IS NEAR-TERM QUANTUM SUPREMACY FEASIBLE? A RUNTIME-BASED INVESTIGATION

Why this project, and why now? Quantum computers attract headlines because, in principle, they can solve some problems much faster than today's machines. But the practical question is simpler and more important: if we count everything that matters in real use: setup, waiting in queues, device errors, and the quality of the answer, do current quantum devices actually beat the best classical computers? The answer determines where public and private investment should go: quantum hardware, classical accelerators, or clever hybrids that combine both. This project replaces hype with hard numbers and clear, fair tests, so that science, industry, and policy can move forward on solid ground.

What is the goal? We will deliver a clear, evidence-based assessment of whether near-term quantum devices can provide **faster, good-enough answers** than finely tuned classical systems under operationally meaningful conditions. Concretely, we will define fair benchmarks and measure performance end-to-end, from job submission to usable results, including energy use.

For each task we will build **strong classical baselines** on modern hardware: multi-core CPUs, powerful GPUs, reconfigurable chips called **FPGAs**, and wafer-scale processors that place an entire silicon wafer on one giant chip. In parallel, we will run the best available **quantum** approaches on cloud-accessible devices.

A key innovation is our use of **end-to-end, user-centric metrics**. Rather than timing only what happens inside the chip, we measure how long it takes to obtain an answer of a given quality from the moment a job is submitted until the result is ready to use. This includes preparation, compilation, queueing, execution, and post-processing. We will also measure **energy per answer**, because an "advantage" that requires excessive electricity is not practical.

Why attempt this now? Three trends make this research urgent:

- **Quantum devices are improving** but still face noise, overheads, and limited scale.
- **Classical methods keep leaping forward**: new algorithms and specialized chips have repeatedly matched or erased earlier quantum wins.
- **Decisions depend on clarity**: funders, policymakers, and industry need reliable guidance on where quantum truly helps today and what must change for it to help tomorrow.

In short, the field is at a tipping point. Honest, reproducible, apples-to-apples evidence is essential.

How will we measure success? We will report results in plain, decision-relevant terms:

- **Time-to-useful-answer**: how long to reach an agreed quality level.
- **Reliability under a time budget**: chance of success within, say, one minute or one hour.
- **Fidelity for sampling tasks**: how close the outputs are to what they should be.
- **Energy per answer**: electricity cost to reach that quality.

All results will include uncertainty estimates and be **fully reproducible** (open code, datasets, and configuration logs). We will pre-register benchmarks and publish the evaluation tools to set a community standard others can adopt.

In a sentence: This project will show exactly when quantum (alone or in hybrid form) can beat the best classical systems, and when it cannot, using fair, end-to-end measurements that decision makers can trust.