

Amplitude Amplification Algorithms*

Mithilesh Kumar,^{1,†} Yusuf Tahir,^{1,‡} and Varun Daiya^{1,§}

¹*Krea University, Sri City, India*

(Dated: May 11, 2026)

This survey traces the evolution of quantum amplitude amplification from its geometric origins in Grover’s algorithm to its algebraic formulation within the Quantum Singular Value Transformation (QSVT) framework. We develop a unified mathematical formalism intended to reduce the historical fragmentation of these algorithms. By establishing a consistent notational foundation early on, framing unstructured search as rotation within an invariant two-dimensional subspace, and following the progression through oblivious, fixed-point, and distributed amplification schemes, we aim to clarify how these methods can be interpreted within the broader language of singular-value polynomial synthesis. We additionally introduce a variational template for amplitude amplification based on parameterized phase rotations, connecting amplification primitives to hybrid quantum-classical algorithms suitable for near-term quantum architectures. To support the pedagogical objectives of this survey, we provide a companion website hosting interactive visualizations and implementation details at <https://quantumamplification.github.io/QuantumAmplitudeAmplification/>.

CONTENTS

I. Introduction	2	E. The Equivalence to Unitary Block Encodings	11
A. The Unstructured Search Paradigm	2	VII. Controlled Quantum Amplification	12
B. The Invariant Subspace Formalism	3	A. From Detection to Finding	12
C. Survey Roadmap	3	B. Controlled Architecture and the Principal Eigenvector	12
II. Problem Formulation and the Quantum Query Model	4	C. The Core Insight	12
III. Historical Background and Survey Scope	4	VIII. Fixed-Point Oblivious Amplitude Amplification	13
IV. Standard Amplitude Amplification and Unstructured Search	4	A. Linear Combination of Unitaries (LCU)	13
A. Initialization and the Uniform Superposition	4	B. The algorithm	14
B. The Reflection Geometry: Oracle and Diffusion	5	IX. Distributed Amplitude Amplification	15
C. The High-Density Edge Case ($M > N/2$)	7	A. Motivation: NISQ Hardware Constraints	15
V. Fixed-Point Amplitude Amplification	7	B. Partitioning the Global Search Space	15
A. Generalized Grover Iterate	7	C. Tensor Decomposition of the Initialization Operator	16
B. Success Probability and Chebyshev Mapping	7	D. Local Probability Obstruction	16
C. Phase Schedule Derivation	9	E. Distributed Fixed-Point Amplification	17
D. Recursive Nesting	9	F. Global Success Guarantee	18
E. From Geometry to Algebraic Synthesis	9	G. Complexity Analysis	18
VI. Oblivious Amplitude Amplification	9	H. Empirical Validation	18
A. Generalized Algebraic Setup	9	I. Hardware Scaling	19
B. Roadblocks: State Collapse and Reflection Obstruction	10	J. Architectural Advantages	19
C. The Purified Setting and Ancilla Padding	10	K. Limitations and Relation to QSVT	19
D. Synthesizing the Oblivious Reflection	11	L. Distributed Exact Quantum Amplitude Amplification (DEQAAA)	19
		X. Variable-Time Amplitude Amplification	21
		A. Variable-Time Framework	22
		B. Variable-Time Amplification Limit	23
		C. Recursive Filtering Mechanics	23
		D. Application to Linear Systems and Bridge to QSVT	23
		XI. Quantum Singular Value Transformation	24
		A. The general framework	24
		B. Subsuming Amplitude Amplification	25

* These authors have contributed equally

[†] mithilesh.kumar@krea.edu.in

[‡] yusuf_tahir.sias25@krea.ac.in

[§] varun_daiya.sias25@krea.ac.in

C. Selected Advanced QSVT Instantiations	26
XII. Applications of the Unified Framework	26
XIII. Limits and Optimality	28
XIV. Variational template for amplification	29
A. QAOA-Inspired Search and the Transverse Field	30
XV. Conclusion and Outlook	30
References	31

I. INTRODUCTION

The discovery of algorithms capable of outperforming classical computational models is a foundational pillar of quantum computer science. Among the core primitives enabling these polynomial and exponential speedups, quantum amplitude amplification stands as one of the most versatile. First introduced by Lov Grover as a discrete, geometric reflection mechanism to solve the unstructured database search problem [10], the algorithm was quickly proven to be asymptotically optimal [2, 3, 28]. The core mechanism was later rigorously generalized to accommodate arbitrary state-preparation unitaries [6]. This generalization elevated amplitude amplification from a specific search protocol to a universal subroutine utilized in myriad complex algorithms, ranging from solving linear systems of equations [16] to simulating Hamiltonian dynamics [4, 5].

Despite its ubiquity, the pedagogical and review literature often documents the progression of amplitude amplification as a collection of isolated algorithmic tricks. A major historical hurdle in standard amplification was its sinusoidal instability, known colloquially as the “soufflé” problem, where excessive iterations cause the target state probability to overshoot and degrade. This limitation spurred the development of fixed-point amplitude amplification, which guarantees a robust high-success final state over a prescribed passband without requiring exact knowledge of the target state’s initial density [11, 22, 27]. Parallel to these developments, alternative mathematical frameworks such as Duality Quantum Computing were established to simulate the linear combination of unitaries [12–14]. The convergence of these theoretical tracks eventually enabled the extension of fixed-point bounds to oblivious scenarios where exact initial reflections are intractable [7, 25], culminating in robust algorithms like the fixed-point oblivious quantum amplitude amplification [26].

Simultaneously, the pursuit of practical quantum advantage in the Noisy Intermediate-Scale Quantum (NISQ) era [23] has driven adaptations of amplitude amplification toward hardware-aware implementations. For instance, techniques for amplitude estimation that avoid the deep circuits required by traditional quantum phase estimation

have become crucial for near-term applications [9, 24]. Furthermore, the physical constraints of contemporary quantum processors have motivated the recent design of scalable, distributed architectures [15], resulting in distributed quantum amplitude amplification schemes that parallelize the global search space to mitigate strict circuit-width and coherence bottlenecks [17, 29].

Today, the historical fragmentation of these various spatial rotation techniques is elegantly unified under the algebraic frameworks of Quantum Signal Processing (QSP) [20, 21] and Quantum Singular Value Transformation (QSVT) [8]. This survey aims to provide a cohesive exposition of amplitude amplification by tracing its trajectory from basic geometric reflections to this modern operator polynomial synthesis. By establishing a consistent notational foundation early on, specifically framing search and amplification as targeted rotations within an invariant two-dimensional subspace, we systematically connect standard, fixed-point, oblivious, duality-based, and distributed schemes. Ultimately, this survey clarifies how the evolution of amplitude amplification can be fundamentally understood as the development of an encompassing algebraic language for singular-value shaping. Implementation of the algorithms is provided in the Qiskit based `ampamp` library [19].

A. The Unstructured Search Paradigm

There is a geometric structure to search. The input is encoded as a n -qubit vector in an N -dimensional Hilbert space \mathcal{H} . The items to be searched form a *good* subspace and the rest a *bad* subspace. This motivates us to partition \mathcal{H} into two mutually exhaustive and strictly orthogonal subspaces:

1. The **target subspace** (or “good” subspace) $\mathcal{H}_{\text{Good}}$, encompassing the span of all valid solution states.
2. The **orthogonal subspace** (or “bad” subspace) \mathcal{H}_{Bad} , encompassing all non-target components.

Accordingly, we define the corresponding orthogonal projection operators Π_{Good} and Π_{Bad} , satisfying $\Pi_{\text{Good}}^2 = \Pi_{\text{Good}}$ and $\Pi_{\text{Good}} + \Pi_{\text{Bad}} = I$, where I is the identity operator on \mathcal{H} .

The problem is initialized by the action of an efficiently preparable unitary algorithm \mathcal{A} , which acts on a fiducial zero state $|0\rangle^{\otimes n}$ to produce an initial coherent trial state $|\text{All}\rangle$:

$$|\text{All}\rangle = \mathcal{A}|0\rangle^{\otimes n}. \quad (1)$$

In the absence of any amplification technique, directly measuring $|\text{All}\rangle$ in the computational basis yields a target solution with an initial success probability p , defined strictly as the expectation value of the target projector:

$$p = \langle \text{All} | \Pi_{\text{Good}} | \text{All} \rangle = \|\Pi_{\text{Good}} |\text{All}\rangle\|^2. \quad (2)$$

When p is overwhelmingly small (e.g., $p \sim \mathcal{O}(1/N)$ in typical unstructured search scenarios), classical randomized sampling necessitates $\mathcal{O}(1/p)$ queries to succeed with high probability. Quantum amplitude amplification seeks to coherently boost this success probability to order unity using only $\mathcal{O}(1/\sqrt{p})$ applications of the unitary \mathcal{A} and its associated target-recognizing oracle, thereby yielding the familiar quadratic improvement in the oracle model.

B. The Invariant Subspace Formalism

To formalize the action of amplitude amplification computationally, it is useful to adopt a standard mathematical formalism that remains globally consistent, independent of whether the initial state is known, oblivious, or dispersed across variable runtimes.

We decompose the initial trial state $|\text{All}\rangle$ explicitly along the orthogonal direct sum $\mathcal{H} = \mathcal{H}_{\text{Good}} \oplus \mathcal{H}_{\text{Bad}}$. We define the strictly normalized **target state** $|\text{Good}\rangle$ and **non-target state** $|\text{Bad}\rangle$ as:

$$|\text{Good}\rangle = \frac{1}{\sqrt{p}} \Pi_{\text{Good}} |\text{All}\rangle, \quad (3)$$

$$|\text{Bad}\rangle = \frac{1}{\sqrt{1-p}} \Pi_{\text{Bad}} |\text{All}\rangle. \quad (4)$$

By construction, these components satisfy absolute orthogonality, $\langle \text{Good} | \text{Bad} \rangle = 0$. The initial success probability is parameterized by an angle θ : $\sin^2(\theta_0) = p$. The initial state can be written as:

$$|\text{All}\rangle = \sin(\theta_0) |\text{Good}\rangle + \cos(\theta_0) |\text{Bad}\rangle. \quad (5)$$

The plane spanned by $\{|\text{Good}\rangle, |\text{Bad}\rangle\}$ serves as the invariant geometric setting for amplitude amplification. One of the recurring themes of this survey is that the evolution of amplitude amplification, from Grover's discrete reflections to the generalized functional viewpoint underlying QSVT, can often be understood as a sequence of increasingly flexible polynomial transformations acting on invariant two-dimensional subspaces analogous to the one defined in Equation (5).

Symbol	Meaning / Standing Assumption
$\mathcal{H}_{\text{Good}}, \mathcal{H}_{\text{Bad}}$	target/non-target subspaces
$\Pi_{\text{Good}}, \Pi_{\text{Bad}}$	orthogonal projectors onto those subspaces
$ \text{All}\rangle = \mathcal{A} 0\rangle^{\otimes n}$	prepared input state before amplification
$p = \ \Pi_{\text{Good}} \text{All}\rangle\ ^2$	initial success probability
$N = 2^n, M$	search-space size and marked count
$\sin^2(\theta_0) = p$	angle parameter used across sections
Oracle model	complexity measured in oracle/query calls

TABLE I. Core notation and assumptions used throughout the survey.

C. Survey Roadmap

The remainder of this survey is structured to construct the conceptual edifice of amplitude amplification systematically:

Section II formalizes the unstructured search problem in the black-box setting, defines the oracle/query-cost model used throughout this survey, and states the asymptotic target and lower-bound landscape.

Section III establishes the historical timeline and scope of this survey, positioning each algorithmic development in the broader transition from geometric search to algebraic signal processing.

Section IV presents the canonical Grover algorithm, re-derived firmly within the two-subspace notation introduced above, highlighting the core geometrical reflection mechanism and the instability inherent in unitarity.

Section V generalizes this geometry into Fixed-Point Amplitude Amplification (FPAA), resolving the critical instability issue when the initial success probability p is strictly unknown, by relying on Chebyshev polynomial phase synthesis.

Section VI introduces Oblivious Amplitude Amplification, a primitive indispensable for embedding subroutines when reflections about a known initial state cannot be synthesized.

Section VII then studies Controlled Quantum Amplification, which converts detection-style search procedures into finding procedures without asymptotic overhead, especially in quantum-walk settings.

Sections VIII and IX address advanced modifications geared toward robustness and hardware scaling, namely Fixed-Point Oblivious amplification and Distributed Amplitude Amplification, the latter formulated explicitly to mitigate structural bottlenecks in Noisy Intermediate-Scale Quantum (NISQ) devices.

Section X details Variable-Time Amplitude Amplification (VTAA), optimizing overhead when discrete sub-components terminate unpredictably.

Finally, **Section XI** places the preceding modules in the language of Quantum Singular Value Transformation. By subsuming the initial state preparation unitary into a broader block-encoding architecture, we show how the geometric intuition over $\{|\text{Good}\rangle, |\text{Bad}\rangle\}$ extends to an algebraic calculus over the singular values of a general matrix. **Sections XII, XIII** then survey applications of this framework, discuss the polynomial lower bounds relevant to its optimality, and outline several directions for fault-tolerant and automated quantum algorithm design. In **Section XIV**, we present a variational template for amplification algorithms.

II. PROBLEM FORMULATION AND THE QUANTUM QUERY MODEL

We consider unstructured black-box search over the Boolean hypercube. Let

$$f : \{0, 1\}^n \rightarrow \{0, 1\}$$

be the indicator of marked inputs, with total search-space size $N = 2^n$. Define $All = \{0, 1\}^n$ and let

$$M = |f^{-1}(1)|$$

denote the number of marked items. A uniformly random classical guess succeeds with probability $p = \frac{M}{N}$.

Oracle access is given by the standard bit-flip unitary \mathcal{O}_f acting on an input register and one ancilla qubit:

$$\mathcal{O}_f |x\rangle |q\rangle = |x\rangle |q \oplus f(x)\rangle. \quad (6)$$

Preparing the ancilla in $|-\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$ yields the equivalent phase-kickback form

$$\mathcal{O}_f |x\rangle |-\rangle = (-1)^{f(x)} |x\rangle |-\rangle. \quad (7)$$

For the remainder of this survey, we assume access to this induced phase oracle and write it as $O : |x\rangle \mapsto (-1)^{f(x)} |x\rangle$.

Complexity is measured in *query complexity*: the dominant cost is the number of oracle calls, while intermediate non-oracle unitary operations are treated as lower-order overhead unless explicitly stated otherwise.

Under this model, classical randomized search requires $\mathcal{O}(N/M)$ oracle queries to find a marked item with constant success probability, while amplitude amplification achieves $\mathcal{O}(\sqrt{N/M})$ [6, 10]. Crucially, Bennett, Bernstein, Brassard, and Vazirani proved that $\Omega(\sqrt{N/M})$ queries are necessary in the unstructured oracle model [3], so Grover-type scaling is asymptotically optimal. We return to the formal lower-bound mechanisms in Section XIII.

III. HISTORICAL BACKGROUND AND SURVEY SCOPE

The theoretical evolution of amplitude amplification is characterized by a steady abstraction from rigid geometric rotations into generalized algebraic transformations. Over the past three decades, algorithmic generalizations have systematically removed the operational constraints of the original unstructured search problem, such as the requirements of a known initial state, a known target density, or a monolithic coherent hardware architecture.

Table II outlines the chronological milestones that define this lineage.

Survey Scope: Rather than presenting these developments as isolated algorithmic tricks, a central objective of this survey is to reconstruct the mathematical bridge connecting them. We restrict our scope to the algorithmic

lineage that carries standard geometric amplification into modern polynomial signal processing. By formalizing a single invariant subspace notation early on, we traverse fixed-point stabilization, oblivious generalizations, temporal variations, and spatial partitioning, and then relate these paradigms to Quantum Singular Value Transformation (QSVT).

IV. STANDARD AMPLITUDE AMPLIFICATION AND UNSTRUCTURED SEARCH

We now formalize the core geometric mechanism that underpins standard amplitude amplification, originally discovered by Lov Grover in 1996 for the unstructured database search problem [10]. The objective is to construct an iterative operator capable of monotonically rotating the initial state vector $|All\rangle$ toward the target subspace $\mathcal{H}_{\text{Good}}$, thereby driving the success probability toward unity. While Grover's original formulation relied strictly on the uniform superposition generated by the Walsh-Hadamard transform, it was subsequently generalized by Brassard, Høyer, Mosca, and Tapp [6] to accommodate any arbitrary state-preparation unitary \mathcal{A} , formally elevating the technique from a specific search algorithm to a universal amplitude amplification subroutine. In their operator formalism, the generalized iterate is written

$$Q = -\mathcal{A}S_0\mathcal{A}^{-1}S_\chi, \quad (8)$$

with $S_0 = I - 2|0\rangle\langle 0|^{\otimes n}$ and $S_\chi = I - 2\Pi_{\text{Good}}$ [6].

To maintain algebraic continuity with the unified formalism established in Section I, we derive the algorithm entirely within the invariant two-dimensional subspace spanned by $\{|Good\rangle, |Bad\rangle\}$.

