FAMU-FSU College of Engineering

Data Enabled Computational Engineering and Applied Quantum Computing

Workshop on Quantum Singular Value Transformation (QSVT)

Smahane El-Halouy

FAMU-FSU College of Engineering

Tallahassee, Florida, USA

Why quantum computing?

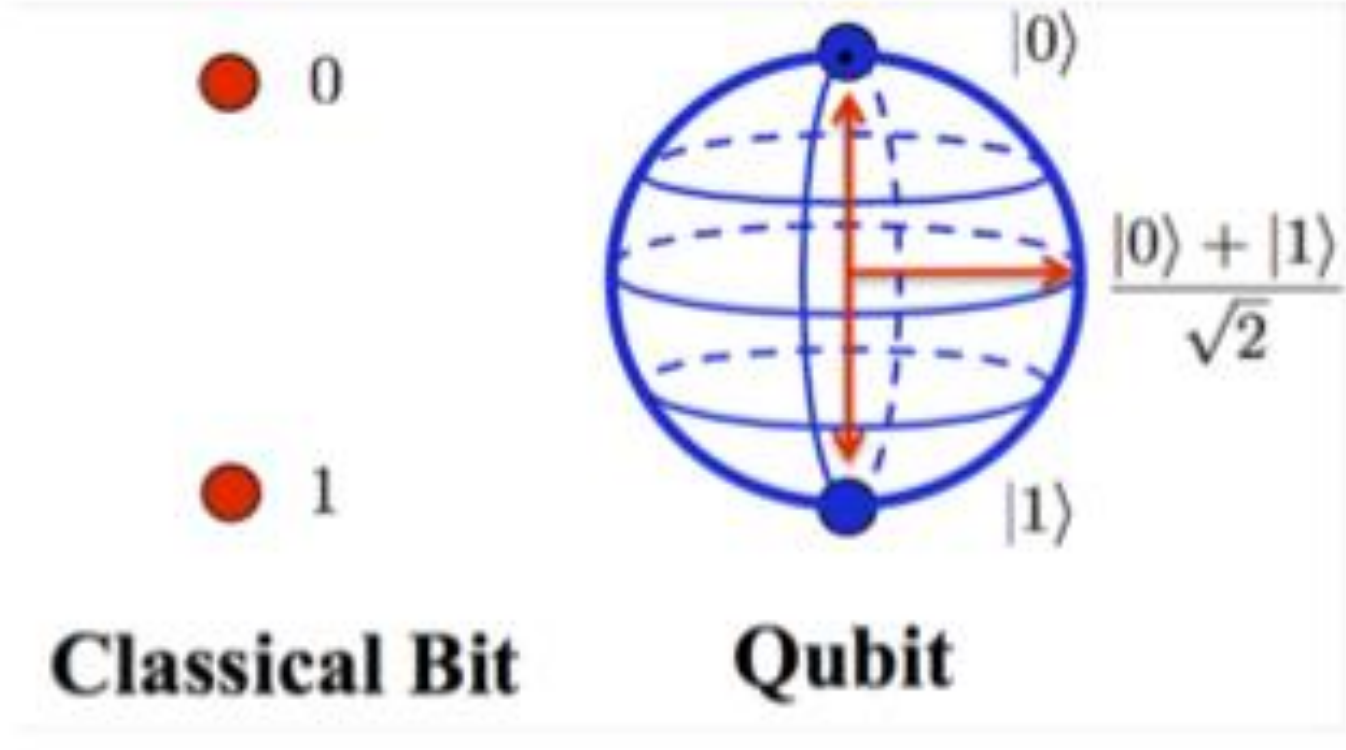

Superposition and entanglement

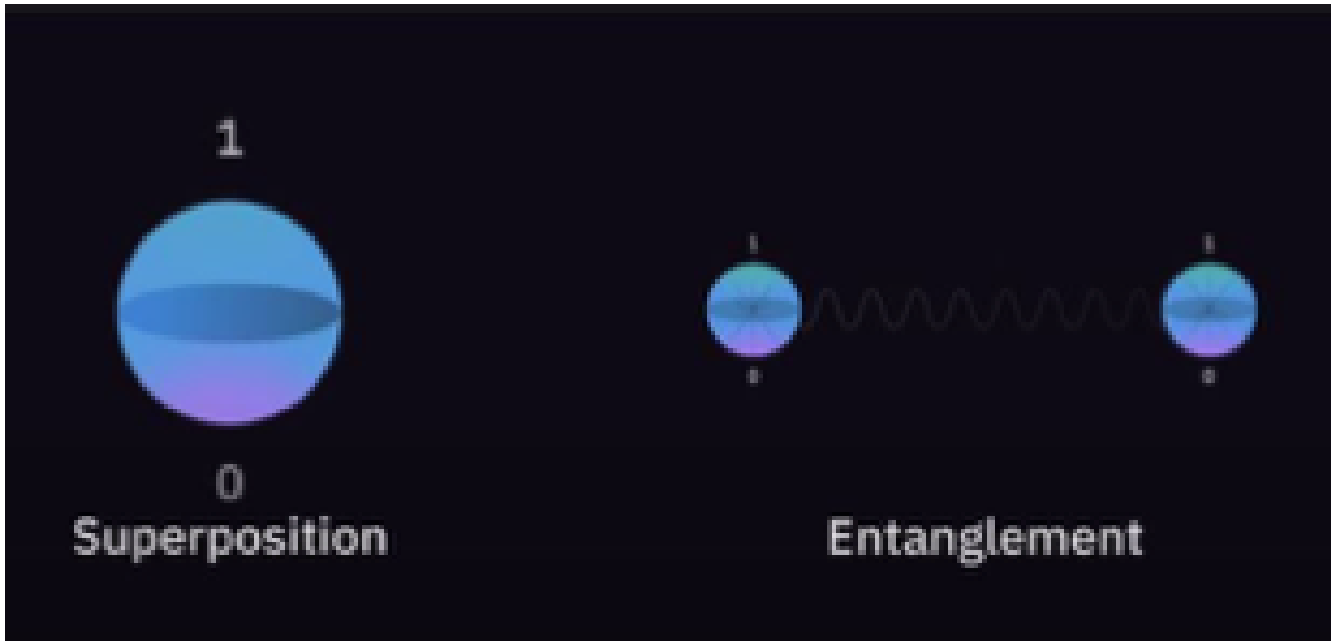

Postulates of Quantum mechanics

- State space: The set of all quantum state of a quantum system forms a Hilbert space.
  For instance, a qubit (quantum bit) forms a Hilbert space $\mathcal{H} \cong \mathbb{C}^2$.

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Postulates of Quantum mechanics

- State space: The set of all quantum state of a quantum system forms a Hilbert space.
  For instance, a qubit (quantum bit) forms a Hilbert space $\mathcal{H} \cong \mathbb{C}^2$.

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad |1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

- The evolution of a quantum state: Through a unitary

Postulates of Quantum mechanics

- Quantum measurement: To extract classical information about the quantum state

Quantum gates

Pauli Transformation Gates are four kind of gates which take one qubit and are 2×2 matrix

$$I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, Y = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

For any qubit $|\psi\rangle = \alpha \mid 0\rangle + \beta \mid 1\rangle$, the qubit is transformed as follows.

1. $I \mid \psi\rangle = \alpha \mid 0\rangle + \beta \mid 1\rangle$
2. $X \mid \psi\rangle = \beta \mid 0\rangle + \alpha \mid 1\rangle$
3. $Y \mid \psi\rangle = -\beta i \mid 0\rangle + \alpha i \mid 1\rangle$
4. $Z \mid \psi\rangle = \alpha \mid 0\rangle - \beta \mid 1\rangle$

