

Quantum Algorithms for Connectivity and Related Problems

MICHAEL JARRET¹, STACEY JEFFERY², SHELBY KIMMEL³, AND ALVARO PIEDRAFITA²

¹*Perimeter Institute*

²*Qusoft, CWI*

³*Middlebury College*

Abstract

An important family of span programs, *st*-connectivity span programs, have been used to design quantum algorithms in various contexts, including a number of graph problems and formula evaluation problems. The complexity of the resulting algorithms depends on the largest *positive witness size* of any 1-input, and the largest *negative witness size* of any 0-input. Belovs and Reichardt first showed that the positive witness size is exactly characterized by the effective resistance of the input graph, but only rough upper bounds were known previously on the negative witness size. We show that the negative witness size in an *st*-connectivity span program is exactly characterized by the *capacitance* of the input graph. This gives a tight analysis for algorithms based on *st*-connectivity span programs on any set of inputs.

We use this analysis to give a new quantum algorithm for estimating the capacitance of a graph. We also describe a new quantum algorithm for deciding if a graph is connected, which improves the previous best quantum algorithm for this problem if we're promised that either the graph has at least $\kappa > 1$ components, or the graph is connected and has small *average resistance*, which is upper bounded by the diameter. We also give an alternative algorithm for deciding if a graph is connected that can be better than our first algorithm when the maximum degree is small. Finally, using ideas from our second connectivity algorithm, we give an algorithm for estimating the *algebraic connectivity* of a graph, the second largest eigenvalue of the Laplacian.

1 Introduction

Span programs are an algebraic model of computation first developed by Karchmer and Wigderson [KW93] to study classical logspace complexity, and introduced to the study of quantum algorithms by Reichardt and Špálek [RŠ12]. In [Rei11, Rei09], Reichardt used the concept of span programs to prove that the general adversary bound gives a tight lower bound on the quantum query complexity of any given decision problem, thus showing the deep connection between span programs and quantum query algorithms.

Given a span program, a generic transformation compiles it into a quantum algorithm, whose query complexity is analyzed by taking the geometric mean of two quantities: the largest *positive witness size* of any 1-input; and the largest *negative witness size* of any 0-input. Thus, in order to analyze the query complexity of an algorithm obtained in this way, it is necessary to characterize, or at least upper bound, these quantities. Moreover, there is always a span program based algorithm with asymptotically optimal quantum query complexity [Rei11, Rei09].

The relationship between quantum query algorithms and span programs is potentially a powerful tool, but this correspondence alone is not a recipe for finding such an algorithm, and actually producing an optimal (or even good) span program for a given problem is generally difficult. Despite this difficulty, a number of span programs have been found for important problems such as *k*-distinctness [Bel12a], formula evaluation [RŠ12, Rei10], and *st*-connectivity [BR12]. The latter span program is of particular importance, as it has been applied to a number of graph problems [CMB16], to generic formula evaluation problems [JK17], and underlies the learning graph framework [Bel12b]. The *st*-connectivity based algorithms are also of interest because, unlike with generic span program algorithms, it is often possible to analyze not only query complexity, but also the time complexity.

While span program algorithms are universal for quantum query algorithms, it can also be fruitful to analyze the unitaries used in these algorithms in ways that are different from how they appear in the standard

span program algorithm. For example, Ref. [IJ16] presents an algorithm to estimate span program witness sizes based on techniques from the standard span program algorithm. We will take a similar approach in this paper, deriving new algorithms based on unitaries from the span program algorithm for st -connectivity.

The problems of st -connectivity and connectivity will be considered in this paper. For a family of undirected graphs G on N edges, for $N \in \mathbb{N}$, and vertex set containing s and t , the problem st -CONN $_G$ is the following: Given $x \in \{0, 1\}^{E(G)}$, decide if there is a path from s to t in $G(x)$, where $G(x)$ is the subgraph of G obtained by including an edge e if $x_e = 1$ ¹. Similarly, the problem of CONN $_G$ is the following: Given $x \in \{0, 1\}^{E(G)}$, determine if every vertex in $G(x)$ is connected to every other vertex in $G(x)$.

1.1 Contributions

In all of the following problems, we assume we have access to a black box unitary O_x that tells us about the presence or absence of edges in a graph G , and the query complexity refers to the number of uses of O_x to solve a problem with high probability.