A. Initialization and the Uniform Superposition

Because the database is completely unstructured, the algorithm initializes from the zero state $|0\rangle^{\otimes n}$ and applies the tensor-product Hadamard operation $H^{\otimes n}$ to explicitly define the initial trial state $|All\rangle$:

$$|All\rangle = H^{\otimes n} |0\rangle^{\otimes n} = \frac{1}{\sqrt{N}} \sum_{x \in All} |x\rangle. \quad (9)$$

We define the normalized geometric components representing the non-solutions and solutions across the entire database. If there are exactly M target solutions $|\{x \in All : f(x) = 1\}| = M$, we define the orthogonal bases:

$$|Good\rangle = \frac{1}{\sqrt{M}} \sum_{f(x)=1} |x\rangle, \quad (10)$$

$$|Bad\rangle = \frac{1}{\sqrt{N-M}} \sum_{f(x)=0} |x\rangle. \quad (11)$$

Year	Algorithmic Milestone	Primary Contribution
1996	Grover's Algorithm [10]	Unstructured search via geometric reflection
2000	Amplitude Estimation [6]	Extracting target probabilities via phase estimation
2002	General Amplitude Amp. [6]	Generalizing search to arbitrary unitary algorithms
2010	Variable-Time AA [1]	ℓ_2 -averaged scaling for early-halting branches
2014	Fixed-Point AA [27]	Monotonic convergence via Chebyshev polynomials
2017	Controlled Quantum Amplification [7]	Detection-to-finding conversion for walk-based search
2017	QAOA for Grover [18]	Near-optimal search using a single-qubit transverse-field mixer
2017	Qubitization & QSP [20, 21]	Optimal polynomial transforms on $SU(2)$ subspaces
2018	QSVT [8]	Algebraic unification via singular value transforms
2022	Fixed-Point Oblivious [26]	Resolving subnormalization via LCU damping
2025	Distributed AA [17]	Hardware-aware parallelization for NISQ devices
2026	Distributed Exact AA [29]	Exact distributed amplification for arbitrary amplitude distributions

TABLE II. Chronological development of the amplitude amplification framework, tracing the transition from basic geometric search primitives to generalized operator polynomials and modern distributed architectures.

Algorithm	Problem Solved	Representative Complexity
Grover / Standard AA	Unstructured search with known calibration	$\mathcal{O}\left(\sqrt{N/M}\right) = \mathcal{O}(1/\sqrt{p})$
FPAA	Unknown lower-bounded p , fixed-point no-overshoot amplification	$\mathcal{O}\left(\frac{1}{\sqrt{p}} \log \frac{2}{\delta}\right)$
OAA	Subroutine amplification without direct input-state reflection	$\mathcal{O}(1/\sqrt{p})$
CQAA	Detection-to-finding conversion for abstract search operators	$\Theta(\text{QHT}(A, \phi_{\text{initial}}\rangle))$
Fixed-Point OAA	Oblivious setting with robustness to unknown/drifted p	$\mathcal{O}(1/\sqrt{p})$ (up to constants)
VTAA	Variable stopping times across computational branches	$\mathcal{O}\left(\frac{1}{\sqrt{p}} \sqrt{\sum_j p_j T_j^2} \text{polylog}(T_m)\right)$
DEQAAA	Exact distributed amplification for arbitrary amplitude distributions	$\mathcal{O}(1/\sqrt{p})$

TABLE III. Survey-level comparison of core amplitude amplification variants.

The uniform initialization state partitions into these two parameters:

$$|\text{All}\rangle = \cos(\theta/2) |\text{Bad}\rangle + \sin(\theta/2) |\text{Good}\rangle, \quad (12)$$

where $\sin(\theta/2) = \sqrt{M/N}$. Note that establishing the angle parameter as $\theta/2$ (as opposed to θ_0 in Section I B) intentionally yields a rotational step size of θ , a standard convention derived directly from Grover's specific construction. The two parameterizations are equivalent under the relation $\theta = 2\theta_0$, with $p = \sin^2(\theta_0) = M/N$.

B. The Reflection Geometry: Oracle and Diffusion

With explicit respect to the initialization angle definitions, the action of the phase oracle O acts as a discrete reflection specifically about the $|\text{Bad}\rangle$ vector axis:

$$O|\text{All}\rangle = \cos(\theta/2) |\text{Bad}\rangle - \sin(\theta/2) |\text{Good}\rangle. \quad (13)$$

As an operator in projection notation, the oracle is written $O = I - 2\Pi_{\text{Good}}$, where Π_{Good} is the projector onto the target subspace $\mathcal{H}_{\text{Good}}$ established in Section I. This form is used throughout all subsequent sections.

The oracle O is combined with another reflection operator R , constructed specifically to reflect about the uniform $|\text{All}\rangle$ initialization state:

$$R = 2|\text{All}\rangle\langle\text{All}| - I. \quad (14)$$

This reflection is implemented utilizing $H^{\otimes n}(2|0\rangle^{\otimes n}\langle 0|^{\otimes n} - I)H^{\otimes n}$. The intermediate operator $2|0\rangle^{\otimes n}\langle 0|^{\otimes n} - I$ acts as the complement of a multi-controlled Z operation, mechanically negating the sign when at least one computational bit evaluates to 1.

Together these operators define the complete Grover operator G :

$$G = RO \quad (15)$$

With elementary rotation arithmetic, evaluating this composite operator precisely in the $\{|\text{Good}\rangle, |\text{Bad}\rangle\}$ basis yields the unitary matrix:

$$G = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}. \quad (16)$$

That is, a vector deterministically rotates towards $|\text{Good}\rangle$ by a constant angle θ upon each sequential application of G . Evaluating the first application demonstrates this rotation explicitly:

$$G|\text{All}\rangle = \cos(3\theta/2) |\text{Bad}\rangle + \sin(3\theta/2) |\text{Good}\rangle \quad (17)$$

After k repeated applications of the G operator, the state




FIG. 1. Grover's Search




FIG. 2. **Geometric progression of the Grover iterate.** The state vector precesses from an initial angle of $\theta/2$ to $3\theta/2$ in a single iteration. The rigid nature of this rotation leads to the “soufflé problem,” where over-iteration past the $|\text{Good}\rangle$ axis causes the success probability to decrease sinusoidally.

systematically evolves to:

$$\begin{aligned}
 |\psi_k\rangle &= G^k |\text{All}\rangle \\
 &= \cos\left(\frac{(2k+1)\theta}{2}\right) |\text{Bad}\rangle \\
 &\quad + \sin\left(\frac{(2k+1)\theta}{2}\right) |\text{Good}\rangle
 \end{aligned} \tag{18}$$

The computational state $|\psi_k\rangle$ points entirely along the $|\text{Good}\rangle$ target axis when iteration numbers k and angular parameters θ satisfy the specific continuous alignment

condition:

$$\frac{(2k+1)\theta}{2} = \frac{\pi}{2} \tag{19}$$

$$\Rightarrow k = \frac{\pi}{2\theta} - \frac{1}{2} \tag{20}$$

Because k must be an integer, one chooses the nearest integer to this continuous optimum. In the standard amplitude-amplification notation with $a = M/N$, a common choice is

$$k = \left\lfloor \frac{\pi}{4 \arcsin(\sqrt{M/N})} - \frac{1}{2} \right\rfloor \tag{21}$$

which yields constant success probability [6]. In the sparse-marked regime $M \ll N$, one may use $\theta \approx 2\sqrt{M/N}$, giving the familiar asymptotic estimate

$$k \approx \frac{\pi}{4} \sqrt{\frac{N}{M}},$$

and hence the $\mathcal{O}(\sqrt{N/M})$ query scaling.

Algorithm 1: Quantum Search Algorithm (Grover's)

Input : (1) An oracle \mathcal{O} such that
 $\mathcal{O}|x\rangle|q\rangle = |x\rangle|q \oplus f(x)\rangle$
(2) n qubits initialized to $|0\rangle$ and one ancilla qubit to $|1\rangle$
Output : The index x such that $f(x) = 1$
// Initialization and Superposition
 $|0\rangle^{\otimes n}|1\rangle \xrightarrow{H^{\otimes n} \otimes H} \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle \left[\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right]$
// Grover Iterations
for $i = 1$ **to** $R \approx \frac{\pi}{4} \sqrt{N/M}$ **do**
 $|\psi\rangle \xrightarrow{\mathcal{O}} (-1)^{f(x)} |\psi\rangle$ // Apply Oracle Phase
 Kickback
 $|\psi\rangle \xrightarrow{H^{\otimes n}} \dots \xrightarrow{2|0\rangle\langle 0| - I} \dots \xrightarrow{H^{\otimes n}} |\psi'\rangle$
 // Diffusion Operator
// Measurement
Measure the first n qubits $\xrightarrow{\text{Collapse}} x$
return x

Intuition Box: Standard Amplitude Amplification

The set of all inputs to the search problem is partitioned into solutions and non-solutions. A starting state vector is chosen such that it has non-trivial overlap with solutions. As this overlap can be small, the amplification is about rotating this initial vector towards the solution vector.

C. The High-Density Edge Case ($M > N/2$)

The exact geometric derivation of the optimum iteration count does not require a sparsity assumption; the condition in Equation (5) and the alignment equation above remain valid for all $0 < p < 1$. Sparsity ($M \ll N$) is used only to obtain the asymptotic estimate $k \approx (\pi/4)\sqrt{N/M}$ via the small-angle approximation. However, in scenarios where the exact solution density is exceptionally high ($M > N/2$), the initial separation angle $\theta/2$ evaluates broadly. Under these geometric conditions, a single computational application of the Grover iterate G can forcibly overshoot the target $|\text{Good}\rangle$ subspace, plunging the immediate success probability precipitously.

While classical random sampling already handles such high-density boundary problems efficiently, the coherent quantum setting requires some mitigation of this over-rotation. A standard device is to append a single auxiliary qubit, thereby doubling the effective database size $N \rightarrow 2N$ without introducing additional marked items. This dilution halves the relative target density to $M/2N \leq 1/2$, reducing the risk of immediate overshoot.

This vulnerability to overshooting is a consequence of the rigid unitary rotation structure. Avoiding such over-rotations when the target density is variable or unknown is one of the central motivations for Fixed-Point Amplitude Amplification.

V. FIXED-POINT AMPLITUDE AMPLIFICATION

Grover's iterate G is a *continuous rotation* in the span $\{|\text{Good}\rangle, |\text{Bad}\rangle\}$ plane. This creates the “soufflé problem”: the success probability rises up to unity and then inevitably falls again. Thus, if the number of marked items M (equivalently the initial success probability $p = M/N$) is unknown, there is no principled stopping time.

Furthermore, as highlighted by the high-density edge case where $p > 1/2$ at the conclusion of Section IV, this continuous rotation can overshoot the target $\mathcal{H}_{\text{Good}}$ subspace even after a single query. Fixed-point amplitude amplification (FPAA) addresses both this sinusoidal instability and the high-density overshooting problem by designing a *sequence* of phase-shifted Grover iterates whose final success probability remains above $1 - \delta^2$ over a broad range of marked fractions, with $\delta \in (0, 1)$ a user-chosen error tolerance. More precisely, the fixed-point guarantee is a statement about the completed synthesized sequence; intermediate prefixes of the sequence need not themselves be monotone in success probability.

A. Generalized Grover Iterate

Fix the target subspace projector $\Pi_{\text{Good}} = \sum_{x: f(x)=1} |x\rangle\langle x|$ onto $\mathcal{H}_{\text{Good}}$, and the uniform initial state $|\text{All}\rangle$. Let $p = \|\Pi_{\text{Good}} |\text{All}\rangle\|^2$ denote the initial success probability.

Introduce generalized phase reflections about the target and start state:

$$S_{\text{Good}}(\beta) = I - (1 - e^{i\beta}) |\text{Good}\rangle\langle \text{Good}| \quad (22)$$

$$S_{\text{All}}(\alpha) = I - (1 - e^{-i\alpha}) |\text{All}\rangle\langle \text{All}| \quad (23)$$

The generalized Grover iterate is then defined as:

$$G(\alpha, \beta) = -S_{\text{All}}(\alpha) S_{\text{Good}}(\beta) \quad (24)$$

Standard Grover search is recovered as the special case where all generalized phases equal π .

B. Success Probability and Chebyshev Mapping

The key design goal in FPAA is to explicitly replace the rigid sinusoidal dependence of standard Grover amplification by a synthesized polynomial map exhibiting a flat high-success “passband.” While standard Grover amplification creates a sharp, oscillating sine wave of success




FIG. 3. Rotation of vectors in FPAA is better represented on the Bloch sphere.

probability, the Chebyshev sequence effectively “flattens” the peak of that wave into a wide, stable plateau.

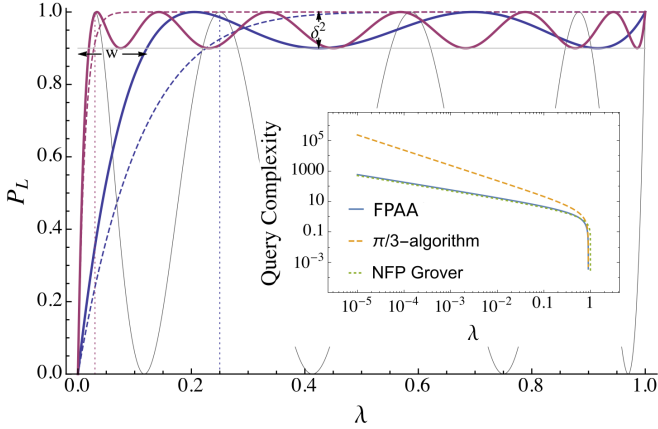


FIG. 4. Chebyshev Polynomials passband graph

This ensures that once the state vector enters the target region, the variable phase shifts dampen the rotation, actively preventing it from leaving. Following [27], for an odd integer $L = 2l + 1$ one uses a palindromic sequence of l generalized iterates

$$S_L = G(\alpha_l, \beta_l) \cdots G(\alpha_2, \beta_2) G(\alpha_1, \beta_1), \quad (25)$$

with query complexity $L - 1 = 2l$. This sequence acts on the two-dimensional invariant subspace so that the final success probability can be written exactly as a polynomial transformation $P_L(p)$ of the initial success probability p .

In the construction of Yoder, Low, and Chuang (2014) [27], one obtains the explicit Chebyshev-form map

$$P_L(p) = 1 - \delta^2 T_L^2\left(T_{1/L}(1/\delta) \sqrt{1-p}\right) \quad (26)$$

where T_L is the L -th Chebyshev polynomial of the first kind and $T_{1/L}(\cdot)$ denotes the functional inverse of $T_L(\cdot)$ defined strictly on $[1, \infty)$ (equivalently calculated as $T_{1/L}(x) = \cosh(\frac{1}{L} \operatorname{arcosh} x)$ for $x \geq 1$).

Intuition Box: Fixed-Point Amplitude Amplification

Instead of using fixed rotations in 2D-space of Grover Search, use a set of angles obtained via Chebyshev polynomials that start slow, accelerate and decelerate. This behavior forces amplitudes to oscillate within a small band near 1. Geometrically, the $e^{-i\alpha}, e^{i\beta}$ factors force the intermediate vectors to leave the 2D-plane of rotation.

Algorithm 2: Fixed-Point Amplitude Amplification (FPAA)

Input: (1) State preparation unitary \mathcal{A} producing $|\text{All}\rangle$ with unknown success probability $p \geq w$

(2) Generalized oracles $S_{\text{Good}}(\beta)$ and $S_{\text{All}}(\alpha)$

Output: State with target subspace probability at least $1 - \delta^2$

// Initialization

$|\psi\rangle \leftarrow \mathcal{A} |0\rangle^{\otimes n};$

Determine sequence length

$$L = 2l + 1 = \mathcal{O}\left(\frac{\log(2/\delta)}{\sqrt{w}}\right);$$

Derive palindromic phase schedule $\{(\alpha_j, \beta_j)\}_{j=1}^l$ from Chebyshev polynomials;

// Generalized Iterations

for $j = 1$ to l do

$|\psi\rangle \leftarrow -S_{\text{All}}(\alpha_j) S_{\text{Good}}(\beta_j) |\psi\rangle;$ // Apply generalized iterate $G(\alpha_j, \beta_j)$

end

return $|\psi\rangle;$

We define the passband edge as the threshold variable w :

$$w = 1 - T_{1/L}(1/\delta)^{-2} \quad (27)$$

Then for all initial probabilities $p \geq w$, the sequence satisfies

$$P_L(p) \geq 1 - \delta^2 \quad (28)$$

The exact finite- L fixed-point requirement is

$$1 - T_{1/L}(1/\delta)^{-2} \leq p \quad (29)$$




FIG. 5. **Fixed-Point Amplitude Amplification (FPAA) Circuit.** The algorithm applies a palindromic sequence of generalized Grover iterates $G(\alpha_j, \beta_j) = -S_{\text{All}}(\alpha_j)S_{\text{Good}}(\beta_j)$ with phase angles chosen according to a Chebyshev schedule. The first and last iterates are expanded to illustrate the underlying phase-shifted reflections about the target subspace and the initial state $|\text{All}\rangle$. This symmetric construction suppresses overshooting and produces a monotonic amplification of success probability. Gates are applied left-to-right in circuit order.

