# Quantum Singular Value Transformation

Explain singular-value filtering with explicit polynomial, domain and success assumptions.

Research 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Identify polynomial boundedness and parity requirements.
- Calculate a cubic spectral filter on a diagonal example.
- Include conditioning, access and output costs in claims.

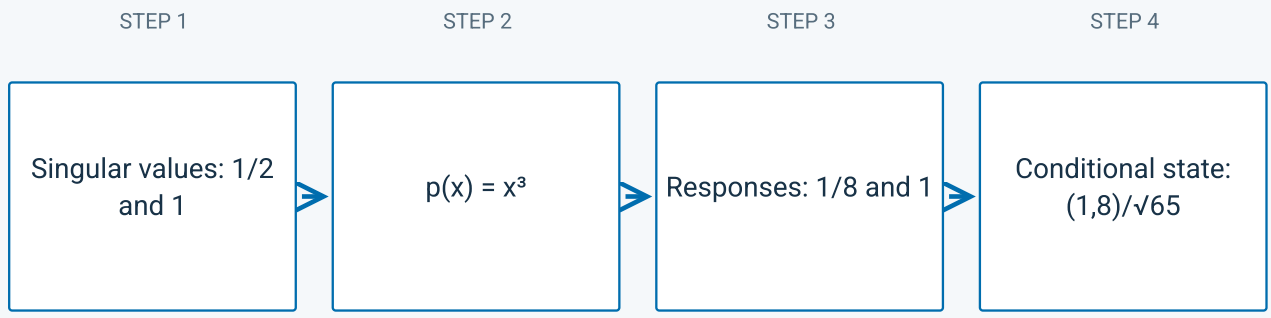
## Apply a polynomial to the important part of a matrix

Quantum singular value transformation, or QSVT, is a framework for using a block encoding to implement polynomial transformations of singular values. Singular values describe how a matrix stretches directions. Transforming them can suppress small components, approximate an inverse or implement other useful spectral filters. The framework unifies several algorithmic constructions without making arbitrary data access free.

Prerequisites: complex matrices, block encoding and the rotation picture of qubitization. Begin with the diagonal example, where singular vectors are simply basis states, before considering a general singular-value decomposition. Gilyén and colleagues.

WORKED DIAGRAM · A28

### Cubic singular-value filter




Starting from |+⟩, the cubic filter gives success probability 65/128 and conditional probabilities 1/65 and 64/65.

## State the polynomial setup

Assume access to a suitable block encoding of A divided by alpha, its inverse and required ancilla-phase operations. Choose a polynomial satisfying the boundedness, parity and degree conditions of the chosen QSVT construction. A

degree-d sequence uses a number of block accesses proportional to d, together with ancillary phase rotations. Approximation quality must be specified over the relevant singular-value interval.

Not every desired function is bounded on the full unit interval. An inverse diverges near zero, so an inversion algorithm needs a promise excluding very small singular values or a regularisation rule. The polynomial must approximate a scaled inverse on that restricted interval while respecting implementability elsewhere.

## Worked example: a cubic filter

Take the positive diagonal matrix with entries one half and one. For this Hermitian toy case, applying the odd polynomial p(x) equal to x cubed gives diagonal entries one eighth and one. Start from plus. The unnormalised transformed state has amplitudes one eighth and one, with the common square-root-of-two denominator.

After conditioning on the appropriate encoded output and normalising, the state is

$$|\psi_{\text{filtered}}\rangle = \frac{|0\rangle + 8|1\rangle}{\sqrt{65}}.$$

Its Z probabilities are one sixty-fifth and sixty-four sixty-fifths. The cubic has strongly suppressed the smaller singular direction. The transformation has not measured which direction was present; it has coherently changed their relative weights. Its success mechanism and normalisation remain part of the physical procedure.

This calculation explains the target response, not a complete phase-synthesis recipe. Constructing a QSVT phase sequence is another algorithmic step. A scalable implementation must provide those phases, their precision and a tested circuit. A tiny direct matrix multiplication is a useful reference for validating the resulting sequence.

Show the unnormalised response before the conditional state, so its probability remains visible.

$$p(x) = x^3, \quad A = \operatorname{diag}(1/2, 1), \qquad p(A)|+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1/8 \\ 1 \end{pmatrix}, \qquad P_{\text{selected}} = \|p(A)|+\rangle\|^2 = \frac{65}{128}.$$

This assumes a selected block implementing the stated response with unit scale in the positive diagonal example. Conditioning divides by the branch norm, giving the existing amplitudes $1/\sqrt{65}$ and $8/\sqrt{65}$. General non-Hermitian QSVT still needs left/right singular-vector bookkeeping.

