




Research Faculty Summit 2018

Systems | Fueling future disruptions



An HPC Systems Guy's View of Quantum Computing


Torsten Hoefler

ETH Zurich, Switzerland (Professor)

Microsoft Quantum, Redmond (Visiting Researcher)



Who is this guy and what is he doing here?



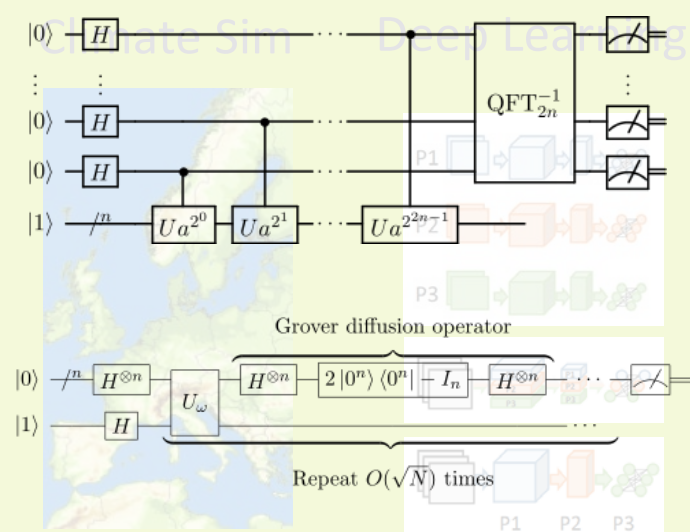


SPCL
1 professor, 6 scientific staff, 13 PhD students

ETH zürich
6.5k staff, 20k students, focus on research





Applications



Programming Systems

Q# **MPI**
LIQ $|i\rangle$

Accelerator Hardware



[1] M. Besta, TH: Slim Fly: A Cost Effective Low-Diameter Network Topology, IEEE/ACM SC14, best student paper


What is a qubit and how do I get one?

“I don't like it, and I'm sorry I ever had anything to do with it.”

Schrödinger (about the probability interpretation of quantum mechanics)

$$|\Psi\rangle = \alpha_0 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \alpha_1 \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad |\alpha_0|^2 + |\alpha_1|^2 = 1$$

$$\text{For example: } |+\rangle = \frac{1}{\sqrt{2}} |0\rangle + \frac{1}{\sqrt{2}} |1\rangle$$



One qubit can include a lot of information in α_0 and α_1 but **can only sample one bit while losing all**


(encoding n bits takes $\Omega(n)$ operations)

RESTRICTED ACCESS

NO COPY


n qubits live in a vector space of 2^n complex numbers (all combinations + entanglement)

$$|\Psi_n\rangle = \sum_{i=0..2^n-1} \alpha_i |i\rangle \quad \text{e.g., } |\Psi_2\rangle = \alpha_0 |00\rangle + \alpha_1 |01\rangle + \alpha_2 |10\rangle + \alpha_3 |11\rangle$$


