Quantum Linear Algebra and its applications

András Gilyén

Alfréd Rényi Institute of Mathematics
Budapest, Hungary

Summer school on QA and QEC, CWI, Amsterdam
2026 August 24-28

Quantum algorithm design – a unifying perspective

A bird's eye view on quantum linear algebra

Motivating example - the quantum matrix inversion (HHL) algorithm

We want to solve large systems of linear equations

$$Ax = b.$$

A quantum computer can nicely work with exponential sized matrices!
Given $|b\rangle$, we can prepare a solution $\propto A^{-1}|b\rangle$.

Matrix arithmetic on a quantum computer using block-encoding

$$\text{Input matrix}: A; \ \text{Implementation}: U = \begin{bmatrix} A & \cdot \\ \cdot & \cdot \end{bmatrix}; \ \text{Algorithm}: U' = \begin{bmatrix} f(A) & \cdot \\ \cdot & \cdot \end{bmatrix}.$$

In HHL $f(x) = \frac{1}{x}$. Use Singular Value Transformation to approximate it!

More examples

- *Optimal Hamiltonian simulation [Low et al.], quantum walks [Szegedy]*
- *Fixed point [Yoder et al.] and oblivious amplitude amplification [Berry et al.]*
- *HHL, regression [Chakraborty et al.], SDPs & LPs [Brandão et al.], ML [Kerendis et al.]*

Block-encoding

A way to represent large matrices on a quantum computer efficiently

$$U = \begin{bmatrix} A & \cdot \\ \cdot & \cdot \end{bmatrix} \quad \Longleftrightarrow \quad A = (\langle 0|^a \otimes I)\, U\, (|0\rangle^a \otimes I)\,.$$

Any complex matrix A with operator norm $\|A\| \leq 1$ can be block-encoded.

One can efficiently construct block-encodings of

- an efficiently implementable unitary U,
- a sparse matrix with efficiently computable elements,
- a matrix stored in a clever data-structure in a QRAM,
- a density operator ρ given a unitary preparing its purification.
- a POVM operator M given we can sample from the rand.var.: $\mathrm{Tr}(\rho M)$,

Example: Block-encoding sparse matrices

Suppose that A is s-sparse and $|A_{ij}| \leq 1$ for all i, j indices. Given "sparse-access " we can efficiently implement unitaries preparing "rows"

$$R\colon |0\rangle|0\rangle|i\rangle \to |0\rangle \sum_k \frac{(\sqrt{A_{ik}})^*}{\sqrt{s}}|i\rangle|k\rangle + |1\rangle|i\rangle|\text{garbage}\rangle,$$

and "columns"

$$C\colon |0\rangle|0\rangle|j\rangle \to |0\rangle \sum_\ell \frac{\sqrt{A_{\ell j}}}{\sqrt{s}}|\ell\rangle|j\rangle + |2\rangle|j\rangle|\text{garbage}\rangle,$$

$R^\dagger C$ then forms a block-encoding of $A/s = (\langle 0|\langle 0| \otimes I)R^\dagger C(|0\rangle|0\rangle \otimes I)$:

$$\langle 0|\langle 0|\langle i|R^\dagger C|0\rangle|0\rangle|j\rangle = (R|0\rangle|0\rangle|i\rangle)^\dagger \cdot (C|0\rangle|0\rangle|j\rangle) = \left(\sum_k \frac{(\sqrt{A_{ik}})^*}{\sqrt{s}}|i\rangle|k\rangle\right)^\dagger \left(\sum_\ell \frac{\sqrt{A_{\ell j}}}{\sqrt{s}}|\ell\rangle|j\rangle\right) = \frac{A_{ij}}{s}$$

Efficient matrix arithmetics

Implementing arithmetic operations on block-encoded matrices

- Given block-encodings A_j we can implement convex combinations.
- Given block-encodings A, B we can implement block-encoding of AB.