Superposition and entanglement

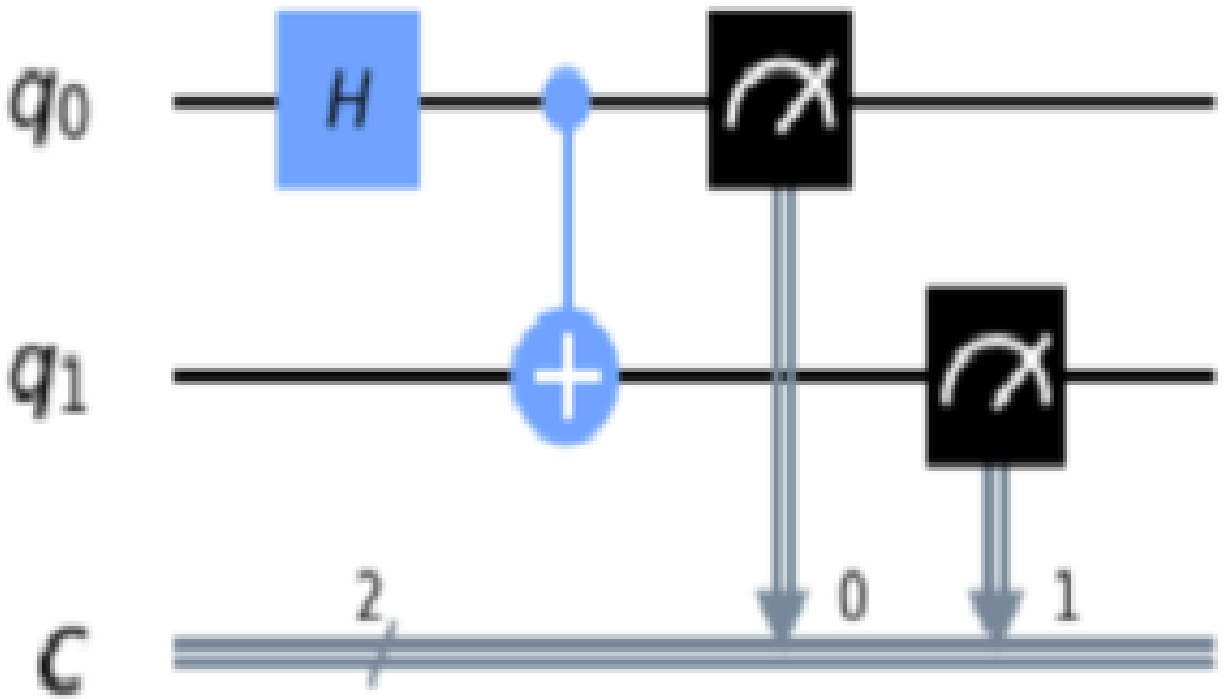

Quantum circuits

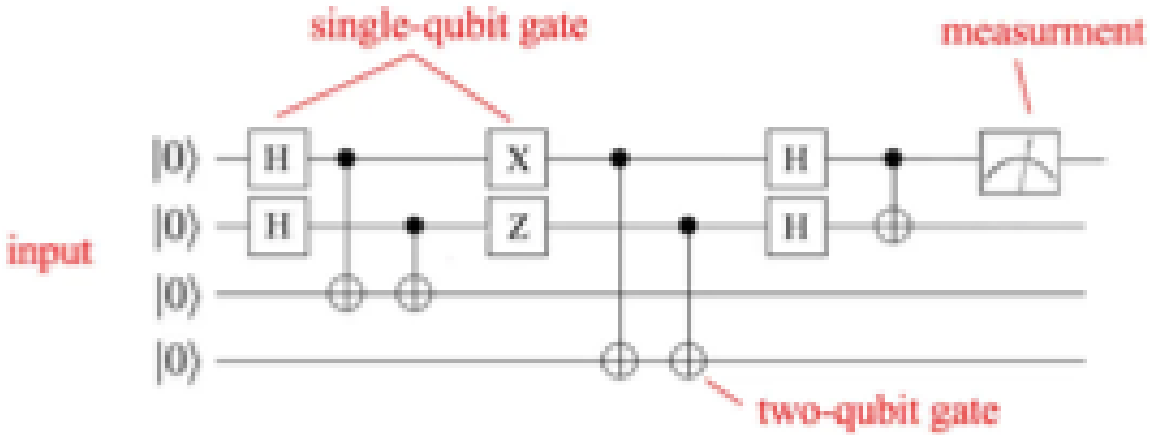

Quantum algorithms

- Grover's algorithm



Box 0 Box 1

Quantum algorithms

- Grover's algorithm



Box 0 Box 1

Mathematically, consider a boolean function $f : \{0,1\} \longrightarrow \{0,1\}$. The question is whether $f(0) = f(1)$ or $f(0) \neq f(1)$? A quantum fruit-checker assumes the access to f via the following quantum oracle:

$$U_f |x, y\rangle = |x, y \oplus f(x)\rangle$$

Quantum algorithms

- Grover's algorithm

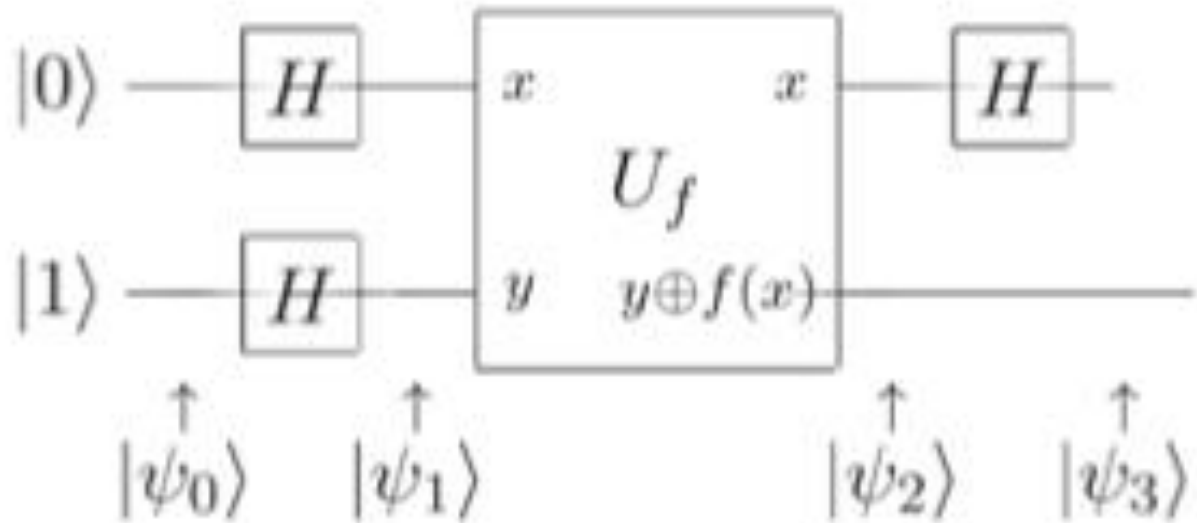

Figure: Quantum circuit implementing Deutsch's algorithm

$$|\Psi_3\rangle = \pm |f(0) \oplus f(1)\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}}\right)$$

Back to classical computing

How a computer stores the input

$$0.6_{10} = 0.100110011\ldots_2$$

$$\text{float}(0.6) = 0.5999999999999977\ldots$$

A floating-point number uses a finite collection of bits.
Most decimal fractions therefore need a small rounding.