Characterizing the negative witness of st -connectivity span programs An important span program for solving solving st -connectivity in subgraphs of complete graphs without edge weights was presented in Ref. [BR12]. When generalized to subgraphs of arbitrary weighted graphs², this span program was applied to develop the learning graph framework [Bel12b], and quantum algorithms for formula evaluation [JK17].

In Ref. [BR12], Belovs and Reichardt gave a tight characterization of the positive witness size as the *effective resistance* between s and t in the input graph. However, for the negative witness size of an input in which s and t are not connected, they gave only a rough upper bound of n^2 , which is refined in [Bel12b] to the total weight of an st -cut, which is still not a tight bound. In Ref. [JK17], it was shown that when the parent graph is planar and s and t are on the same face, the negative witness can be characterized exactly as the effective resistance of a graph related to the planar dual of the parent graph. In particular, this allowed for a tight analysis of st -connectivity-based span program algorithms for formula evaluation in [JK17].

In this work, we bring the story to its conclusion, by showing that the negative witness size of the st -connectivity span program is exactly characterized by the *effective capacitance* of the input graph (Theorem 17). At a high-level, this well-studied electrical network quantity is a measure of the potential difference that the network could store between the component containing s and the component containing t . The more, shorter paths between these two components in the graph $G \setminus G(x)$, the greater the capacitance. This characterization tells us that quantum algorithms can quickly decide st -connectivity on graphs that are promised to have either small effective resistance or small effective capacitance.

Quantum algorithm for estimating st -capacitance As one immediate application, we get a new quantum algorithm for estimating the capacitance of an input graph $G(x)$ to multiplicative error ϵ , with query complexity $\tilde{O}(\epsilon^{-3/2} \sqrt{C_{s,t}(G(x)) p})$, where $C_{s,t}(G(x))$ is the st -capacitance of $G(x)$, and p is the length of the longest self-avoiding st -path in G (Corollary 20). This follows from Ref. [IJ16], which shows that given any span program, there is a quantum algorithm that, on input x , outputs an estimate of the witness size of x .

New quantum algorithm for connectivity We use this tighter analysis of the negative witness to analyze a new algorithm for graph connectivity. This problem was first studied in the context of quantum algorithms by Dürr, Høyer, Heiligman and Mhalla [DHHM06], who gave an optimal $\tilde{O}(n^{3/2})$ upper bound on the time complexity. An optimal span-program-based quantum algorithm was later presented by Āriņš [Āri16], whose algorithm also uses only $O(\log n)$ space.

Since a graph is connected if and only if every pair of vertices $\{u, v\}$ are connected, we propose an algorithm that uses the technique of [NTS95, JK17] to convert the conjunction of $\binom{n}{2}$ st -connectivity span programs into a single st -connectivity span program: take $n(n-1)/2$ copies of $G(x)$, one for each pair of distinct vertices $\{u, v\}$ with $u < v$, and call u the source and v the sink of this graph. Connect these graphs

¹We consider more complicated ways of associating edges with input variables in Section 2.2, but the basic idea is captured by this simpler picture.

²Assigning positive weights to edges does not change whether or not a graph is st -connected, but rather, the weights should be considered as parameters of the span program that affect its complexities.

in series, in any order, by identifying the sink of one to the source of the next. Call the source of the first graph s , and the sink of the last graph t . See Figure 1 for an example when G is a triangle. In this way we have created a graph (which we denote $\mathcal{G}(x)$) that is st -connected if and only if $G(x)$ is connected. In other words, for any $x \in \{0, 1\}^{E(G)}$, $\text{CONN}_G(x) = st\text{-CONN}_{\mathcal{G}}(x)$.




Figure 1: The graph \mathcal{G} is st -connected if and only if G is connected.

For a graph $G(x)$, define the average resistance as $R_{\text{avg}}(G(x)) = \frac{1}{n(n-1)} \sum_{s,t:s \neq t} R_{s,t}(G(x))$. We consider the case where we are promised that if $G(x)$ is connected, then $R_{\text{avg}}(G(x)) \leq R$, and if $G(x)$ is not connected it has at least $\kappa > 1$ components. By analyzing the effective resistance and capacitance of \mathcal{G} , we show that when G is a subgraph of a complete graph, meaning it has no multi-edges, CONN_G under this promise can be solved in query complexity $O(n\sqrt{R/\kappa})$ (Theorem 25), and time complexity $\tilde{O}(n\sqrt{R/\kappa}U)$, where U is the cost of implementing one step of a quantum walk on G (Corollary 26). For the case when G has multi-edges, we get an upper bound of $O(n^{3/4}\sqrt{Rd_{\max}(G)}/\kappa^{1/4})$ on the query complexity, where $d_{\max}(G)$ is the maximum degree of any vertex in the graph.

