, 2015. https://sdgs.un.org/goals