A polynomial replaces the function

$$\cos x = 1 - \frac{x^2}{2!} + \frac{x^4}{4!} - \frac{x^6}{6!} + \cdots$$

$$p_2(x) = 1 - \frac{x^2}{2}, \qquad p_4(x) = 1 - \frac{x^2}{2} + \frac{x^4}{24}$$

A polynomial needs only additions and multiplications.
We will use the simple approximation $p_2(x)$ in our quantum circuit.

More terms improve this approximation

Reference: $\cos(0.6) \approx 0.8253356149$

Approximation at x = 0.6	Value	Absolute error
1	1.000000000	1.75e-01
$1 - x^2/2$	0.820000000	5.34e-03
$1 - x^2/2 + x^4/24$	0.825400000	6.44e-05
Previous expression $- x^6/720$	0.825335200	4.15e-07

A qubit stores two amplitudes

$$|\psi\rangle = a|0\rangle + b|1\rangle, \qquad |a|^2 + |b|^2 = 1$$

$$\Pr(0) = |a|^2, \qquad \Pr(1) = |b|^2$$

The amplitudes can be complex. Their squared magnitudes sum to one.
Measurement returns 0 or 1 with the corresponding probability.

Encoding one real number

$$W(x) = \begin{pmatrix} x & i\sqrt{1-x^2} \\ i\sqrt{1-x^2} & x \end{pmatrix}$$

$$W(0.6)|0\rangle = 0.6|0\rangle + 0.8i|1\rangle$$

For $-1 \leq x \leq 1$, place x in the amplitude of $|0\rangle$.
The other amplitude supplies the remaining probability.

Two signal gates already create a polynomial

$$\langle 0|W(x)^2|0\rangle = x^2 + (i\sqrt{1-x^2})^2$$

$$= 2x^2 - 1$$

$$x = 0.6: \qquad 2(0.6)^2 - 1 = -0.28$$

Multiplying the two gate matrices creates cross terms.
Their interference changes the amplitude.

Phase rotations control the coefficients

$$Z(\phi) = e^{i\phi Z} = R_z(-2\phi)$$

$$Z(\phi_0)\, W(x)\, Z(\phi_1)\, W(x)\, Z(\phi_2)$$

Quantum signal processing, or QSP, inserts phase gates between uses of W(x). The phase angles describe the polynomial and stay fixed as x changes.

A circuit for 1 − x²/2

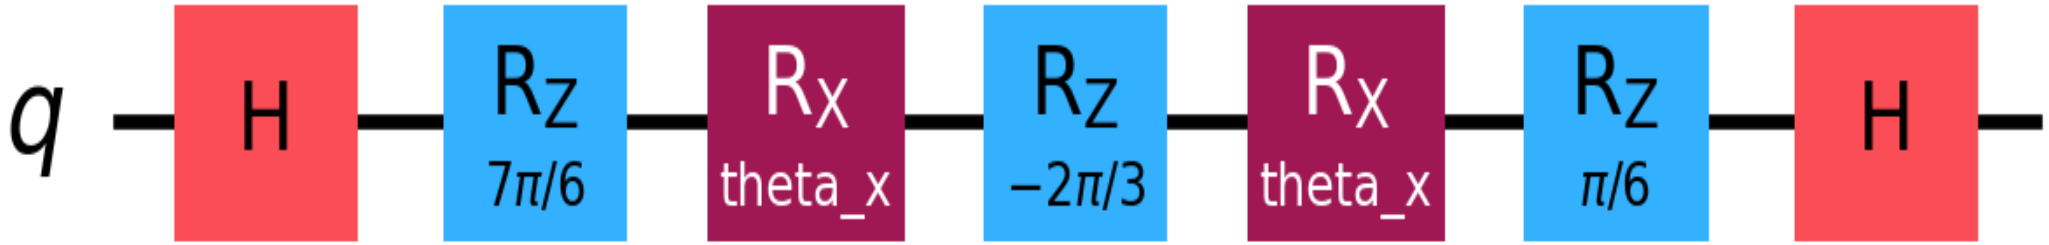

The H gates choose the input and output basis.
Each signal rotation has angle $\theta_x = -2 \arccos(x)$.

Reading a result requires measurements

$$\widehat{P}_0 = \frac{\text{number of zeros}}{\text{number of shots}}$$

$$\widehat{p_2(x)} = \sqrt{\widehat{P}_0} \quad \text{because } p_2(x) \geq 0$$

At x = 0.6, the ideal fraction of zeros is 0.6724.
Repeated measurements estimate that probability.