For asymptotically small values of p , this yields the sufficient scaling estimate

$$L = \mathcal{O}\left(\frac{\log(2/\delta)}{\sqrt{p}}\right) \quad (30)$$

and the query complexity is $L - 1$. Thus FPAA preserves Grover's optimal quadratic scaling while providing fixed-point robustness against unknown initial densities.

C. Phase Schedule Derivation

The Chebyshev design in (26) can be implemented *exactly* by choosing an analytic phase schedule $\{(\alpha_j, \beta_j)\}_{j=1}^L$ for odd $L = 2l + 1$. Let

$$\gamma^{-1} = T_{1/L}(1/\delta), \quad \text{so that} \quad \gamma \in (0, 1]. \quad (31)$$

The discrete phases are then given by

$$\begin{aligned} \alpha_j &= 2 \cot^{-1}\left(\tan(2\pi j/L) \sqrt{1 - \gamma^2}\right), \\ \beta_{l-j+1} &= -\alpha_j, \end{aligned} \quad (32)$$

for $j = 1, \dots, l$, so the schedule is palindromic.

D. Recursive Nesting

An important structural feature of the FPAA family is *recursive nesting*. In the exact Yoder–Low–Chuang construction, if an outer sequence of length L_2 targets final error bound δ , then the inner sequence of length L_1 must be assigned the retuned bound

$$\delta_1 = T_{1/L_2}(1/\delta)^{-1}. \quad (33)$$

With that choice, composing the two fixed-point blocks yields

$$T_{L_2}(T_{L_1}(x)) = T_{L_1 L_2}(x), \quad (34)$$

so one obtains a new fixed-point sequence of effective computational length $L_1 L_2$ without restarting from the

fiducial $|\text{All}\rangle$ state. Operationally, the state output of the first FPAA block becomes the input to the next block, and the polynomial degree multiplies under composition. This recursive architecture provides a natural way to broaden the target passband (or iteratively reduce δ) without incurring restart overhead.

E. From Geometry to Algebraic Synthesis

The structural formulation of Fixed-Point Amplitude Amplification marks an important shift in perspective: it moves from *pure geometry* (rigid spatial rotations defined by fixed vectors) toward *algebraic synthesis* (designing operator sequences to approximate bounded classical polynomials). The idea of interleaving generalized target reflections with state reflections to shape a transformation polynomial provides a useful conceptual precursor to Quantum Singular Value Transformation (QSVT), which we discuss in Section XI.

VI. OBLIVIOUS AMPLITUDE AMPLIFICATION

Section V removed the stopping-time instability of Grover rotations when p is unknown. However, both Grover and FPAA still assume the ability to reflect about a known start state $|\text{All}\rangle$. Oblivious amplification removes that assumption and studies the regime where the input is unknown [25].

In this setting, the input is an arbitrary pure state $|\psi\rangle$ rather than a known basis-derived initialization. The key question is whether amplitude amplification can still be implemented without explicit knowledge of $|\psi\rangle$. This subsection formalizes the obstruction and then identifies the additional structure needed for oblivious amplification to become possible.

A. Generalized Algebraic Setup

Let $|\psi\rangle$ be an arbitrary, unknown n -qubit pure state. We are given black-box access to a Boolean predicate

$f : \{0, 1\}^n \rightarrow \{0, 1\}$ representing the target index, and an initialization unitary \mathcal{A} such that $\mathcal{A}|\psi\rangle = \sum_x \alpha_x |x\rangle$.

Define the success probability analytically as $p(\psi) = \sum_{f(x)=1} |\alpha_x|^2$. The normalized good and bad components are parameterized as:

$$\begin{aligned} |\text{Good}(\psi)\rangle &= \frac{1}{\sqrt{p(\psi)}} \sum_{f(x)=1} \alpha_x |x\rangle \\ |\text{Bad}(\psi)\rangle &= \frac{1}{\sqrt{1-p(\psi)}} \sum_{f(x)=0} \alpha_x |x\rangle \end{aligned} \quad (35)$$

The prepared state (with non-zero component with $f(x) = 1$) can be written as

$$\mathcal{A}|\psi\rangle = \sqrt{1-p(\psi)} |\text{Bad}(\psi)\rangle + \sqrt{p(\psi)} |\text{Good}(\psi)\rangle.$$

Unlike the standard setting, both $p(\psi)$ and the associated two-dimensional amplification plane depend on the unknown input. The rotation subspace is therefore not known a priori.

B. Roadblocks: State Collapse and Reflection Obstruction

In standard amplitude amplification, failed runs can be discarded and the process restarted from $|0\rangle^{\otimes n}$. In the oblivious setting, we typically have only coherent access to $|\psi\rangle$. Measuring a success flag collapses the system, and the no-cloning theorem prevents us from making backup copies of the unknown input state. Consequently, the standard repeat-until-success strategy is unavailable.

A more fundamental obstruction is geometric. In standard implementations, the diffusion operator is $R = \mathcal{A}(I - 2|0\rangle^{\otimes n}\langle 0|^{\otimes n})\mathcal{A}^{-1}$. This works because $|0\rangle^{\otimes n}$ is a known, fixed state about which we can easily reflect.

In the oblivious model, the analogous reflection would be $R_{|\psi\rangle} = I - 2|\psi\rangle\langle\psi|$. Implementing this for an unknown $|\psi\rangle$ is generally impossible without state tomography, and a universal construction would violate the linearity of quantum mechanics across non-orthogonal inputs. Thus, the standard Grover-style mechanism cannot be directly applied. Even if the success probability p is constant for a family of states, oblivious amplification still requires reflections that depend on the input state. This becomes possible only if we impose additional structure, such as the purified ancilla or block-encoding layouts discussed later.

Algorithm 3: Oblivious Amplitude Amplification (OAA)

Input: (1) Unknown arbitrary pure state $|\phi\rangle$
 (2) Black-box operation \mathcal{A} acting on $|0\rangle^{\otimes l}|\phi\rangle$
 (3) Subspace projector Π_{Good} (4) A known lower bound or calibration for the relevant success amplitude

Output: Amplified target state without knowledge of $p(\psi)$ or $|\phi\rangle$

// Subnormalized Block Encoding Form

$|\psi_{\text{init}}\rangle \leftarrow \mathcal{A}|0\rangle^{\otimes l}|\phi\rangle$;

Define clean-ancilla reflection

$$R_{\text{clean}} = 2(|0\rangle\langle 0|^{\otimes l} \otimes I) - I;$$

Define target reflection $S_{\text{Good}} = I - 2\Pi_{\text{Good}}$;

// Oblivious Iterations

for $i = 1$ **to** $\mathcal{O}(1/\sqrt{p})$ **do**

$|\psi_{\text{init}}\rangle \leftarrow \mathcal{A}R_{\text{clean}}\mathcal{A}^\dagger S_{\text{Good}}|\psi_{\text{init}}\rangle$;

end

return $|\psi_{\text{init}}\rangle$;

C. The Purified Setting and Ancilla Padding

General oblivious amplification fails because both restart and unknown-state reflections are unavailable. The standard workaround is a stricter model: the *purified setting with independent initial weight*.

Assume a tensor split $\mathcal{H} = \mathcal{H}_{\text{ancilla}} \otimes \mathcal{H}_{\text{data}}$ with dimensions 2^l and 2^m ($n = l + m$), and promise inputs of the form

$$|\psi\rangle = |0\rangle^{\otimes l}|\phi\rangle,$$

where $|\phi\rangle$ is arbitrary but the ancilla prefix is known and clean. The requirement $l > 0$ provides the fixed reference subspace needed for oblivious constructions and is precisely the structural ingredient used in block-encoding pipelines.

The initial probability of success can be obtained as:

$$p = \left\| \Pi_{\text{Good}} \mathcal{A} |0\rangle^{\otimes l} |\phi\rangle \right\|_2^2 \quad (36)$$

It so happens that amplification is feasible only if p is independent of $|\phi\rangle$ on the relevant data subspace, i.e., $p(\phi) \equiv p$.

To analyze this condition, define the Hermitian operator

$$M = \mathcal{A}^\dagger \Pi_{\text{Good}} \mathcal{A}$$

The success probability is given by $p = \langle 0^{\otimes l} | \langle \phi | M | 0^{\otimes l} \rangle | \phi \rangle$. By restricting the auxiliary register to the zero-state $|0\rangle^{\otimes l}$, one can isolate a specific block of M . Partitioning M according to the computational basis yields:

$$M = \begin{bmatrix} M_{\text{TL}} & M_{\text{TR}} \\ M_{\text{BL}} & M_{\text{BR}} \end{bmatrix} \quad (37)$$

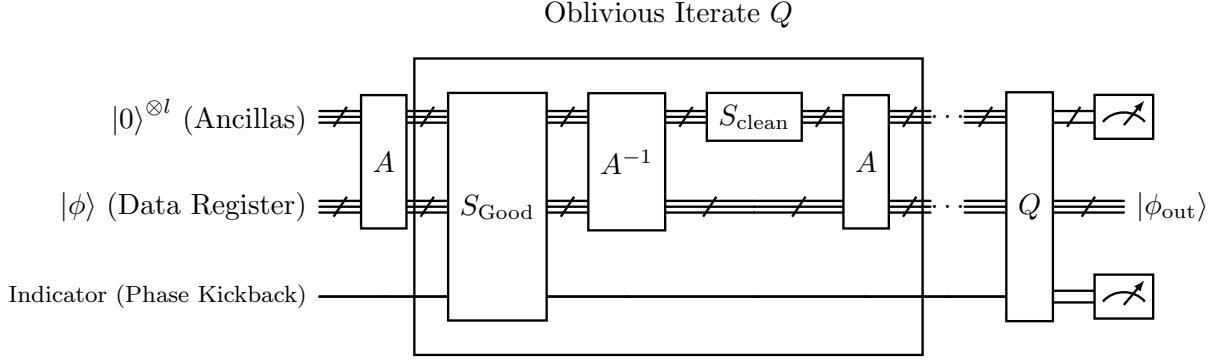


FIG. 6. **Oblivious Amplitude Amplification (OAA) Circuit Architecture.** OAA enables the amplification of an unknown state $|\phi\rangle$ within a purified setting using l clean ancillas. Unlike standard search, the data register is preserved (not measured) so the amplified output can be reused coherently inside larger pipelines such as block-encoding and QSVT constructions. Each iterate Q implements the sequence $AS_{\text{clean}}A^{-1}S_{\text{Good}}$. A successful measurement of the clean ancilla state $|0\rangle^{\otimes l}$ and the success indicator heralds that the data register has collapsed into the amplified target state $|\phi_{\text{out}}\rangle$. Gates are applied left-to-right in circuit order.

where M_{TL} corresponds to the subspace defined by the $0^{\otimes l}$ auxiliary prefix. Since the initial state $|0\rangle^{\otimes l}|\phi\rangle$ lies entirely within this sector, the probability reduces to $p = \langle\phi|M_{\text{TL}}|\phi\rangle$. Because this expectation value must be constant for any state $|\phi\rangle$, every eigenvalue of M_{TL} must equal p . This implies the structural identity:

$$M_{\text{TL}} = pI \quad (38)$$

Thus, on the subspace where the ancilla qubits are zero, $\mathcal{A}^\dagger \Pi_{\text{Good}} \mathcal{A}$ acts as a scalar multiple of the identity. In the case $l = 0$, this simplifies to $M = pI$, representing the limit where oblivious amplification cannot be performed.

Intuition Box: Oblivious Amplitude Amplification

Core mechanism

Given an unknown state $|\phi\rangle$, pad it with ancilla $|0\rangle^{\otimes l}$. Apply reflections about this ancilla subspace for amplification.

Geometric interpretation

State reflections are replaced by subspace reflections. The geometry is still similar to the standard amplitude amplification.

Algorithmic consequence

Yields optimal $\mathcal{O}(1/\sqrt{p})$ scaling directly on unknown data structures, meaning an algorithm can securely boost internal probabilistic success rates without intermediate collapse or explicit copying.

Structural insight

Ancilla padded state amplification is equivalent to unitary block-encoding that sets stage for QSVT.

D. Synthesizing the Oblivious Reflection

The identity block $M_{\text{TL}} = pI$ replaces the unavailable unknown-state reflection. Let

$$C = |0\rangle^{\otimes l}\langle 0|^{\otimes l} \otimes I_{\text{data}}$$

From $\mathcal{A}^\dagger \Pi_{\text{Good}} \mathcal{A} = M$, projection gives

$$CM|0\rangle^{\otimes l}|\phi\rangle = p|0\rangle^{\otimes l}|\phi\rangle$$

uniformly over $|\phi\rangle$. Therefore the reflection about the clean-ancilla sector is

$$R_{\text{clean}} = 2C - I$$

which plays the role needed inside the effective two-dimensional amplification subspace. The resulting oblivious iterate

$$Q = \mathcal{A}(2C - I)\mathcal{A}^\dagger(I - 2\Pi_{\text{Good}})$$

recovers Grover-type rotation with $\mathcal{O}(1/\sqrt{p})$ scaling, without tomography, cloning, or measurement-driven restarts.

E. The Equivalence to Unitary Block Encodings

This setting is equivalent to the block-encoding framework used in modern quantum algorithms, such as Hamiltonian simulation and QSVT. A unitary \mathcal{A} is a block-encoding of an operator B if its matrix representation takes the form:

$$\mathcal{A} = \begin{bmatrix} B & * \\ * & * \end{bmatrix} \quad (39)$$

Restricting the input to the zero-ancilla state isolates the transformation B . If \mathcal{A} block-encodes $B = \sqrt{p}U$, then for any state $|\phi\rangle$:

$$\mathcal{A}|0\rangle^{\otimes l}|\phi\rangle = \sqrt{p}|0\rangle^{\otimes l}U|\phi\rangle + |\perp(\phi)\rangle$$

where $|\perp(\phi)\rangle$ is an orthogonal state with a non-zero ancilla. Here, the success probability p is independent of $|\phi\rangle$, satisfying the purified invariance condition. Conversely, if we define the success projection as the clean-ancilla sector ($\Pi_{\text{Good}} = |0^{\otimes l}\rangle\langle 0^{\otimes l}| \otimes I$), then:

$$M = \mathcal{A}^\dagger \Pi_{\text{Good}} \mathcal{A}$$

The top-left block of M is given by $M_{\text{TL}} = \mathcal{A}_{\text{TL}}^\dagger \mathcal{A}_{\text{TL}}$. Since the purified condition requires $M_{\text{TL}} = pI$, it follows that $\mathcal{A}_{\text{TL}}^\dagger \mathcal{A}_{\text{TL}} = pI$. This implies $\mathcal{A}_{\text{TL}} = \sqrt{p}U$ for some unitary U , proving that the two formalisms are equivalent.

Final Verdict. The fundamental relational dynamic proving the equivalence:

$$\begin{array}{c} \text{Purified Amplification Setting} \\ \Longleftrightarrow \\ \text{Unitary Block Encoding} \end{array}$$

establishes OAA as the coherent amplification primitive for subnormalized block encodings, and prepares the transition to Section XI.

VII. CONTROLLED QUANTUM AMPLIFICATION

A. From Detection to Finding

Standard Amplitude Amplification (SAA) is designed so that the initial state $|\psi\rangle$ is rotated within a 2D plane until it has a large overlap with the target state $|\phi\rangle$. However, in quantum-walk search, one often utilizes a search operator:

$$A = WS_\phi$$

where W is the update/walk unitary whose $(+1)$ -eigenvector is the raw starting state, and $S_\phi = I - 2|\phi\rangle\langle\phi|$ is the reflection about the target state $|\phi\rangle$.

Spectral information about A may suffice to *detect* that a target exists (e.g., via phase estimation), but it does not automatically imply that repeated use of A will *produce* $|\phi\rangle$ with constant probability upon measurement. Controlled Quantum Amplitude Amplification (CQAA) resolves this detection-versus-finding bottleneck by wrapping the search dynamics in a one-qubit control register. This framework converts abstract detection into a concrete finding procedure while preserving the same asymptotic quantum hitting time.