## Why parity and singular vectors matter

For a general non-Hermitian matrix, singular-value transformations act between or within left and right singular-vector spaces according to the construction. An odd polynomial can produce a transformation with the same left-right structure as A. An even polynomial generally acts in one associated space. It is therefore unsafe to describe every QSVT polynomial as simply taking an ordinary matrix power.

The diagonal positive example hides that distinction because the left and right singular vectors coincide. Use it to learn the scalar filter, then retain the full singular-vector bookkeeping for general matrices. Eigenvalues and singular values are also different objects when a matrix is not Hermitian or normal.

## From filtering to inversion

To approximate an inverse on singular values between one over kappa and one, choose a bounded scaled reciprocal. The scale prevents the polynomial from exceeding unit magnitude. Resolving the transition near the excluded small-value region costs degree, and the desired accuracy also affects degree. The condition number therefore remains a resource parameter rather than disappearing inside a matrix box.

Regularised inversion replaces a dangerous reciprocal near zero with a controlled response. That changes the mathematical task and can improve stability. State the regularisation and its bias alongside approximation error, sampling uncertainty and data-model error. A small circuit error does not imply a small application error for an ill-conditioned problem.

## Step-by-step algorithm design

Define the requested output property. Establish the matrix-access model and block scale. Identify spectral promises such as a gap or condition bound. Choose a function and a polynomial approximation with the correct parity and boundedness. Synthesize phases, implement the sequence and validate its response on a small spectrum. Finally include state preparation, success amplification and observable estimation in the total cost.

For the cubic example, test both basis inputs and plus. The basis tests check the response values one eighth and one, while plus checks relative weights and coherence. Include an input with an imaginary component to expose a sign or conjugation error. A test based only on normalised output probabilities can miss an incorrect success amplitude.

## Separate filtering strength from useful output

For the cubic toy filter, the selected branch probability on plus is the average of one sixty-fourth and one, or sixty-five one-hundred-twenty-eighths. Its conditional one probability is sixty-four sixty-fifths. The latter is close to one, but the former is only about 0.508. A sharper-looking conditional distribution does not mean every attempt yields it.

If the input were concentrated entirely in the smaller singular direction, the success probability would be one sixty-fourth while the conditional direction would remain that same basis state. Filtering cannot manufacture an absent large-singular-value component. Initial overlap and success amplitude are therefore resources in state-preparation applications.

The calculation is simple enough to validate independently of phase synthesis. After implementing an actual sequence, compare both its block amplitudes and leakage with these values. Checking only the renormalised successful output could hide a wrong overall response factor that substantially changes the algorithm’s cost.

## Complexity and limits

QSVT can yield strong query-complexity results for particular matrix tasks, but the advantage depends on efficient block access, manageable scale and conditioning, and a useful compact output. Loading an arbitrary dense matrix or printing an entire solution can dominate. Compare against classical algorithms that exploit the same structure and access assumptions.

This framework is powerful precisely because it makes reusable assumptions explicit. It does not imply that every matrix function or every machine-learning application becomes easy. Degree lower bounds, precision requirements and output costs still constrain the full computation.

## Exercises and worked answers

**Exercise 1.** Is p(x) equal to two x bounded by one over the entire interval from negative one to one?

Check the worked answer

No. Its magnitude reaches two at the endpoints. It cannot be inserted directly as the bounded response of the standard construction without rescaling or changing the promised domain and implementation.

**Exercise 2.** Why does an inverse filter need information about small singular values?

Check the worked answer

The reciprocal grows without bound near zero, making both implementation and the underlying solution sensitive. A spectral lower bound or regularisation determines the interval, scaling and approximation degree needed for a meaningful algorithm.

## Check your understanding

PAUSE & THINK

### What is the conditional probability of $|1\rangle$ after the worked cubic filter?

- [ ] 1/2
- [ ] 8/9
- [ ] 64/65

**Answer:** 64/65. The normalised amplitudes are $1/\sqrt{65}$ and $8/\sqrt{65}$, so the second probability is 64/65.

## Sources

- Gilyén and colleagues