Linear combination of (non-)unitary matrices [Childs and Wiebe '12, Berry et al. '15]

Suppose that $U = \sum_i |i\rangle\langle i| \otimes U_i$, and $P : |0\rangle \mapsto \sum_i \sqrt{p_i}|i\rangle$ for $p_i \in [0, 1]$.
Then $(P^\dagger \otimes I)U(P \otimes I)$ is a block-encoding of $\sum_i p_i U_i$.
In particular if $(\langle 0| \otimes I)U_i(|0\rangle \otimes I) = A_i$, then it is a block-encoding of

$$\sum_i p_i A_i.$$

Quantum Singular Value Transformation (QSVT)

Our main theorem about QSVT

Let $P\colon [-1,1] \to [-1,1]$ be a degree-d odd polynomial map. Suppose that

$$U = \begin{bmatrix} A & \cdot \\ \cdot & \cdot \end{bmatrix} = \begin{bmatrix} \sum_i \varsigma_i |w_i\rangle\langle v_i| & \cdot \\ \cdot & \cdot \end{bmatrix} \Longrightarrow U_\Phi = \begin{bmatrix} \sum_i P(\varsigma_i) |w_i\rangle\langle v_i| & \cdot \\ \cdot & \cdot \end{bmatrix},$$

where $\Phi(P) \in \mathbb{R}^d$ is efficiently computable and U_Φ is the following circuit:

Alternating phase modulation sequence $U_\Phi :=$

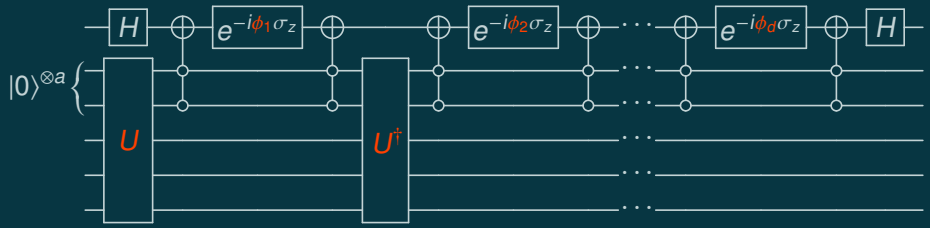

Simmilar result holds for even polynomials.

Direct implementation of HHL / the (pseudo)inverse

Singular value decomposition and (pseudo)inverse

Suppose $A = W\Sigma V^\dagger$ is a singular value decomposition.
Then the (pseudo)inverse of A is $A^{-1} = V\Sigma^{-1}W^\dagger$, (note $A^\dagger = V\Sigma W^\dagger$).
For pseudoinverse: Σ^{-1} contains the inverses of the non-zero elements of Σ.

Implementing the (pseudo)inverse using QSVT

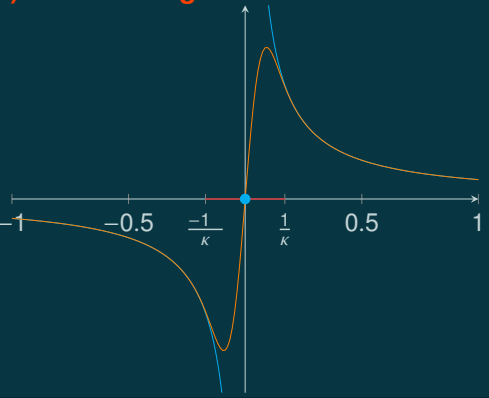

Solution: $|x\rangle \propto A^{-1}|b\rangle = (\langle 0| \otimes I)U_{A^{-1}}|0\rangle|b\rangle$ Degree / complexity: $\mathcal{O}\left(\kappa \log\left(\frac{1}{\varepsilon}\right)\right)$

Quantum walks and Hermitian matrices

Connection to Szegedy quantum walks

Szegedy defined (2004) quantisation of a symmetric Markov chain M via a product of two reflection operators. We can understand his algorithm as

$$\text{Markov chain}: M; \ \text{Updates}: U = \begin{bmatrix} M & \cdot \\ \cdot & \cdot \end{bmatrix}; \ \text{Walk}: W^n = \begin{bmatrix} T_{2n}(M) & \cdot \\ \cdot & \cdot \end{bmatrix}.$$

(T_d is the d-th Chebyshev polynomial of the first kind.)
If we choose $\phi_j = \frac{\pi}{2}$ for all $j \in \{1, \ldots, d\}$, we get $P = \pm T_d$ in QSVT.

Quantum Fast-Forwarding Markov Chains [Apers & Sarlette (2018)]

Simulate t classical steps using $\propto \sqrt{t}$ quantum operations. I.e., implement

$$U' = \begin{bmatrix} M^t & \cdot \\ \cdot & \cdot \end{bmatrix}.$$

Proof: x^t can be ε-apx. on $[-1, 1]$ with a degree-$\sqrt{2t \ln(2/\varepsilon)}$ polynomial.