In the worst case, when $R = n$ and $\kappa = 2$, our algorithm for the case when G is a subgraph of the complete graph achieves the optimal upper bound of $\tilde{O}(n^{3/2})$. Like the algorithm of Ref. [Ari16], our algorithm uses only $O(\log n)$ space. It is also the first connectivity algorithm that applies to the case where G is not necessarily the complete graph, although the other algorithms can likely be adapted to the more general case.

The algorithm of Ariš can be seen as similar to ours, except that rather than connecting copies of $G(x)$ for each $\{u, v\}$ pair, his algorithm only considers pairs $\{1, v\}$ for $v \neq 1$. In contrast, our algorithm is symmetric in the vertex set, which makes a detailed analysis more natural.

Alternative quantum algorithms for connectivity In Section 5, we present an alternative approach to deciding graph connectivity. It is based on phase estimation of a particular unitary that is also used in the st -connectivity span program, but applied to a different initial state.

We first show that the quantum query complexity of deciding CONN_G is $O(\sqrt{nd_{\max}(G)/(\kappa\lambda)})$, when we're promised that if $G(x)$ is connected, the second smallest eigenvalue of the Laplacian of $G(x)$, $\lambda_2(G(x))$, is at least λ , and otherwise, $G(x)$ has at least $\kappa > 1$ connected components (Corollary 27). In the unweighted worst case, $\lambda_2(G(x)) \geq 2/n^2$ and $d_{\max} = n - 1$, which gives a sub-optimal $O(n^2)$ algorithm. However, for some classes of inputs, this algorithm performs better than our first algorithm. Neglecting constants, and using the fact that $R_{\text{avg}}(G(x)) \leq 1/\lambda_2(G(x))$, our first algorithm has query complexity (in the case of no multi-edges)

$$T_1 = n\sqrt{R/\kappa} \leq n/\sqrt{\kappa\lambda} \quad (1)$$

whereas our second algorithm has query complexity

$$T_2 = \sqrt{nd_{\max}(G)/\kappa\lambda}. \quad (2)$$

When G is a complete graph, our second algorithm can only be worse, since in that case $d_{\max}(G) = n - 1$. However, when G is a Boolean hypercube so that $d_{\max} = \log n$, our second algorithm may be significantly better.

However, thus far we have only been considering the query complexity of our second algorithm. This algorithm also requires an initial state of a particular form, and while this state is independent of the input, it may generally not be time efficient to produce such a state. We are able to give time-efficient versions of our second algorithm in two contexts.

First, in Theorem 28, we show that for any G , under the promise that if $G(x)$ is connected, then $\lambda_2(G(x)) \geq \lambda$, and otherwise $G(x)$ has at least $\kappa > 1$ connected components, we can solve CONN_G in time complexity

$$\tilde{O} \left(\sqrt{\frac{nd_{\text{avg}}(G)}{\kappa\lambda_2(G)}} \left(S + \sqrt{\frac{d_{\text{max}}(G)}{\lambda}} U \right) \right), \quad (3)$$

where U is the complexity of implementing a step of a quantum walk on G , S is the cost generating a quantum state corresponding to the stationary distribution of a random walk on G , and $d_{\text{avg}}(G)$ is the average degree of the vertices of G . This time complexity might generally be significantly worse than the query complexity, but has the advantage of applying to all G .

Second, in Theorem 29, we give an improved time complexity for the case where G is a Cayley graph of degree d , showing an upper bound of

$$\tilde{O} \left(\sqrt{\frac{nd}{\kappa\lambda}} U + \sqrt{\frac{nd}{\kappa\lambda_2(G)}} \Lambda \right), \quad (4)$$

where, as above, U is the cost of implementing a step of a quantum walk on G , and Λ is the cost of computing the eigenvalues of G ; that is, given $g \in V(G)$, computing λ_g , the corresponding eigenvalue. For example, this gives an upper bound of $\tilde{O}(n/\sqrt{\lambda\kappa})$ when G is a complete graph, and $\tilde{O}(\sqrt{n/(\lambda\kappa)})$ when G is a Boolean hypercube.

