

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
MARCH 2023	JOURNAL ARTICLE (POST PRINT)	START DATE JUNE 2022	END DATE JULY 2022
4. TITLE AND SUBTITLE GAUSSIAN AMPLITUDE AMPLIFICATION FOR QUANTUM PATHFINDING			
5a. CONTRACT NUMBER		5b. GRANT NUMBER N/A	
5d. PROJECT NUMBER PROGRAM: CQC		5f. WORK UNIT NUMBER	
6. AUTHOR(S) Daniel Koch, Massimiliano Cutugno, Samuel Karlson, Saahil Patel, Laura Wessing, and Paul M. Alsing			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITQ 525 Brooks Road Rome NY 13441-4505			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITQ 525 Brooks Road Rome NY 13441-4505			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. PA# AFRL-2022-2385 Date Cleared: May 2022			11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-RI-RS-TP-2023-007
13. SUPPLEMENTARY NOTES © 2022 The Authors. This work is copyrighted. One or more of the authors is a U.S. Government employee working within the scope of their Government job; therefore, the U.S. Government is joint owner of the work and has the right to copy, distribute, and use the work. All other rights are reserved by the copyright owner. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).			
14. ABSTRACT We study an oracle operation, along with its circuit design, which combined with the Grover diffusion operator boosts the probability of finding the minimum or maximum solutions on a weighted directed graph. We focus on the geometry of sequentially connected bipartite graphs, which naturally gives rise to solution spaces describable by Gaussian distributions. We then demonstrate how an oracle that encodes these distributions can be used to solve for the optimal path via amplitude amplification. And finally, we explore the degree to which this algorithm is capable of solving cases that are generated using randomized weights, as well as a theoretical application for solving the Traveling Salesman problem.			
15. SUBJECT TERMS Quantum computing; quantum algorithms; amplitude amplification			
16. SECURITY CLASSIFICATION OF:		17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	SAR
19a. NAME OF RESPONSIBLE PERSON PAUL M. ALSING			19b. PHONE NUMBER (Include area code) N/A

Article

Gaussian Amplitude Amplification for Quantum Pathfinding

Daniel Koch ^{1,*} , Massimiliano Cutugno ¹, Samuel Karlson ², Saahil Patel ¹, Laura Wessing ¹ and Paul M. Alsing ¹

¹ Air Force Research Lab, Information Directorate, Rome, NY 13441, USA; massimiliano.cutugno@us.af.mil (M.C.); saahil.patel@us.af.mil (S.P.); laura.wessing.1@us.af.mil (L.W.); paul.alsing@us.af.mil (P.M.A.)

² Air Force Academy, Colorado Springs, CO 80840, USA; c22samuel.karlson@afacademy.af.edu

* Correspondence: daniel.koch.13@us.af.mil

Abstract: We study an oracle operation, along with its circuit design, which combined with the Grover diffusion operator boosts the probability of finding the minimum or maximum solutions on a weighted directed graph. We focus on the geometry of sequentially connected bipartite graphs, which naturally gives rise to solution spaces describable by Gaussian distributions. We then demonstrate how an oracle that encodes these distributions can be used to solve for the optimal path via amplitude amplification. And finally, we explore the degree to which this algorithm is capable of solving cases that are generated using randomized weights, as well as a theoretical application for solving the Traveling Salesman problem.

Keywords: quantum computing; quantum algorithms; amplitude amplification



Citation: Koch, D.; Cutugno, M.; Karlson, S.; Patel, S.; Wessing, L.; Alsing, P.M. Gaussian Amplitude Amplification for Quantum Pathfinding. *Entropy* **2022**, *24*, 963. <https://doi.org/10.3390/e24070963>

Academic Editors: GuiLu Long, Shijie Wei and Heng Fan

Received: 1 June 2022

Accepted: 2 July 2022

Published: 11 July 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The use of quantum computers for tackling difficult problems is an exciting promise, but not one without its own set of challenges. Qubits allow for incredible parallelism in computations via superposition states, but reliably pulling out a single answer via measurements is often difficult. In 1996, Grover demonstrated one of the first mechanisms overcoming this weakness [1], later shown to be optimal [2,3], and has since been refined into a broader technique in quantum algorithms known as ‘amplitude amplification’ [4–10]. In this study, we seek to extend the capabilities of amplitude amplification as a means of pathfinding on a directed graph with weighted edges.

The success of Grover’s algorithm can be boiled down to two primary components: the oracle operation U_G and diffusion operation U_S . While U_S is typically considered a straightforward mathematical operation—achieving a reflection about the average amplitude—critics of Grover’s algorithm often point to U_G as problematic [11–14]. Nielsen and Chuang elegantly describe the dilemma of implementing U_G as differentiating between an operation in which *knows* the desired marked state, versus a true blackbox oracle which can *recognize* the answer [15]. Only an oracle of the latter case can truly be considered a speedup for quantum, otherwise, the solution to the unstructured search problem is already encoded into U_G , defeating the purpose of using a quantum computer in the first place. We note this specific issue with Grover’s algorithm because it is exactly the problem we aim to address in this study, specifically for the gate-based model of quantum computing. In this study, we demonstrate an alternative to the standard Grover oracle, which we refer to as a ‘cost oracle’ U_P , capable of solving weighted directed graph problems.

Beyond the specific geometry used to motivate U_P and build its corresponding quantum circuit, much of this study is aimed at formulating a deeper understanding of amplitude amplification. The idea of using an oracle that applies phases different from the standard U_G was first investigated by Long and Hoyer [16–18] and later others [19–21], who showed the degree to which a phase other than π on the marked state(s) could still be used for probability boosting. Here, we study a U_G replacement which affects *all* states with

unique phases, not just a single marked state. Consequently, the effect of U_P is analogous to a cost function, whereby U_P acting on any state results in a phase proportional to that state's representative weighted path. The advantage of quantum is to utilize superposition, evaluating all costs simultaneously, and ultimately boosting the probability of measuring the solution to the optimization problem. Using U_P results in an amplitude amplification process that is more complex than standard Grover's, but still achieves high probabilities under ideal conditions. And most importantly, we demonstrate the degree to which probability boosting is possible under randomized conditions which one would expect from realistic optimization problems [22–24].

After demonstrating results for the success of U_P , the final topic of this study is a theoretical application of cost oracles for solving the Traveling Salesman Problem (TSP) [25], or all-to-all connected directed graphs. Notable strategies thus far for a quantum solution to the TSP are based on phase estimation [26], backtracking [27], and adiabatic quantum computing [28–31]. Here we approach the problem from an amplitude amplification perspective, continuing an idea that goes back over a decade [32]. However, to realize the appropriate quantum states for this application of U_P , we must look beyond binary superposition states provided by qubits, in favor of a mixed qudit quantum computer which more naturally suits the problem. Although still in their technological infancy compared to qubits, the realization of qudit technologies [33–36], qudit-based universal computation [37,38], their fundamental quantum circuits [39–44], and algorithm applications [45,46] have all seen significant advancements over the last decade, making now an exciting time to consider their use for future algorithms.

Layout

Section 2 begins with an alternative oracle to Grover's U_G , which we use to introduce fundamental features of amplitude amplification and oracle operations. The progression of this study then revolves around a specific directed graph problem, where the underlying characteristics of each graph's solution space are describable by the Central Limit Theorem [47] and the Law of Large Numbers [48], resulting in solution space distributions which resemble a Gaussian function [49]. Section 3 covers specifics of this weighted directed graph problem, a graphical representation of all possible paths, and a proposed classical solving speed based on arguments of information access. Sections 4 and 5 show how each graph can be represented as a pathfinding problem, translated into quantum states, and ultimately solved using a modified Grover's algorithm. In Section 6 we present results from simulated perfect Gaussian distributions, providing insight into fundamental properties of optimization problems that are viable for amplitude amplification. In Section 7, we explore the viability of using a cost oracle to solve optimization problems involving randomness. Section 8, explores a theoretical application of U_P for solving the Traveling Salesman Problem, and, finally, Section 9, concludes with a summary of our findings and discussions of future research.

2. Gate-Based Grover's

Shown below in Equation (1) is U_s , known as the diffusion operator, which is the driving force behind amplitude amplification. The power of this operation lies in its ability to reflect every state in the quantum system about the average amplitude *without* computing the average itself.

$$U_s = 2|s\rangle\langle s| - \mathbb{I} \quad (1)$$

In order to make use of this powerful geometric operation, we must pair it with an oracle operator in order to solve interesting problems. For clarity, in order for an operator to qualify as an oracle, we require that the probability of measuring each state in the system must be the same before and after applying the oracle. This requirement excludes any and all operations which cause interference, leaving only one type of viable operator: phase

gates. Thus, it is the aim of this study to investigate viable oracle operations which encode the information of problems into phases and solve them using amplitude amplification.

2.1. Optimal Amplitude Amplification

It is important to understand why the standard Grover oracle U_G , given in Equation (2) and used in Algorithm 1, is optimal in the manner in which it boosts the probability of the marked state to nearly 1 [2,3]. Geometrically, this is because the entire amplitude amplification process takes place along the real axis in amplitude space (i.e., at no point does any state acquire an imaginary amplitude component). Consequently, the marked state, origin, mean amplitude point, and non-marked states are all linearly aligned, which ensures that the marked state receives the maximal reflection of the average (mean point) at each step. This property holds true for not only the real axis, but any linear axis that runs through the origin, so long as the marked and non-marked states differ in phase by π as a result of the oracle operation.

$$U_G|\Psi\rangle = \begin{cases} \text{marked} & , e^{i\pi}|\Psi_i\rangle \\ \text{non-marked} & , |\Psi_i\rangle \end{cases} \quad (2)$$

Algorithm 1 Grover's Search Algorithm

- 1: Initialize Qubits: $|\Psi\rangle = |0\rangle^{\otimes N}$
 - 2: Prepare Equal Superposition: $H^{\otimes N}|\Psi\rangle = |s\rangle$
 - 3: **for** $k \approx \frac{\pi}{4}\sqrt{2^N}$ **do**
 - 4: Apply $U_G|\Psi\rangle$ (Oracle)
 - 5: Apply $U_s|\Psi\rangle$ (Diffusion)
 - 6: Measure
-

We note the optimality of U_G because it is directly tied to the nature of the problem which it solves, namely an unstructured search [1]. The power of amplitude amplification using U_G goes hand-in-hand with the rigidity of the operator. Thus, if we want to expand the capabilities of amplitude amplification on gate-based quantum computers to more interesting problems, we must explore more flexible oracle operators, and consequently expect probability boosting that is less than optimal.

2.2. Alternate Two-Marked Oracle

Here we present an example analogous to Grover's search algorithm with two marked states, but with an oracle operator of our own design. The purpose of this exercise is to illustrate several key ideas that will be prominent throughout the remainder of this study. Firstly, the general idea of a multi-phased oracle operation [50,51], or 'non-boolean' oracles [52]. Secondly, to demonstrate that the success of amplitude amplification can be directly traced back to the inherent mathematical properties of an oracle. And finally, to introduce terminology and features of amplitude amplification on discrete systems which will apply to later oracles. All of the following results were verified using IBM's Qiskit simulator as well as our own python-based simulator.

$$U'_{G2}|\Psi\rangle = \begin{cases} |0\rangle^{\otimes N} & , |0\rangle^{\otimes N} \\ |1\rangle^{\otimes N} & , e^{i\pi}|1\rangle^{\otimes N} \\ |\Psi_i\rangle \in |G_\theta\rangle & , e^{i\theta}|\Psi_i\rangle \\ |\Psi_i\rangle \in |G_{-\theta}\rangle & , e^{-i\theta}|\Psi_i\rangle \end{cases} \quad (3)$$

where

$$\begin{aligned} |G_\theta\rangle &\equiv \sum_i |\Psi_i\rangle = |0\rangle \otimes |\psi\rangle, \text{ for } |\psi\rangle \neq |0\rangle^{\otimes N-1} \\ |G_{-\theta}\rangle &\equiv \sum_i |\Psi_i\rangle = |1\rangle \otimes |\psi\rangle, \text{ for } |\psi\rangle \neq |1\rangle^{\otimes N-1} \end{aligned} \quad (4)$$

We begin with the mathematical definition of our oracle function in Equation (3) above, which we shall refer to as U'_{G2} , as well as its quantum circuit composition in Figure 1. Contrary to Equation (2), we now have an oracle operation with four distinct outcomes depending on which state $|\Psi_i\rangle$ it is acting on. Additionally, U'_{G2} has a free parameter θ , controlled by the experimenter, which dictates how the states $|0\rangle^{\otimes N}$ and $|1\rangle^{\otimes N}$ boost in probability. Altogether, the effect of U'_{G2} can be seen in Figure 2, which displays the position of each state in amplitude space (the complex plane) after the first application: $U'_{G2}|s\rangle$.