The special case of Hermitian matrices

Singular value transf. = eigenvalue transf. [Low & Chuang (2017)]

Let $P\colon [-1,1] \to [-1,1]$ be a degree-d even/odd polynomial map.
If H is Hermitian, then $P(H)$ coincides with the singular value transform.

Removing parity constraint for Hermitian matrices

Let $P\colon [-1,1] \to [-\frac{1}{2},\frac{1}{2}]$ be a degree-d polynomial map. Suppose that U is an a-qubit block-encoding of a Hermitian matrix H. We can implement

$$U' = \begin{bmatrix} P(H) & \cdot \\ \cdot & \cdot \end{bmatrix},$$

using d times U and $U^\dagger$, 1 controlled U, and $O(ad)$ extra two-qubit gates.

Proof: let $P_{\text{even}}(x) := P(x) + P(-x)$ and $P_{\text{odd}}(x) := P(x) - P(-x)$ then

$$P(H) = \frac{1}{2}(P_{\text{even}}(H) + P_{\text{odd}}(H)) \quad \text{implement using QSVT + LCU}$$

Singular vector transformation and projection

Fixed-point and oblivious amplitude ampl. [Yoder et al., Berry et al.]

Amplitude amplification problem: Given U such that

$$U|0\rangle = \sqrt{p}|0\rangle\left|\psi_{\text{good}}\right\rangle + \sqrt{1-p}|1\rangle|\psi_{\text{bad}}\rangle, \quad \text{prepare } \left|\psi_{\text{good}}\right\rangle.$$

Note that $(|0\rangle\langle 0| \otimes I)U(|0\rangle\langle 0|) = \sqrt{p}|0,\psi_{\text{good}}\rangle\langle 0|$; we can apply QSVT. $P : [\sqrt{p_0}, 1] \mapsto [1-\varepsilon, 1]$

Generalization: Singular vector transformation

Given a unitary U, and projectors $\widetilde{\Pi}, \Pi$, such that

$$A = \widetilde{\Pi}U\Pi = \sum_{i=1}^{k} \varsigma_i |\phi_i\rangle\langle\psi_i|$$

is a singular value decomposition. Transform one copy of a quantum state

$$|\psi\rangle = \sum_{i=i}^{k} \alpha_i|\psi_i\rangle \quad \text{to} \quad |\phi\rangle = \sum_{i=i}^{k} \alpha_i|\phi_i\rangle.$$

If $\varsigma_i \geq \delta$ for all $0 \neq \alpha_i$, we can ε-apx. using QSVT with compl. $\mathcal{O}\left(\frac{1}{\delta}\log\left(\frac{1}{\varepsilon}\right)\right)$.

Optimal block-Hamiltonian simulation

Suppose that H is given as an a-qubit block-encoding, i.e., $U = \begin{bmatrix} H & \cdot \\ \cdot & \cdot \end{bmatrix}$.

Complexity of block-Hamiltonians simulation [Low & Chuang (2016)]

Given $t, \varepsilon > 0$, implement a unitary U', which is ε close to e^{itH}. Can be achieved with query complexity

$$O\left(t + \log(1/\varepsilon)\right).$$

Gate complexity is $O(a)$ times the above.

Proof sketch

Approximate to ε-precision $\sin(tx)$ and $\cos(tx)$ with polynomials of degree as above. Then use QSVT and combine even/odd parts.

Optimal complexity

$$\Theta\left(t + \frac{\log(1/\varepsilon)}{\log(e + \log(1/\varepsilon)/t)}\right)$$ cf. density matrix exp. $\Theta(t^2/\varepsilon)$ Lloyd et al., Kimmel et al.]

An intuitive lower bound

Lower bound on eigenvalue transformation

Suppose that U is a block-encoding of a Hermitian matrix H from a family of operators. Let $f\colon [-1,1] \to \mathbb{C}$, then implementing a block-encoding of $f(H)$ requires at least $\left\| \frac{df}{dx} \right\|_I$ uses of U, if $I \subseteq [-\frac{1}{2}, \frac{1}{2}]$ is an interval of potential eigenvalues of H.

Proof sketch

The proof is based on an elementary argument about distinguishability of unitary operators.