B. Controlled Architecture and the Principal Eigenvector

Dohotaru and Høyer [7] introduce a control qubit together with a specific rotation R that defines a rotated basis:

$$\begin{aligned} |\tilde{0}\rangle &= R|0\rangle = \cos(\tilde{\theta})|0\rangle + \sin(\tilde{\theta})|1\rangle \\ |\tilde{1}\rangle &= R|1\rangle = -\sin(\tilde{\theta})|0\rangle + \cos(\tilde{\theta})|1\rangle \end{aligned}$$

Algorithm 4: Controlled Quantum Amplification (CQAA)

Input: (1) A search unitary $A = WS_\phi$ where W is an update walk and $S_\phi = I - 2|\phi\rangle\langle\phi|$
(2) A tuned control parameter $\tilde{\theta}$
Output: State with constant overlap on the marked set $|\tilde{1}, g\rangle$
// **Initialization and Wrapper Control**
Prepare logical control qubit in rotated basis $\{|\tilde{0}\rangle, |\tilde{1}\rangle\}$;
Define conditional search circuit: $U = (|0\rangle\langle 0| \otimes W + |1\rangle\langle 1| \otimes I)(|\tilde{0}\rangle\langle\tilde{0}| \otimes S_\phi + |\tilde{1}\rangle\langle\tilde{1}| \otimes I)$;
// **Principal Eigenvector Amplification**
Construct the unique $(+1)$ -eigenvector $|U_0\rangle$ that links both initial and marked states;
Apply the controlled search iterate until the principal eigenspace contributes constant overlap to the marked configuration;
return $|\tilde{1}, g\rangle$ *component of the measurement*;

The controlled search circuit U is the product of two conditional operators:

1. **Controlled Reflection:** $|\tilde{0}\rangle\langle\tilde{0}| \otimes S_\phi + |\tilde{1}\rangle\langle\tilde{1}| \otimes I$
2. **Controlled Update:** $|0\rangle\langle 0| \otimes W + |1\rangle\langle 1| \otimes I$

Let the initial overlap be $\sin(\theta) = |\langle\phi|\psi\rangle|$, with initial success probability $\epsilon = \sin^2 \theta$. We define the orthogonalized search-plane start vector (the “bad” component):

$$|\phi^\perp\rangle = \frac{|\psi\rangle - \sin(\theta)|\phi\rangle}{\cos(\theta)}$$

In the sparse-target regime $\epsilon < 1/2$, we tune the control rotation such that:

$$\sin(\tilde{\theta}) = \frac{\sin(\theta)}{\cos(\theta)} = \tan(\theta)$$

With this specific tuning, the circuit U acquires a distinguished **$(+1)$ -eigenvector** $|U_0\rangle$ on the joint Hilbert space:

$$|U_0\rangle = \frac{1}{\sqrt{2}} \left(|0, \phi^\perp\rangle - |\tilde{1}, \phi\rangle \right)$$

C. The Core Insight

The principal eigenspace of U now simultaneously contains the initial configuration and the marked configuration with constant overlap.

In standard AA, the amplitude moves from $|\phi^\perp\rangle$ to $|\phi\rangle$ within the search register. In CQAA, the evolution occurs across the joint system. The state $|U_0\rangle$ serves as a stationary bridge: by applying the controlled iterate for a number of steps governed by the corresponding quantum hitting time, the system evolves so that a measurement

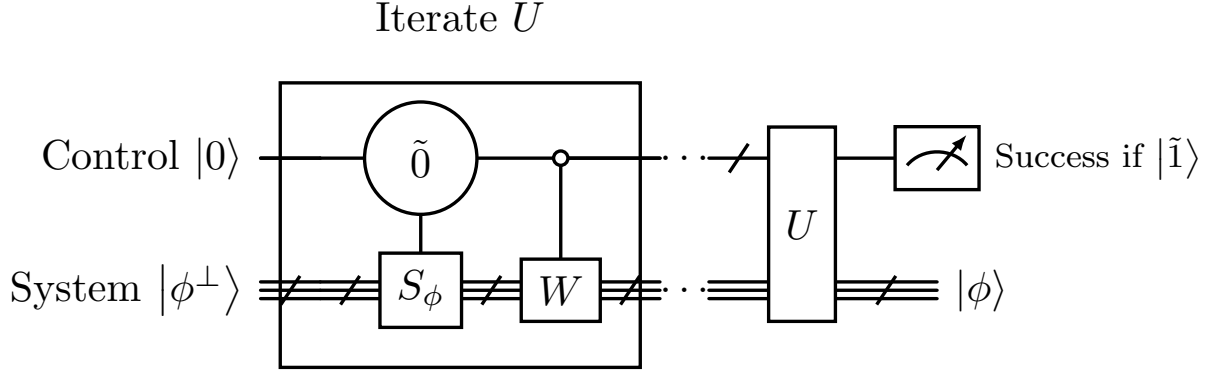


FIG. 7. **Controlled Quantum Amplification (CQAA) Circuit Architecture.** CQAA embeds search dynamics within a one-qubit control register. Following Dohotaru and Høyer, the iterate U is built from two sequential conditional operations: a controlled reflection S_ϕ and a controlled update walk W . Success is heralded by the control-qubit component aligned with $|\tilde{1}\rangle$.

of the control qubit in the $|\tilde{1}\rangle$ state collapses the search register into the target state $|\phi\rangle$.

In this way, the controlled wrapper converts an abstract walk/detection process into a robust finding process on the enlarged Hilbert space.

Intuition Box: Controlled Quantum Amplification

Core mechanism.

Embeds a quantum walk search sequence inside an auxiliary control qubit, redefining hitting times entirely through global unitary evolution.

Geometric interpretation.

Standard walks often only detect a target, oscillating past it in spectrum. By shifting into the parameterized basis $|\tilde{0}\rangle, |\tilde{1}\rangle$, CQAA forces the target and the starting configuration to coexist statically inside the newly formed principal eigenstate of the modified walk operator.

Structural insight.

Shows that detection-to-finding conversions do not inherently require iterative backtracking; one can absorb the detection dynamics natively into an interpolated walk structure using simple one-qubit controls.

A. Linear Combination of Unitaries (LCU)

Since the set of unitary operators spans the set of linear operators, any operator H can be expressed as

$$H = \sum_{l=0}^{L-1} \beta_l U_l \quad (40)$$

where $C = \sqrt{\sum_l \beta_l^2} = 1$. C is not necessarily 1 for every H , but for implementation purposes, every H is assumed to be normalized. The main steps of this implementation are as follows:

Wave division: Start with an ancillary state

$$|\psi_i\rangle = \sum_l \beta_l |l\rangle$$

which contains the coefficients of H . The target system is assumed to be $|\psi_0\rangle$. Then the combined system starts in:

$$|\psi_1\rangle = |\psi_i\rangle |\psi_0\rangle = \sum_l \beta_l |l\rangle |\psi_0\rangle$$

Controlled U operations: Now use $|l\rangle$ as control to apply U_l to $|\psi_0\rangle$:

$$|\psi_2\rangle = \sum_l \beta_l |l\rangle U_l |\psi_0\rangle$$

Wave combination: To obtain $H |\psi_0\rangle = \sum_l \beta_l U_l |\psi_0\rangle$, we need a recombination step. This reintegration happens in the ancillary qubits in $|0\rangle^{\otimes m}$.

$$|\psi_2\rangle = \frac{1}{\sqrt{L}} |0\rangle^{\otimes m} \sum_l \beta_l U_l |\psi_0\rangle = \frac{1}{L} H |\psi_0\rangle$$

Measurement: Measure the ancilla qubits, if the output is $|0\rangle^{\otimes m}$, then process succeeds with probability $\|H |\psi_0\rangle\|^2 / L$.

VIII. FIXED-POINT OBLIVIOUS AMPLITUDE AMPLIFICATION

As established in Section VI, OAA removes dependence on known-input reflections. But it still inherits Grover-style overshoot when the subnormalization (equivalently, effective success probability p) is unknown or drifting [26]. Fixed-Point Oblivious Quantum Amplitude Amplification (FOQA) addresses this by introducing a damping schedule that suppresses oscillatory overshoot and drives the process toward a stable high-success regime.

As is clear for above steps is that the wave combination step is the most crucial step. Now we describe the LCU operation. It starts with a rotation about y of the ancilla qubit.

$$\begin{aligned} |\bar{0}\rangle &= V_n |0\rangle = \cos \frac{\alpha_n}{2} |0\rangle + \sin \frac{\alpha_n}{2} |1\rangle \\ |\bar{1}\rangle &= V_n |1\rangle = -\sin \frac{\alpha_n}{2} |0\rangle + \cos \frac{\alpha_n}{2} |1\rangle \end{aligned}$$

This is followed by $-Z$ controlled by index qubit ($|0\rangle$). At this point, apply oblivious operator U^\dagger to index + content register. It is followed by controlled Z on the index qubit by the ancilla qubit ($|\bar{0}\rangle$). Finally, V_n^\dagger is applied to ancilla and U to index + content register.

B. The algorithm

The algorithm contains three registers: ancilla (for combination step), index (for oblivious step) and content (containing $|\varphi\rangle$). Just like oblivious amplitude amplification, we assume existence of a unitary operator U that prepares the initial index + content register state.

Algorithm 5: Fixed-point Oblivious Quantum Amplitude-amplification (FOQA)

Input : Initial state $|\varphi\rangle$, Oracle U , Damping parameters $\{\alpha_1, \alpha_2, \dots, \alpha_N\}$
Output : Amplified target state $V|\varphi\rangle$ in the content register

// Initialization
 Prepare registers: $|a\rangle|i\rangle|c\rangle \leftarrow |0\rangle|0\rangle|\varphi\rangle$
 Apply database generation: $|\Psi_0\rangle = U(|0\rangle|0\rangle|\varphi\rangle)$
 // Initial state is
 $\sin(\theta)|0\rangle V|\varphi\rangle + \cos(\theta)|1\rangle|\phi\rangle$

for $n \leftarrow 1$ **to** N **do**
 // Step 1: Quantum Wave Division (QWD)
 Apply rotation $V_n = e^{-i\alpha_n Y/2}$ to ancilla $|a\rangle$
 // Step 2: Entanglement Generation (Part A)
 Apply $(-Z)$ to ancilla $|a\rangle$ conditioned on index register $|i\rangle = |0\rangle$
 // Step 3: Work Register Transformation
 Apply U^\dagger to work registers (Index and Content)
 // Step 4: Entanglement Generation (Part B)
 Apply Z to index register $|i\rangle$ conditioned on ancilla $|a\rangle = V_n|0\rangle$
 // Step 5: Quantum Wave Combination (QWC)
 Apply V_n^\dagger to ancilla $|a\rangle$ and U to work registers
 // Step 6: Success Verification
 Measure ancilla register $|a\rangle$
if *Measurement result is* $|0\rangle$ **then**
 | **return** target state in register $|c\rangle$
 | **break**
end
else
 | Continue to iteration $n + 1$
end
end

$$U|0\rangle|\varphi\rangle = \sin \theta |0\rangle V|\varphi\rangle + \cos \theta |1\rangle|\phi\rangle \quad (41)$$

where $V|\varphi\rangle$ is the target state for searching and the corresponding index state is $|0\rangle$. The LCU circuit amplifies the amplitude of the target state. The angles are chosen according to the fixed-point amplification mechanism. At each stage, the ancillary register is measured. The algorithm terminates if the outcome is $|0\rangle$, otherwise LCU operation is repeated.




FIG. 8. **FOQA Circuit Architecture with Linear Combination of Unitaries (LCU)**. Internal LCU iterations are mediated through an ancilla so the oblivious-amplification step can be implemented without explicit reflection about an unknown input state. The vertical control lines show the two entangling controls: first, a $-Z$ gate on the ancilla conditioned on the index register being in $|0\rangle$; second, a Z gate on the index conditioned on the ancilla being in the rotated state $|\bar{0}\rangle$. The ancilla measurement then heralds whether the fixed-point oblivious amplification step has succeeded.

Intuition Box: Fixed-Point Oblivious Amplification

Fixed-Point Oblivious Quantum Amplification (FOQA) overcomes the reliance on precise state reflections in traditional $SU(2)$ Grover search by implementing a Linear Combination of Unitaries (LCU) to handle unknown initial states. It utilizes an ancillary qubit to apply a rotation $V_n = e^{-i\alpha_n Y/2}$, where the angle α_n acts as a damping coefficient that suppresses Grover-style over-rotation while preserving coherent amplification in the oblivious setting. By determining the sequence of damping angles via the Mizel fixed-point method, the algorithm drives the success probability toward a stable high-success regime rather than the oscillatory behavior of fixed-angle Grover search. Structurally, this alternating application of ancillary rotations and the work unitary U acts as a precursor to Quantum Singular Value Transformation (QSVT), functioning as a polynomial transformation that shapes the singular values of the oblivious operator to isolate the target signal.

IX. DISTRIBUTED AMPLITUDE AMPLIFICATION

A. Motivation: NISQ Hardware Constraints

The amplitude amplification procedures discussed in the previous sections assume the availability of a single coherent quantum processor capable of executing circuits on an n -qubit Hilbert space of dimension 2^n . While theoretically elegant, this assumption becomes increasingly unrealistic on present-day Noisy Intermediate-Scale Quantum (NISQ) hardware. Current devices suffer from severe

constraints on qubit counts, limited coherence times, calibration drift, and connectivity restrictions. This motivates distributed amplitude amplification architectures as developed recently in [17].

Distributed Quantum Amplitude Amplification (DQAA) addresses these architectural limitations by partitioning the global search space across multiple smaller quantum processors. Rather than executing the search algorithm on a monolithic n -qubit device, the algorithm distributes subproblems across 2^j processors, each operating on only $n - j$ qubits. The processors execute amplitude amplification locally, and the classical controller subsequently aggregates the measurement outcomes.

It is crucial to emphasize that this approach does not alter the fundamental asymptotic complexity of quantum search. The optimal query complexity $\mathcal{O}(\sqrt{N/M})$ established for Grover's algorithm remains unchanged. Instead, DQAA should be interpreted as an *architectural redistribution* of the computation, trading hardware width for reduced per-device qubit requirements.

B. Partitioning the Global Search Space

Let the global search function be

$$f : \{0, 1\}^n \rightarrow \{0, 1\},$$

where $f(x) = 1$ indicates a marked solution.

We partition the input register into a prefix of j qubits and a suffix of $n - j$ qubits,

$$x = (i_k, y), \quad i_k \in \{0, 1\}^j, \quad y \in \{0, 1\}^{n-j}.$$

The prefix i_k indexes the processor responsible for evaluating the corresponding subspace. Each processor therefore receives a restricted oracle

$$f_k(y) = f(i_k y)$$




FIG. 9. **Distributed Quantum Amplitude Amplification (DQAA) Architecture.** The DQAA framework partitions a global n -qubit search space across 2^j independent processors. Each node operates on a reduced $n - j$ qubit Hilbert space by evaluating a local sub-function $f_k(y) = f(i_k y)$. By assuming a separable initialization operator ($\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$), the processors execute fixed-point iterations locally without quantum communication. The “Lucky Node Theorem” (Theorem 3) guarantees that at least one processor \hat{k} starts with a local success probability $p_{\hat{k}} \geq p$, ensuring the global target is found with probability at least $1 - \epsilon^2$.

Thus the global function decomposes into 2^j independent sub-functions

$$f_k : \{0, 1\}^{n-j} \rightarrow \{0, 1\}$$

Each processor performs amplitude amplification within the Hilbert space

$$\mathcal{H}_k = \text{span}\{|y\rangle : y \in \{0, 1\}^{n-j}\}$$

C. Tensor Decomposition of the Initialization Operator

A key structural requirement for distributed amplification is that the global initialization unitary admits a strict tensor decomposition

$$\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$$

where \mathcal{A}_1 acts on the prefix register and \mathcal{A}_2 acts on the suffix register. The global initial state therefore factorizes as

$$\mathcal{A}|0\rangle^{\otimes n} = (\mathcal{A}_1|0\rangle^{\otimes j}) \otimes (\mathcal{A}_2|0\rangle^{\otimes(n-j)})$$

Within processor k , the local suffix state is therefore

$$|\psi_k\rangle = \mathcal{A}_2|0\rangle^{\otimes(n-j)}$$

Using the decomposition introduced in Section IV, the local Hilbert space separates into

$$\mathcal{H}_k = \mathcal{H}_{\text{Good},k} \oplus \mathcal{H}_{\text{Bad},k}$$

Let

$$p_k = \langle \psi_k | \Pi_{\text{Good},k} | \psi_k \rangle$$

denote the local success probability. Equivalently, one may write

$$|\psi_k\rangle = \sqrt{p_k} |\text{Good}_k\rangle + \sqrt{1 - p_k} |\text{Bad}_k\rangle,$$

where $|\text{Good}_k\rangle \in \mathcal{H}_{\text{Good},k}$ and $|\text{Bad}_k\rangle \in \mathcal{H}_{\text{Bad},k}$ are normalized and orthogonal.

a. Limitations of the Tensor Assumption. The requirement $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$ is highly restrictive. In many modern quantum algorithms the initialization procedure generates significant entanglement across registers. In particular, block-encoded operators and Quantum Singular Value Transformation (QSVD) pipelines rely on entangled system-ancilla registers that generally do not admit simple tensor factorizations. Consequently, distributed amplitude amplification is primarily applicable to algorithms whose initial states are separable across the partition boundary.