Figure 1. Quantum circuit for implementing U'_{G2} . Boxes with θ and π are phase gates, both single and controlled. For the controlled operations, black dots indicate a $|1\rangle$ control state, and similarly white dots for $|0\rangle$.




Figure 2. An illustration of $U'_{G2}|s\rangle$. A unit circle of radius $1/\sqrt{2^N}$ is shown by the blue-dashed line, along with the point of average amplitude with a red 'X'. The parameter θ controls the phase acquired by the cluster of states $|G_\theta\rangle$ and $|G_{-\theta}\rangle$, which in turn dictates the location of the mean point along the real axis.

Before revealing how this alternate two-marked oracle performs at amplitude amplification, note the red 'X' located along the real axis of Figure 2. This 'X' marks the mean point, or average amplitude, where every state in the system will be reflected about after the first diffusion operator U_s . Because $2^N - 2$ states are evenly distributed between $|G_\theta\rangle$ and $|G_{-\theta}\rangle$, this initial mean point can be made to lie anywhere along the real axis between $(-1/\sqrt{2^N}, 1/\sqrt{2^N})$ as θ ranges from 0 to π . Shown in Figure 3 below is the relation between θ and the resulting probability boosts for $|0\rangle^{\otimes N}$ and $|1\rangle^{\otimes N}$.

We define the metric P_M , shown as the y-axis in Figure 3, to be the peak probability achievable through amplitude amplification as defined in Algorithm 1 for a given state. Here we track P_M for the states $|0\rangle^{\otimes N}$ and $|1\rangle^{\otimes N}$ as a function of θ , for the case of $N = 20$. Firstly, note the two extremes of θ : 0 and π , for which the resulting amplitude amplification processes are *exactly* equal to standard Grover's for $|1\rangle^{\otimes N}$ and $|0\rangle^{\otimes N}$ respectively. This

is in agreement with the geometric picture of U'_{G2} outlined in Figure 2, whereby all of the states of $|G_\theta\rangle$ and $|G_{-\theta}\rangle$ receive phases of 0 or π , isolating a single state to be π phase different from the remaining $2^N - 1$ states.




Figure 3. A plot of θ vs. peak probability P_M for the states $|0\rangle^{\otimes N}$ (orange-dashed) and $|1\rangle^{\otimes N}$ (blue-solid). Approximate forms for the two plots are given in Equations (5) and (6).

While U'_{G2} is able to reproduce U_G at the θ bounds, it is the intermediate θ values which are more revealing about the capabilities of amplitude amplification. For sufficiently large N , the mean point produced from U'_{G2} is dominated by the states making up $|G_\theta\rangle$ and $|G_{-\theta}\rangle$, approximately equal to $\approx 1/\sqrt{2^N} \cdot \cos(\theta)$ (the real axis). We note this $\cos(\theta)$ dependance because it also describes the two P_M plots shown in Figure 3, given by Equations (5) and (6) below.

$$P_M(|1\rangle^{\otimes N}) \approx \frac{1}{2}(\cos(\theta) + 1) \quad (5)$$

$$P_M(|0\rangle^{\otimes N}) \approx \frac{1}{2}(\cos(\theta - \pi) + 1) \quad (6)$$

The emphasis here is that we have a one-to-one correlation between a property of U'_{G2} , specifically θ , and the resulting peak probabilities P_M achievable through amplitude amplification. But more accurately, θ is just a parameter for controlling the mean amplitude point produced by U'_{G2} , which is the more fundamental indicator of successful amplitude amplification. This is evidenced by the $\cos(\theta)$ relation found in both P_M plots here, as well as properties of oracle operators to come in this study, which can similarly be directly linked to the initial mean points they produce.

3. Pathfinding Geometry

While the U'_{G2} oracle is useful for gaining insight into non-boolean amplitude amplification processes, ultimately it does not correspond to a meaningful problem we would ideally look to a quantum computer to solve. In particular, we want an oracle operation that boosts a quantum state unknown to the experimenter beforehand, yielding the answer to some unsolved problem. To this end, we now introduce one such optimization problem which can be encoded as an oracle and ultimately solved through amplitude amplification.

3.1. Graph Structure

Shown in Figure 4 is the general structure of the problem which will serve as the first primary focus for this study: a series of sequentially connected bipartite graphs with weighted edges, for which we are interested in finding the path of least or greatest resistance through the geometry. More formally, we seek the solution to a weighted directed graph optimization problem. Each geometry can be specified by two variables, N and L , which represent the number of vertices per column and the total number of columns, respectively. Throughout this study, we refer to vertices as ‘nodes’, and each complete set of nodes in a

vertical column as a ‘layer’. For example, Figure 4’s geometry represents a 4-layer system ($L = 4$), with 3 nodes per layer ($N = 3$).

Given the geometric structure shown above, we now assign a complete set of weights ω_i , one for each of the total $N^2 \cdot (L - 1)$ edges throughout the geometry. These weights are one-directional, as we only consider solutions that span the full geometry from layer S to F in Figure 4. In total, there are N^L solutions to the directed graph, which we refer to as ‘paths’. For clarity, a single path P_j is defined as the collection of edges that span from the leftmost to rightmost layers (S to F), touching exactly one node in every layer (see Figure 7 for an $N = 2$ example).

$$\omega_i \in [0, R], \quad \omega_i \in \mathbb{Z} \quad (7)$$

$$W_j = \sum_{i \in P_j} \omega_i \quad (8)$$

$$\mathbb{P} \equiv \{P_1, P_2, \dots, P_{N^L}\} \equiv \text{All Paths} \quad (9)$$




Figure 4. A geometry composed of sequentially connected bipartite directed graphs with weighted edges for which we are interested in finding the optimal path from layer S to layer F, touching exactly 1 node per layer. N denotes the number of nodes per layer, while L is the total number of layers. With full connectivity between nearest neighboring layers, each geometry has a total of $N^2 \cdot (L - 1)$ edges, yielding N^L possible paths from layer S to F.

For each path P_j , there is a cumulative weight W_j that is obtained by summing the individual weighted edges that make up the path (Equation (8)). The goal is to find the optimal solution path with a cumulative weight of either $W_{\min}$ or $W_{\max}$:

$$\mathbb{W} \equiv \{W_1, W_2, \dots, W_{N^L}\} \equiv \text{All Solutions} \quad (10)$$

$$W_{\min} = \min(\mathbb{W}) \quad (11)$$

$$W_{\max} = \max(\mathbb{W}) \quad (12)$$

For simplicity, we consider problems where each edge ω_i is an integer number between 0 and some max R . This will allow for a clearer picture when visualizing solution spaces $\mathbb{W}$ later on. However, we note that all the results which follow are equally applicable to the continuous case $\omega_i \in \mathbb{R}$ (set of real numbers), which we discuss in Section .

3.2. Classical Solving Speed

As outlined in Equations (7)–(12), we are interested in finding the path (collection of weighted edges) which corresponds to the smallest or largest W_i value within the set $\mathbb{W}$. However, the cumulative values W_i are assumed to be initially unknown and must be computed from a given directed graph like in Figure 4. Importantly, this means that the base amount of information given to either a classical or quantum computer is the set of ω_i weights and their locations, for which either computer must then find an optimal solution. For graphs defined according to Figure 4, yielding $N^2 \cdot (L - 1)$ total weights, we argue that

the optimal classical solving speed is of this order. Figure 5 below is an example of how a classical algorithm solves the pathfinding problem one layer at a time, checking each weighted edge exactly once.




Figure 5. A layer by layer example of a classical approach to finding $W_{\min}$ or $W_{\max}$, for the case of $N = 3$ and $L = 4$. The blue-dashed, green-solid, and red-dotted lines each represent possible solutions for the optimal path ending on each of the three nodes per layer.

The steps illustrated in Figure 5 can be summarized as the recursive process given in Algorithm 2. The general strategy is to work through the graph one layer at a time, checking all N^2 edges between layers, and continually updating a list (labeled OP in Algorithm 2) of possible optimal paths as one moves through the geometry. Importantly, the classical algorithm only needs to check each weighted edge one time to determine the optimal path. At each layer of the algorithm, N candidate paths are stored in memory (the blue, red, and green lines in Figure 5) and used to compute the next N^2 possible paths (grey solid lines), repeating this process up to the final layer.

Algorithm 2 Classical Pathfinding

```

1: OP = { 0, 0, ..., 0 } (length N)
2: for L - 1 do
3:   for N2 do
4:     Check each edge OPk + wi
5:     if OPk + wi is optimal then
6:       Update OPk
7: Wmin/max = min/max OP

```

The algorithm shown above has an $O(N^2 \cdot (L - 1))$ query complexity, which we will later compare with quantum. However, this speed is specifically for directed graphs defined according to Figure 4 and Equations (7)–(12). And while quantum will offer a speedup for certain N and L ranges, this particular speedup is not the primary interest of this study. As we demonstrate next, these sequential bipartite graphs were chosen to illustrate a problem with an efficient quantum circuit construction for the oracle. Different graph structures will have varying classical speeds for quantum to compete against, but not all graph structures are easily encoded into quantum states and solvable using amplitude amplification.

4. Quantum Cost Oracle

Having now outlined the problem of interest, as well as a classical solving speed, in this section we present the quantum strategy for pathfinding. We begin by outlining the

manner in which all N^L possible paths are uniquely assigned a quantum state, with the goal of encoding each total path weight W_i via phases. Then later in Section 4, we show how these phases can be used for amplitude amplification to solve for $W_{\min}$ or $W_{\max}$.

4.1. Representing Paths in Quantum

For qubit-based quantum computing, the methodology put forth in this section is most naturally suited to problem sizes where $N = 2^n$ (nodes per layer). This is because N dictates how many quantum states are needed for encoding a layer, for which 2^n is achievable using qubits. We begin by presenting two example cases in Figure 6 of size $N = 2$ and $N = 4$, both $L = 4$. Accompanying each graph are the qubit states needed to represent each node per layer.




Figure 6. (top) An example geometry of size $N = 2$, $L = 4$. For the case of $N = 2$, a single qubit is sufficient for representing all possible node choices per layer via the states $|0\rangle$ and $|1\rangle$. (bottom) An example geometry of size $N = 4$, $L = 4$, requiring two qubits for representing the nodes in each layer.

Because we are interested in solving a quantum pathfinding problem, the manner in which the qubits' orthogonal basis states $|0\rangle$ and $|1\rangle$ are used needs to reflect this fact. A final measurement at the end of the algorithm will yield a state $|P_i\rangle$, comprised of all $|0\rangle$'s and $|1\rangle$'s, from which the experimenter must then extrapolate its meaning as the path P_i . We achieve this by encoding each individual qubit state (or group of qubits) as the location of a particular node in the geometry. Using $\sqrt{N}$ qubits allows us to identify each of the N nodes per layer (for problem sizes $N = 2^n$), for a total of $\sqrt{N} \cdot L$ qubits representing a complete graph. For problems of size $N > 2$, multiple qubits are grouped together in order to represent all possible nodes per layer, such as in Figure 6 (two qubits for representing four nodes).

Figure 7 shows an example path for $N = 2$, and its corresponding state $|P_i\rangle$. For this particular graph size there are a total of 16 possible paths, which can be exactly encoded using the basis states $|0\rangle$ and $|1\rangle$ of four qubits. Conversely for an $N = 4$ geometry, two qubits are necessary for representing the four possible nodes per layer (states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$). This yields a total of 8 qubits for the complete graph ($N = 4$, $L = 4$), for a Hilbert space of size 2^8 , which is exactly equal to the total number of possible paths 4^4 . With quantum states encoded in this manner, the goal of the algorithm is to measure $|P_{\min}\rangle$ or $|P_{\max}\rangle$, which will yield the answer $W_{\min}$ or $W_{\max}$ upon classically checking the path.