Optimality of pseudoinverse implementation

$$\text{Let } I := \left[\frac{1}{\kappa}, \frac{1}{2}\right] \text{ and let } f(x) := \frac{1}{\kappa x}, \text{ then } \left.\frac{df}{dx}\right|_{\frac{1}{\kappa}} = -\kappa.$$

Thus our implementation is optimal up to the $\log(1/\varepsilon)$ factor.

Summarizing the various speed-ups

Speed-up	Source of speed-up	Examples of algorithms
Exponential	Dimensionality of the Hilbert space	Hamiltonian simulation
	Precise polynomial approximations	Improved HHL algorithm
Quadratic	Singular value $=$ square root of probability	Grover search
	Singular values are easier to distinguish	Amplitude estimation
	Close-to-1 singular values are more flexible	Quantum walks

Some more applications

- Quantum walks, fast QMA amplification, fast quantum OR lemma
- Quantum Machine learning: PCA, principal component regression
- “Non-commutative measurements” (for ground state preparation)
- Sample and gate efficient metrology, fractional queries
- $\vdots$

Summary of some applications of QSVT

$\sin(tx), \cos(tx)$:

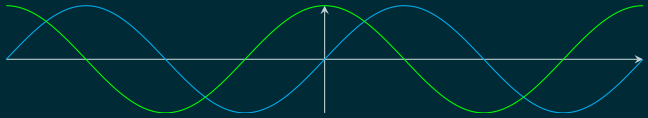

Hamiltonian simulation

$\exp(-\beta x)$:

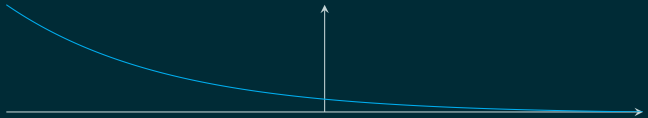

Brute-force Gibbs sampling

$T_n(x)$:

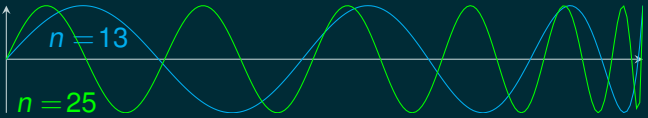

Grover search
Ampl. ampl.
Quantum walks

$\approx$ Heaviside(x) :

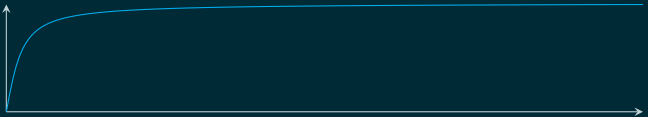

"Fixed-point" ampl. ampl.
Ground state prep.

Jordan's quantum algorithm for gradients (2004)

A generalization of the Bernstein-Vazirani algorithm

- Given a function $f\colon \mathbb{Z}_N^d \to \mathbb{Z}_N$ so that $f(x) = s \cdot x \pmod N$; find $s \in \mathbb{Z}_N^d$.
- The function is given as a phase oracle $U_f\colon |x\rangle \mapsto e^{\frac{2\pi i}{N} f(x)}|x\rangle = e^{2\pi i \frac{sx}{N}}|x\rangle$.

Recall: $QFT_N\colon |j\rangle \mapsto \frac{1}{\sqrt{N}} \sum_{\ell=0}^{N-1} e^{-2\pi i \frac{j\ell}{N}} |\ell\rangle$

Faster quantum gradient descent!

