

Quantum Inspired

EE599-001 & EE699-010, Spring 2026

Hank Dietz

<http://aggregate.org/hankd/>

Why?

- Quantum computers might be much faster to solve various problems, but are they?
 - The disappointing scale of NISQ hardware
 - What are the basic quantum operations?
 - The search for quantum supremacy
- Quantum computing is re-thinking some basic assumptions of computing... maybe re-thinking conventional computation is good too?

Quantum-Inspired (QI) Algorithms

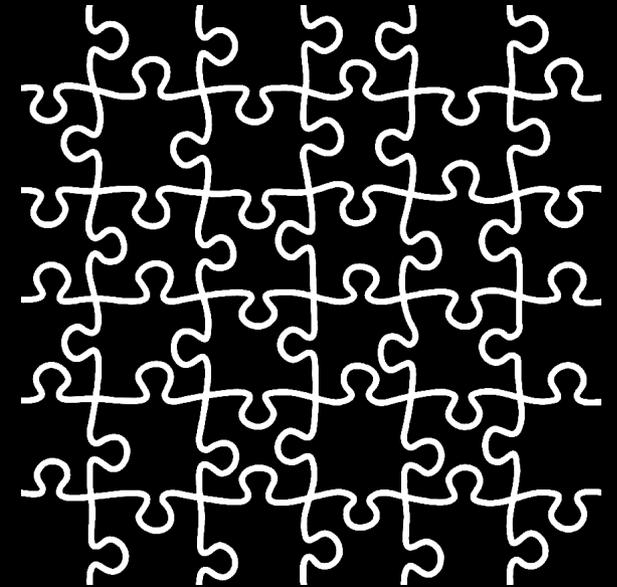
- Quantum algorithms operate on entangled superpositions, often using interference and amplitude amplification...
- Can approximate quantum qubit operations by linear algebra on exponentially large vectors
- QI exploits structure (e.g., sparsity or sampling) to achieve sub-linear or poly-log runtime
- Is quantum speedup obtained if conventional code adopts a quantum algorithm's approach?

Quantum Supremacy

- Quantum supremacy is measured against an *assumed classical “hardness”* of the *problem*
- Converting a quantum algorithm into code that can execute conventionally, “*dequantization*,” doesn’t invalidate; it clarifies where quantum advantage lies
- Understanding QI algorithms provides better insight into future quantum hardware claims

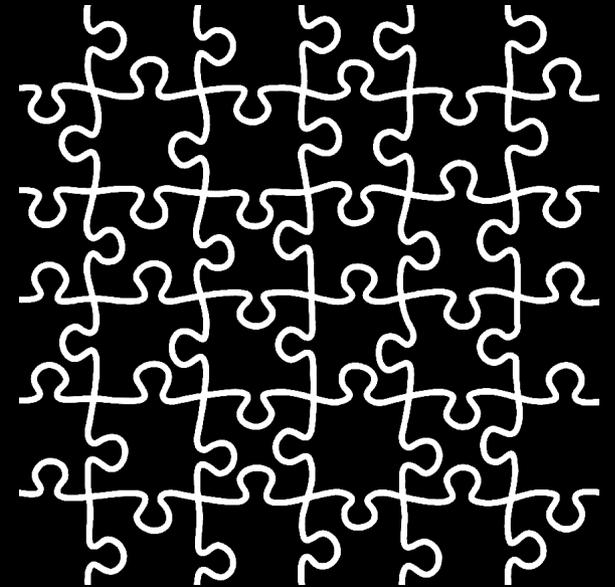
A Benchmark Problem

- Solving a jigsaw puzzle was proposed as a benchmark
- Each piece has 4 edges with uniquely numbered shapes
- Use *any algorithm* to match piece edges with same number
- A variant of this is **Quantum Awesomeness**, intended to show quantum supremacy



A Problematic Benchmark

- Solving a jigsaw puzzle was proposed as a benchmark
- Each piece has 4 edges with uniquely numbered shapes
- Use *any algorithm* to match piece edges with same number
- Hashing on a conventional machine, *n-piece puzzle is solved in $O(1)$ time and $\sim O(n)$ effort*



QRAM: Quantum RAM

- QRAM model: $O(1)$ quantum query for any amplitude within a superposed vector
 - Think Grover's algorithm...
 - Initialize a superposition with a vector?
- Analogous to SIMD parallel vector indexing, where each $PE_{i\text{proc}}$ accesses $a[f(i\text{proc})]$...
TMC CM implemented by $\Sigma((j==i\text{proc})*a[j])$
- The hard part for quantum is initializing $a[]$, not indexing (which is basically Grover's alg.)