D. Local Probability Obstruction

Suppose the global algorithm has success probability

$$p = \sum_{x:f(x)=1} |\alpha_x|^2$$

Although this quantity may be known or bounded, its distribution across the 2^j partitions is generally unknown. The local success probabilities p_k may vary dramatically between nodes. Consequently, the standard Grover stopping rule cannot be applied locally because the required iteration count depends on p_k .

The following theorem addresses this obstruction.

Theorem 1 (Lucky Node Theorem). *Let p denote the global success probability and p_k the success probability associated with partition k . Then there exists at least one node \hat{k} such that*

$$p_{\hat{k}} \geq p$$

Proof. Using $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$, write the initialized global state as

$$|\Psi\rangle = \sum_{k=0}^{2^j-1} \sum_{y \in \{0,1\}^{n-j}} \alpha_{k,y} |i_k\rangle |y\rangle$$

For each partition k , define

$$a_k = \sum_{y: f_k(y)=1} |\alpha_{k,y}|^2, \quad b_k = \sum_{y: f_k(y)=0} |\alpha_{k,y}|^2$$

so $w_k := a_k + b_k$ is the total probability weight of partition k and $\sum_k w_k = 1$. If $w_k > 0$, define local success probability

$$p_k = \frac{a_k}{w_k}, \quad 0 \leq p_k \leq 1$$

The global success probability is

$$p = \sum_{k=0}^{2^j-1} a_k = \sum_{k=0}^{2^j-1} w_k p_k$$

Assume for contradiction that $p_k < p$ for all k with $w_k > 0$. Multiplying by $w_k \geq 0$ and summing gives

$$\sum_k w_k p_k < \sum_k w_k p = p \sum_k w_k = p$$

This implies $p < p$, a contradiction. Therefore there exists at least one \hat{k} such that $p_{\hat{k}} \geq p$. \square

This theorem shows that there exists a *lucky node* whose initial success probability is at least as large as the global value.

Algorithm 6: Distributed Quantum Amplitude Amplification (DQAA-2025)

Input: (1) Global search space partitioned across 2^j quantum processors
 (2) Initialization operator decomposed tightly via $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$
Output: Results aggregated classically achieving high likelihood of isolating a valid target
 // **Parallel Partitioning**
 Node k uniquely addresses a specific problem slice mapped via local oracle f_k ;
 // **Local Amplification via Lucky Node guarantees**
for every processor $k \in \{0, \dots, 2^j - 1\}$ *in parallel*
do
 Apply Fixed-Point Amplitude Amplification (FPAA) using fixed global lower bound parameter p ;
 Measure localized output;
end
 // **Classical Aggregation**
 Classical controller merges localized data prefixes against successfully generated suffixes to identify target elements;

E. Distributed Fixed-Point Amplification

Because the Lucky Node Theorem implies that $p_k \geq p$ for at least one node, the Fixed-Point Amplitude Amplification (FPAA) procedure of Section V can be applied at each processor using the global lower bound p .

Each processor executes the local iteration operator

$$Q_k = -\mathcal{A}_2 S_0 \mathcal{A}_2^\dagger S_{f_k}$$

where

$$S_{f_k} = I - 2\Pi_{\text{Good},k}$$

is the phase oracle for the local function f_k , and

$$S_0 = I - 2|0\rangle^{\otimes(n-j)}\langle 0|^{\otimes(n-j)}$$

This is the reflection about the all-zero state of the suffix register.

The phase sequence used in these reflections follows the Chebyshev fixed-point schedule derived previously in Section V. Each processor therefore executes the same FPAA sequence optimized for the lower bound p .

Intuition Box: Distributed Amplitude Amplification

Distributed Amplitude Amplification segments a large amplification problem across 2^j smaller NISQ processors by relying on a tensor product constraint ($\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$) to reduce circuit width. Because target items distribute non-uniformly, the computing nodes isolate independent Hilbert spaces where "Lucky Node" bounds guarantee that at least one local subspace begins with an enhanced target density. Consequently, each physical computer executes shallower circuits on fewer qubits, utilizing fixed-point techniques to suppress overshooting caused by unpredictable target weights, thereby improving practical hardware feasibility despite maintaining the same asymptotic depth. However, the method's primary limitation is its strict reliance on separability; the independent partitioning breaks down if the algorithm generates system-ancilla entanglement across registers, such as during block encoding.

F. Global Success Guarantee

Theorem 2. *Let each processor execute the FPAA schedule with l generalized iterates satisfying*

$$l \geq \frac{1}{2\sqrt{p}} \log \left(\frac{2}{\epsilon} \right)$$

Equivalently, the corresponding odd sequence length satisfies

$$L = 2l + 1 = \mathcal{O} \left(\frac{1}{\sqrt{p}} \log \frac{2}{\epsilon} \right)$$

Then the distributed algorithm succeeds with probability at least $1 - \epsilon^2$

Proof. By the Lucky Node Theorem there exists a node \hat{k} with $p_{\hat{k}} \geq p$. Running FPAA with

$$l \geq \frac{1}{2\sqrt{p}} \log \left(\frac{2}{\epsilon} \right)$$

therefore also satisfies the amplification requirement for node \hat{k} , so its final success probability obeys

$$P_{\hat{k}} \geq 1 - \epsilon^2$$

Let P_k denote the final success probability at node k . Since node executions are independent, the global success probability is

$$P_{\text{tot}} = 1 - \prod_{k=0}^{2^j-1} (1 - P_k)$$

Hence

$$P_{\text{tot}} \geq P_{\hat{k}} \geq 1 - \epsilon^2$$

Therefore the distributed algorithm succeeds with probability at least $1 - \epsilon^2$. \square

G. Complexity Analysis

The parallel runtime of the distributed algorithm is

$$\mathcal{O} \left(\frac{1}{\sqrt{p}} \log \left(\frac{2}{\epsilon} \right) \right),$$

which matches fixed-point amplitude amplification up to the same logarithmic error factor.

However, the aggregate computational work across all processors is

$$2^j \cdot \mathcal{O} \left(\frac{1}{\sqrt{p}} \log \left(\frac{2}{\epsilon} \right) \right).$$

Thus DQAA preserves Grover optimality while trading increased hardware parallelism for reduced per-device circuit size.

H. Empirical Validation

To validate the theoretical guarantees in a concrete setting, we performed a numerical simulation on a small distributed search instance.

The global problem used an $n = 6$ qubit search space ($N = 64$). We chose partition parameter $j = 2$, yielding $2^j = 4$ independent processors, each operating on a suffix register of $n - j = 4$ qubits. Each processor therefore ran fixed-point amplitude amplification over a local space of dimension $2^4 = 16$.

Two distinct behaviors were observed:

- **Unlucky nodes** ($p_k = 0$): if a partition contained no marked element, fixed-point iteration could not amplify absent target amplitude. These nodes produced near-uniform measurement outcomes over the local suffix basis, i.e., random-noise-like behavior.
- **Lucky nodes** ($p_k > 0$): partitions containing marked states converged monotonically toward the target subspace. The success probability increased without Grover-style overshoot and exceeded the bound $1 - \epsilon^2$ in the tested runs.

This experiment supports the practical content of the Lucky Node Theorem: under uneven partition statistics, at least one processor can still be amplified reliably to high success probability.

I. Hardware Scaling

Let n denote the global qubit count and j the partition parameter.

Quantity	Value
Processors	2^j
Qubits per processor	$n - j$
Total qubits	$2^j(n - j)$

Three scaling regimes are particularly relevant:

- $j = \mathcal{O}(1)$: modest hardware parallelism with minimal overhead.
- $j = \mathcal{O}(\log n)$: polynomial processor growth.
- $j = \Theta(n)$: exponential processor requirements.

Practical implementations therefore restrict j to small values.

J. Architectural Advantages

Distributed amplification possesses several practical advantages:

- **Reduced qubit requirements.** Each processor requires only $n - j$ qubits.
- **Parallel execution.** All processors operate simultaneously.
- **Zero quantum communication.** Nodes operate independently and communicate only classically after measurement by concatenating local suffix outcomes with their corresponding prefixes.

These properties make DQAA particularly attractive for heterogeneous quantum networks composed of multiple small NISQ devices.

K. Limitations and Relation to QSVT

Despite its architectural advantages, DQAA exhibits important limitations. The strict tensor decomposition requirement $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$ prevents its direct application to algorithms whose initialization circuits generate significant entanglement across registers.

This restriction is substantial in modern algorithmic settings. Many ansatz-based state preparations, block-encoded operators, and Quantum Singular Value Transformation (QSVT) pipelines rely on correlated system-ancilla structure rather than separable initialization. In such settings, the partition boundary generally cuts across entangling operations, so the local processors no longer inherit an autonomous amplification problem of the form

assumed here. Consequently, extending distributed architectures to QSVT-based or block-encoding-based algorithms remains an open research direction rather than a direct corollary of the present framework.

This limitation marks the boundary between the geometric amplitude amplification techniques developed in the preceding sections and the algebraic operator transformations that characterize modern quantum algorithm design.

L. Distributed Exact Quantum Amplitude Amplification (DEQAAA)

The 2025 DQAA framework shows that amplitude amplification can be parallelized across multiple nodes, but it does so under three significant restrictions: the state-preparation unitary must factor as $\mathcal{A} = \mathcal{A}_1 \otimes \mathcal{A}_2$, the number of nodes is constrained to 2^j , and the resulting procedure remains approximate rather than exact. Zhou et al. (2026) [29] address these three bottlenecks simultaneously by proposing a Distributed Exact Quantum Amplitude Amplification Algorithm (DEQAAA) for quantum states with arbitrary amplitude distributions.

Historically, the paper presents DEQAAA as the next step in a distributed-search lineage that runs from earlier distributed Grover-style constructions through Hua and Qiu’s 2025 distributed amplitude amplification framework. In that progression, the 2025 result established that amplification could be spread across many small processors, while the 2026 result replaces the earlier fixed-point, tensor-factorized architecture by a distributed exact-amplification architecture built for arbitrary amplitude distributions [29].

The architectural change is substantial. Instead of requiring a power-of-two node count, DEQAAA supports any number of computing nodes t satisfying $2 \leq t \leq n$. If node j is assigned n_j qubits with $\sum_{j=0}^{t-1} n_j = n$, then the largest local memory requirement is

$$\max(n_0, n_1, \dots, n_{t-1})$$

rather than the $2^j(n - j)$ total-qubit overhead of the 2025 construction. The comparison table in [29] further emphasizes that DEQAAA removes the stringent tensor-product assumption on \mathcal{A} and supports exact multi-target amplification in settings where DQAA does not.

The algorithm proceeds in two phases. First, the global target set X_g is partitioned into local target sets by the substring map

$$h(x, s, l) = x_s x_{s+1} \cdots x_{s+l-1}$$

which induces local Boolean subfunctions f_j and node-wise success probabilities p_j [29]. Each node then performs a local exact amplitude amplification step using

$$EQ_j = \mathcal{A}_{|\phi_j\rangle} R_{|0\rangle}^{\phi_j} \mathcal{A}_{|\phi_j\rangle}^\dagger R_{f_j}^{\phi_j}$$

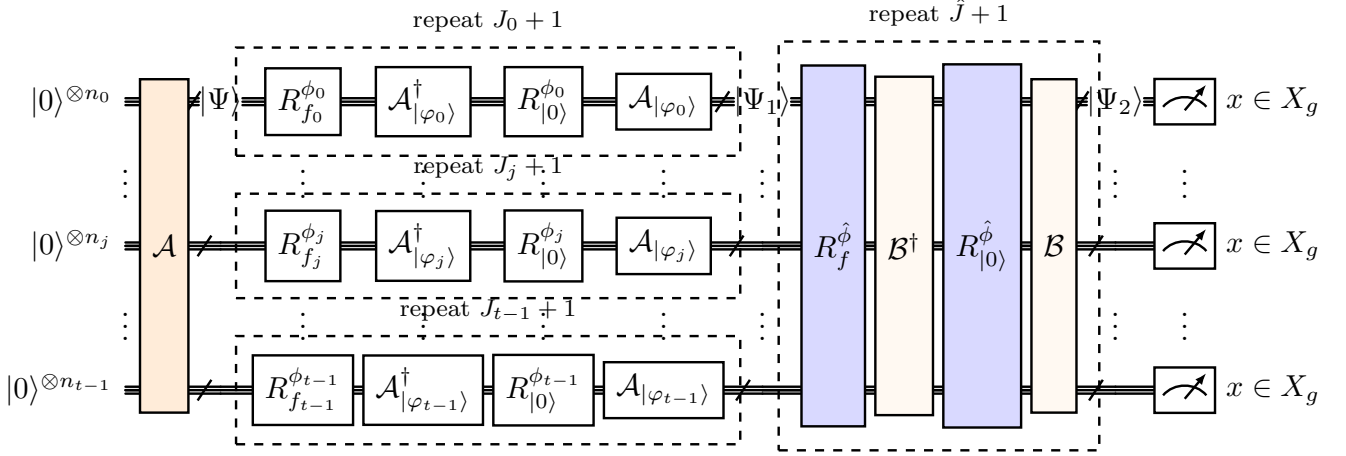


FIG. 10. **DEQAAA Architecture.** This two-phase framework achieves exact amplitude amplification by first parallelizing rotations locally across t nodes and then performing a global synchronization step. It removes the 2^j node constraint and the rigid $\mathcal{A}_1 \otimes \mathcal{A}_2$ assumption of previous approximate distributed models.

repeated $J_j + 1$ times, where

$$J_j = \left\lfloor \frac{\pi}{4 \arcsin(\sqrt{p_j})} - \frac{1}{2} \right\rfloor$$

and the phase is

$$\phi_j = 2 \arcsin \left(\frac{\sin(\pi/(4J_j + 6))}{\sqrt{p_j}} \right)$$

If the resulting global success probability after this local phase already satisfies $p'_g = 1$, the algorithm terminates. Otherwise, the full first-phase unitary is bundled into a composite operator

$$\mathcal{B} = \left[\bigotimes_{j=0}^{t-1} (EQ_j^{J_j+1}) \right] \mathcal{A}$$

and a second, global exact amplification step is applied:

$$\widehat{EQ} = \mathcal{B} R_{|0>^{\otimes n}}^{\hat{\phi}} \mathcal{B}^\dagger R_f^{\hat{\phi}}, \quad \hat{J} = \left\lfloor \frac{\pi}{4 \arcsin(\sqrt{p'_g})} - \frac{1}{2} \right\rfloor$$

with

$$\hat{\phi} = 2 \arcsin \left(\frac{\sin(\pi/(4\hat{J} + 6))}{\sqrt{p'_g}} \right)$$

and the global exact operator is applied $\hat{J} + 1$ times. This two-phase construction is the mechanism by which the paper proves exact final amplification for the global target set [29].

It is also worth emphasizing that the algorithm is not merely “distributed Grover” with a corrected stopping rule. Structurally, both phases are distributed instances of Exact Quantum Amplitude Amplification (EQAAA): the

first phase executes local exact amplification on each node, while the second phase applies a global exact-amplification step using \mathcal{B} as the effective state-preparation unitary. In that sense, DEQAAA should be read as a distributed realization of EQAAA rather than as a fixed-point search method. This marks an important contrast with the 2025 DQAA construction, whose distributed correctness relied on Yoder–Low–Chuang fixed-point schedules precisely because the local success amplitudes were not known in advance.

From the perspective of this survey, DEQAAA is especially notable because it shifts distributed amplification away from the lucky-node logic of the 2025 model and toward explicit phase synthesis. Standard distributed Grover-style constructions depend on rigid sign flips; by contrast, DEQAAA replaces those reflections with parameterized phase-rotation operators R_f^{ϕ} and $R_{|0>^{\otimes n}}^{\phi}$, bringing the distributed setting conceptually closer to the phase-engineering viewpoint that culminates in QSVT.

The cost model in [29] is also different from the query-centric analysis used earlier in this survey. The paper studies decomposed circuit depth and proves

$$\text{dep}(\text{DEQAAA}) = \begin{cases} \text{dep}(\text{First phase}), & p'_g = 1, \\ \text{dep}(\text{First phase}) + \text{dep}(\text{Second phase}) & \text{otherwise.} \end{cases} \quad (42)$$

$$\begin{aligned} \text{dep}(\text{First phase}) &= \text{dep}(\mathcal{A}) \\ &+ \max_j \left\{ \text{dep}(j\text{-th node}, p_j, |X_j|) \right\} \end{aligned} \quad (43)$$

This is paired with an explicit decomposition of the multi-controlled phase-shift gate $C^{n-1}PS(\phi)$ into elementary gates. In the paper’s MindSpore Quantum experiments,

the resulting distributed circuits achieve more than a 97% reduction in both gate count and circuit depth in the 10-qubit example relative to centralized QAAA and EQAAA implementations [29].