We remark that our alternative connectivity algorithms apply for any choice of edge weights on G : $d_{\text{avg}}(G)$ and $d_{\text{max}}(G)$ should be interpreted as the average and maximum *weighted* degrees in G , and $\lambda_2(G)$ and $\lambda_2(G(x))$ the second-smallest eigenvalue of the *weighted* Laplacian of G and $G(x)$ respectively. The choice of weights may also impact the costs U and S .

Estimating the algebraic connectivity We give an algorithm to estimate the *algebraic connectivity* of $G(x)$, $\lambda_2(G(x))$, when G is a complete graph. The algebraic connectivity is closely related to the inverse of the mixing time, which is known to be small for many interesting families of graphs such as expander graphs. We give a protocol that with probability at least $2/3$ outputs an estimate of $\lambda_2(G(x))$ up to multiplicative error ε in time complexity $\tilde{O} \left(\frac{1}{\varepsilon} \frac{n}{\sqrt{\lambda_2(G(x))}} \right)$ (Theorem 30).

1.2 Open Problems

Our work suggests several directions for new research. Since *st*-connectivity is fairly ubiquitous, it seems that our approach may, in turn, help analyze applications of *st*-connectivity. Additionally, we provide two algorithms for deciding connectivity, in Section 4.2 and Section 5. At least naively, it seems like our two algorithms are incomparable, even though they are based on similar unitaries. It would be worthwhile to understand whether the two approaches are fundamentally different. Another open question is to determine how to set the weights of edges in our graphs; these weights can have a significant effect on query complexity. Finally, it would be interesting to see whether one can extend our algorithm for estimating algebraic connectivity to accept more general parent graphs than the complete graph.

1.3 Organization

The rest of this paper is organized as follows. In Section 2, we introduce the necessary background on which we build our results, including basic notation (Section 2.1), graph theory (Section 2.2), quantum algorithms (Section 2.3) and span programs (Section 2.4). In Section 3, we show that the negative witness size of an *st*-connectivity span program is the effective capacitance of the graph, then in Section 4, we give two applications of this observation: The first is a quantum algorithm for estimating the effective capacitance of a graph (Section 4.1); and the second is our first quantum algorithm for deciding connectivity, as a composition of *st*-connectivity span programs (Section 4.2). Finally, in Section 5, we give our second algorithm for deciding connectivity, based on estimating the second-smallest eigenvalue of the Laplacian, and also give an algorithm for estimating the algebraic connectivity of a graph.

2 Preliminaries

2.1 Linear Algebra Notation

For a subspace V of some inner product space, we let Π_V denote the orthogonal projector onto V .

For a linear operator A , we will let $\sigma_{\min}(A)$ denote its smallest non-zero singular value, and $\sigma_{\max}(A)$ its largest singular value. We let $\ker A$ denote the kernel of A , $\text{row}(A)$ denote the rowspace of A , and $\text{col}(A)$ the columnspace of A . We let A^+ denote the Moore-Penrose pseudoinverse of A . If A has singular value decomposition $A = \sum_i \sigma_i |a_i\rangle\langle b_i|$, (for left singular vectors $\{|a_i\rangle\}$ and right singular vectors $\{|b_i\rangle\}$) then $A^+ = \sum_i 1/\sigma_i |b_i\rangle\langle a_i|$. Then we have $AA^+ = \Pi_{\text{col}(A)}$ and $A^+A = \Pi_{\text{row}(A)}$.

For a unitary U with eigenvalues $e^{i\theta_1}, \dots, e^{i\theta_N}$ for $\theta_1, \dots, \theta_N \in (-\pi, \pi]$, let $\Delta(U) = \min\{|\theta_i| : \theta_i \neq 0\}$ denote the *phase gap* of U .

2.2 Graph Theory

Multigraphs We will consider multigraphs, which may have multiple edges between a pair of vertices. Thus, to each edge, we associate a unique identifying label ℓ . We refer to each edge in the graph using its endpoints and the label ℓ , as, for example: $(\{u, v\}, \ell)$. The label ℓ uniquely specifies the edge, but we include the endpoints for convenience. Let $\vec{E}(G) = \{(u, v, \ell) : (\{u, v\}, \ell) \in E(G)\}$ be the directed edges of G . Furthermore, for any set of edges E , we let $\vec{E} = \{(u, v, \ell) : (\{u, v\}, \ell) \in E\}$ represent the corresponding set of directed edges. We will sometimes write (u, v, ℓ) for an undirected edge, but when talking about undirected edges, we have $(u, v, \ell) = (v, u, \ell)$.