A rotation gives the QSP signal convention

$$W(a) = e^{i \arccos(a) X} = \begin{pmatrix} a & i\sqrt{1-a^2} \\ i\sqrt{1-a^2} & a \end{pmatrix}$$

$$W(a) = R_x(-2 \arccos a), \qquad \langle 0|W(a)|0\rangle = a$$

We use W(a), whose off-diagonal entries carry a factor i.
This choice connects directly to Qiskit's RX gate.

Repeating the signal creates Chebyshev polynomials

$$W(a)^d = e^{id\theta X}, \qquad \theta = \arccos a$$

$$\langle 0|W(a)^d|0\rangle = \cos(d\theta) = T_d(a)$$

$$d = 3: \quad T_3(a) = 4a^3 - 3a$$

The rotation angle adds when identical rotations multiply.
The top-left entry is a degree-d polynomial in a.

The polynomial circuit

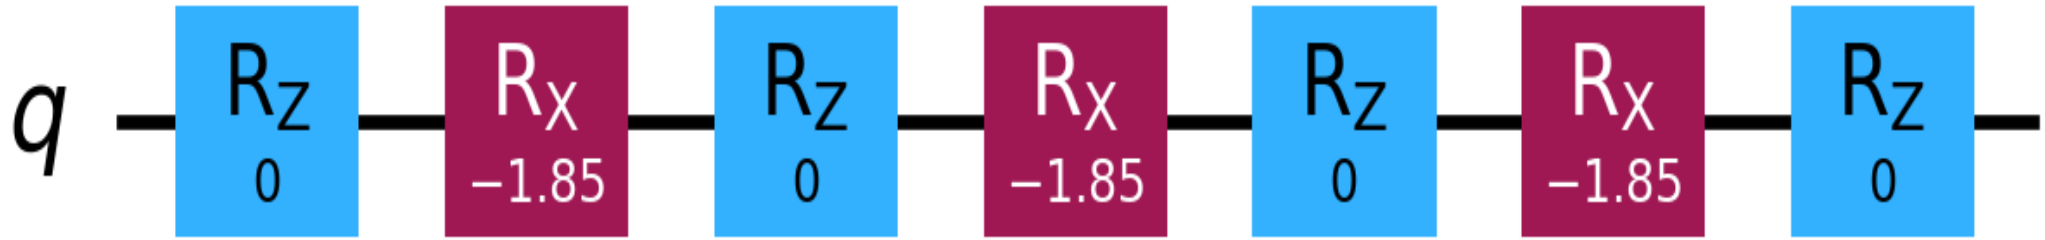

Three uses of W(a) implement $T_3(a)$ in the zero-to-zero amplitude.
The circuit diagram reads from left to right.

QSP inserts tunable phase rotations

$$Z(\phi) := e^{i\phi Z} = R_z(-2\phi)$$

$$U_\Phi(a) = Z(\phi_0)W(a)Z(\phi_1)\cdots W(a)Z(\phi_d)$$

$$\Phi = (\phi_0, \ldots, \phi_d), \qquad d \text{ signal uses}$$

The scalar a enters through W.
The phases determine the polynomial and can be chosen in advance.

The polynomial pair inside a QSP circuit

$$U_\Phi(a) = \begin{pmatrix} P(a) & i\sqrt{1-a^2}\,Q(a) \\ i\sqrt{1-a^2}\,Q^*(a) & P^*(a) \end{pmatrix}$$

$$\deg P \le d, \qquad \deg Q \le d-1$$

$$|P(a)|^2 + (1-a^2)|Q(a)|^2 = 1$$

The diagonal amplitude is a polynomial P.
The off-diagonal amplitude contains $\sqrt{(1-a^2)}$ times another polynomial Q.

From scalar to matrices and why?

Polynomials are the first matrix functions

$$p(t) = \sum_{k=0}^{d} c_k t^k \qquad \Longrightarrow \qquad p(A) = \sum_{k=0}^{d} c_k A^k, \quad A^0 = I$$

$$p(A)b = c_0 b + c_1 Ab + \cdots + c_d A^d b$$

A polynomial needs only addition, scalar multiplication and repeated matrix action.
This gives an executable definition before we discuss diagonalization.

A polynomial acts on each eigenvalue

$$A^k = (Q\Lambda Q^\dagger)^k = Q\Lambda^k Q^\dagger$$

$$p(A) = \sum_k c_k Q\Lambda^k Q^\dagger = Q\,p(\Lambda)\,Q^\dagger$$

$$p(A)q_j = p(\lambda_j)q_j$$

The powers of $Q^\dagger Q$ cancel.
Every mode keeps its eigenvector while its eigenvalue changes through p.