At the same time, DEQAAA introduces a different hidden cost. The local operators $\mathcal{A}_{|\phi_j\rangle}$ must prepare substates derived from the global amplitude distribution, so the burden of tensor-factorizing \mathcal{A} is replaced by the burden of synthesizing local substates. Moreover, the paper explicitly notes that exact amplification presumes access to the exact probability distribution P of the input state. In realistic experiments one instead estimates an approximate distribution \tilde{P} from repeated measurements and quantifies the mismatch via the Kullback–Leibler divergence

$$D_{\text{KL}}(\tilde{P}\|P) = \sum_x \tilde{P}(x) \times \log \frac{\tilde{P}(x)}{P(x)}$$

The mapping from a local probability distribution P_j to a substate $|\phi_j\rangle$ is also not unique; the paper explicitly allows any substate consistent with P_j and points to concrete synthesis routes such as Top-Down amplitude encoding, uniformly controlled rotations, and QGAN-based preparation circuits. Thus DEQAAA should be read not as eliminating all state-preparation bottlenecks, but as relocating them: it removes the rigid tensor-product restriction of DQAA, while replacing it with a more flexible but still nontrivial requirement of probability-aware local substate synthesis.

The empirical story in the paper is likewise slightly more nuanced than the headline percentages alone suggest. In the smallest undecomposed examples, DEQAAA does not automatically outperform centralized QAAA and EQAAA in raw gate count or depth. The practical advantage emerges after multi-controlled-gate decomposition, and it becomes markedly stronger as the qubit count grows. This is consistent with the survey-level interpretation advanced here: the principal gain is architectural and compilation-level scalability, rather than an asymptotic query-speedup over single-node exact amplification [29].

Algorithm 7: Distributed Exact Quantum Amplitude Amplification Algorithm (DEQAAA-2026)

Input: (1) An arbitrary array of t connected processor nodes (where $2 \leq t \leq n$)
(2) An arbitrarily distributed global problem requiring exact amplification
Output: Algorithm outputs the globally valid item exclusively ($p'_g = 1$)

// Phase 1: Local Exact Amplification
Partition the global target space into t localized subproblems with dynamically analyzed local target likelihoods p_j ;
for each node j in parallel **do**
 Run EQ_j strictly parameterized by ϕ_j to
 forcefully rotate local targets exactly onto axis;
end
if Total global success probability reaches unity **then**
 Terminate successfully;
end
// Phase 2: Global Exact Verification
Composite resulting aggregated unitary into \mathcal{B} ;
Run unified step \widehat{EQ} exactly $\hat{J} + 1$ times under unified parameter $\hat{\phi}$;

Intuition Box: Distributed Exact Amplitude Amplification
Core mechanism.

Moves away from approximate fixed-point limits to calculate and synthesize custom multi-controlled target amplitudes across multiple variable-cost nodes structurally in a robust two-stage process.

Algorithmic consequence.

Bypasses the restrictive tensor factorization restriction entirely, accommodating nearly arbitrary initial probability shapes. Drastically compresses required overall circuit layout while delivering exact amplitude mapping.

Structural insight.

It acts as the truest bridge in distributed topologies; relocating state preparation burden (requiring intense local probability evaluation) while physically embracing the phase mapping logic foundational to QSP and QSVT.

X. VARIABLE-TIME AMPLITUDE AMPLIFICATION

Standard amplitude amplification models the state-preparation routine \mathcal{A} as a rigid black box. If \mathcal{A} internally branches so that some computational paths halt early while others continue much longer, coherent amplification

Feature	DQAA (2025)	DEQAAA (2026)
Exact solution	No	Yes
Stringent assumption on \mathcal{A}	Yes	No
Arbitrary initial amplitude distribution	No	Yes
Number of computing nodes	2^j	$2 \leq t \leq n$
Maximum qubits at one node	$n - j$	$\max(n_0, \dots, n_{t-1})$

TABLE IV. Key contrasts between the 2025 distributed amplitude amplification framework and the 2026 DEQAAA extension, adapted from [17, 29].

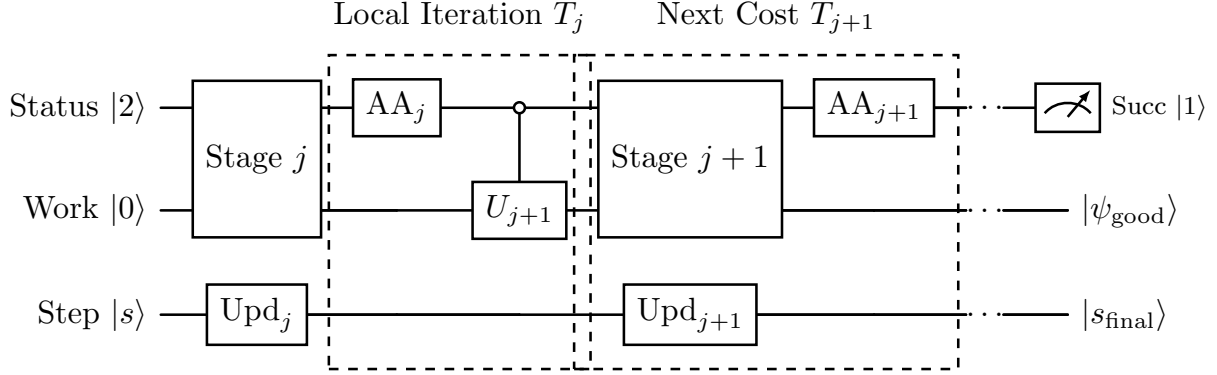


FIG. 11. **Variable-Time Amplitude Amplification (VTAA) Filtering Architecture.** The circuit operates in m discrete stages. At each stage j , the algorithm performs evolution and local amplitude amplification (AA) for cost T_j . The status register $|O\rangle$ flags branches as successful ($|1\rangle$), failed ($|0\rangle$), or continuing ($|2\rangle$). The open-circle control (\circ) ensures that subsequent unitary evolutions U_{j+1} act exclusively on the “continuing” superposition, preventing halted branches from incurring the worst-case cost T_m . The step register $|s\rangle$ prevents interference between branches resolved at different times.

still requires padding all branches to the worst-case runtime. Denoting the maximal branch cost by T_m and the global success probability by p , the resulting complexity remains

$$\mathcal{O}\left(\frac{T_m}{\sqrt{p}}\right).$$

This worst-case synchronization overhead is the bottleneck resolved by Variable-Time Amplitude Amplification (VTAA) [1]. The obstruction is especially acute in linear-systems solvers such as HHL, where branch-dependent phase-estimation precision induces runtime spread controlled by the condition number κ .

A. Variable-Time Framework

Let a variable-time quantum procedure be specified by a sequence of stage-evolution unitaries

$$U_1, U_2, \dots, U_m,$$

and define the prefix algorithm

$$\mathcal{A}_j := U_j U_{j-1} \cdots U_1,$$

with cumulative query cost T_j satisfying $T_1 < T_2 < \cdots < T_m$. The state after stage j is $|\Psi_j\rangle = \mathcal{A}_j |0\rangle^{\otimes n}$.

To formalize early halting coherently, introduce an auxiliary status space

$$\mathcal{H}_{\text{status}} = \text{span}\{| \text{succ} \rangle, | \text{fail} \rangle, | \text{cont} \rangle\}.$$

At stage j , the joint state admits the orthogonal decomposition

$$|\Psi_j\rangle = \sqrt{p_{s,j}} | \text{succ} \rangle | \text{Good}_j \rangle + \sqrt{p_{f,j}} | \text{fail} \rangle | \text{Bad}_j \rangle + \sqrt{p_{c,j}} | \text{cont} \rangle | \Phi_j \rangle, \quad (44)$$

where $|\Phi_j\rangle$ contains branches that must continue to U_{j+1} , and $p_{s,j} + p_{f,j} + p_{c,j} = 1$.

Extending the Section I decomposition, the target space is a stage-resolved direct sum

$$\mathcal{H}_{\text{Good}} = \bigoplus_{j=1}^m \mathcal{H}_{\text{Good},j},$$

where $\mathcal{H}_{\text{Good},j}$ contains valid solutions that halt exactly at stage j . Consequently, the final prepared state can be written as

$$| \text{All} \rangle = \sum_{j=1}^m \sqrt{p_j} | \text{Good}_j \rangle + \sqrt{1-p} | \text{Bad} \rangle,$$

with

$$p = \sum_{j=1}^m p_j$$

B. Variable-Time Amplification Limit

Theorem 3 (Variable-Time Amplification Limit [11]). *For a variable-time algorithm with stage probabilities $\{p_j\}_{j=1}^m$, cumulative costs $\{T_j\}_{j=1}^m$, and global success probability $p = \sum_{j=1}^m p_j$, there exists a variable-time amplification procedure that boosts success probability to $\Omega(1)$ using query complexity*

$$\mathcal{O}\left(T_m \log T_m + \frac{1}{\sqrt{p}} \sqrt{\sum_{j=1}^m p_j T_j^2} \cdot \log^{1.5} T_m\right).$$

The cost term $\sqrt{\sum_j p_j T_j^2}$ is the ℓ_2 -average stopping time appearing in Ambainis's analysis. When most success amplitude terminates at early stages ($T_j \ll T_m$ for dominant p_j), the second term can be substantially smaller than the worst-case synchronized cost $\mathcal{O}(T_m/\sqrt{p})$, although the additive $T_m \log T_m$ term remains.

The multiplicative $\text{polylog}(T_m)$ overhead arises from the precision requirements of coherent filtering subroutines (implemented via phase-estimation-style or polynomial filters) that separate $|\text{succ}\rangle$ from $|\text{cont}\rangle$ at each recursive level without collapsing superposition.

C. Recursive Filtering Mechanics

VTAA operates recursively by amplifying early-halting success components first. At stage $j = 1$, one applies standard amplification to increase overlap with the $|\text{succ}\rangle$ branch associated with $\mathcal{H}_{\text{Good},1}$. A coherent quantum filter then isolates branches already terminated from those in $|\text{cont}\rangle$, after which amplification is applied recursively to the continuing component at stage $j + 1$. Repeating this stage by stage prevents amplitudes that have already resolved at cost T_1, \dots, T_j from being forced through the full worst-case depth T_m , which is precisely the mechanism behind the ℓ_2 runtime gain.

D. Application to Linear Systems and Bridge to QSVT

Ambainis's original motivation was the Harrow–Hassidim–Lloyd (HHL) quantum linear systems algorithm [1]. In HHL, the required simulation/phase-estimation precision is strongly controlled by the condition number κ . Replacing standard amplitude amplification with VTAA reduces the dependence from $\mathcal{O}(\kappa^2 \log N)$ to $\mathcal{O}(\kappa \log^3 \kappa \log N)$, providing a concrete historical proof that variable-time structure yields genuine asymptotic gains.

Conceptually, VTAA still remains inside the Grover-reflection paradigm: it optimizes temporal inhomogeneity by recursive composition of geometric amplification primitives. The next step is more radical: replacing this

reflection geometry with direct polynomial transformations of singular values, leading to Section XI on Quantum Singular Value Transformation.

Algorithm 8: Variable-Time Amplitude Amplification (VTAA)

Input: A sequence of unitaries U_1, \dots, U_m that halt computations nonuniformly
Output: State matching target criteria bypassing worst-case sequence execution maximums
// Recursive Branch Filtering
for stage $j = 1$ **to** m **do**
 Execute U_j ;
 Partition space into $|\text{succ}_j\rangle$, $|\text{fail}_j\rangle$ and $|\text{cont}_j\rangle$;
 // Coherent Polynomial Subroutine
 Apply quantum amplitude amplification strictly bounded locally at cost sequence T_j ;
 Pass unhalting vectors through filtering operation allowing them exclusively independent extension towards U_{j+1} ;
end
Measure and identify correctly extracted state via $|\text{succ}\rangle$ projection branches;

Intuition Box: Variable-Time Amplitude Amplification

Core mechanism.

Stops rapid successful quantum states from waiting in the queue for unhalting states. Applies localized amplification iteratively checking for branch failures step-by-step.

Geometric interpretation.

Instead of projecting over a single axis, VTAA constructs a staircase of sub-spaces. Halted states are siphoned carefully into a secure plane using quantum polynomial filtration, effectively freezing their geometric decay while active states continue rotating.

Algorithmic consequence.

Substitutes an extremely detrimental maximum branch time (T_m) dependency with a much faster average weighted ℓ_2 -norm metric. Crucial to resolving linear systems solver overheads.

Structural insight.

The need for highly sensitive coherent phase filters connecting sequential computational layers logically begs for a single operation acting mathematically on variable matrix ranks at once, which is exactly what QSVT mathematically achieves via polynomials.

Amplification Framework	Suffers Overshoot?	Requires Exact p ?	Asymptotic Query Complexity
Standard Grover (Sec. IV)	Yes	Yes	$\mathcal{O}\left(\frac{1}{\sqrt{p}}\right)$
Fixed-Point AA (Sec. V)	No	No ^a	$\mathcal{O}\left(\frac{1}{\sqrt{p}} \log \frac{2}{\delta}\right)$
Oblivious AA (Sec. VI)	Yes	Yes	$\mathcal{O}\left(\frac{1}{\sqrt{p}}\right)$
Fixed-Point Oblivious (Sec. VIII)	No	No	$\mathcal{O}\left(\frac{1}{\sqrt{p}}\right)$ (average iterations, up to constants)
Distributed AA (Sec. IX)	No	No ^a	$\mathcal{O}\left(\frac{1}{\sqrt{p}} \log \frac{2}{\epsilon}\right)$ (Parallel Runtime)
Distributed Exact AA (Sec. IX)	No	Yes	exact two-phase distributed depth
Variable-Time AA (Sec. X)	No	No	$\mathcal{O}\left(T_m \log T_m + \frac{1}{\sqrt{p}} \sqrt{\sum p_j T_j^2} \log^{1.5} T_m\right)$

^a Requires only a reliable global lower bound, not the exact probabilistic value.

TABLE V. Comparison of the core amplitude amplification frameworks developed prior to QSVT. The evolution demonstrates a systematic elimination of the "soufflé problem" (overshoot) and the relaxation of prior-knowledge constraints, culminating in optimized parallel and temporally-inhomogeneous bounds.




FIG. 12. **Quantum Singular Value Transformation (QSVT) Architecture.** The circuit implements a degree- d polynomial transformation of the singular values of an operator A embedded in a block-encoding U . The phase qubit controls a sequence of projector-controlled phase shifts $\Pi_{\phi_k} = e^{i\phi_k(2\Pi - I)}$, interleaved with alternating applications of the unitary U and its adjoint U^\dagger . This unified framework generalizes amplitude amplification, Hamiltonian simulation, and matrix inversion through the choice of the phase vector $\vec{\phi}$.

XI. QUANTUM SINGULAR VALUE TRANSFORMATION

Amplitude amplification can be seen as *amplification* of singular values of an operator. Let us present how. We have seen that alternating sequence of reflections and rotations computes some polynomial of the amplitudes. Starting with the state

$$|All\rangle = \sqrt{p}|Good\rangle + \sqrt{1-p}|Bad\rangle \quad (45)$$

the success amplitude is \sqrt{p} . The objective of amplification is to increase this amplitude to 1. This translates to mapping the amplitudes via some polynomial such that the success amplitudes are mapped to 1. Consider this problem algebraically. Ideal objective of amplification is to apply the projector

$$A = |Good\rangle\langle All| \quad (46)$$

whose only non-zero singular value is \sqrt{p} . Since A is not unitary, we cannot build a circuit that applies it to the starting state directly. In addition, without amplification of $|All\rangle$, the probability of success is p . It is better to have another operator $P(A)$ for some polynomial $P(\cdot)$ such that

$P(\sqrt{p}) \approx 1$. The objective of QSVT is to transform the state by $P(A)$ for problem-specific polynomial P .

A. The general framework

In the case of amplification, the task reduces to applying the operator A . As the probability of favorable outcome is p (which can be very small), we want to apply a function $P(A)$ that preserves the favorable outcomes while increasing the success probability. This can be achieved by considering the singular value decomposition of A .

$$A = \sum_k \sigma_k |w_k\rangle\langle v_k|$$

Any function of A is applied directly to the singular values.

$$P(A) = \sum_k P(\sigma_k) |w_k\rangle\langle v_k|$$

Which polynomial to use will depend on the problem. Let us define *bounded parity polynomials*.

Definition 1. A polynomial $P(x)$ is called bounded parity polynomial if it satisfies the following conditions:

1. **Boundedness:** $|P(x)| \leq 1$ for all $x \in [-1, 1]$
2. **Parity:** $P(x)$ must have the same parity as its degree d , i.e.
 - If d is odd, then $P(x)$ must be an odd function, $P(x) = -P(-x)$
 - If d is even, then $P(x)$ must be an even function, $P(x) = P(-x)$

Intuition Box: Quantum Singular Value Transformation

Quantum Singular Value Transformation (QSVT) unifies matrix processing by applying alternating sequential phase variations ($\vec{\phi}$) to implement generalized polynomial transformations. Geometrically, rather than operating within a single two-dimensional reflection plane, QSVT constructs and simultaneously manipulates decoupled 2D singular sub-planes spanning the state space. Consequently, diverse operations, such as matrix inversion, unstructured search, and Hamiltonian simulation, are executed within an identical quantum circuit architecture simply by selecting different defining polynomials. Structurally, QSVT abstracts geometric amplitude dynamics into linear algebra, shaping the polynomials that act directly on block-encoded subnormalized matrices.