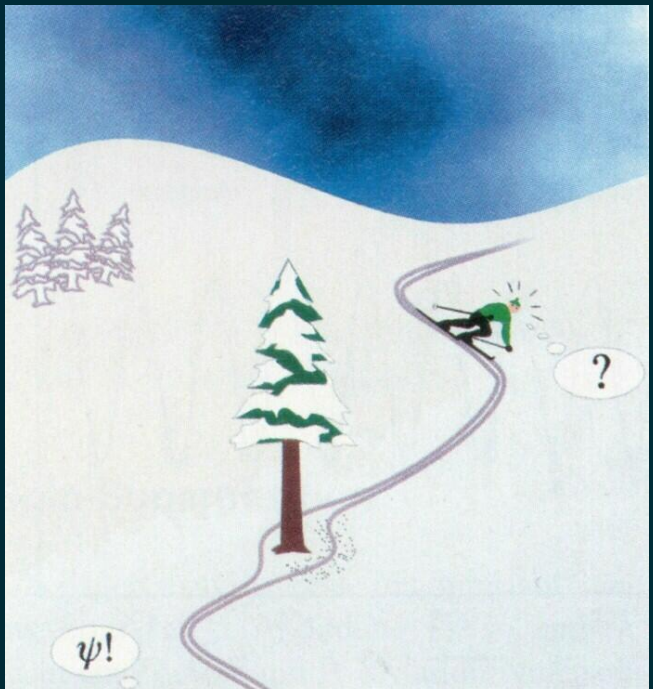

Quantum Gradient Computation Algorithm

Quantum Fourier Transform - extracting linear phase factors

Let $\varepsilon = \frac{1}{N}$ be the precision we want to achieve, and set

$$G_N := \left\{ \frac{0}{N}, \frac{1}{N}, \ldots, \frac{N-1}{N} \right\}.$$

Suppose $x, k \in G_N$ are quantum (basis) states, then

$$\sum_{x \in G_N} |x\rangle \frac{e^{2\pi i (Nxk)}}{\sqrt{N}} \xrightarrow{QFT_N} |k\rangle.$$

Quantum Gradient Computation Algorithm

Gradient computation - S. Jordan's algorithm (2004)

Input: phase oracle $O_f : |\vec{x}\rangle \to |\vec{x}\rangle e^{2\pi i f(\vec{x})}$, where $\vec{x} \in G_N^d$
Output: gradient with (hopefully) $\varepsilon = 1/N$ coordinate-wise precision

Assumption: $f(\vec{x}) \approx f(\vec{0}) + \vec{x}\nabla f(\vec{0})$, then

$$\sum_{\vec{x}\in G_N^d} \frac{|\vec{x}\rangle}{N^{\frac{d}{2}}} \xrightarrow[N\times]{O_f} \sum_{\vec{x}\in G_N^d} |\vec{x}\rangle \frac{e^{2\pi i N f(\vec{x})}}{N^{\frac{d}{2}}} \approx e^{2\pi i N f(\vec{0})} \sum_{\vec{x}\in G_N^d} |\vec{x}\rangle \frac{e^{2\pi i \left(N\vec{x}\nabla f(\vec{0})\right)}}{N^{\frac{d}{2}}} \xrightarrow[\otimes d]{QFT_N} \approx \left|\nabla f(\vec{0})\right\rangle.$$

$$\sum_{\vec{x}\in G_N^d} \frac{|\vec{x}\rangle}{N^{\frac{d}{2}}} e^{2\pi i \left(N\vec{x}\nabla f(\vec{0})\right)} = \bigotimes_{i=1}^{d} \sum_{x_i \in G_N} \frac{|x_i\rangle}{\sqrt{N}} e^{2\pi i N x_i \nabla_i f(\vec{0})}$$

Exponential speed-up?

- If we have a circuit computing f gradient computation introduces small overheads.
- "Cheap gradient principle": $\leq 4\times$ overhead for classical gradient computation

Application to purified mixed state tomography

Input model and problem statement

- Suppose we are given purified state preparation circuit
$$V\colon |\bar{0}\rangle \to |\psi\rangle_{AB}$$
such that $\mathrm{Tr}_A(|\psi\rangle\langle\psi|) = \rho$.
- We wish to estimate ρ to precision ε in trace distance

Idea: consider the linear function $f\colon X \in [-1,1]^{d\times d} \to \mathbf{Tr}(X\rho)/\sqrt{d}$

- Need to build U_f, but worst case $\|X\| = d$ (all ones matrix)
- Suppose the matrix elements of X are uniformly random in $[-1,1]$
- Apart from exponentially small probability: $\|X\| = O\left(\sqrt{d}\right)$ (matrix Chernoff bound)
- We can build block-encoding of $\mathrm{diag}(\mathrm{Tr}(X\rho)/\sqrt{d}) = \mathrm{diag}(f(X))$ for most X.
- With $1/\varepsilon$ uses of U_f we get ε coordinate-wise (independent) estimates of the gradient $\frac{\rho}{\sqrt{d}}$
- If the estimator is unbiased we very likely get $\varepsilon\sqrt{d}$ estimate of $\frac{\rho}{\sqrt{d}}$ in $\|\cdot\|$
- Implies $\varepsilon r d$ estimate of ρ in trace norm ($r = \mathrm{rank}(\rho)$) $\quad\Rightarrow\quad$ $\widetilde{O}(dr/\varepsilon)$ complexity!