Figure 7. An example path (red-dashed) for a graph of size $N = 2$, $L = 4$. The quantum state $|0100\rangle$ represents the path shown in red, using the single qubit states $|0\rangle$ and $|1\rangle$ for bottom and top row nodes respectively.

4.2. Cost Oracle U_P

The four qubit state shown in Figure 7 corresponds to a single path, but a superposition state is capable of representing all 2^4 solutions simultaneously (and more generally any N^L). In order to use these states for finding the optimal path, we now need a mechanism for assigning each path state $|P_i\rangle$ its unique path weight W_i . To achieve this, we implement an operator U_P , which we refer to as a ‘cost oracle’, capable of applying the cumulative weights W_i of each path through phases:

$$\begin{aligned} U_P|0100\rangle &= (e^{i\omega_1} \cdot e^{i\omega_2} \cdot e^{i\omega_3})|0100\rangle \\ &= e^{i(\omega_1+\omega_2+\omega_3)}|0100\rangle \\ &= e^{iW_{0100}}|0100\rangle \end{aligned} \quad (13)$$

In Equation (13) above, we’ve used the numerical weights ω_i from Figure 7 as an example, where each edge is directly translated into a phase contribution. In practice, however, a scaling factor p_s is necessary for meaningful results (which we discuss in Sections 4 and 5). The reason we refer to U_P as a cost oracle is because the manner in which it affects quantum states is analogous to that of a cost function. More specifically, applying U_P to any state $|P_i\rangle$ will cause the state to pick up a phase proportional to its cumulative weight W_i . However, it is more accurate to call U_P an oracle because the exact manner in which phases are distributed throughout the quantum system is unknown to the experimenter. That is to say, the experimenter is unaware of which $|P_i\rangle$ state is receiving the desired phase proportional to $W_{\min}$ or $W_{\max}$ until the conclusion of the algorithm. The matrix representation of U_P has the form of Equation (14) below, where each phase ϕ_i is a scalar of the form $p_s \cdot W_i$. (the role of p_s is discussed later). The matrix for U_P has dimensions $N^L \times N^L$, equal to the total number of possible solutions, with each path’s phase along the main diagonal.

$$U_P|\Psi\rangle = \begin{bmatrix} e^{i\phi_1} & 0 & 0 & \cdot & \cdot \\ 0 & e^{i\phi_2} & 0 & & \\ 0 & 0 & e^{i\phi_3} & & \\ \cdot & & & \cdot & \\ \cdot & & & & \cdot \end{bmatrix} \begin{bmatrix} |P_1\rangle \\ |P_2\rangle \\ |P_3\rangle \\ \cdot \\ \cdot \end{bmatrix} \quad (14)$$

It is important to note that the matrix shown in Equation (14) is *not* necessary for the quantum circuit implementation of U_P . In particular, computing all N^L phases is already slower than the $O(N^2 \cdot (L - 1))$ approach laid out in Section 2. Thus, as we demonstrate in the next subsection, a viable quantum approach needs to implement U_P *without* calculating any total path lengths W_i .

4.3. Quantum Circuit

Having now seen the desired effect from U_P (Equation (14)), here we present a qubit-based quantum circuit design that efficiently achieves all N^L unique phases, with no *a priori* classical computations of any W_i . Here we will focus on the case $N = 2$ for simplicity, leaving the general case for the next section. We begin by defining the operator U_{ij} shown

below in Equation (15), and its corresponding quantum circuit in Figure 8. The operator U_{ij} encodes all of the phases contained between layers i and j , from which we can build up to the full U_P .

$$U_{ij} \equiv \begin{bmatrix} e^{i\phi_{00}} & 0 & 0 & 0 \\ 0 & e^{i\phi_{01}} & 0 & 0 \\ 0 & 0 & e^{i\phi_{10}} & 0 \\ 0 & 0 & 0 & e^{i\phi_{11}} \end{bmatrix} \quad (15)$$




Figure 8. (top) Illustration of layers i and j for an $N = 2$ graph, and the four weighted edges shared between them. (bottom) Quantum circuit for achieving the U_{ij} operation outlined in Equation (15).

The circuit shown in Figure 8 applies a unique phase to each of the 2-qubit basis states $|Q_i Q_j\rangle$, one for each of the four edges connecting layers i and j . The complete information of all weighted edges connecting layers i and j is achieved with exactly one phase gate (controlled) per edge, which is a property that holds true for all geometry sizes. Importantly, from a qubit connectivity viewpoint, the qubits which make up layer i only need to interact with the qubits making up layers $i \pm 1$. This in turn can be used to significantly reduce circuit depth, demonstrated below in Figure 9.




Figure 9. The complete circuit design for U_P , for the case of $N = 2$. Each U_{ij} operation applies the four ϕ_i phases corresponding to the ω_i weights connecting layers i and j . Because of the way in which phases add exponentially, the order in which a total weight W_i is applied to a state $|P_i\rangle$ can be done in two sets of parallel operations, shown by the dashed-grey line.

$$U_P \equiv \prod_{i=1}^{L-1} U_{i,i+1} \quad (16)$$

Let us now compare the desired effect of U_P from Equation (14), with its layer-by-layer construction shown in Figures 8 and 9. Each U_{ij} operation applies phases proportional to the locally weighted edges connecting layers i and j , involving only the qubits representing those layers. Also, due to the way in which phases add exponentially (Equation (13)), the full path weight W_i for each $|P_i\rangle$ state is achieved from the product of U_{ij} operations, shown above in Equation (16). Importantly, note that nowhere in U_P 's construction do we compute a single W_i value. As mentioned earlier, this is a necessary requirement of U_P in order to truly consider it an oracle operation. Here we have achieved exactly that

by splitting U_P up into localized U_{ij} operations for each section of the graph. For results on how an $N = 2$ U_P operation performs on IBM's superconducting qubits, please see Appendix A.

We would like to stress that the structure of Figure 9 is general for all geometry sizes, which is one of the motivations for studying these sequential bipartite graphs. The parameter N dictates the number of quantum states per layer, which in turn determines the dimensionality of U_{ij} . But for all graphs, the parameter L has no impact on circuit depth, as the complete implementation of U_P can always be achieved through two sets of parallel U_{ij} operations, shown by the dashed grey line in Figure 9.

4.4. Qudit Quantum Circuit

To compliment the results from the previous section for constructing U_P on a qubit-based quantum computer, here we shall briefly mention how qudits can be used to greatly expand beyond simply $N = 2^n$ sized graphs, as well as further reduce circuit depth. Since we will be interested in using qudits again in Section 8, let us now introduce the notation for a general d -level quantum bit:

$$|Q\rangle_d \equiv \sum_{i=0}^{d-1} \alpha_i |i\rangle_d \quad (17)$$

As shown in Equation (17), the quantum state for any d -dimensional qudit can be expressed as a superposition of orthogonal basis states, spanning $|0\rangle$ through $|d-1\rangle$. Experimentally, the realization of qudits has been steadily progressing over the past decade [33–36], which makes it an exciting time to start considering their applications for quantum algorithms. Here, the use of qudits allows us to represent graphs beyond $N = 2^n$. For example, a qutrit-based computer ($d = 3$) can encode graphs of size $N = 3^n$. Better still, a mixed qudit computer grants us the ability to encode graphs with a different N at each layer, such as in Figure 10.




Figure 10. (top) A sequential bipartite graph of varying N at each layer. (bottom) A mixed qudit quantum state capable of representing all possible paths through the geometry.

Note that it is still possible to create a varying N graph using qubits, so long as every layer has $N = 2^n$ nodes. However, even for geometry sizes that are implementable using qubits, the use of qudits is still advantageous for several reasons. Consider the two quantum circuits shown below in Figure 11, which both achieve a U_{ij} operation connecting two $N = 4$ layers, applying the same 16 phases in total.

The primary issue with using qubits is that there is a hidden resource cost when constructing higher-order control operations. In order to achieve an N -control phase gate, the true quantum circuit requires N additional ancilla qubits to serve as intermediate excited states [53]. This is because the qubit operations from which we build up higher-order control-phase gates are $P(\theta)$ (single-qubit phase), CX (control-X), and CCX (Toffoli). The significant advantage that the $d = 4$ qudit circuit has is the absence of Toffoli gates, as all 16 control-phase operations only need to occur between the two qudits. Thus, the qudit circuit is advantageous in both resource cost (two qudits vs. seven qubits) and circuit depth (reduction of four Toffoli gates per each of the 16 phase operations). Of course, the trade-off

is that qudit technologies are still primitive compared to the more popular qubit, and as such would be expected to come with much higher error rates. Nevertheless, we will return to the use of qudits again in Section 8, as the Hilbert space sizes they offer will be necessary for unlocking meaningful problems to solve.




Figure 11. Quantum circuits for U_{ij} connecting two layers of $N = 4$ nodes. **(top)** A qubit-based quantum circuit **(bottom)** A $d = 4$ qudit-based quantum circuit. More information on qudit unitary operations and circuits can be found in the review study by Wang et al. [54], such as the X_d^- operator shown here.

5. Gaussian Amplitude Amplification

With the construction of U_P outlined in Section 4, here we discuss how this cost oracle operator can be used to solve for $W_{\min}$ or $W_{\max}$. Because U_P applies phases to every quantum state, substituting it for U_G in Grover's algorithm has dramatic consequences on the way in which the amplitude amplification process plays out.

5.1. Solution Space Distributions

The motivation for studying directed graphs according to Figure 4 is only partially due to their circuit implementation (Figures 8–11). Additionally, these sequential bipartite graphs possess a second important quality necessary for the success of the algorithm: their $\mathbb{W}$ distributions. In Equation (7) we restricted each edge weight ω_i to be an integer value, for a reason that we will now discuss. By forcing each ω_i to be an integer within $[0, R]$, we can create directed graphs that have a high likelihood of repeat W_i values. Consequently, two independent paths $|P_i\rangle$ and $|P_j\rangle$ will both yield the same cumulative weights $W_i = W_j$, from different contributing ω_i 's. As we let N and L increase, these repeat values lead to $\mathbb{W}$ distributions which become describable by a gaussian function, given in Equation (18), where the majority of W_i values cluster around the expected mean $\mu \approx \frac{R}{2}(L - 1)$.

$$G(x) = \alpha e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (18)$$

Figure 12 illustrates a few example problem sizes for various N and L , and their resulting $\mathbb{W}$ histogram distributions. These distributions represent the range of expected outcomes from picking a path through the directed graph at random and seeing what W_i value one gets. The odds of picking the optimal path are 1 in N^L , while the most probable W_i corresponds to the peak of the gaussian. Importantly, the tail ends of the distribution represent our desired solutions $W_{\min}$ and $W_{\max}$ (top left plot in Figure 12), which are always maximally distanced from the cluster of states around the mean. Also note that letting ω_i be continuous within $[0, R]$ still produces the same gaussian effect, but discrete bin sizes

are necessary for viewing the resulting $\mathbb{W}$ histogram distributions, hence our choice to let ω_i be integers only.

$$Y_i = \text{pop.}(W_i) \quad Y'_i = G(W_i)$$

$$R_{\text{corr}} = \sqrt{\frac{\sum_i (Y'_i - Y_i)^2}{N^L}} \quad (19)$$

Shown above in Figure 13 is an example distribution and accompanying gaussian best-fit. This particular distribution was derived from a graph of size $N = 6$, $L = 10$, in anticipation of results later to come (Figures 16, 22 and 24). With an R_{corr} value of approximately 3.98, given by Equation (19), it is clear that the gaussian approximation for this example is not perfect. Even for a problem such as this one, composed of over 60 million in possible solutions, the resulting $\mathbb{W}$ distribution still has non-negligible deviations from a perfect Gaussian, which will be a primary focus of Section 7. Nevertheless, these approximate Gaussian profiles are sufficient for the success of the algorithm.




Figure 12. Histograms of W_i for randomly generated graphs of various N and L sizes, with $R = 100$. As N and L increase while keeping R constant, the profile of these $\mathbb{W}$ distributions approach perfect gaussians, given by Equation (18).