Definition: a function of a normal matrix

$$A = Q\Lambda Q^\dagger \quad \Longrightarrow \quad f(A) := Q\ \mathrm{diag}\big(f(\lambda_1), \ldots, f(\lambda_N)\big)Q^\dagger$$

$$f(A)b = \sum_j f(\lambda_j)\langle q_j, b\rangle q_j$$

Nonnormal matrices

$$A = X\Lambda X^{-1}: \quad \|p(A) - f(A)\|_2 \leq \kappa_2(X) \max_j |p(\lambda_j) - f(\lambda_j)|$$

$$J = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}: \quad f(J) = \begin{pmatrix} f(1) & f'(1) \\ 0 & f(1) \end{pmatrix}$$

The singular value decomposition

$$A = L\Sigma R^\dagger = \sum_j \sigma_j u_j v_j^\dagger, \qquad \sigma_j \geq 0$$

$$Av_j = \sigma_j u_j, \qquad A^\dagger u_j = \sigma_j v_j$$

$$A^\dagger A = R\Sigma^2 R^\dagger, \qquad AA^\dagger = L\Sigma^2 L^\dagger$$

Each right singular vector maps to its matching left singular vector.
The singular value gives the nonnegative scale of that map.

Singular-value functions remember orientation

$$p^{\mathrm{SV}}_{\mathrm{odd}}(A) = L\, p(\Sigma) R^{\dagger}, \qquad p^{\mathrm{SV,R}}_{\mathrm{even}}(A) = R\, p(\Sigma) R^{\dagger}$$

$$p(x) = x^3 : \; p^{\mathrm{SV}}(A) = A A^{\dagger} A$$

$$p(x) = x^2 : \; p^{\mathrm{SV,R}}(A) = A^{\dagger} A$$

An odd transformation maps input singular directions to output directions.
An even transformation can stay in the right singular-vector space.

Encoding a matrix inside a unitary

$$W_A = \begin{pmatrix} A & i\sqrt{I - A^2} \\ i\sqrt{I - A^2} & A \end{pmatrix}$$

$$(\langle 0| \otimes I) W_A (|0\rangle \otimes I) = A$$

One signal qubit identifies the desired block.
One system qubit carries the two-component vector.

Definition of a matrix block encoding

$$(\langle 0^a| \otimes I)\, U_A\, (|0^a\rangle \otimes I) = B = \frac{A}{\alpha}, \qquad \alpha \geq \|A\|_2$$

$$U_A(|0^a\rangle|b\rangle) = |0^a\rangle B|b\rangle + |\mathrm{fail}\rangle$$

The signal-zero block contains a scaled matrix.
The normalization α must be large enough for the block to be a contraction.

Singular vectors define paired input and output spaces

$$\mathcal{E}_j^{\text{in}} = \text{span}\{|0\rangle v_j, |1\rangle u_j\}, \quad \mathcal{E}_j^{\text{out}} = \text{span}\{|0\rangle u_j, |1\rangle v_j\}$$

$$\mathcal{U}_A : \ \mathcal{E}_j^{\text{in}} \longrightarrow \mathcal{E}_j^{\text{out}}, \qquad \begin{pmatrix} \sigma_j & \sqrt{1-\sigma_j^2} \\ \sqrt{1-\sigma_j^2} & -\sigma_j \end{pmatrix}$$

The matrix acts between two bases.
In the coordinates of a singular pair, the dilation is the scalar reflection R(σ).

Alternating U and U† keeps the spaces aligned

$$d \text{ odd}: \quad L\,p(\Sigma)R^\dagger$$

$$d \text{ even}: \quad R\,p(\Sigma)R^\dagger$$

$$\mathcal{U}_A,\ \mathcal{U}_A^\dagger,\ \mathcal{U}_A,\ldots \quad \text{in chronological order}$$

A forward use moves from the right space to the left space.
An adjoint use returns to the right space.
Parity records which space remains at the end.

The QSVT theorem

$$\|A\|_2 \le 1, \quad p \in \mathbb{R}[x], \quad \deg p = d, \quad |p(x)| \le 1 \text{ on } [-1,1]$$

$$\widetilde{\Pi}\, U_p\, \Pi = p^{\mathrm{SV}}(A) \quad \text{with the parity-dependent orientation}$$

Signal and extraction conventions may require an extra ancilla.

A bounded real polynomial of fixed parity can transform singular values of a block-encoded matrix.
The number of block-encoding calls scales linearly with polynomial degree.

Small example

Qiskit labels basis states as |q4 q3 q2 q1 q0⟩. Success means q4 q3 q2 = 000, so state[:4] contains the successful vector.

$$g = \frac{c}{\|b\|} A^{-1} b, \qquad p_{\text{success}} = \|g\|^2, \qquad |x\rangle = \frac{g}{\sqrt{p_{\text{success}}}}$$

Accepted states: 00000, 00001, 00010, 00011

Simulator recovery: x = ‖b‖ g / c
Hardware postselection prepares the normalized state |x⟩.