We will be concerned with certain subgraphs of a graph G , associated with bit strings of length N . We denote by $G(x)$ the subgraph associated with the string $x \in \{0, 1\}^N$. In particular, each edge in G is associated with a variable x_i or its negation \bar{x}_i , called a *literal*, and is included in $G(x)$ if and only if the associated literal evaluates to 1. Here x_i is the i^{th} bit of x . For example, an edge (u, v, ℓ) associated with the literal \bar{x}_i is in $G(x)$ if and only if x_i takes value 0, as in Figure 2. Precisely how this association of edges and literals is chosen depends on the problem of interest, so we will leave the description implicit, and often assume for simplicity that there is a one-to-one mapping between the edges and positive literals.




Figure 2: Example of how each edge in G is associated with a bit of x and a value of that bit. For edges labeled by x_i we include the edge in $G(x)$ if $x_i = 1$, while for edges labeled by \bar{x}_i , we include the edge in $G(x)$ if $x_i = 0$.

Networks A *network* $\mathcal{N} = (G, c)$ consists of a graph G combined with a positive real-valued *weight* function $c : E(G) \rightarrow \mathbb{R}^+$. Since c is a map on undirected edges, we can easily extend it to a map on directed edges such that $c(u, v, \ell) = c(v, u, \ell)$, and we overload our notation accordingly. We will often assume that some c is implicit for a graph G and let

$$\mathcal{A}_G = \sum_{(u, v, \ell) \in E(G)} c(u, v, \ell) (|u\rangle\langle v| + |v\rangle\langle u|) \quad (5)$$

denote its weighted adjacency matrix. Note that \mathcal{A}_G only depends on the total weight of edges from u to v , and is independent of the number of edges across which this weight is distributed. Let $d_G(u) =$

$\sum_{v,\ell:(u,v,\ell)\in E(G)} c(u,v,\ell)$ denote the weighted degree of u in G , under the implicit weight function c , and let $d_{\max}(G) = \max_{u\in V(G)} d_G(u)$. Let

$$\mathcal{D}_G = \sum_{u\in V(G)} d_G(u)|u\rangle\langle u| \quad (6)$$

denote the weighted degree matrix, and let

$$L_G = \mathcal{D}_G - \mathcal{A}_G \quad (7)$$

denote the Laplacian of G . The Laplacian is always positive semidefinite, so its eigenvalues are real and non-negative. For $|\mu\rangle = \sum_{u\in V(G)} |u\rangle$, it is always the case that $L_G|\mu\rangle = 0$, so the smallest eigenvalue of L_G is 0. Let $\lambda_2(G)$ denote the second smallest eigenvalue of L_G , including multiplicity. This value is called the *algebraic connectivity* or the *Fiedler value* of G , and it is non-zero if and only if G is connected.

Electric networks Consider a graph G with specially labeled vertices s and t that are connected in G . One can consider a fluid that enters a graph G at s , flows along the edges of the graph, and exits the graph at t . The fluid can spread out along some number of the st -paths in G . An st -flow is any linear combination of st -paths. More precisely:

Definition 1 (Unit st -flow). *Let G be an undirected graph with $s, t \in V(G)$, and s and t connected. Then a unit st -flow on G is a function $\theta : \vec{E}(G) \rightarrow \mathbb{R}$ such that:*

1. For all $(u, v, \ell) \in \vec{E}(G)$, $\theta(u, v, \ell) = -\theta(v, u, \ell)$;
2. $\sum_{v,\ell:(s,v,\ell)\in\vec{E}(G)} \theta(s, v, \ell) = \sum_{v,\ell:(v,t,\ell)\in\vec{E}(G)} \theta(v, t, \ell) = 1$; and
3. for all $u \in V(G) \setminus \{s, t\}$, $\sum_{v,\ell:(u,v,\ell)\in\vec{E}(G)} \theta(u, v, \ell) = 0$.