Figure 13. (black circles/blue lines) A histogram of $\mathbb{W}$ for a randomly generated graph with parameters: $N = 6$, $L = 10$, $R = 100$. (red dash) A best-fit gaussian plot of the form given in Equation (18), minimizing Equation (19) ($R_{\text{corr}} \approx 3.981$), with gaussian parameter values reported in the top-right.

5.2. Mapping to 2π

When using the cost oracle as defined in Equation (14), one must be mindful that U_P does not only mark the states corresponding to $W_{\min}$ and $W_{\max}$, but *all* states uniquely. This is quite different from the standard Grover oracle U_G , which *only* marks the state(s) of interest. For this reason, the use of U_P for amplitude amplification can be viewed as less flexible than U_G . While U_G can in principle be used to boost any of the N^L quantum states in $|\Psi\rangle$, U_P on the other hand is better suited for boosting a much smaller percentage of states. However, the states which U_P is effective at boosting are $|P_{\min}\rangle$ and $|P_{\max}\rangle$, perfect for solving a directed graph problem.

In viewing the $\mathbb{W}$ histograms in Figure 12, let us now consider the effect of applying U_P from Equation (14) on an equal superposition state $|s\rangle \equiv H^{\otimes n}|0\rangle^{\otimes n}$. Each point along the x-axis corresponds to a particular path length W_j , while the y-axis represents the total number of quantum states which will receive a phase proportional to that weight: $e^{i\phi_j}|P_j\rangle$. Thus, the net result of U_P will apply all N^L phases in a gaussian-like manner, with the majority of states near the mean receiving similar total phases (from different contributing ω_i 's). And in order to capitalize on this distribution of phases, we will introduce a phase scaling constant p_s into the oracle operation, which affects all states equally:

$$U_P(p_s)|\Psi\rangle = \sum_j^{N^L} e^{i(p_s \cdot W_j)} |P_j\rangle \quad (20)$$

The scaling constant p_s in Equation (20) is a value which must be multiplied into every cumulative W_j phase throughout the oracle. This can be achieved by setting each individual phase in U_{ij} to $p_s \cdot \omega_i$, such that the cumulative operation of U_P is equal to Equation (20). The phase p_s can be thought of as simply the translation of any problem's $\mathbb{W}$, for any scale of numbers used, into a regime of phases which can be used for boosting. More specifically, a range of phases $[x, x + 2\pi]$ for which the state $|P_{\min}\rangle$ or $|P_{\max}\rangle$ is optimally distanced from the majority of states in amplitude space (complex plane). See Figure 14 for an illustrated example, and note the location of the red 'x' corresponding to $|\Psi\rangle$'s collective mean after U_P .




Figure 14. (left) An example histogram of all W_i paths for the case of $N = 4$, $L = 10$, $R = 100$. (right) The same distribution mapped to a complete 2π cycle of phases via the cost oracle U_P acting on the equal superposition state $|s\rangle$. Additionally, the resulting mean (red 'X') and $|P_{\min}\rangle/|P_{\max}\rangle$ states (blue diamond) are shown. An accompanying color scale is provided on the far right, illustrating the percentile distribution of states for both plots.

Without p_s , the numerical W_i values from a given directed graph have no guarantee of producing any meaningful amplitude amplification. However, when scaled properly with an optimal p_s (which is discussed in Sections 6 and 7), U_P can be made to distribute phases like shown in Figure 14, where the phases picked up by $|P_{\min}\rangle$ and $|P_{\max}\rangle$ form a range of $[x, x + 2\pi]$. This in turn ensures that the majority of states will cluster near $x + \pi$, pulling the amplitude mean (red 'X') away from $|P_{\min}\rangle$ and $|P_{\max}\rangle$ (blue diamond).

5.3. U_G vs. U_P Diffusion

As with the standard Grover search Algorithm [1], the U_P oracle operation in isolation is not enough to solve for $W_{\min}$ or $W_{\max}$. A second mechanism for causing interference is necessary in order to boost the probability of measuring the desired state. For this, we once again use the standard Grover diffusion operator U_s , given in Equation (1). With U_P distributing phases to each state, and U_s causing reflections about the average, we now have sufficient tools for quantum pathfinding, shown in Algorithm 3.

As noted previously, the algorithm outlined here is identical to that of Grover's search algorithm, with U_G swapped out for U_P . However, this replacement significantly changes the way in which the states of $|\Psi\rangle$ go through amplitude amplification, illustrated in Figure 15. For comparison, the amplitude space when using the standard U_G is also shown.




Figure 15. Examples of amplitude amplification, comparing the use of U_P vs. U_G for five iterations, both with the same number of total states $N' = 24,000$. In both plots, the origin (0,0) (black '+'), the mean point (red 'x'), the desired boosted state (blue diamond), and all other points (black circles) are shown. For scale, the radius of the equal superposition state $|s\rangle$ ($1/\sqrt{N'}$), as well as the probability of measuring the blue diamond state (which can be used to infer distance to the origin).

Algorithm 3 Quantum Pathfinding

- 1: Initialize Qubits: $|\Psi\rangle = |0\rangle^{\otimes N}$
- 2: Prepare Equal Superposition: $H^{\otimes N}|\Psi\rangle = |s\rangle$
- 3: **for** $k \approx \frac{\pi}{4}\sqrt{N^L}$ **do**
- 4: Apply $U_P(p_s)|\Psi\rangle$ (Phase Oracle)
- 5: Apply $U_s|\Psi\rangle$ (Diffusion)
- 6: Measure

Step 1 of Figure 15 shows the effect of using the diffusion operator U_s immediately following the first application of U_P (see Figure 14). The location of the mean point (red 'X') causes states near $|P_{\min}\rangle$ and $|P_{\max}\rangle$ (blue diamond) to reflect further than those around the mean of the gaussian. However, when compared with the lower plots using U_G , this increase in probability is always smaller than that of standard Grover's. Geometrically, this is a consequence of having states with phases spread out over a 2π range, resulting in a mean amplitude point that is closer to the origin (similar to U'_{G2} from Section 2).

What follows after step 1 for the case of U_P is a process with no simple mathematical description. As illustrated in steps 2–5, repeat applications of U_P and U_S result in quantum superposition states which exhibit a ‘spiraling’ effect around the mean point, which itself is also moving around the complex plane. Although quite clearly different from standard Grover’s, two key elements remain the same: (1) the distance between the mean point and the origin gradually decreases with each step, while (2) the distance between $|P_{\min}\rangle/|P_{\max}\rangle$ and the origin increases (i.e., incremental probability gains with each step). Just like Grover’s, both of these statements hold true for $O(\sqrt{N^L})$ iterations, after which the process begins to rebound.

Shown above in Figure 16 is a step-by-step comparison of probabilities for standard Grover’s versus gaussian amplitude amplification (i.e., amplitude amplification using a 2π gaussian distribution of phases), both for problem sizes of 6^{10} quantum states ($N = 6$, $L = 10$). The blue-dashed line tracks the probability of measuring the marked state as it approaches 1, while the red-solid line represents the probability of measuring $|P_{\min}\rangle$. Notably, the probability of $|P_{\min}\rangle$ achieves a lower peak P_M , and at a later step count. This is the trade-off for using U_P versus U_G : a lower boost in probability, but a solution to an inherently different problem (unstructured search vs. weighted directed graph). Importantly, however, the combination of iterations and peak probability for $|P_{\min}\rangle$ is still high enough for a potential quantum speedup under certain conditions, which we discuss in the next two sections.




Figure 16. A comparison of probability boosting using U_G (blue-dashed) vs. U_P (red-solid) as a function of steps (oracle + diffusion iterations), both acting on a quantum system of 6^{10} states. For U_G we track the probability of the marked state, while the U_P case tracks the probability of measuring $|P_{\min}\rangle$.

6. Simulating Gaussian Amplitude Amplification

Much like the analysis of U'_{G2} from Section 2, here we present results which illustrate the capacity for successful amplitude amplification one can expect from a gaussian distribution of phases encoded by U_P . To do this, we use a classical python-based simulator, capable of mimicking the amplitude amplification process outlined in Algorithm 3, allowing us to track quantum states and probabilities throughout. Results from various simulations are provided in the coming subsections, as well as their significance for identifying properties of problems that are viable for amplitude amplification.

6.1. Modeling Quantum Systems

As illustrated in Figure 14, amplitude amplification is viable for solving optimization problems with naturally gaussian solution spaces $\mathbb{W}$, scaled down to a 2π range of phases via p_s . In the next section, we address the challenges of finding p_s , while here we will focus solely on how the amplitude amplification process performs under ideal conditions.

$$G(\theta) = \alpha e^{-\frac{(\theta-\pi)^2}{2\sigma^2}}, \quad \theta \in [0, \pi] \quad (21)$$

Let us now outline our methodology for creating and simulating discrete U_P 's derived from Equation (21), shown in Figure 17. In step 1, we begin with a normalized gaussian ($\alpha = 1$) centered at π , with σ (standard deviation) as the only free parameter. Next we discretize the gaussian by using (x, y) points along the function ($x = \theta$, $y = G(\theta)$), taken in evenly spaced intervals of θ based on how many unique phases we want to model between 0 and 2π . These x and y values are then stored in two vectors: G_x and G_y . At this stage, G_x represents the various phases encoded by some U_P , but together with G_y they do not represent a valid oracle yet. This is because the values of G_y need to model a histogram of states, which means: (1) every value in G_y must be an integer, and (2) the sum of G_y is the Hilbert Space size of the quantum system. Analogous to the histograms shown throughout this study, G_x represents the space of possible W_i solutions, while G_y represents how many states will receive a phase proportional to W_i . Thus, a viable U_P operator is finally achieved in step 3 of Figure 17, after all the values of G_y are multiplied by a constant factor and rounded to integers (preserving σ from step 1).




Figure 17. (1–3) Illustrations of how our python-based simulator creates gaussian $\mathbb{W}$ distributions for testing. In step 1, we pick a standard deviation σ and create a continuous gaussian from 0 to 2π , with $\alpha = 1$ and $\mu = \pi$. In step 2 we select how many unique W_i phases we want to model, and use this number to discretize the continuous gaussian into two discrete arrays G_x and G_y . In step 3 we select a target Hilbert space size N to model, and scale all of the values in G_y up to integers, such that the sum(G_y) is as close to N as possible. And finally in step 4 we simulate amplitude amplification using G_x and G_y , tracking the probability of $|P_{\min}\rangle$.

For each simulation according to Figure 17, the full construction of U_P is based upon three free parameters of our choosing: σ , size(G_x), and the sum(G_y), shown in steps 1, 2, and 3 respectively. The motivation for these three parameters is based on their direct ties to the quantities N , L , and R from Equations (7)–(12). For example, the combination of N and L determines the Hilbert space size of the quantum system needed to represent all possible paths, which we can control with the sum(G_y). Simultaneously, L and R together dictate the maximum number of possible W_i weights: $[0, R \cdot (L - 1)]$, which we can model with the size(G_x). And finally, σ is impacted by all three parameters together, and as we show next, has the strongest correlation to whether or not amplitude amplification is viable.

6.2. Long Tail Model

Using the methodology put forth in Figure 17, there is still one important choice that impacts the nature of the quantum system we are modeling, namely rounding. In step 3 of Figure 17, we must implement a rounding protocol to meet the requirement that all G_y values be integers. For phases near the central region of the gaussian, the choice in rounding is practically inconsequential for the amplitude amplification process, but not for the tails where $W_{\min}$ and $W_{\max}$ lie. This can be seen in the two $U_P|s\rangle$ plots in Figure 18, where in one case all G_y values are rounded up to the nearest integer (left), and one where all values are rounded down (right).




Figure 18. (top) An example distribution created from our simulator, before rounding in stage 3, with properties of the distribution given on the left. (bottom) Two different U_P interpretations of the distribution shown on top. (left) The long tail model, whereby all values of G_y are rounded up to the nearest integer. Grey dashes indicate the region where $\text{pop}(W_i) = 1$. (right) The short tail model where all values are rounded down, causing $\text{pop}(W_i)$ values near the tails to be zero for small σ .

In this subsection, we shall focus on simulated distributions according to the left $U_P|s\rangle$ encoding in Figure 18, which we refer to as the ‘long tail’ model. Compared to the randomly generated distributions in Figure 12, this turns out to be an unrealistic model for problems where we expect $W_{\min}$ to be larger than the theoretical minimum. Nevertheless, this long tail model will serve to illustrate the most ideal case for Gaussian amplitude amplification. In particular, it allows us to simulate the theoretical limit of a gaussian distribution as σ goes to zero, for which the resulting amplitude amplification process is most nearly a replication of standard Grover’s.