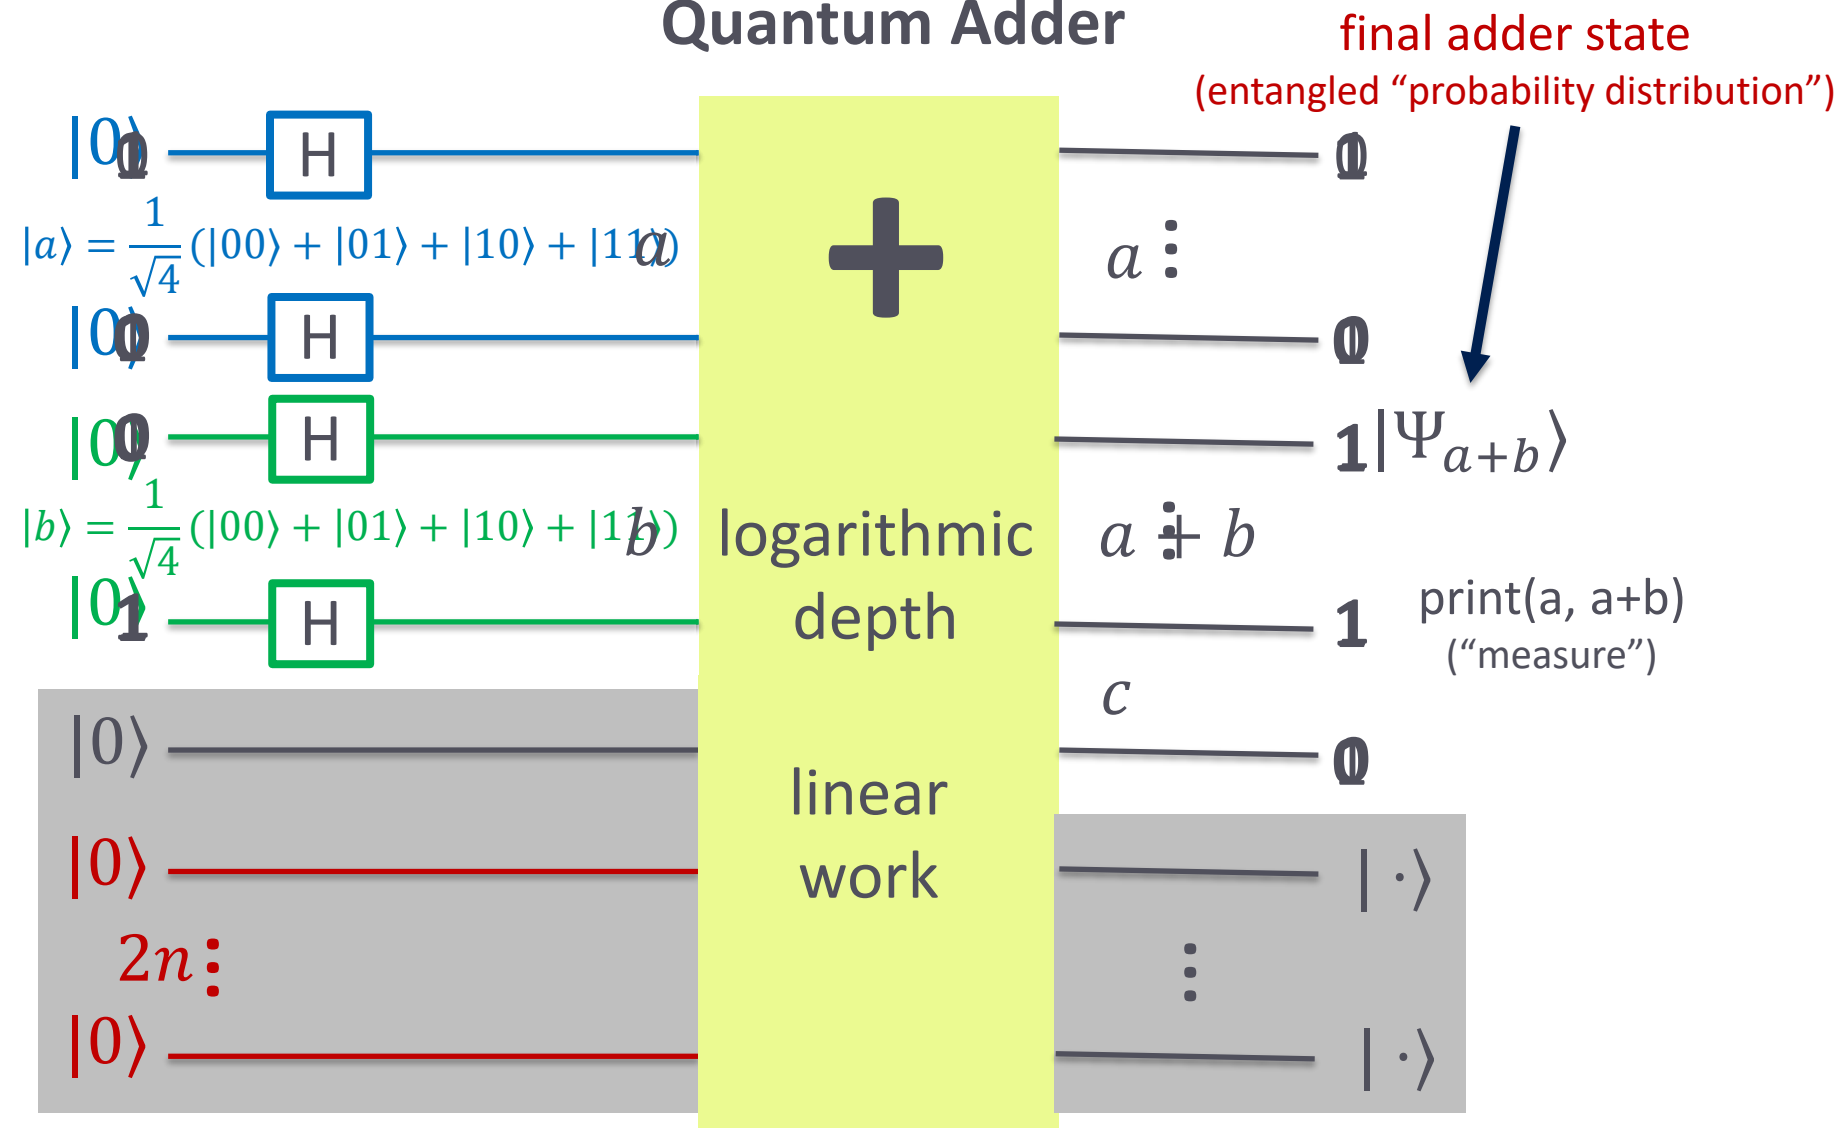


Example: adding 2^n numbers in $O(\log n)$ time

Reminder: Classical Adder



Quantum Adder





Example: adding 2^n numbers in $O(\log n)$ cycles

We add all 2^n numbers in parallel but only recover n classical bits!

A Corollary to Holevo's Theorem (1973): ***at most n classical bits can be extracted from a quantum state with n qubits even though that system requires $2^n - 1$ complex numbers to be represented!***

My corollary: *practical quantum algorithms read a linear-size input and modify an exponential-size quantum state such that the correct (polynomial size) output is likely to be measured.*



Question: Are quantum algorithms good at solving problems where a solution is verifiable efficiently (polynomial time)? Answer: Kind of 😊




So quantum computers can solve NP-complete problems!?

A problem is in NP if a solution can be verified deterministically in polynomial time.

- Even quantum computers may not solve NP-hard problems (limited by linearity of operators). But since quantum is at least as powerful as classic, we do not know!
- New complexity class: **B**ounded-**Q**uantum **P**olynomial time (BQP)




Hardware and software architecture for quantum computing



Classical Logic

Logical layer 	qubits Quantum circuits	bits Instruction stream + data	abstraction Q# programming language Q intermediate representation	SW
Gate synthesis 	Build Gates (T, Rotation, multi-control, ...)	Factory control and qubit routing	Microcoded instructions	MW
Q error correction 	QEC Codes (Steane etc., 1:n mapping)	Control for QEC (varies with code)	Microcoded instructions	MW
Physical control 	Physical quantum state	Analog pulse generators	Qubit control pulses	HW




Full Example: Grover's search



allocate $\lceil \log_2 |D| \rceil$ qubits

- Task: find $x \in D$ for which $f(x) = y$ (invert $f(x)$)
 - Classical requires $O(|D|)$ queries
 - Quantum requires $O(\sqrt{|D|})$ queries



```
operation GroverSearch(n_searchQubits: Int): (Result[]) {
  body {
    mutable resultElement = new Result[n_searchQubits];


    using (qubits = Qubit[n_searchQubits]) {
      ApplyToEachCA(H, qubits); // qubits to uniform superposition

      let n_iterations = Floor(0.25 * PI()