Definition 2 (Unit Flow Energy). *Given a graph G with implicit weighting c and a unit st -flow θ on $G(x)$, the unit flow energy of θ on $E' \subseteq E(G(x))$, is*

$$J_{E'}(\theta) = \frac{1}{2} \sum_{e\in E'} \frac{\theta(e)^2}{c(e)}. \quad (8)$$

Definition 3 (Effective resistance). *Let G be a graph with implicit weighting c and $s, t \in V(G)$. If s and t are connected in $G(x)$, the effective resistance of $G(x)$ between s and t is $R_{s,t}(G(x)) = \min_{\theta} J_{E(G(x))}(\theta)$, where θ runs over all unit st -unit flows of $G(x)$. If s and t are not connected in $G(x)$, $R_{s,t}(G(x)) = \infty$.*

Intuitively, $R_{s,t}$ characterizes “how connected” the vertices s and t are in a network. The more, shorter paths connecting s and t , and the more weight on those paths, the smaller the effective resistance.

The effective resistance has many applications. For example, $R_{s,t}(G) \left(\sum_{e\in E(G)} c(e) \right)$ is equal to the *commute time* between s and t , or the expected time a random walker starting from s takes to reach t and then return to s [CRR⁺96]. If $\mathcal{N} = (G, c)$ models an electrical network in which each edge e of G is a $1/c(e)$ -valued resistor and a potential difference is applied between s and t , then $R_{s,t}(\mathcal{N})$ corresponds to the resistance of the network, which determines the ratio of current to voltage in the circuit (see [DS84]). Thus, the values $c(e)$ can be interpreted as conductances.

For a connected graph G , we can define the *average resistance* by:

$$R_{\text{avg}}(G) := \frac{1}{n(n-1)} \sum_{s,t\in V:s\neq t} R_{s,t}(G).$$

Now that we have a measure of the connectedness of s and t in a graph G , we next introduce a measure of how disconnected s and t are, in the case that we are considering a subgraph $G(x)$ of G where s and t are not connected.

Definition 4 (Unit st -potential). Let G be an undirected weighted graph with $s, t \in V(G)$, and s and t connected. For $G(x)$ such that s and t are not connected, a unit st -potential on $G(x)$ is a function $\mathcal{V} : V(G) \rightarrow \mathbb{R}^+$ such that $\mathcal{V}(s) = 1$ and $\mathcal{V}(t) = 0$ and $\mathcal{V}(u) = \mathcal{V}(v)$ if $(u, v, \ell) \in E(G(x))$.

Note that this is a different definition from the typical potential function. Usually, if we have a flow from a vertex s to a vertex t , we define the potential difference between u and v for an edge (u, v, ℓ) to be the amount of flow across that edge divided by the weight of the edge. In our definition, the potential difference across all edges in $E(G(x))$ is zero, and we have potential difference across edges that are in $E(G) \setminus E(G(x))$.

A unit st -potential is a witness of the disconnectedness of s and t in $G(x)$, in the sense that it is a generalization of the notion of an st -cut. (An st -cut is a unit potential that only takes values 0 and 1.)

Definition 5 (Unit Potential Energy). Given a graph G with implicit weighting c and a unit st -potential \mathcal{V} on $G(x)$, the unit potential energy of \mathcal{V} on $E' \subseteq E(G)$ is defined

$$\mathcal{J}_{E'}(\mathcal{V}) = \frac{1}{2} \sum_{(u,v,\ell) \in \vec{E}'} (\mathcal{V}(u) - \mathcal{V}(v))^2 c(u, v, \ell). \quad (9)$$

Definition 6 (Effective capacitance). Let G be a graph with implicit weighting c and $s, t \in V(G)$. If s and t are not connected in $G(x)$, the effective capacitance between s and t of $G(x)$ is $C_{s,t}(G(x)) = \min_{\mathcal{V}} \mathcal{J}_{E(G)}(\mathcal{V})$, where \mathcal{V} runs over all unit st -potentials on $G(x)$. If s and t are connected, $C_{s,t}(G(x)) = \infty$.

In physics, capacitance is a measure of how well a system stores electric charge. The simplest capacitor is a set of separated metal plates at a fixed distance. To see how a capacitor works, we imagine one terminal of a battery attached to each plate. The battery acts as a sort of pump that moves negative charges from one plate to the other, against their natural tendencies. Because like charges repel, as the plates become increasingly polarized, it requires more energy to move additional charges. In other words, as additional charge accumulates, the voltage difference between the plates increases and, correspondingly, the voltage that it takes to overcome the differential increases. At equilibrium, the voltage difference between the plates will be equal to the voltage difference of the terminals of the battery. The ratio of the amount of charge moved to the voltage difference created between the plates is a property of the capacitor itself and depends only on its geometry. This ratio is called its effective capacitance.