Shown in Figure 19 are results from simulated amplitude amplifications for quantum systems of size $N \approx 60 \cdot 10^6$ ($\text{sum}(G_y)$). Each U_P oracle represents 700 unique weights W_i ($\text{size}(G_x)$) scaled to a 2π range, for σ values ranging from $[0, 1.2]$. The top plot shows the peak probabilities P_M achievable for the $|P_{\min}\rangle$ state, while the bottom plot shows the corresponding number of needed $U_P U_S$ iterations S_M .

Beginning with $\sigma = 0$, we note how close the results from Figure 19 are to that of standard Grover’s: P_M is ~ 0.997 vs. ~ 1 , and S_M is 6089 vs. 6083. For this σ , we are modeling an oracle where $N - 699$ states all receive π phase, $|P_{\min}\rangle$ receives a phase of 0, and the remaining 698 states all receive phases of varying $\pi/350$ multiples. If instead these 698 states were also set to receive phases of π , then U_P would be exactly U_G . But by having them evenly spread out over a full 2π range, their impact on the amplitude amplification process can be seen in P_M and S_M .

While the special case of $\sigma = 0$ can be thought of as the theoretical limit where U_P approaches U_G , the remaining results shown in Figure 19 illustrate how gaussian amplitude amplification performs for σ values which represent more realistic optimization problems. As one might expect, the top plot shows a steadily decreasing trend in P_M as σ increases, accompanied by similar incremental increases in S_M . These trends continue smoothly up to approximately $\sigma \approx 0.64$, which we shall refer to as σ_{cutoff} , at which point both plots change dramatically. The critical difference between the quantum systems we are modeling above and below σ_{cutoff} is that beyond this point the Gaussian distributions of U_P are so wide that they begin to populate multiple states with the value $W_{\min}$. Consequently, if there are M states all with the same $W_{\min}$, then they will all share $1/M^{\text{th}}$ of the probability boosting from amplitude amplification. For this reason, we’ve included the red-dashed line in the top plot of Figure 19, which multiplies each peak P_M by the $\text{pop.}(W_{\min})$. Thus, the red-dashed

line is a more accurate representation of the relation between P_M and σ for this particular Hilbert space size, independent of how many $W_{\min}$'s are present in the system.




Figure 19. Results for simulated gaussian distributions of Hilbert space size $N = 60 \cdot 10^6$, following the long tail model, as a function of standard deviation σ . **(top)** Black data points indicate the highest achievable probabilities P_M for $|P_{\min}\rangle$, while the red-dashed line shows $P_M \cdot \text{pop}(W_{\min})$ for cases with multiple $W_{\min}$ solutions. **(bottom)** The number of required iterations S_M in order to reach P_M .

The value σ_{cutoff} can be interpreted as the limit where a particular optimization problem is expected to have more than one optimal solution. For sequential bipartite graphs, we can manipulate the odds of getting multiple $W_{\min}$ paths by increasing N while simultaneously decreasing L and R . Importantly, the presence of multiple $W_{\min}$'s does not detract from a U_P 's aptitude for boosting states, as evidenced by the red dashed line which represents the shared probability across all $|P_{\min}\rangle$ states. However, it does significantly impact the expected optimal number of iterations S_M , which can be seen in the bottom plot of Figure 19. Having multiple states share the optimal phase is analogous to a result from 1998 [5], where the step count for Grover's search algorithm is reduced from $O(\frac{\pi}{4}\sqrt{N})$ to $O(\frac{\pi}{4}\sqrt{N/M})$ for M marked states. Here the same effect can be observed in the S_M plot, where each increase in the $\text{pop.}(W_{\min})$ results in a fractional reduction to S_M .

6.3. Short Tail Model

One important trend from long-tail model and Figure 19, which will continue throughout this study, is the inverse relation between the standard deviation σ of a problem's solution space $\mathbb{W}$, and U_P 's ability to boost $|P_{\min}\rangle$. Thus, the ideal optimization problem for amplitude amplification is one with a naturally small σ , and $W_{\min}$ as distanced from the mean as possible (i.e., long tails). More realistically, however, these two conditions are contradictory to each other: the smaller σ is for a given problem, the *closer* we expect $W_{\min}$ to be to the mean.

Returning now to the bottom right $U_S|s\rangle$ plot of Figure 18, here we present results from our simulator which model problems more akin to Figure 12. We refer to these $\mathbb{W}$ distributions as the 'short tail' model, by which we mean the expected number of solutions where $\text{pop.}(W_i) = 1$ is small, and the expected number of solutions where $\text{pop.}(W_i) = 0$ increases as σ decreases. Unlike the long tail model, this represents an optimization problem

where $W_{\min}$ is unknown (changing as a function of σ), making it more difficult to find an effective p_s scaling factor, such as Equation (24) below.

$$W_{\min} \cdot p_s = x \quad (22)$$

$$W_{\max} \cdot p_s = x + 2\pi \quad (23)$$

$$p_s = \frac{2\pi}{W_{\max} - W_{\min}} \quad (24)$$

Because we have full information of the quantum systems we are modeling, both $W_{\min}$ and $W_{\max}$ are known for every simulation so we are able to use Equation (24) to find the optimal p_s for each U_P . In the long tail model no p_s scaling was necessary, whereas here it is required in order to align $|P_{\min}\rangle$ for optimal boosting. Shown below in Figure 20 is an illustration of this rescaling process, analogous to Figure 14.




Figure 20. (top left) An example distribution created from our simulator following the short tail model, causing $W_{\min}$ and $W_{\max}$ to be located away from 0 and 2π . (top right) The same distribution scaled by p_s to a full 2π range. (bottom) Below each histogram distribution is an amplitude space plot of $U_P |s\rangle$, tracking the location of $|P_{\min}\rangle$ (blue diamond) and the mean point (red 'X').

The process shown in Figure 20 takes place in our simulations immediately following step 3 of Figure 17, before simulating amplitude amplification for P_M and S_M . The consequence of this rescaling can be seen in the statistics of the top right distribution, resulting in new σ' and $\text{size}(G_x)$ values from the original. This rescaled $\text{size}(G_x)$ value comes from the number of $G_y \neq 0$ states ($\text{pop.}(W_i) = 0$) in the system, which have no impact on the amplitude amplification process. Consequently, the boosting of $|P_{\min}\rangle$ is driven by a new effective standard deviation σ' , which notably is always $\sigma' \geq \sigma$.

Shown in Figure 21 are results of simulated amplitude amplification for the short tail model, for a range of initial σ values $[0, 0.8]$ and initial $\text{size}(G_x) = 700$. In all four plots there are three sets of data for various Hilbert space sizes: $N = 60 \cdot 10^6$ (blue), $N = 10 \cdot 10^6$ (orange), and $N = 2 \cdot 10^6$ (green). In contrast to the long tail model results of Figure 19, Figure 21 illustrates a different trend for P_M vs. σ up to σ_{cutoff} . The highest P_M achievable for $N = 60 \cdot 10^6$ at $\sigma = 0$ was previously ~ 0.997 , now only ~ 0.917 following the short tail model. However, if we look at the top right plot of σ' vs. σ , we can see where this lower P_M value comes from. Over the range of initial σ values $[0, \sigma_{\text{cutoff}}]$, the consequence of rescaling with p_s are σ' values between $[0.54, 0.59]$. Comparing these σ' values with Figure 19, the long tail model predicts P_M values around $0.89 \sim 0.92$, which is exactly what we find for P_M 's reported in Figure 21.




Figure 21. Results for simulated gaussian distributions of various Hilbert space sizes (blue = $60 \cdot 10^6$, orange = $10 \cdot 10^6$, and green = $2 \cdot 10^6$), following the short tail model, as a function of initial standard deviation σ . **(left)** P_M and S_M plots for boosting $|P_{\min}\rangle$. **(top right)** The standard deviation σ' after rescaling each distribution by the p_s value which maximizes P_M (see Figure 20). **(bottom right)** The total number of unique W_i phases modeled by each distribution.

To explain this new relation between σ' and P_M , we must note the two additional Hilbert sizes N (orange and green data points) shown in Figure 21. For any given initial σ , all three simulation sizes were derived from the same normalized gaussian in step 1 of Figure 17. Yet due to their differing N values, each system size populates a different number of unique W_i states, shown in the bottom right plot of $\text{size}(G_X)$ vs. σ . For each σ , the largest Hilbert space $N = 60 \cdot 10^6$ always results in the biggest $\text{size}(G_X)$ after rounding, which consequently yields the largest distance between $W_{\min}$ and $W_{\max}$. This distance dictates the necessary amount of rescaling by p_s (Equation (24)), resulting in different σ' values, which in turn determine achievable P_M 's for $|P_{\min}\rangle$.

To summarize, the findings presented here for the long and short tail models demonstrate the range of success that gaussian amplitude amplification can produce. For any optimization problem, we must consider not only the solution space $\mathbb{W}$'s natural σ , but how the distribution of W_i 's can be mapped to a 2π range of phases for U_P . This was the motivation for introducing σ' via the short tail model, which demonstrated that problem size N is just as important as σ . Even for a problem that may possess a naturally small σ , if N isn't sufficiently large enough to probabilistically produce $W_{\min}/W_{\max}$ solutions away from the mean, then the problem may not be viable for a quantum solution. Conversely, if we are able to encode large optimization problems into U_P oracles, then we can expect successes analogous to the long tail model with small σ .

7. Algorithmic Viability

The hope of quantum computers isn't to solve artificially created ideal scenarios, but problems that arise naturally with inherent difficulties. Following the simulated Gaussian amplitude amplification results from the previous section, we now ask how reliable this boosting mechanism is for $\mathbb{W}$ distributions with imperfections that one would expect from realistic problems. What follows in the coming subsections are observations and techniques for applying the quantum pathfinding Algorithm 3 to randomly generated $\mathbb{W}$ distributions according to Equations (7)–(12).

7.1. Finding an Optimal p_s

In order to achieve a successful gaussian amplitude amplification on $|P_{\min}\rangle/|P_{\max}\rangle$, for a $\mathbb{W}$ distribution with deviations from a perfect gaussian, the key lies in finding an

optimal scaling parameter p_s . In Section 5.2 we introduced p_s as a necessary means for translating the full range of $\mathbb{W}$ down to $[x, x + 2\pi]$, and again in the short tail model for Section 6.3.

The approach outlined in Equation (24) is a way of ensuring $|P_{\min}\rangle$ and $|P_{\max}\rangle$ form a complete 2π range, but it is not necessarily the optimal p_s for amplitude amplification. Firstly, it causes the states $|P_{\min}\rangle$ and $|P_{\max}\rangle$ to share the boosting effect equally, which is not ideal for problems where we are interested in finding only one or the other. But more importantly, randomness in $\mathbb{W}$ means that the overall distribution of phases from U_P is very likely to be not symmetric. This means that the optimal p_s for boosting $|P_{\min}\rangle$ will differ from the optimal p_s for $|P_{\max}\rangle$. These different p_s 's correspond to values which best align $|P_{\min}\rangle/|P_{\max}\rangle$ with a π phase difference from the mean. Figure 22 illustrates an example of this, as well as the margin for error in finding the optimal p_s value before accidentally boosting an unintended state.




Figure 22. A plot of p_s vs. achievable probabilities via amplitude amplification, for the $\mathbb{W}$ distribution shown in Figure 13. The state $|P_{\min}\rangle$ represents the solution to the pathfinding problem $W_{\min}$, while $|P'_{\min}\rangle$ corresponds to the next smallest W_i .

Derived from the same directed graph used to produce Figure 13, the two plots shown in Figure 22 were created by carefully simulating Algorithm 3 over the range of p_s values shown along the x-axis, for $|P_{\min}\rangle$ as well as the second best solution state $|P'_{\min}\rangle$. It is clear by the two spikes in probability, and the space in between, that the role of p_s for unlocking successful amplitude amplifications cannot be ignored. For this particular example, using a scaling factor of $p_s \approx 0.008957$ causes the state $|P_{\min}\rangle$ to reach a peak probability of about 80.37%, while using $p_s \approx 0.008982$ causes $|P'_{\min}\rangle$ to boost to about 80.47%. Thus, a margin of error on the order of $\sim 3 \cdot 10^{-5}$ in p_s is enough to change what state gets boosted.