Once the right polynomial is identified, the next step is to embed the operator A inside a unitary U , often as the first block. Reusing the purified formalism from Section VI, let A act on s qubits and U on $a + s$ qubits. Define

$$\Pi := |0\rangle^{\otimes a} \langle 0|^{\otimes a} \otimes I_s. \quad (47)$$

Definition 2. U is an (α, a, ϵ) -block encoding of A if $\alpha \geq 1$ and

$$\left\| (\langle 0|^{\otimes a} \otimes I_s) U (|0\rangle^{\otimes a} \otimes I_s) - \frac{A}{\alpha} \right\| \leq \epsilon. \quad (48)$$

Equivalently, in the ancilla computational basis,

$$U = \begin{bmatrix} A/\alpha & * \\ * & * \end{bmatrix} \quad \text{up to error } \epsilon.$$

The parameter α is the subnormalization factor, a is the ancilla width, and ϵ is the synthesis error.

With the polynomial P and the unitary U identified, the next step is to obtain a sequence of angles $\vec{\phi} = (\phi_0, \dots, \phi_d)$ that will be needed to approximate P using U .

Define the projector-controlled phase unitary

$$\Pi_{\phi} := e^{i\phi(2\Pi - I)}. \quad (49)$$

For a phase vector $\vec{\phi} = (\phi_0, \dots, \phi_d)$, define

$$U_{\vec{\phi}} = \Pi_{\phi_0} \prod_{j=1}^d \left(U^{(\pm)} \Pi_{\phi_j} \right), \quad (50)$$

where $U^{(\pm)}$ alternates between U and U^\dagger .

Theorem 4 (QSVT Singular Value Transformation (Gilyén et al. [8])). *Let U be an (α, a, ϵ) -block encoding of*

$$A = \sum_k \sigma_k |w_k\rangle \langle v_k|,$$

and let P be a degree- d bounded parity polynomial. Then there exists $\vec{\phi}$ such that for odd d ,

$$\left\| (\langle 0|^{\otimes a} \otimes I) U_{\vec{\phi}} (|0\rangle^{\otimes a} \otimes I) - \sum_k P(\sigma_k/\alpha) |w_k\rangle \langle v_k| \right\| \leq \mathcal{O}(d\epsilon). \quad (51)$$

and for even d ,

$$\left\| (\langle 0|^{\otimes a} \otimes I) U_{\vec{\phi}} (|0\rangle^{\otimes a} \otimes I) - \sum_k P(\sigma_k/\alpha) |v_k\rangle \langle v_k| \right\| \leq \mathcal{O}(d\epsilon). \quad (52)$$

The parity condition determines output geometry: odd polynomials map right singular vectors to left singular vectors, while even polynomials act within the right-singular space.

Algorithm 9: Quantum Singular Value Transformation (QSVT)

Input: (1) A block-encoding unitary U containing target matrix A/α parameterized by a ancillae
 (2) A polynomial phase sequence vector $\vec{\phi} = (\phi_0, \dots, \phi_d)$
Output: A coherent operator block-encoding $P(A/\alpha)$

// Quantum Signal Processing Pipeline
 Initialize the signal register together with a ancilla qubits;
 Apply the initial phase operator $\Pi_{\phi_0} = e^{i\phi_0(2\Pi - I)}$;
for $k = 1$ **to** d **do**
 Apply the block-encoding unitary: U if k is odd, and U^\dagger if k is even;
 Apply the next phase operator $\Pi_{\phi_k} = e^{i\phi_k(2\Pi - I)}$;
end
 Measure or project onto the ancilla $|0^a\rangle$ subspace to obtain the transformed signal block implementing $P(A/\alpha)$;

B. Subsuming Amplitude Amplification

QSVT reproduces the phase-synthesized amplification frameworks from Sections IV–VI and VIII–X as polynomial-design instances.

1. **Standard Grover (Section IV).** For odd degree $d = 2r + 1$, the induced polynomial is

$$P_d(x) = \sin(d \arcsin x), \quad (53)$$

so $P_d(\sin \theta) = \sin(d\theta)$ (the trigonometric Chebyshev form). Amplification to constant success requires $d = \Theta(1/\sqrt{p})$.

2. **Fixed-Point Amplification (Section V).** The phase schedules derived in Section V are exact synthesis solutions for bounded minimax sign-like approximation with controlled ripple δ . The fixed-point no-overshoot property is exactly the bounded passband property of the synthesized polynomial.
3. **Oblivious and Fixed-Point Oblivious Amplification (Sections VI and VIII).** In the purified model, the signal block can be written as $\sqrt{p}V$ with $V^\dagger V = I$. In block-encoding language one may either regard this as a $(1, a, 0)$ -encoding of $\sqrt{p}V$ or, equivalently, as a $(1/\sqrt{p}, a, 0)$ -encoding of V . Oblivious amplification may therefore be viewed as QSVT acting on the same singular-value parameter \sqrt{p} , uniformly amplifying the encoded isometry without state-dependent reflections. From this perspective, the Linear Combination of Unitaries (LCU) damping mechanism used in Fixed-Point Oblivious Amplification (Section VIII) can be interpreted as a precursor to the more general polynomial synthesis available in QSVT.
4. **Variable-Time Amplification (Section X).** VTAA's recursive filtering is algebraically equivalent to composing polynomials that approximate a steep Heaviside selector over stage-resolved singular sectors, isolating early-halting amplitude while preserving coherence; the precision requirement yields the $\text{polylog}(T_m)$ overhead.

C. Selected Advanced QSVT Instantiations

Beyond the direct unification statements above, several high-impact algorithmic primitives become immediate once amplitude amplification is interpreted as polynomial synthesis over singular values:

1. **Spectral thresholding and filtering.** By synthesizing bounded approximants to step-like functions, QSVT implements coherent singular-value filtering, enabling selective amplification or suppression of spectral bands without intermediate measurement [8].
2. **Matrix inversion from polynomial reciprocals.** Choosing $P(x) \approx 1/x$ on a gapped interval $[1/\kappa, 1]$ recovers linear-system subroutines in the same algebraic framework, clarifying why condition-number dependence is governed by approximation degree rather than bespoke circuit geometry [1, 8].

3. **Stable amplification inside block-encoding pipelines.** The fixed-point intuition of Sections V and VI lifts to robust singular-value shaping: instead of tuning iteration counts branch-by-branch, one compiles a single bounded polynomial with explicit passband and stopband behavior.

XII. APPLICATIONS OF THE UNIFIED FRAMEWORK

By shifting the focus from geometric state rotations to singular-value polynomial transforms, the amplification framework extends naturally to other quantum computational primitives. Because the structure applies bounded polynomials to matrix singular blocks, it connects amplitude amplification to broader classes of algorithms beyond unstructured search.

Linear Systems. Choosing $P(x) \approx 1/x$ over an interval $[1/\kappa, 1]$ yields quantum linear-system solvers. This recovers the condition-number scaling improvements discussed in the context of the HHL algorithm [1, 16]. Both amplifying marked amplitude and inverting spectral amplitude reduce to synthesizing a bounded polynomial with the requisite approximation profile.

Hamiltonian Simulation. Simulating quantum time evolution, e^{-iHt} , reduces to applying trigonometric polynomials $P(x) = \cos(tx)$ and $P(x) = \sin(tx)$ to a block-encoded Hamiltonian H . If H is given via an (α, a, ϵ) -block encoding, this framework yields the near-optimal scaling $\mathcal{O}(\alpha t + \log(1/\epsilon))$ without relying on Trotterization or truncated Taylor series expansions.

Ground State Preparation (Beyond VQE and QAOA). In the current Noisy Intermediate-Scale Quantum (NISQ) era, ground state preparation and optimization are dominated by variational heuristics such as the Variational Quantum Eigensolver (VQE) and the Quantum Approximate Optimization Algorithm (QAOA). While useful for near-term hardware, these heuristic algorithms lack rigorous asymptotic performance guarantees. In the fault-tolerant setting, QSVT offers a more systematic alternative to these methods. By applying a polynomial approximation of the sign function or the exponential function to a block-encoded Hamiltonian, one can project an initial trial state toward the ground state with provable, near-optimal query complexity.

Amplitude Estimation and IQAE. Quantum Amplitude Estimation (QAE), introduced by Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp in the same foundational work that formalized generalized amplitude amplification [6], is the natural twin primitive for estimating success probabilities rather than only amplifying them. The original construction combines Grover-style iterations with Shor-style Fourier analysis: for the generalized iterate Q with eigenvalues $e^{\pm i2\theta_a}$ (where $a = \sin^2 \theta_a$ is the target success probability), QAE applies controlled powers of Q and the size- M Fourier transform F_M to estimate θ_a , then recovers a . This coherent “probability counting”

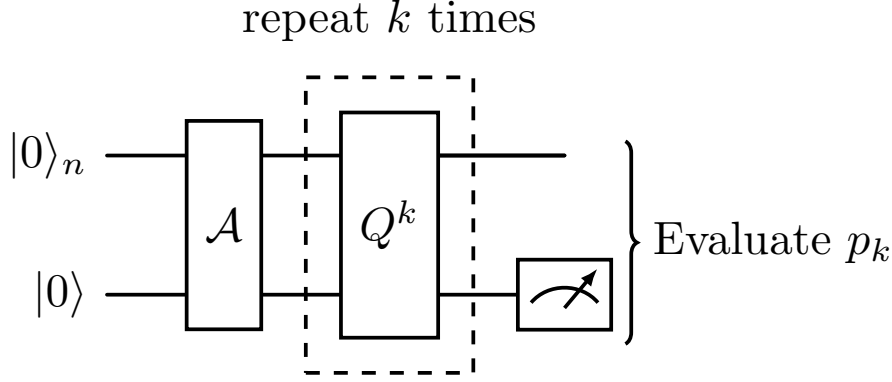


FIG. 13. **Iterative Quantum Amplitude Estimation (IQAE) Refined Architecture.** IQAE eliminates the Quantum Fourier Transform (QFT) by executing a series of circuits using only one auxiliary qubit. Each circuit prepares the state $Q^k \mathcal{A}|0\rangle_{n+1}$ to estimate the success probability $p_k = \sin^2((2k+1)\theta_a)$. This measurement provides the mathematical bridge to the classical interval update through the identity $p_k = \frac{1 - \cos(K\theta_a)}{2}$, where $K = 4k + 2$. By managing the rotation scale K classically, the algorithm refines a confidence interval for θ_a while maintaining unique cosine inversion within a single half-plane.

capability made QAE a central subroutine in quantum Monte Carlo methods for finance and in quantum machine-learning workflows. The main implementation bottleneck is the explicit Quantum Fourier Transform (QFT), which is costly in ancilla and depth.

Algorithm 10: Iterative Quantum Amplitude Estimation (IQAE)

Input: (1) State preparation unitary \mathcal{A} with unknown success probability a
 (2) Grover iterate $Q = -\mathcal{A}S_0\mathcal{A}^{-1}S_\chi$
 (3) Target accuracy ε , failure probability α , and shots per iteration N_{shots}
Output: An interval $[a_l, a_u]$ containing a with confidence at least $1 - \alpha$

```
// Initialization
 $[\theta_l, \theta_u] \leftarrow [0, \pi/2]$ ;
 $k_0 \leftarrow 0, K_0 \leftarrow 2$ ;
// Iterative interval refinement
while  $\theta_u - \theta_l > 2\varepsilon$  do
  Choose the largest  $K_i = 4k_i + 2$  such that
   $[K_i\theta_l, K_i\theta_u] \bmod 2\pi$  lies entirely in one
  half-plane;
  Prepare  $Q^{k_i}\mathcal{A}|0\rangle^{\otimes n}|0\rangle$  and measure the final
  qubit  $N_{\text{shots}}$  times;
  Construct a confidence interval  $[p_i^{\min}, p_i^{\max}]$  for
   $p_i = (1 - \cos(K_i\theta_a))/2$ ;
  Invert the cosine on the known branch and
  update  $[\theta_l, \theta_u]$ ;
end
return  $[\sin^2(\theta_l), \sin^2(\theta_u)]$ ;
```

Iterative Quantum Amplitude Estimation (IQAE) [9] is a useful modern representative of the QPE-free direction. Let

$$\mathcal{A}|0\rangle^{\otimes n}|0\rangle = \sqrt{1-a}|\psi_0\rangle|0\rangle + \sqrt{a}|\psi_1\rangle|1\rangle \quad (54)$$

where $a = \sin^2(\theta_a)$. As in ordinary amplification, the Grover iterate $Q = -\mathcal{A}S_0\mathcal{A}^{-1}S_\chi$ acts only inside the two-dimensional plane generated by the bad and good components:

$$Q^k \mathcal{A}|0\rangle^{\otimes n}|0\rangle = \cos((2k+1)\theta_a)|\psi_0\rangle|0\rangle + \sin((2k+1)\theta_a)|\psi_1\rangle|1\rangle \quad (55)$$

Intuition Box: Iterative Quantum Amplitude Estimation

Objective: Estimate $a = \sin^2(\theta_a)$ to precision ϵ with high confidence.

Approach Insight: IQAE functions as a multi-scale measurement process. Each Grover power k scales the phase to $(2k+1)\theta_a$, increasing the sensitivity of the success probability to the underlying parameter. The algorithm connects these scales by using the confidence interval from a lower k (coarse but unique) to disambiguate the periodic signal of a higher k (precise but aliased). This iterative "handover" allows for exponential precision gain while maintaining a valid estimate.

Advantages: Achieves the optimal $\mathcal{O}(\epsilon^{-1})$ query complexity with only 1 auxiliary qubit. By replacing the QFT with classical interval arithmetic, it reduces circuit depth by $\approx 10\times$ compared to QPE-based methods, enabling high-precision results on NISQ hardware.

Thus measuring the final qubit after k Grover powers estimates

$$p_k = \sin^2((2k+1)\theta_a) = \frac{1 - \cos(K\theta_a)}{2} \quad (56)$$

where $K = 4k + 2$. IQAE maintains a classical confidence interval $[\theta_l, \theta_u]$ for θ_a , chooses the largest feasible

K for which $[K\theta_l, K\theta_u] \bmod 2\pi$ stays inside one cosine half-plane, samples p_k , and then safely inverts the cosine on that known branch. It returns an interval $[a_l, a_u]$ with $a_u - a_l \leq 2\varepsilon$ and failure probability at most α , using

$$N_{\text{oracle}} = \mathcal{O}\left(\frac{1}{\varepsilon} \log\left(\frac{1}{\alpha} \log \frac{1}{\varepsilon}\right)\right) \quad (57)$$

applications of Q . This preserves the near-quadratic Monte Carlo improvement without the QPE/QFT register, making amplitude estimation more favorable for NISQ and early fault-tolerant architectures [8, 9, 20, 24].

Quantum Machine Learning. Matrix inversion and projection are central bottlenecks in classical machine learning. Applying appropriate thresholding or inverse polynomials to a data matrix block-encoding provides a native mechanism for quantum principal component analysis (PCA) and regression.

XIII. LIMITS AND OPTIMALITY

Every algorithm described so far is bound to use $\mathcal{O}(\sqrt{N/M})$ oracle queries. It turns out that when restricted to oracles that only distinguish marked inputs via phase flip, any algorithm is bound to use at least these many queries.

Theorem 5 (BBBV Lower Bound [3]). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a black-box Boolean oracle and $M = |f^{-1}(1)| \geq 1$ marked items. Any quantum query algorithm that finds a marked input with success probability at least $2/3$ (equivalently, any fixed constant bounded away from 0) must use*

$$\Omega\left(\sqrt{N/M}\right)$$

oracle queries in the worst case.

Proof. Let $S \subseteq [N]$ denote the marked set, with $|S| = M$, and define the phase oracle

$$O_S |i\rangle |z\rangle = (-1)^{[i \in S]} |i\rangle |z\rangle,$$

where the second register $|z\rangle$ contains arbitrary workspace. A T -query algorithm has the form

$$U_T O_S U_{T-1} O_S \cdots U_1 O_S U_0 |0\rangle,$$

where the U_t are oracle-independent unitaries.

We first reduce search to decision. Suppose an algorithm \mathcal{A} finds a marked item with probability at least $2/3$ using T queries whenever $|S| = M$. Construct a decision algorithm \mathcal{D} with $T+1$ queries as follows: run \mathcal{A} , measure its output candidate i , query the oracle once on i , and accept iff $i \in S$. Then:

1. if $|S| = M$, the acceptance probability of \mathcal{D} is at least $2/3$;
2. if $S = \emptyset$, the acceptance probability of \mathcal{D} is 0.

Therefore it suffices to prove that any algorithm distinguishing the empty oracle (a database with no marked items) from every oracle with exactly M marked items and constant bias (probability of success at least $2/3$) must use $\Omega(\sqrt{N/M})$ queries.