Now consider a capacitor that corresponds to the graph $G(x)$ in which a 0-resistance wire is connected between vertices whenever there is an edge in $G(x)$, and a $c(e)$ -unit capacitor is connected between vertices whenever there is an edge $e \in E(G) \setminus E(G(x))$. If s and t are not connected in $G(x)$, it is as though s and t are on separate “plates” (with some complicated geometry) that can accumulate charge relative to each other. Then the effective capacitance given in Definition 6 is precisely the ratio of charge (accumulated on the plates corresponding to s and t) to voltage (on those plates) that is achieved when electrical energy is stored in this configuration. We make the connection between Definition 6 and the standard definition of effective capacitance, as well as effective conductance, more explicitly in Appendix A.

Definitions 3 and 6 may seem unwieldy for actually calculating the effective resistance and effective capacitance, so we now recall that when calculating effective resistance, $R_{s,t}$ (respectively effective capacitance $C_{s,t}$), one can use the rule that for edges in series, or more generally, graphs connected in series, resistances add (resp. inverse capacitances add). Edges in parallel, or more generally, graphs connected in parallel, follow the rule that inverse resistances add (resp. capacitances add). That is:

Claim 7. Let two networks (G_1, c_1) and (G_2, c_2) each have connected nodes s and t . Let $G(x_1)$ and $G(x_2)$ be subgraphs of G_1 and G_2 respectively. Then we consider a new graph G by identifying the s nodes and the t nodes of G_1 and G_2 (i.e. connecting the graphs in parallel) and define $c : E(G) \rightarrow \mathbb{R}^+$ by $c(e) = c_1(e)$ if $e \in E(G_1)$ and $c(e) = c_2(e)$ if $e \in E(G_2)$. Similarly, we set $G(x)$ to be the subgraph of G that includes the corresponding edges e such that $e \in E(G_1(x_1))$ or $e \in E(G_2(x_2))$. Then

$$\frac{1}{R_{s,t}(G(x))} = \frac{1}{R_{s,t}(G_1(x_1))} + \frac{1}{R_{s,t}(G_2(x_2))}, \quad C_{s,t}(G(x)) = C_{s,t}(G_1(x_1)) + C_{s,t}(G_2(x_2)) \quad (10)$$

If we create a new graph G by identifying the t node of G_1 with the s node of G_2 , relabeling this node $v \notin \{s, t\}$ (i.e. connecting the graphs in series) and define c and $G(x)$ as before, then

$$R_{s,t}(G(x)) = R_{s,t}(G_1(x)) + R_{s,t}(G_2(x)), \quad \frac{1}{C_{s,t}(G(x))} = \frac{1}{C_{s,t}(G_1(x_1))} + \frac{1}{C_{s,t}(G_2(x_2))}. \quad (11)$$

2.3 Quantum Algorithms

We will consider problems parametrized by a parent graph G , by which we more precisely mean a family of graphs $\{G_n\}_{n \in \mathbb{N}}$ where G_n is a graph on n vertices. We will generally drop the subscript n .

A graph is connected if there is a path between every pair of vertices. For a family of graphs G , and a set $X \subset \{0, 1\}^N$, let $\text{CONN}_{G,X}$ denote the *connectivity problem*, defined by $\text{CONN}_{G,X}(x) = 1$ if $G(x)$ is connected (see Section 2.2 for description of $G(x)$), and $\text{CONN}_{G,X}(x) = 0$ if $G(x)$ is not connected, for all $x \in X$.

Similarly, for $s, t \in V(G)$, defined $st\text{-CONN}_{G,X}$ by $st\text{-CONN}_{G,X}(x) = 1$ if there is a path from s to t in $G(x)$, and $st\text{-CONN}_{G,X}(x) = 0$ otherwise, for all $x \in X$.

We will consider CONN and $st\text{-CONN}$ in the edge-query input model, meaning that we have access to a standard quantum oracle O_x , defined $O_x|i\rangle|b\rangle = |i\rangle|b \oplus x_i\rangle$, where x_i is the i^{th} bit of i . Since every edge of G is associated with an input variable, as described in Section 2.2, for any edge in G , we can check if it is also present in $G(x)$ using one query to O_x .