Additional notables from Figure 22 are as follows: (1) Despite a single optimal p_s for boosting $|P_{\min}\rangle$, the plot shows a range of p_s values around the optimal case for which the algorithm can still be successful. (2) The range of p_s values between the two peaks can be regarded as a ‘dead zone’, where no state in the system receives a meaningful probability boost. (3) Because states near $|P_{\min}\rangle$ are also able to receive meaningful amplitude amplifications ($|P'_{\min}\rangle$), this suggests that the algorithm may be viable for a heuristic technique. (4) From an experimental viewpoint, the scale of precision shown for p_s must be achievable via phase gates, which means the size of implementable problems will be dictated by the technological limits of state-of-the-art quantum devices.

7.2. Single vs. Multiple p_s

The two plots shown in Figure 22 represent potential amplitude amplification peaks, where a single p_s scaling factor is used for every iteration of $U_s U_P$. However, in principle this is not necessarily the optimal strategy for boosting $|P_{\min}\rangle$, as p_s could theoretically be

different with each iteration. The choice of p_s at each step is an extra degree of freedom available to the experimenter, which we explore here as a potential tool for overcoming randomness in $\mathbb{W}$.

In order to better quantify the advantage a step-varying p_s approach has to offer, let us first define our metric for a successful amplitude amplification in Equation (28) below. We refer to this metric as ‘probability of success’, labeled P_{succ} , which combines an amplitude amplification’s peak probability and step count into a single number, quantifying the probability of a quantum speedup over classical.

$$C_{\text{steps}} = N^2 \cdot (L - 1) \quad (25)$$

$$r = \lfloor C_{\text{steps}} / Q_{\text{steps}} \rfloor \quad (26)$$

$$P_M = \text{Prob.}(|P_{\min}\rangle) \quad (27)$$

$$P_{\text{succ}} = 1 - (1 - P_M)^r \quad (28)$$

To summarize the components making up Equation (28): C_{steps} is the number of classical steps needed to find $W_{\min}$ (equal to the total number of edges), Q_{steps} is the number of $U_s U_p$ iterations needed in order to reach the peak probability P_M , and r is the number of allowable amplitude amplification attempts to measure $|P_{\min}\rangle$ before exceeding C_{steps} . Altogether, P_{succ} represents the probability that $|P_{\min}\rangle$ will be successfully measured within r attempts. Using dice as a simple example, the probability of success that one will roll a 1–5 in four attempts is $P_{\text{succ}} = 1 - (1 - \frac{5}{6})^4 \approx 99.92\%$.

The quantity P_{succ} is a simplified way of comparing quantum vs. classical speeds, more specifically query complexity, which ignores many of the extra complicating factors of a more rigorous speed comparison (classical CPU speeds, quantum gate times, quantum decoherence and error correction, etc.). Here, we are simplifying one step in classical as the processing of information from a single weighted edge ω_i (steps 4–6 in Algorithm 2), versus one step in quantum as a single iteration of $U_s U_p$ (steps 4 & 5 in Algorithm 3). This is the typical manner in which Grover’s search algorithm is considered a quadratic speedup, and is sufficient for our study’s purpose.

With P_{succ} now defined, we return to the question of whether a step-varying approach to p_s can improve gaussian amplitude amplification. For details on how an optimal p_s can be computed at each step of the algorithm, please see Appendix B for our technique. To summarize, we simulate a range of p_s values at each step such that the distance in amplitude space between $|P_{\min}\rangle$ and the mean point is maximized, resulting in the largest reflection about the average from U_s per step. Figure 23 shows an example for the case $N = 30$, $L = 4$, and resulting P_M & P_{succ} .

As evidenced by the accompanying numbers in Figure 23, a step-varying approach to p_s is indeed advantageous for getting the maximal peak probability P_M out of a given $\mathbb{W}$. However, it is also clear that the exact sequence of p_s values (bottom plot) are non-trivial, and likely unpredictable from an experimental perspective when dealing with randomized data. Although the majority of p_s ’s are near a single value, there are constant sharp fluctuations at every step, some small while others are quite large. These fluctuations can be understood as a signature of the $\mathbb{W}$ distribution, unique to every problem, actively counteracting the randomness of the graph’s weighted edges at every step.

The result shown in Figure 23 for improving P_M was found to be very consistent. More specifically, *every* randomly generated graph that was studied, for all N and L , could always be optimized to produce a higher P_M using a step-varying p_s approach versus only a single p_s . However, in some cases it was found that the larger P_M value did not directly translate to a better P_{succ} , as the resulting higher Q_{step} count caused P_{succ} to be lower (fewer attempts to measure $|P_{\min}\rangle$). In general, our tests found the step-varying p_s approach to be most effective at improving P_M and P_{succ} for smaller problem sizes. But these smaller cases oftentimes produced p_s vs. step plots (bottom of Figure 23) which were highly chaotic and irregular from problem to problem, even for the same N and L . Conversely, as problem

sizes increase, the difference between the single vs. step-varying approaches became more negligible, with much more regular and stable p_s vs. step plots.




Figure 23. (top) An example $\mathbb{W}$ histogram distribution for the case $N = 30$, $L = 4$, $R = 200$. (bottom) A plot of all p_s values used at each step in order to optimized the probability of measuring $|P_{\min}\rangle$. Note the small black arrow, marking the p_s value at step 1. To the right of each plot are accompanying details about the success of each amplitude amplification process for each approach.

7.3. Statistical Viability

While the results from the previous subsection can be regarded as a more theoretical strategy for optimizing P_M , here we address the issue of finding p_s from a more practical perspective. In any realistic optimization problem, it is fair to assume that the experimenter has limited information about $\mathbb{W}$. Consequently, using a strategy for finding a suitable p_s such as Equation (24) may be impossible, which then begs the question: how feasible is gaussian amplitude amplification when used blindly? To help answer this question we conducted a statistical study, shown in Figure 24. The general idea is to imagine a scenario in which the experimenter needs to solve the same sized directed graph problem numerous times, with randomized but similar values each time (for example, optimal driving routes throughout a city can change hourly due to traffic patterns). Under these conditions, we are interested in whether a quantum strategy can use information from past directed graphs in order to solve future ones.

The results shown in Figure 24 illustrate the varying degrees of success one can expect using three different p_s approaches. The Figure showcases 100 randomly generated directed graphs of size $N = 6$, $L = 10$, $R = 100$, and their resulting peak P_M probabilities. Optimal P_M values for each graph were found through simulating amplitude amplification using (1) (light blue) a step-varying p_s approach, (2) (green) a single optimal p_s , and (3) (dark red) an average p_s . For the average p_s , this value was computed by averaging together the 100 single optimal p_s values: ~ 0.0083478 .

Two notables from Figure 24 are as follows: (1) Even for this appreciably large problem size (over 60 million paths), about 15% of the $\mathbb{W}$ distributions studied could not be optimized for P_M values over 50%. We found this to be of interest for a future study: what is it about these $\mathbb{W}$ distributions and their randomness that makes them inherently difficult to boost? (2) The large discrepancy between the single optimal and average p_s plots can be seen quite clearly across the 100 trials. However, returning to the question posed at the top of the subsection, the average P_M of these blind attempts is roughly 20% (top right corner of Figure 24). If a quantum computer could reliably be trusted to find $|P_{\min}\rangle$ 20% (or more) of the time using a single p_s , this could be a viable use case for quantum, used in conjunction with a classical computer for a hybrid approach.




Figure 24. Results from testing on 100 randomly generated $\mathbb{W}$ distributions, for $N = 6$, $L = 10$, $R = 100$. For each trial, we report the highest P_M probability found for the state $|P_{\min}\rangle$ using (light blue) a step-varying p_s approach, (green) a single optimal p_s approach, and (dark red) an average p_s approach. Reported on the right side of the figure are the averages found for all three approaches.

8. The Traveling Salesman

As the final topic of this study, here we present results for a theoretical application of gaussian amplitude amplification as a means to solve the Traveling Salesman problem [25] (TSP). Solving the TSP in this manner is an idea that goes back to 2012 [32], which we build upon here using the new insights gained from this study, particularly Sections 6 and 7. Because the adaptation of U_P discussed here relies on qudit technologies, which we will not explicitly cover, we encourage interested readers to see [54] for an overview of unitary operations and quantum circuits for qudits.

8.1. Weighted Graph Structure

Let us begin by defining the exact formalism of the Traveling Salesman Problem that we seek to solve using amplitude amplification. Shown in Figure 25 is an example TSP for the case of $N = 8$, where N corresponds to the total number of cities (nodes). Just as with the sequential bipartite graphs from Sections 3–6, a TSP can be represented as a weighted directed (or undirected) graph. Here we are interested in the most general case, an asymmetric TSP, where each edge has two unique weights w_{ij} and w_{ji} , one for traveling in either direction across the edge.




Figure 25. (left) Geometric structure for the Traveling Salesman Problem, for the case $N = 8$. Each edge contains a weighted value w_{jk} , where j and k are the two connected nodes. (right) An example path, touching each node exactly once. Each path P_i is defined by a unique ordering of all N nodes ($N!$ in total), with W_i corresponding to the sum of all weighted edges composing the path.

Once again, the solution we seek is $W_{\min}$ or $W_{\max}$, given in Equations (29) and (30). For clarity, here we are defining a path P_i as shown in Figure 25, traversing every node in the graph exactly once (and not returning to the starting node). In total, this produces a

solution space of $N!$ unique path permutations for a given TSP (for a symmetric TSP the number of permutations is the same, but the number of unique solutions is halved). We will continue to denote the set of all possible paths as $\mathbb{P}$, and similarly the set of all possible solutions as $\mathbb{W}$.

$$\omega_{jk} \in [0, R] \quad (29)$$

$$W_i = \sum_{jk \in P_i} \omega_{jk} \quad (30)$$

8.2. Encoding Mixed Qudit States

In order to realize a Hilbert space of size $N!$ such that every possible path P_i can be encoded as a quantum state $|P_i\rangle$, we require a mixed qudit quantum computer. Given in Equation (17) is the quantum state of a d -dimensional qudit, capable of creating superposition states spanning $|0\rangle_d$ through $|d-1\rangle_d$. When using qudits of different dimensions together, their combined Hilbert space size is the product of each qudit's dimensionality, as shown in Equation (31) below. If one is restricted to a quantum computer composed of a single qudit size d , then only quantum systems of size d^n are achievable. Thus, a single d -qudit computer can never produce the needed $N!$ Hilbert space size (unless $d = N!$, which is impractical) for solving the TSP.

$$|\Psi\rangle_{24} = |Q\rangle_4 |Q'\rangle_3 |Q''\rangle_2 = \sum_{i=0}^3 \sum_{j=0}^2 \sum_{k=0}^1 \alpha_{ijk} |i\rangle_4 |j\rangle_3 |k\rangle_2 \quad (31)$$

The quantum state shown above in Equation (31) is the mixed qudit composition which can encode an $N = 4$ TSP, capable of creating a superposition of $4! = 24$ states. These 24 states span every combination from the lowest energy state $|0\rangle|0\rangle|0\rangle$, up to the highest energy level for each qudit $|3\rangle|2\rangle|1\rangle$. Each of these basis states will serve as a $|P_i\rangle$, receiving a phase proportional to its total path weight W_i via the oracle U_P . See Figure 26 for an $N = 4$ TSP example.




Figure 26. (left) Geometric illustrations for 12 of the possible solution paths for an $N = 4$ TSP weighted graph. (right) Quantum state representations for the 12 paths shown, plus 12 additional states with opposite direction.

The quantum states shown in Figure 26 are meant to be symbolic, representing the information needed to specify each of the 24 unique paths (order of nodes traversed). For the realization of U_P however, we must encode the information of these 24 paths into

the orthogonal basis states $|i\rangle|j\rangle|k\rangle$ via phases. But unlike the convention used in Figure 7, where individual qubit states represent a single node in the graph, where we cannot use qudits in the same manner. To understand why, it is helpful to visualize the problem from a different geometric perspective, as shown in Figure 27.