The circuit recovers the Poisson solution

$x = (0.8, 0.6, 0.4, 0.2)^T$
Success probability = 4.9536%

$$|x\rangle = \frac{4|00\rangle + 3|01\rangle + 2|10\rangle + |11\rangle}{\sqrt{30}}$$

Normalized amplitudes
0.73029674
0.54772256
0.36514837
0.18257419

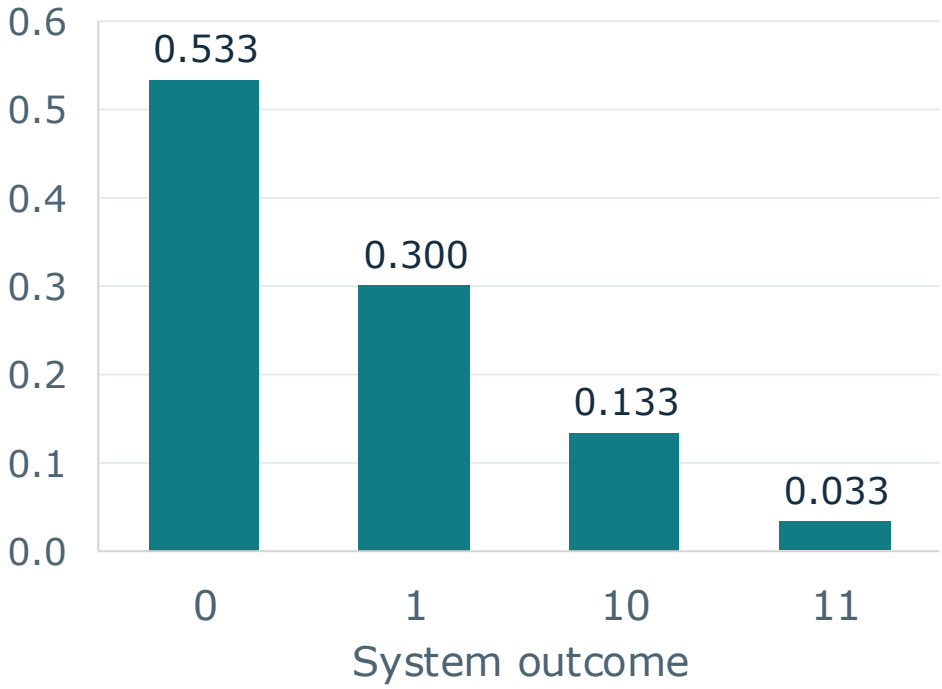

A larger problem needs interval approximation

$$B = A/\alpha, \qquad \delta = \sigma_{\min}(A)/\alpha, \qquad \widetilde{\kappa} = 1/\delta$$

$$p(z) \approx \gamma/z \quad \text{on } [-1, -\delta] \cup [\delta, 1]$$

$$d = O(\widetilde{\kappa} \log(\widetilde{\kappa}/\varepsilon)) \quad \text{for standard bounded inverse constructions}$$

For a general normalized matrix, the smallest singular value sets a gap δ.
A reciprocal approximation must remain bounded and respect parity.
Its degree and normalization determine both circuit cost and success.