        * Sqrt(ToDouble(2^n_searchQubits)));

      // Grover iteration
      for (nonce in 1..n iterations) {
        OracleAND(qubits); // flips phase of desired state

        // apply Grover diffusion operator
        ApplyToEachCA(H, qubits);
        ApplyToEachCA(X, qubits);
        (Controlled Z)(qubits[1..n_searchQubits-1], qubits[0]);
        ApplyToEachCA(X, qubits);
        ApplyToEachCA(H, qubits);
      }
      set resultElement = MultiM(qubits);
    }
    return (resultElement);
  }
}
```

Q# code



Quadratic speedup? Grover on a real machine

Performance estimates must be understood to be believed (inspired by Donald Knuth's "An algorithm must be seen to be believed")

1. Query complexity model – how algorithms are developed

- $T = \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil$ queries ($|D| = 2^n$ - represented by n bits)

2. Express (oracle and diffusion operator) as n -bit unitary

- Assuming O n -bit operations for oracle!
- $T = O \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil$ n -bit operations - $T_t = \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil$

3. Decompose unitary into two-bit (+arbitrary rotation) gates

- $T = O_2 \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil \cdot 2(n-1)$ elementary operations - $T_t = \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil \cdot 4(n-1)$

4. Design approximate implementations in discrete gate set (using HTHT...)







- $T = O_2 \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil \cdot 2(n-1)$ discrete T gate operations - $T_t = \left\lceil \frac{\pi}{4} \sqrt{2^n} \right\rceil \cdot 48(n-1)$

5. Mapping to real hardware (swaps and teleport)

- Not so simple to model, depends on oracle – potentially $\Theta(\sqrt{2^n})$ slowdown

6. Quantum error correction

- Not so simple, depends on quality of physical bits and circuit depth, huge constant slowdown




Quadratic speedup? Grover on a real machine

Performance estimates must be understood to be believed (inspired by Donald Knuth's "An algorithm must be seen to be believed")

1. Query complexity model – how algorithms are developed

Quantum computer with logical error rates $\leq 10^{-24}$ and gate times of 10^{-6} s vs. classical at 1 teraop/s.