Figure 27. Spanning tree representation of all possible paths for an $N = 4$ Traveling Salesman problem.

The spanning-tree representation shown in Figure 27 is equivalent to the weighted directed graph in Figure 25, with the same solution $W_{\min}$. The motivation for looking at the problem in this manner is to highlight the decreasing number of possible choices with each successive layer. Returning now to Equation (31), $|\Psi\rangle_{24}$'s mixed qudit composition was chosen to exactly mimic the dimensionality of choices at each layer in Figure 27. For example, the largest qudit $|Q\rangle_4$ in the system has four available states, one to represent each of the four possible starting nodes. Similarly, the next largest qudit $|Q'\rangle_3$ provides three possible states, one for each of the remaining untouched nodes, and so forth until the final qubit. However, while the four states of $|Q\rangle_4$ can all be exactly assigned to one of the four starting nodes, the same cannot hold true for the states of $|Q'\rangle_3$ and $|Q''\rangle_2$.

If we want to repeat the strategy for labeling $|P_i\rangle$ states like in Figure 7, then we require N total $d = N$ qudits, such that each $|i\rangle_d$ basis state can be uniquely specified as a particular node in the graph. However, this leads to a Hilbert space size of N^N , which is more than the number of total possible paths (for $N = 4$, this is 256 states for only 24 paths). These extra states are problematic because they represent invalid solutions to the TSP we want to solve, i.e., paths that traverse a single node more than once. Thus, in order to solve an $N!$ sized problem, we must use a Hilbert space created from a mixed qudit approach like in Equation (31).

Our solution to this $N!$ path/state encoding problem is outlined in Figure 28, for the case $N = 5$. The strategy for identifying each basis state of $|\Psi\rangle$ as a particular $|P_i\rangle$ follows from two rules: (1) initially label all nodes in the TSP graph with a unique $|i\rangle_d$ basis state for the $d = N$ largest qudit (leftmost graph). (2) For subsequent $d < N$ qudits, each $|j\rangle_d$ basis state corresponds to one of the remaining untraversed nodes, ordered clockwise from the position of the *previous* qudit state. See Figure 28 for two example paths, where possible qudit states at each step are shown in blue, and previous qudit states in black.

The two rules specified above are enough to guarantee every $|P_i\rangle$ is unique, even though the meaning of individual qudit states are not. While this encoding is sufficient, we note that other encodings are equally valid as well. So long as U_P is able to apply each phase $p_s \cdot W_i$ to the correct basis state $|P_i\rangle$, then the amplitude amplification results of the following subsection are applicable.

8.3. Simulated TSP Results

To conclude this discussion of the Traveling Salesman problem, here we present results which demonstrate how amplitude amplification performs as a function of N . To do this, we analyzed each problem size from two approaches: (1) Analogous to Figure 24, find the optimal single p_s for randomly generated graphs of each size, and record P_M values. (2) Compare these results against our simulator from Section 6.3 by gathering average statistics for $W_{\min}$, $W_{\max}$, and σ' , and use these along with $N!$ to predict expected P_M values. Results for method (1) are shown in Figure 29 below.




Figure 28. (leftmost) Initial mapping of an $N = 5$ TSP to the quantum states $|0\rangle$ – $|4\rangle$, and their accompanying city names. (panels 1–4) Step by step outline of two different paths through the geometry, illustrating the ‘clockwise’ nomenclature outlined in this section. At each step, the path thus far is illustrated in solid black lines/states, while potential next nodes are shown in blue arrows/states.

Starting with σ' , indicated by the black dots in Figure 29, we find a trend that is consistent with the sequential bipartite graphs from earlier in this study. As N increases, the rescaled standard deviation σ' of the solution space distribution $\mathbb{W}$ decreases, and consequently we find higher P_M values (blue dots). Accompanying each average P_M are intervals that represent the top 90% of all values found. These bars are in agreement with Figure 24, whereby the average values may be high, but working with randomized data is always subject to occasional $\mathbb{W}$ distributions which are inherently difficult to boost $|P_{\min}\rangle$. Even for $N = 11$, which was the largest size studiable with our computing resources, we still found the effects of randomness to be strong enough to cause P_M values to be under 40%.




Figure 29. Results from using a single optimal p_s for randomly generated TSP weighted graphs as a function of problem size N , $R = 100$. (dots) Average values for σ' (black) and P_M (blue). (bars) Intervals indicating the top 90% of all P_M values found.

Finally, using average $\mathbb{W}$ statistics in our simulator, we found predicted P_M values which were in strong agreement with those shown in Figure 29. For problem sizes $N = 9, 10, 11$, the simulator predicted P_M values which were all within 5% of the averages found experimentally. For smaller N sizes, the resulting $\mathbb{W}$ distributions become less and less resemblant to Gaussian profiles, making their comparison to our perfect gaussian simulator less meaningful. Overall, the two trends shown in Figure 29 are positive for quantum, indicating that as N increases so too does the viability of boosting $|P_{\min}\rangle$.

9. Conclusions

Amplitude amplification is a powerful tool for the future success of quantum computers, but it is not strictly limited to the unstructured search problem proposed by Grover over two decades ago [1]. In this study, we've demonstrated the viability of amplitude amplification as a means for solving a completely different problem type, namely pathfinding through a weighted directed graph. This was made possible by two key factors: (1) a cost oracle capable of encoding all possible solutions via phases, and (2) the gaussian-like manner in which the solution space naturally occurs. It is because of these Gaussian-like distributions that we are able to boost the desired solution state to high probabilities. More specifically, we are able to utilize the central cluster of states around the mean of the gaussian to create an oracle U_P which produces a mean point away from the desired solution state in amplitude space. This in turn allows for reflections about the average at each step via U_S to incrementally increase the probability of the desired solution state up to some maximum (P_M), which can be related to the distribution encoded into U_P . And finally, we've demonstrated that such oracles are implementable for the gate-based model of quantum computing, such that the answer to the optimization problem is not directly encoded into the quantum circuit for U_P .

Future Work

The algorithmic potential for gaussian amplitude amplification presented in this study is a promising first step, but there is still much to be learned. We view the process illustrated in Figure 15 as an open question for a more rigorous mathematical study. Throughout this study, we were able to simulate Gaussian amplitude amplification classically because each Hilbert Space had a finite number of states. However, studying a truncated continuous Gaussian function as it undergoes $U_S U_P$ through many steps is more difficult, but could lead to the improved success of the algorithm. Additionally, studying the same process but with a skewed gaussian could yield highly valuable insight into more realistic problem cases, such as why certain $\mathbb{W}$ distributions in Figure 24 performed better than others.

Much of the discussion in Section 7. was centered around the scaling constant p_s and its role in unlocking successful amplitude amplification. This is arguably the biggest unknown for the future success of the algorithm. We demonstrated that given an optimal p_s the algorithm can solve for the desired solution, but it is still unclear under what circumstances an experimenter can reliably obtain p_s since it changes from problem to problem. We also showed the degree to which an average p_s could be used, which we believe is a viable application for quantum under certain circumstances, requiring further research. Alternatively, it is possible that an optimal p_s could be found through a learning style algorithm, such as QAOA [55,56] or VQE [57], whereby the results of each attempted amplitude amplification are fed back to a classical optimizer.

Finally, the Traveling Salesman oracle in Section 8. is a theoretical application, but with the highest upside for a quantum speedup ($O(\sqrt{N!})$), relying on future qudit technology for realization. Critically, we neglected to provide an efficient quantum circuit for U_P (an inefficient circuit is easy to construct, but too cumbersome to provide a quantum speedup), which is an open question we are still pursuing. Beyond the TSP, however, we plan to investigate more optimization problems which also naturally give rise to gaussian solution space distributions, making them candidates for amplitude amplification.

Author Contributions: The authors of this work contributed in the following ways: conceptualization and preliminary analysis D.K.; investigation and validation, D.K. and M.C.; software design and data collection, D.K. and S.K.; writing—draft preparation, review, and editing, D.K., S.P. and M.C.; computing resources, L.W.; supervision and project administration, L.W. and P.M.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data and code files that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: We gratefully acknowledge support from the Griffiss Institute. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of AFRL.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. U_P Fidelity Results

Here we present experimental results which demonstrate the viability of implementing U_P on IBM's state-of-the-art qubit architectures 'Casablanca' and 'Lagos' [58]. Because U_P only applies phases (which are undetectable through measurements), each experiment consists of an application of U_P followed by $U_P^\dagger$, ensuring that each experiment has a definitive measurement result for calculating fidelity (the state of all $|0\rangle$'s). Equation (A2) below shows the fidelity metric used.

$$|\Psi\rangle = H^{\otimes L} U_P^\dagger U_P H^{\otimes L} |0^{\otimes L}\rangle \quad (\text{A1})$$

$$f = \langle 0^{\otimes L} | \Psi \rangle \quad (\text{A2})$$

Because of the multiplicative nature of fidelities, the actual fidelity of a single U_P application can be estimated as higher than the values shown in Figure A1. Also note the dramatic decrease in fidelity between experiments $L = 2$ and $L = 3$. This drop off can be explained by revisiting Figure 9, and noting the difference in circuit depth for U_P when using 2 versus 3 qubits. For the special case of $L = 2$, we have $U_P = U_{ij}$ (Equation (15)), while for all other cases U_P requires two sets of U_{ij} operations (Figure 9). This difference in circuit depth explains the high fidelity for $L = 2$ versus $L = 3, 4, 5$.




Figure A1. Fidelity results as defined in Equation (A2), for the case $N = 2$, $L \in 2, 3, 4, 5$, performed on IBM's superconducting qubits.

Appendix B. Step-Varying p_s

In order to compute the maximal P_M values displayed in Figures 23 and 24, we used a classical simulation of the quantum state $|\Psi\rangle$ at each step of the amplitude amplification process in order to determine optimal p_s values. At each step of the algorithm we test a range of p_s values when applying U_P , tracking the distance in amplitude space between the state $|P_{\min}\rangle$ and collective mean, given in Equation (A6). Once a maximal D is found at each step, the corresponding p_s value is stored, the diffusion operator U_s is applied to $|\Psi\rangle$, and the resulting probability P_M for $|P_{\min}\rangle$ is recorded. This process is repeated until the simulation finds a P_M value which is smaller than the previous step, signaling the rebound point of the algorithm.

$$|\Psi\rangle = \sum_k^{N^L} \alpha_k |P_k\rangle \quad (\text{A3})$$

$$\text{Dist}(\alpha, \beta) \equiv \sqrt{\text{real}(\alpha - \beta)^2 + \text{imag}(\alpha - \beta)^2} \quad (\text{A4})$$

$$\alpha_{\text{mean}} = \frac{1}{N^L} \sum_k^{N^L} \alpha_k \quad (\text{A5})$$

$$D = \text{Dist}(\alpha_{\text{mean}}, \alpha_{\text{min}}) \quad (\text{A6})$$

Figure A2 illustrates an example $\mathbb{W}$ distribution, along with three p_s values and their effect on $|\Psi\rangle$ after the first application of U_P . In each $U_P|s\rangle$ amplitude plot, the value of D and p_s are shown, along with a line connecting the locations of α_{min} and α_{mean} .




Figure A2. Illustration of the classical simulation technique used to determine the optimal p_s value at each step by maximizing the distance between $|P_{\text{min}}\rangle$ and the mean point.