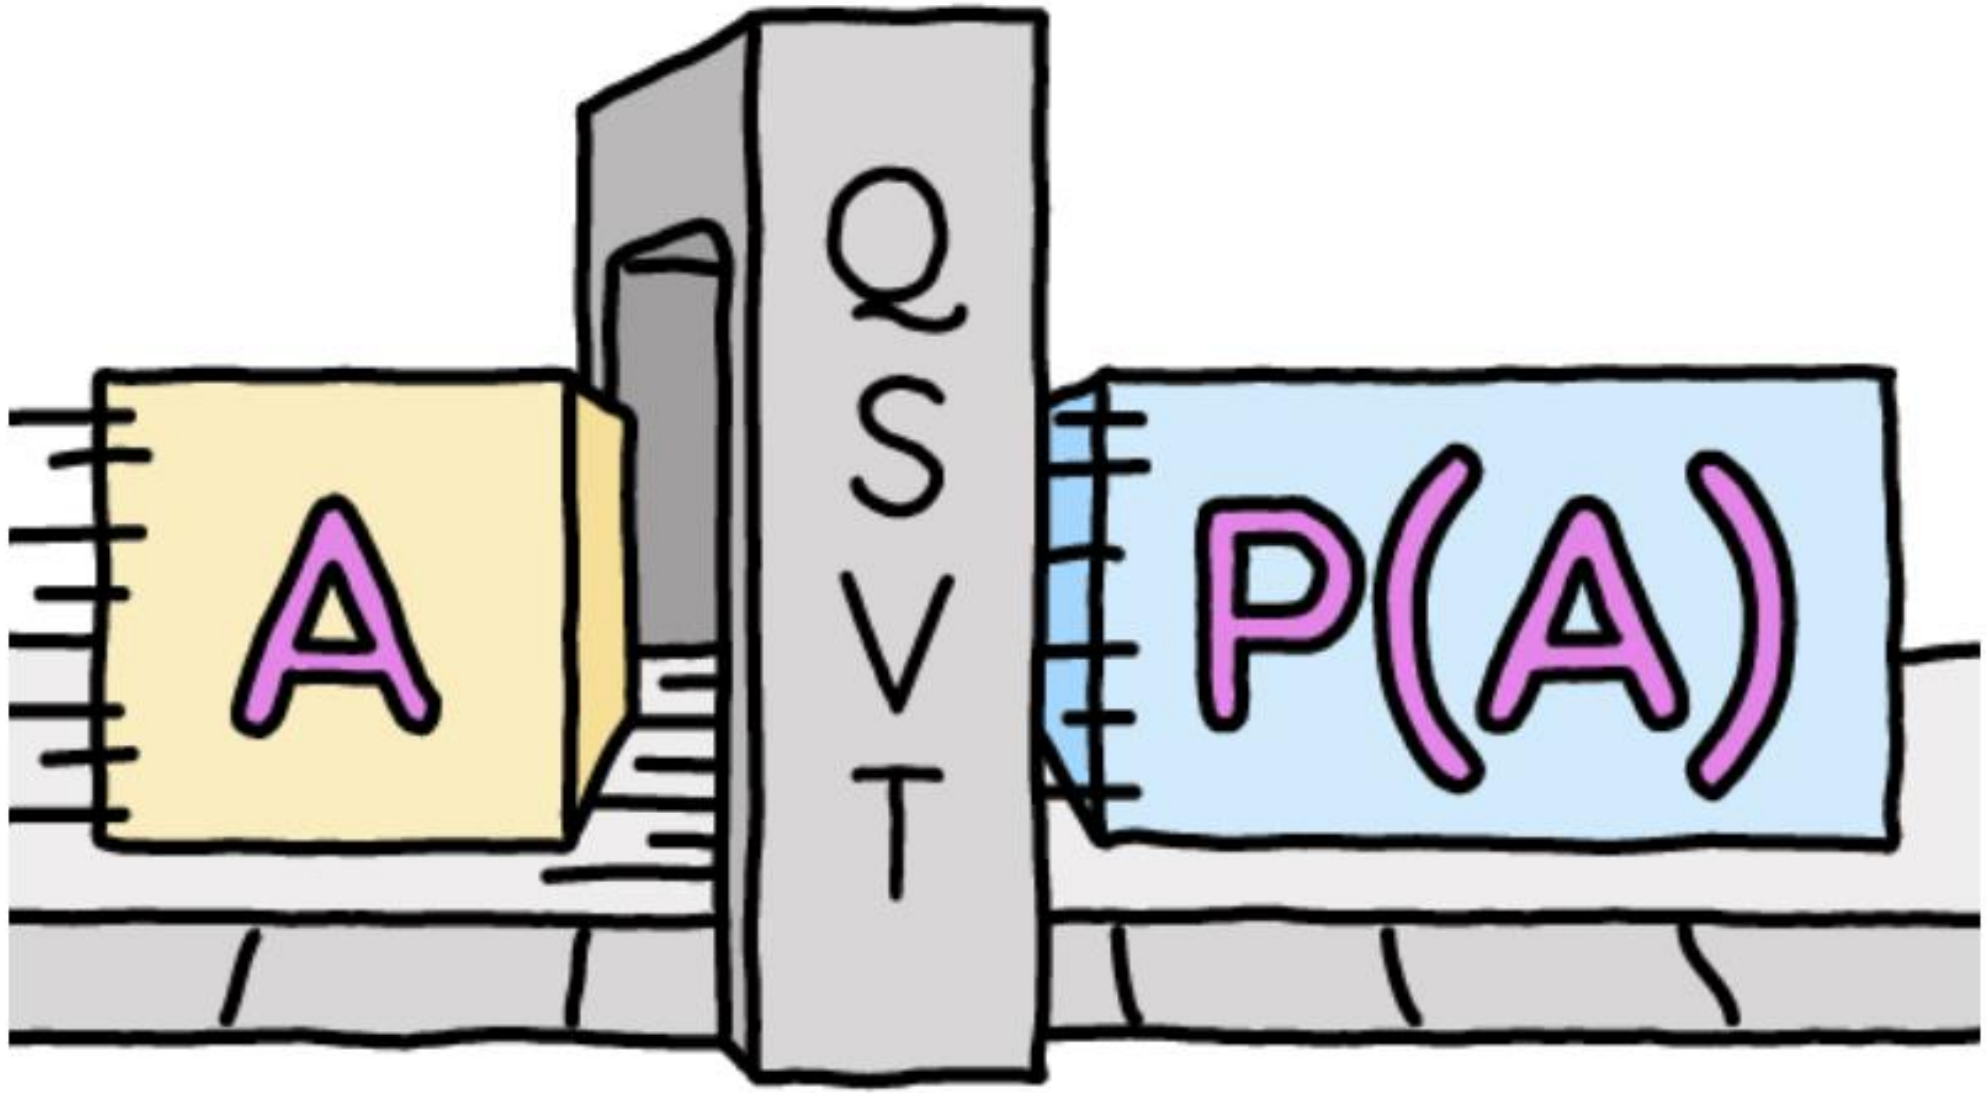
A
QSVT
P(A)