38 billion years



2. Express (oracle and diffusion operator) as n-bit unitary

3. Decompose unitary into (single-qubit rotation) gates

4. Design approximate implementation in discrete gate set


5. Mapping to real hardware (swaps and teleport)

6. Quantum error correction

k	T	#gates Clifford	T	depth overall	#qubits
128	$1.19 \cdot 2^{86}$	$1.55 \cdot 2^{86}$	$1.06 \cdot 2^{80}$	$1.16 \cdot 2^{81}$	2,953
192	$1.81 \cdot 2^{118}$	$1.17 \cdot 2^{119}$	$1.21 \cdot 2^{112}$	$1.33 \cdot 2^{113}$	4,449
256	$1.41 \cdot 2^{151}$	$1.83 \cdot 2^{151}$	$1.44 \cdot 2^{144}$	$1.57 \cdot 2^{145}$	6,681

Table 5. Quantum resource estimates for Grover's algorithm to attack AES- k , where $k \in \{128, 192, 256\}$.



from Grassl et al.: "Applying Grover's algorithm to AES: quantum resource estimates", arXiv:1512.04965



Real applications?

Quantum Chemistry/Physics

- Original idea by Feynman – use quantum effects to evaluate quantum effects
- Design catalysts, exotic materials, ...



Breaking encryption & bitcoin

- Big hype – destructive impact – single-shot (but big) business case
- Not trivial (requires arithmetic) but possible






Your connection is not private



Accelerating heuristical solvers

- Quadratic speedup can be very powerful!
- Requires much more detailed resource analysis → systems problem





Quantum machine learning

- Feynman may argue: “quantum advantage” assumes that circuits cannot be simulated classically → they represent very complex functions that could be of use in ML?

TOM BRIMONTE, BUSINESS 05.19.16 07:00 AM

**GOOGLE, ALIBABA SPAR OVER
TIMELINE FOR 'QUANTUM
SUPREMACY'**




Thanks!



- Special thanks to Matthias Troyer and Doug Carmean
- Thanks to: Thomas Haener, Damian Steiger, Martin Roetteler, Nathan Wiebe, Mike Upton, Bettina Heim, Vadym Kliuchnikov, Jeongwan Haah, Dave Wecker, Krysta Svore
- And the whole MSFT Quantum / QuArC team!

All used images belong to the respective owners – source list in appendix

me on the rocky
path to develop my
intuition for
quantum computation



How does a quantum computer work?

Qubits are arranged on a (commonly 2D) substrate

Reuse big parts of process technology in microelectronics

Quantum systems are most naturally seen as accelerators

Work in close cooperation with a traditional control circuit

Quantum circuits use predication (no control flow)

Circuit view simplifies reasoning but requires classical envelope

Qubits are error prone, need to be highly isolated (major challenge)

Quantum error correction enabled the dream of quantum computers

Commonly limited to neighbor interactions between qubits



Limited range, may require swapping across chip

Operations ("gates") are applied to qubits in place!

As opposed to bits flowing through traditional computers!

Operations ("gates") have highly varying complexity

Some are literally free (classical tracking), some are very expensive




Backup

For the unexpected discussions


Physical qubit implementations

- Photons
 - Polarization encoding, number of photons (Fock state), arrival timing
- Electrons
 - Spin, charge (number),
- Nuclei
 - Spin through NMR
- Josephson junctions/superconducting
 - Charge (incl. transmon), flux, phase
- Quantum dots
 - Spin, electron localization
- Majorana
 - MZMs



Elements of a quantum computation


- Classical control flow
 - Execute outer loop and classical parts of programs ($W=?$, $D=?$)
- Gate application, Measurement, Initialization
 - Clifford gates ($W=1$, $D=1$), T gates ($W=?$, $D=?$), Measurement ($W=?$, $D=?$), Initialization ($W=?$, $D=1$)
 - Qubit control (depends on technology, may be complex as well)
- Error correction
 - Surface code ($W=?$, $D=?$, distribution?), Steane code ($W=?$, $D=?$, distribution)
- Input-specific recompilation
 - (inputs may change during execution due to measurement results, maybe precompile, check algorithms)



Basic components for a universal quantum computer

- State preparation (usually $|0\rangle$)
- Gate application (a universal set is H, CNOT, $\pi/8$ phase rotation)
- Measurement (sufficient in standard basis $|0\rangle, |1\rangle$)
- Wait (apply identity, to sync with operations on other qubits)

- All needs to be performed fault-tolerant!



Thank you!