Let \mathcal{D} be such a decision algorithm using T queries. For each marked set S , define the final state

$$|\psi_T^S\rangle = U_T O_S U_{T-1} O_S \cdots U_1 O_S U_0 |0\rangle,$$

and let

$$|\psi_T^\emptyset\rangle = U_T U_{T-1} \cdots U_1 U_0 |0\rangle$$

be the final state under the empty oracle $O_\emptyset = I$.

To compare these states, introduce the intermediate empty-oracle states

$$|\phi_t\rangle = U_t U_{t-1} \cdots U_0 |0\rangle, \quad t = 0, 1, \dots, T-1$$

Write the query register decomposition of $|\phi_t\rangle$ as

$$|\phi_t\rangle = \sum_{i=1}^N |i\rangle |\eta_{t,i}\rangle$$

where the workspace vectors $|\eta_{t,i}\rangle$ need not be normalized. Define

$$q_t(i) = \|\eta_{t,i}\|^2, \quad \sum_{i=1}^N q_t(i) = 1.$$

Thus $q_t(i)$ is exactly the probability that the t -th query would address index i if all oracle calls were replaced by the identity.

The effect of changing the empty oracle into O_S at this step is

$$(O_S - I) |\phi_t\rangle = -2 \sum_{i \in S} |i\rangle |\eta_{t,i}\rangle$$

so

$$\|(O_S - I) |\phi_t\rangle\| = 2 \sqrt{\sum_{i \in S} q_t(i)}$$

To make the hybrid step explicit, define for $t = 0, 1, \dots, T$ the final hybrid states

$$|\chi_t\rangle = U_T O_S U_{T-1} O_S \cdots U_{t+1} O_S U_t U_{t-1} \cdots U_0 |0\rangle$$

so that $|\chi_0\rangle = |\psi_T^S\rangle$ and $|\chi_T\rangle = |\psi_T^\emptyset\rangle$. Then, for each $t = 0, 1, \dots, T-1$,

$$|\chi_t\rangle - |\chi_{t+1}\rangle = U_T O_S \cdots U_{t+2} O_S U_{t+1} (O_S - I) |\phi_t\rangle$$

Because left multiplication by a unitary preserves norm,

$$\|\chi_t\rangle - \chi_{t+1}\rangle\| = \|(O_S - I) |\phi_t\rangle\|$$

Applying the triangle inequality to the telescoping sum

$$|\psi_T^S\rangle - |\psi_T^\emptyset\rangle = \sum_{t=0}^{T-1} (|\chi_t\rangle - |\chi_{t+1}\rangle)$$

therefore yields

$$\| |\psi_T^S\rangle - |\psi_T^\emptyset\rangle \| \leq \sum_{t=0}^{T-1} \|(O_S - I) |\phi_t\rangle\| = 2 \sum_{t=0}^{T-1} \sqrt{\sum_{i \in S} q_t(i)}.$$

Average this bound over all M -element subsets $S \subseteq [N]$. For each fixed t ,

$$\mathbb{E}_S \left[\sum_{i \in S} q_t(i) \right] = \sum_{i=1}^N q_t(i) \Pr[i \in S] = \frac{M}{N} \sum_{i=1}^N q_t(i) = \frac{M}{N}$$

Because $x \mapsto \sqrt{x}$ is concave, Jensen's inequality gives

$$\mathbb{E}_S \left[\sqrt{\sum_{i \in S} q_t(i)} \right] \leq \sqrt{\mathbb{E}_S \left[\sum_{i \in S} q_t(i) \right]} = \sqrt{\frac{M}{N}}$$

Hence

$$\mathbb{E}_S [\| |\psi_T^S\rangle - |\psi_T^\emptyset\rangle \|] \leq 2T \sqrt{\frac{M}{N}}$$

Therefore there exists at least one marked set S with $|S| = M$ such that

$$\| |\psi_T^S\rangle - |\psi_T^\emptyset\rangle \| \leq 2T \sqrt{\frac{M}{N}}$$

Next we show that $|\psi_T^S\rangle$ and $|\psi_T^\emptyset\rangle$ cannot be arbitrarily close. Let Π be the POVM element corresponding to the algorithm declaring “marked.” By assumption,

$$\langle \psi_T^S | \Pi | \psi_T^S \rangle \geq \frac{2}{3}, \quad \langle \psi_T^\emptyset | \Pi | \psi_T^\emptyset \rangle \leq \frac{1}{3}$$

Define

$$\rho_S = |\psi_T^S\rangle \langle \psi_T^S|, \quad \rho_\emptyset = |\psi_T^\emptyset\rangle \langle \psi_T^\emptyset|$$

For pure states, the bias achievable by any measurement is upper-bounded by their trace distance, and for pure states that trace distance is at most the Euclidean distance:

$$\begin{aligned} |\langle \psi_T^S | \Pi | \psi_T^S \rangle - \langle \psi_T^\emptyset | \Pi | \psi_T^\emptyset \rangle| &\leq \frac{1}{2} \|\rho_S - \rho_\emptyset\|_1 \\ &\leq \| |\psi_T^S\rangle - |\psi_T^\emptyset\rangle \|. \end{aligned} \quad (58)$$

Since the left-hand side is at least $1/3$, we conclude that

$$\frac{1}{3} \leq \| |\psi_T^S\rangle - |\psi_T^\emptyset\rangle \| \leq 2T \sqrt{\frac{M}{N}}$$

Thus

$$T \geq \frac{1}{6} \sqrt{\frac{N}{M}}$$

Recalling the initial reduction from search to decision, any search algorithm must satisfy

$$T + 1 = \Omega \left(\sqrt{\frac{N}{M}} \right),$$

and therefore $T = \Omega(\sqrt{N/M})$ as claimed. \square

XIV. VARIATIONAL TEMPLATE FOR AMPLIFICATION

Formulating amplitude amplification variants within a variational template is a highly effective strategy, particularly for mapping these algorithms onto Noisy Intermediate-Scale Quantum (NISQ) architectures. By treating the amplification steps as a parameterized quantum circuit (PQC), we can unify deterministic algorithms and heuristic hybrid-quantum algorithms under a single mathematical framework.

In standard and fixed-point amplitude amplification, the algorithm proceeds via alternating phase matching operations. We can generalize this into an ansatz parameterized by two vectors of phase angles, $\vec{\alpha} = (\alpha_1, \dots, \alpha_d)$ and $\vec{\beta} = (\beta_1, \dots, \beta_d)$, where d is the circuit depth.

The parameterized trial state is generated as:

$$|\psi(\vec{\alpha}, \vec{\beta})\rangle = \left[\prod_{k=1}^d \mathcal{A} S_0(\alpha_k) \mathcal{A}^\dagger S_{\text{Good}}(\beta_k) \right] \mathcal{A} |0\rangle^{\otimes n}$$

where:

- \mathcal{A} is the state preparation unitary.
- $S_{\text{Good}}(\beta_k) = I - (1 - e^{i\beta_k}) \Pi_{\text{Good}}$ applies a phase shift to the target subspace.
- $S_0(\alpha_k) = I - (1 - e^{i\alpha_k}) |0\rangle \langle 0|^{\otimes n}$ applies a phase shift to the initial zero state (equivalent to the reflection about the $| \text{All} \rangle$ state).

The goal of the classical optimization loop is to maximize the probability of measuring a target state. We define the cost function $C(\vec{\alpha}, \vec{\beta})$ as the complement of the target expectation value:

$$C(\vec{\alpha}, \vec{\beta}) = 1 - \langle \psi(\vec{\alpha}, \vec{\beta}) | \Pi_{\text{Good}} | \psi(\vec{\alpha}, \vec{\beta}) \rangle$$

Table VI summarizes the major amplification variants in terms of their variational cost generators and mixer/reflection evolutions within a unified alternating-operator perspective.

Among these variants, the QAOA-inspired formulation is particularly notable because it replaces Grover's nonlocal diffusion reflection by a physically local transverse-field mixer Hamiltonian, thereby connecting amplitude amplification with variational alternating-operator algorithms.

TABLE VI. Variational interpretation of amplitude amplification algorithms.

Algorithm	Cost Generator	Mixer / Reflection Evolution
Grover / Standard AA	$-\Pi_{\text{Good}}$	$R_{\text{All}} = 2 \text{All}\rangle\langle\text{All} - I$
Fixed-Point AA (FPAA)	$-\Pi_{\text{Good}}$	Generalized phase reflections $S_{\text{All}}(\alpha)$
Oblivious AA (OAA)	$-\Pi_{\text{Good}}$	$R_{\text{clean}} = 2C - I$, with $C = 0\rangle\langle 0 ^{\otimes l} \otimes I$
Controlled Quantum AA (CQAA)	$-\Pi_{\phi} = - \phi\rangle\langle\phi $	Controlled walk evolution W
Fixed-Point Oblivious Quantum AA (FOQA)	$-\Pi_{\text{Good}}$	LCU rotation sequence $V_n = e^{-i\alpha_n Y/2}$
Distributed QAA (DQAA)	$-\Pi_{\text{Good},k}$	Local Grover / FPAA evolutions
QAOA-Inspired Search	$C = - u\rangle\langle u $	$e^{-i\beta B}$, where $B = \sum_{j=1}^n X_j$
QSVT / QSP Amplification	Block-encoded signal operator	Phase-controlled $\text{SU}(2)$ polynomial evolutions

A. QAOA-Inspired Search and the Transverse Field

An important 2017 development re-examined Grover search from the perspective of hardware-efficient alternating-operator circuits. Jiang, Rieffel, and Wang observed that Grover’s diffusion operator is not only conceptually geometric but also hardware-intensive: implementing the reflection about $|\text{All}\rangle$ requires a multi-controlled operation whose standard circuit decomposition uses $\Theta(n)$ two-qubit gates [18]. This cost becomes a practical bottleneck on near-term devices even though the query complexity is optimal.

Their alternative replaces Grover’s diffusion reflection by the transverse-field mixer

$$B = \sum_{j=1}^n X_j$$

where each X_j is a single-qubit Pauli- X operator. The problem Hamiltonian is taken to be

$$C_u = -|u\rangle\langle u|$$

for the unknown marked string u . After relabeling the marked item to $u = 0$, we write

$$C \equiv C_0 = -|0\rangle\langle 0|$$

and the initial state remains the uniform product state

$$|+\rangle^{\otimes n}$$

and the dynamics generated by both B and C preserve the permutation-symmetric subspace of dimension $n+1$ [18].

The algorithm applies oracle and mixer unitaries periodically through the single-period building block

$$W(\gamma) = e^{-i\pi B/n} e^{i\gamma C} e^{-i\pi B/n} e^{-i\gamma C}$$

with $\gamma \in (0, \pi]$ a tunable phase parameter. Jiang et al. show that the choice $\gamma = \pi$ yields the best asymptotic success behavior in their analysis. Although this construction does not use Grover’s Euclidean reflection picture directly, it is still governed by a low-dimensional effective transition: using spin-coherent-state and phase-space methods, the paper shows that repeated application of $W(\pi)$ induces a nearly closed transition between states

having large overlap with the initial uniform state and the marked computational basis state [18].

From the perspective of this survey, the significance of this result is twofold. First, it shows that the quadratic quantum speedup of unstructured search does not depend on implementing Grover’s exact diffusion operator gate by gate. Second, it provides a natural conceptual bridge from circuit-model amplitude amplification to later Hamiltonian- and QAOA-inspired alternating-operator frameworks. The resulting average query complexity is

$$T(n) \simeq \frac{\pi}{2\sqrt{2}} 2^{n/2} = \frac{\pi}{2\sqrt{2}} \sqrt{N}$$

which retains the optimal $\Theta(\sqrt{N})$ scaling while differing from Grover’s absolute optimum only by a small constant factor. In exchange, the mixer layer is built entirely from single-qubit transverse-field rotations, which is substantially more favorable for hardware with costly entangling gates.

XV. CONCLUSION AND OUTLOOK

The progression from geometric reflections to singular value transformations provides a mathematically cohesive framework for amplitude amplification. Algorithm design in this regime shifts from heuristic circuit construction to the identification of bounded polynomials.

While the asymptotic theory within the oracle model is mature, the physical realization of these frameworks presents several active research frontiers:

- **Fault-Tolerant Compilation:** The phase-modulated reflection sequences require deep multi-controlled operations that are costly to compile into fault-tolerant discrete gate sets (e.g., Clifford+T).
- **Distributed Architectures:** As outlined in Section IX, executing massive block-encodings across modular, distributed NISQ networks remains a critical, unsolved hardware challenge.
- **Algorithmic Automation:** Because the framework reduces algorithm design to polynomial synthesis, optimal phase angles can be discovered programmatically via classical optimization algorithms (such as the Remez algorithm), enabling the automated generation of quantum circuits for specific approximation targets.

- **Efficient Data Loading:** Constructing the requisite block-encodings with realistic quantum RAM (QRAM) and memory assumptions remains a practical bottleneck.

Amplitude amplification has evolved from an isolated geometric trick into a standardized application of polynomial operator calculus, shifting the focus of future research from query complexity to hardware implementation.

-
- [1] Andris Ambainis. Variable time amplitude amplification and a faster quantum algorithm for solving systems of linear equations. *arXiv*, 2010. doi:10.1010.4458.
 - [2] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, 2001.
 - [3] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM Journal on Computing*, 26(5):1510–1523, 1997.
 - [4] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma. Exponential improvement in precision for simulating sparse hamiltonians. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, pages 283–292, 2014.
 - [5] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma. Simulating hamiltonian dynamics with a truncated taylor series. *Physical Review Letters*, 114: 090502, 2014.
 - [6] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.
 - [7] Catalin Dohotaru and Peter Høyer. Controlled quantum amplification. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *LIPIcs*, pages 18:1–18:13, 2017.
 - [8] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 193–204, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367059. doi:10.1145/3313276.3316366.
 - [9] D. Grinko, J. Gacon, C. Zoufal, and S. Woerner. Iterative quantum amplitude estimation. *npj Quantum Information*, 7:52, 2021. doi:10.1038/s41534-021-00379-1.
 - [10] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 212–219, 1996.
 - [11] Lov K. Grover. Fixed-point quantum search. *Phys. Rev. Lett.*, 95:150501, Oct 2005. doi: 10.1103/PhysRevLett.95.150501.
 - [12] S. Gudder. Mathematical theory of duality quantum computers. *Quantum Information Processing*, 6:37–48, 2007.
 - [13] L. Gui-Lu. General quantum interference principle and duality computer. *Communications in Theoretical Physics*, 45:825, 2006.
 - [14] L. Gui-Lu and L. Yang. Duality computing in quantum computers. *Communications in Theoretical Physics*, 50: 1303, 2008.
 - [15] L. Gyongyosi and S. Imre. Scalable distributed gate-model quantum computers. *Scientific Reports*, 11:5172, 2021.
 - [16] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103:150502, 2009.
 - [17] Xiao Hua and Daowen Qiu. Distributed quantum amplitude amplification. *arXiv*, 2025. doi:2510.16498.
 - [18] Shanchi Jiang, Eleanor G. Rieffel, and Zhiqiang Wang. Near-optimal quantum circuit for Grover’s unstructured search using a transverse field. *Physical Review A*, 95: 062317, 2017.
 - [19] Mithilesh Kumar, Yusuf Tahir, and Varun Daiya. Quantum Amplification Library, April 2026. URL <https://github.com/QuantumAmplification/ampamp>.
 - [20] Guang Hao Low and Isaac L. Chuang. Optimal Hamiltonian simulation by quantum signal processing. *Physical Review Letters*, 118:010501, 2017.
 - [21] Guang Hao Low and Isaac L. Chuang. Hamiltonian simulation by qubitization. *Quantum*, 3:163, 2019. arXiv:1610.06546.
 - [22] A. Mizel. Critically damped quantum search. *Physical Review Letters*, 102:150501, 2009.
 - [23] J. Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
 - [24] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tsuyoshi Tanaka, Hiroshi Onodera, and Naoki Yamamoto. Amplitude estimation without phase estimation. *Quantum Information Processing*, 19:75, 2020.
 - [25] Dieter van Melkebeek and Brandon Young. Lecture 11: Oblivious amplitude amplification. March 2022.
 - [26] Bing Yan, Shijie Wei, Hong Jiang, Huaiwen Wang, Qiuchi Duan, Zhao Ma, and Gui-Lu Long. Fixed-point oblivious quantum amplitude-amplification algorithm. *Scientific Reports*, 12:14339, 2022.
 - [27] Ted J. Yoder, Guang Hao Low, and Isaac L. Chuang. Fixed-point quantum search with an optimal number of queries. *Physical Review Letters*, 113:210501, 2014.
 - [28] C. Zalka. Grover’s quantum searching algorithm is optimal. *Physical Review A*, 60:2746–2751, 1997.
 - [29] Xing Zhou, Wang Tao, Ke Li, and Shichao Zheng. Distributed exact quantum amplitude amplification algorithm for arbitrary quantum states. *arXiv*, 2026. doi: 2601.09128.