References

1. Grover, L.K. A fast quantum mechanical algorithm for database search. *arXiv* **1996**, arXiv:9605043.
2. Boyer, M.; Brassard, G.; Hoyer, P.; Tapp, A. Tight bounds on quantum searching. *Fortschr. Phys.* **1998**, *46*, 493–506. [\[CrossRef\]](#)
3. Bennett, C.H.; Bernstein, E.; Brassard, G.; Vazirani, U. Strengths and Weaknesses of Quantum Computing. *Siam J. Comput.* **1997**, *26*, 1510–1523. [\[CrossRef\]](#)
4. Farhi, E.; Gutmann, S. Analog analogue of a digital quantum computation. *Phys. Rev. A* **1998**, *57*, 2403. [\[CrossRef\]](#)
5. Brassard, G.; Hoyer, P.; Tapp, A. Quantum Counting. In Proceedings of the LNCS 1443: 25th International Colloquium on Automata, Languages, and Programming (ICALP), Aalborg, Denmark, 13–17 July 1998; pp. 820–831.
6. Brassard, G.; Hoyer, P.; Mosca, M.; Tapp, A. Quantum Amplitude Amplification and Estimation. *Ams Contemp. Math.* **2002**, *305*, 53–74.
7. Childs, A.M.; Goldstone, J. Spatial search by quantum walk. *Phys. Rev. A* **2004**, *70*, 022314. [\[CrossRef\]](#)
8. Ambainis, A. Variable time amplitude amplification and a faster quantum algorithm for solving systems of linear equations. *arXiv* **2010**, arXiv:1010.4458.
9. Singleton, R.L., Jr.; Rogers, M.L.; Ostby, D.L. Grover's Algorithm with Diffusion and Amplitude Steering. *arXiv* **2021**, arXiv:2110.11163.
10. Kwon, H.; Bae, J. Quantum amplitude-amplification operators. *Phys. Rev. A* **2021**, *104*, 062438. [\[CrossRef\]](#)
11. Lloyd, S. Quantum search without entanglement. *Phys. Rev. A* **1999**, *61*, 010301. [\[CrossRef\]](#)
12. Viamontes, G.F.; Markov, I.L.; Hayes, J.P. Is Quantum Search Practical? *arXiv* **2004**, arXiv:0405001.
13. Regev, O.; Schiff, L. Impossibility of a Quantum Speed-up with a Faulty Oracle. *arXiv* **2012**, arXiv:1202.1027.
14. Seidel, R.; Becker, C.K.-U.; Bock, S.; Tcholtchev, N.; Gheorge-Pop, I.-D.; Hauswirth, M. Automatic Generation of Grover Quantum Oracles for Arbitrary Data Structures. *arXiv* **2021**, arXiv:2110.07545.
15. Nielsen, M.A.; Chuang, I.L. *Quantum Computation and Quantum Information*; Cambridge University Press: Cambridge, UK, 2000; p. 249.

16. Long, G.L.; Zhang, W.L.; Li, Y.S.; Niu, L. Arbitrary Phase Rotation of the Marked State Cannot Be Used for Grover's Quantum Search Algorithm. *Commun. Theor. Phys.* **1999**, *32*, 335.
17. Long, G.L.; Li, Y.S.; Zhang, W.L.; Niu, L. Phase matching in quantum searching. *Phys. Lett. A* **1999**, *262*, 27–34. [[CrossRef](#)]
18. Hoyer, P. Arbitrary phases in quantum amplitude amplification. *Phys. Rev. A* **2000**, *62*, 052304. [[CrossRef](#)]
19. Younes, A. Towards More Reliable Fixed Phase Quantum Search Algorithm. *Appl. Math. Inf. Sci.* **2013**, *1*, 10. [[CrossRef](#)]
20. Li, T.; Bao, W.-S.; Lin, W.-Q.; Zhang, H.; Fu, X.-Q. Quantum Search Algorithm Based on Multi-Phase. *Chinese Phys. Lett.* **2014**, *31*, 050301. [[CrossRef](#)]
21. Guo, Y.; Shi, W.; Wang, Y.; Hu, J. Q-Learning-Based Adjustable Fixed-Phase Quantum Grover Search Algorithm. *J. Phys. Soc. Jpn.* **2017**, *86*, 024006. [[CrossRef](#)]
22. Song, P.H.; Kim, I. Computational leakage: Grover's algorithm with imperfections. *Eur. Phys. J. D* **2003**, *23*, 299–303. [[CrossRef](#)]
23. Pomeransky, A.A.; Zhiron, O.V.; Shepelyansky, D.L. Phase diagram for the Grover algorithm with static imperfections. *Eur. Phys. J. D-At. Mol. Opt. Plasma Phys.* **2004**, *31*, 131–135. [[CrossRef](#)]
24. Janmark, J.; Meyer, D.A.; Wong, T.G. Global Symmetry is Unnecessary for Fast Quantum Search. *Phys. Rev. Lett.* **2014**, *112*, 210502. [[CrossRef](#)]
25. Gutin, G.; Punnen, A.P. *The Traveling Salesman Problem and Its Variations*; Springer: New York, NY, USA, 2007.
26. Srinivasan, K.; Satyajit, S.; Behera, B.K.; Panigrahi, P.K. Efficient quantum algorithm for solving travelling salesman problem: An IBM quantum experience. *arXiv* **2018**, arXiv:1805.10928.
27. Moylett, D.J.; Linden, N.; Montanaro, A. Quantum speedup of the traveling-salesman problem for bounded-degree graphs. *Phys. Rev. A* **2017**, *95*, 032323. [[CrossRef](#)]
28. Martoňák, R.; Santoro, G.E.; Tosatti, E. Quantum annealing of the traveling-salesman problem. *Phys. Rev. E* **2004**, *70*, 057701. [[CrossRef](#)]
29. Warren, R.H. Adapting the traveling salesman problem to an adiabatic quantum computer. *Quantum Inf. Process.* **2013**, *12*, 1781–1785. [[CrossRef](#)]
30. Warren, R.H. Solving the traveling salesman problem on a quantum annealer. *SN Appl. Sci.* **2020**, *2*, 75. [[CrossRef](#)]
31. Chen, H.; Kong, X.; Chong, B.; Qin, G.; Zhou, X.; Peng, X.; Du, J. Experimental demonstration of a quantum annealing algorithm for the traveling salesman problem in a nuclear-magnetic-resonance quantum simulator. *Phys. Rev. A* **2011**, *83*, 032314. [[CrossRef](#)]
32. Bang, J.; Yoo, S.; Lim, J.; Ryu, J.; Lee, C.; Lee, J. Quantum heuristic algorithm for traveling salesman problem. *J. Korean Phys. Soc.* **2012**, *61*, 1944. [[CrossRef](#)]
33. Kues, M.; Reimer, C.; Roztocki, P.; Cortés, L.R.; Sciara, S.; Wetzel, B.; Zhang, Y.; Cino, A.; Chu, S.T.; Little, B.E.; et al. On-chip generation of high-dimensional entangled quantum states and their coherent control. *Nature* **2017**, *546*, 622–626. [[CrossRef](#)]
34. Low, P.J.; White, B.M.; Cox, A.A.; Day, M.L.; Senko, C. Practical trapped-ion protocols for universal qudit-based quantum computing. *Phys. Rev. Res.* **2020**, *2*, 033128. [[CrossRef](#)]
35. Yurtalan, M.A.; Shi, J.; Kononenko, M.; Lupascu, A.; Ashhab, S. Implementation of a Walsh-Hadamard gate in a superconducting qutrit. *Phys. Rev. Lett.* **2020**, *125*, 180504. [[CrossRef](#)] [[PubMed](#)]
36. Lu, H.-H.; Hu, Z.; Alshaykh, M.S.; Moore, A.J.; Wang, Y.; Imany, P.; Weiner, A.M.; Kais, S. Quantum Phase Estimation with Time-Frequency Qudits in a Single Photon. *Adv. Quantum Technol.* **2019**, *3*, 1900074. [[CrossRef](#)]
37. Niu, M.Y.; Chuang, I.L.; Shapiro, J.H. Qudit-Basis Universal Quantum Computation Using χ^2 Interactions. *Phys. Rev. Lett.* **2018**, *120*, 160502. [[CrossRef](#)] [[PubMed](#)]
38. Luo, M.-X.; Wang, X.-J. Universal quantum computation with qudits. *Sci. China Phys. Mech. Astron.* **2014**, *57*, 1712–1717. [[CrossRef](#)]
39. Li, B.; Yu, Z.-H.; Fei, S.-M. Geometry of Quantum Computation with Qutrits. *Sci. Rep.* **2013**, *3*, 2594. [[CrossRef](#)]
40. Lanyon, B.P.; Barbieri, M.; Almeida, M.P.; Jennewein, T.; Ralph, T.C.; Resch, K.J.; Pryde, G.J.; O'Brien, J.L.; Gilchrist, A.; White, A.G. Quantum computing using shortcuts through higher dimensions. *Nat. Phys.* **2009**, *5*, 134–140. [[CrossRef](#)]
41. Gokhale, P.; Baker, J.M.; Duckering, C.; Brown, N.C.; Brown, K.R.; Chong, F.T. Asymptotic improvements to quantum circuits via qutrits. In Proceedings of the ISCA '19: 46th International Symposium on Computer Architecture, Phoenix, AZ, USA, 22–26 June 2019; pp. 554–566.
42. Khan, F.S.; Perkowski, M. Synthesis of multi-qudit Hybrid and d-valued Quantum Logic Circuits by Decomposition. *Theor. Comput. Sci.* **2006**, *367*, 336–346. [[CrossRef](#)]
43. Muthukrishnan, A.; Stroud, C.R., Jr. Multi-valued Logic Gates for Quantum Computation. *Phys. Rev. A* **2000**, *62*, 052309. [[CrossRef](#)]
44. Daboul, J.; Wang, X.; Sanders, B.C. Quantum gates on hybrid qudits. *J. Phys. A Math. Gen.* **2003**, *36*, 2525–2536. [[CrossRef](#)]
45. Blok, M.S.; Ramasesh, V.V.; Schuster, T.; O'Brien, K.; Kreikebaum, J.M.; Dahlen, D.; Morvan, A.; Yoshida, B.; Yao, N.Y.; Siddiqi, I. Quantum Information Scrambling on a Superconducting Qutrit Processor. *Phys. Rev. X* **2021**, *11*, 021010. [[CrossRef](#)]
46. Hu, X.-M.; Zhang, C.; Liu, B.-H.; Cai, Y.; Ye, X.-J.; Guo, Y.; Xing, W.-B.; Huang, C.-X.; Huang, Y.-F.; Li, C.-F.; et al. Experimental High-Dimensional Quantum Teleportation. *Phys. Rev. Lett.* **2020**, *125*, 230501 [[CrossRef](#)]
47. Laplace, P.S. Mémoire sur les approximations des formules qui sont fonctions de très grands nombres et sur leur application aux probabilités. In *Mémoires de l'Académie Royale des Sciences de Paris*; Baudouin: Brussels, Belgium, 1810; Volume 10.
48. Bernoulli, J. *Ars Conjectandi*; Thurnisiorum: Basileae, Switzerland, 1713.

49. Gauss, C.F. *Theoria Motus Corporum Coelestium in Sectionibus Conicis Solem Ambientium*; Friedrich Perthes: Hamburg, Geramey; I.H. Besser: Hamburg, Geramey, 1809.
50. Satoh, T.; Ohkura, Y.; Meter, R.V. Subdivided Phase Oracle for NISQ Search Algorithms. *IEEE Trans. Quantum Eng.* **2020**, *1*, 1–15. [[CrossRef](#)]
51. Benchasattabuse, N.; Satoh, T.; Hajdušek, M.; Meter, R.V. Amplitude Amplification for Optimization via Subdivided Phase Oracle. *arXiv* **2022**, arXiv:2205.00602.
52. Shyamsundar, P. Non-Boolean Quantum Amplitude Amplification and Quantum Mean Estimation. *arXiv* **2021**, arXiv:2102.04975.
53. Koch, D.; Wessing, L.; Alsing, P.M. Introduction to Coding Quantum Algorithms: A Tutorial Series Using Qiskit. *arXiv* **2019**, arXiv:1903.04359.
54. Wang, Y.; Hu, Z.; Sanders, B.C.; Kais, S. Qudits and High-Dimensional Quantum Computing. *Front. Phys.* **2020**, *10*, 589504. [[CrossRef](#)]
55. Farhi, E.; Goldstone, J.; Gutmann, S. A Quantum Approximate Optimization Algorithm. *arXiv* **2014**, arXiv:1411.4028.
56. Hadfield, S.; Wang, Z.; O’Gorman, B.; Rieffel, E.G.; Venturelli, D.; Biswas, R. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms* **2019**, *12*, 34. [[CrossRef](#)]
57. Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, Z.-Q.; Love, P.J.; Aspuru-Guzik, A.; O’Brien, J.L. A variational eigenvalue solver on a quantum processor. *Nat. Commun.* **2014**, *5*, 4213. [[CrossRef](#)]
58. IBM 7-Qubit Casablanca and Lagos Architectures. Available online: <https://quantum-computing.ibm.com> (accessed on 10 April 2021).