Quantum Recommendation Systems & Linear Algebra

- **Kerenidis & Prakash (2016)** exponentially faster quantum recommendation systems
 - Sparse user-product preference $m \times n$ matrix
 - Quantum samples k -rank approximation in $O(\text{poly}(k, 1/\epsilon) * \text{polylog}(mn))$ for small k
 - Classical is $\Omega(mn)$
 - Prime example of practical quantum advantage for machine learning

Tang's Dequantization

- **Tang (2018)** proved polynomially-slower conventional version of Kerenidis & Prakash
 - Tree “sample-and-query” (SQ), not QRAM
 - Does ℓ^2 -norm importance sampling
 - Also $O(\text{poly}(k, 1/\epsilon) * \text{polylog}(mn))$ for small k
 - Quantum speedup required better input data than the classical algorithm was given

Generalized Classical SQ

- **Tang (2018)** generalizes to low-rank approximation of any matrix by score sampling
 - Importance of each row/col to top-k
 - Sample rows proportional to ℓ^2 -norm to approx. SVD (singular value decomposition)
 - Many quantum linear algebra algorithms dequantized using similar methods... e.g., for machine learning

The HHL Algorithm

- **Harrow, Hassidim, & Lloyd (2009)** presented quantum solution of linear systems of equations
 - Quantum time $O(\log(n) \cdot K^2 / \epsilon)$, K condition no.
 - Classical is $O(n \cdot \text{poly}(K))$
 - However, quantum output is $|x\rangle$, so it takes at least $O(n)$ repeats to measure

Classical HHL Alternatives

- Various new classical algorithms converge on solutions exponentially quickly
 - **Randomized Kaczmarz (2007)** iteratively projects onto hyperplanes
 - Stochastic gradient descent for least-squares systems converge in $O(n \cdot \log(1/\epsilon) / \lambda_{\min})$
- HHL works for sparse and well-conditioned; for those, classical iterative solvers compete

Classical HHL-Inspired

- Various dequantizations of linear algebra that is used in HHL applications
 - Quantum singular value transformation (QSVT) unifies HHL, phase estimation, and amplitude amplification
 - Chebyshev expansions approximating matrix functions as polynomials behave similarly
 - Solutions with poly-log solution sparsity have polynomial slowdown using conventional

Grover's Search

- Unstructured search on N items in $O(\sqrt{N})$ quantum queries vs. $O(N)$ conventional
 - Uses amplitude amplification
 - Quadratic speedup is “real” because $O(N)$ bound is tight, but it isn't exponential
- Structure is easily imposed in conventional searches...

Quantum Annealing

- Quantum annealing can be exponentially faster than conventional simulated annealing
 - **Simulated Quantum Annealing (SQA)** mimics quantum tunneling in a Monte Carlo method
 - Competitive with D-Wave hardware QA
- An example of quantum thinking giving speedup rather than quantum HW...

QAOA and Tensor Networks

- **Quantum Approximate Optimization Algorithm (QAOA)** is a hybrid combinatorial optimization algorithm for NISQ; e.g., QUBO, MaxCut
 - p alternating cost and mixer Hamiltonians; QAOA at depth p has $2p$ free parameters optimized classically
 - Tensor network contraction simulates QAOA circuits for low entanglement cases
 - For MaxCut on planar/sparse graphs, **Goemans and Williamson SDP (1995)** is faster than QAOA

qPCA

- Quantum Principal Component Analysis (qPCA) by **Lloyd, Mohseni, and Rebentrost (2013)**
 - qPCA is $O(\log n)$ vs. classical PCA is $O(n^2)$
 - Tang's dequantization with SQ-oracle with sampling PCA is also $O(\log n)$
 - qPCA only wins when classical sampling is not efficient (e.g., at higher ranks)

Boson Sampling

- **Aaronson and Arkhipov (2011)** proposed Boson sampling model of linear optical net
 - Requires #P-complete matrix permanents (#P-complete means $O(n^2 2^n)$)
 - Classical approximate simulators compete for small photon counts
 - Classical **tensor network GBS (2024)** basically shows quantum results obscured by noise...

RCS

- Google claimed Sycamore achieved quantum supremacy in 2019 with RCS
 - Original claim: ~10K years on supercomputer
 - 2021 estimate reduced to days on a cluster; then IBM did a classical simulation...
- RCS is still faster on a quantum system, but it isn't really doing anything useful except as a metric of quantum hardware quality

Other Quantum Inspired

- Various Monte Carlo stochastic simulations use amplitude weights or other quantum properties
- Various graph walk algorithms, e.g., **this**
- Quantum Natural Language Processing (QNLP) mapping compositional distributional semantics into tensor networks

Things Without Classical Approximations

- Quantum chemistry: Variational Quantum Eigensolver (VQE) and quantum phase estimation at scale
- Factoring: Shor's algorithm still wins
- Quantum error correction: fault-tolerant quantum computation with constant overhead
- Quantum communication: teleportation and entanglement-based protocols

What Does This Mean?

- Quantum advantage mostly comes from
 - Stochastic algorithms
 - Quantum concepts in classical algorithms
 - Exponential work done with single operations on entangled superpositions
- Data access model is often the limiting factor
- Are quantum phenomena necessary to efficiently operate on entangled superpositions?