Let $f_N : X_N \rightarrow \{0, 1\}$, $X_N \subseteq [q]^N$, for $N \in \mathbb{N}$, be a family of functions. An algorithm *decides* f with *bounded error* if for any $x \in X_N$, the algorithm outputs $f(x)$ with probability at least $2/3$. Let $f_N : X_N \rightarrow \mathbb{R}_{\geq 0}$, $X_N \subseteq [q]^N$, for $N \in \mathbb{N}$, be a family of functions. An algorithm *estimates* f to *relative accuracy* ϵ with *bounded error* if for any $x \in X_N$, on input x , the algorithm outputs \tilde{f} such that $|\tilde{f} - f(x)| \leq \epsilon f(x)$ with probability at least $2/3$. Since all algorithms discussed here will have bounded error, we will omit this description. We will generally talk about a function $f : X \rightarrow \{0, 1\}$, leaving the parametrization over n implicit.

In the remainder of this section, we describe the quantum algorithmic building blocks of the algorithms introduced in Section 5.

Theorem 8 (Phase Estimation [Kit95, CEMM98]). *Let U be a unitary with eigenvectors $|\theta\rangle$ satisfying $U|\theta\rangle = e^{i\theta}|\theta\rangle$ and assume $\theta \in [-\pi, \pi]$. For any $\Theta \in (0, \pi)$ and $\epsilon \in (0, 1)$, there exists a quantum algorithm that makes $O(\frac{1}{\Theta} \log \frac{1}{\epsilon})$ calls to U and, on input $|\theta_j\rangle$ outputs a state $|\theta_j\rangle|w\rangle$ such that if $\theta_j = 0$, then $|w\rangle = |0\rangle$ and if $|\theta_j| \geq \Theta$, $|\langle 0|w\rangle|^2 \leq \epsilon$. If U acts on s qubits, the algorithm uses $O(s + \log \frac{1}{\Theta})$ space.*

We will use the following corollary of Theorem 8, which is a slight generalization of an algorithm introduced in [CKS17], also called Gapped Phase Estimation.

Theorem 9 (Gapped Phase Estimation). *Let U be a unitary with eigenvectors $|\theta\rangle$ satisfying $U|\theta\rangle = e^{i\pi\theta}|\theta\rangle$ and assume $\theta \in [-1, 1]$. Let $\varphi \in (0, 1)$, let $\epsilon > 0$ and let $\delta \in (0, 1 - \varphi)$. Then, there exists a unitary procedure $\text{GPE}(\varphi, \epsilon, \delta)$ making $O(\varphi^{-1} \log \epsilon^{-1})$ queries to U that on input $|0\rangle_C |0\rangle_P |\theta\rangle$ prepares a state $(\beta_0 |0\rangle_C |\gamma_0\rangle_P + \beta_1 |1\rangle_C |\gamma_1\rangle_P) |\theta\rangle$ where $|\gamma_0\rangle$ and $|\gamma_1\rangle$ are some unit vectors, $\beta_0^2 + \beta_1^2 = 1$ and such that*

- if $|\theta| \leq \delta$, then $|\beta_1| \leq \epsilon$, and
- if $\delta + \varphi \leq |\theta|$, then $|\beta_0| \leq \epsilon$.

The registers C and P have 1 and $O(\log(\varphi^{-1}) \log(\epsilon^{-1}))$ qubits respectively. And in addition to the queries to U , the algorithm uses $O\left(\log \frac{1}{\epsilon} (\log^2 \frac{1}{\varphi} + \log \log \frac{1}{\epsilon})\right)$ elementary gates.

Proof. Standard phase estimation [Kit95, CEMM98] (but see, in particular [CEMM98, Appendix C]) on input $|\theta\rangle$ with precision $\varphi/4$ prepares a state $|\tilde{\theta}\rangle|\theta\rangle$ such that upon measuring $|\tilde{\theta}\rangle$, with probability at least c for some $c > 1/2$, we measure some $\tilde{\theta}$ that is within $\varphi/2$ of θ , meaning that if $|\theta| \leq \delta$, then $|\tilde{\theta}| < \delta + \varphi/2$, and if $|\theta| \geq \delta + \varphi$, then $|\tilde{\theta}| > \delta + \varphi/2$.

Let $k = c' \log \frac{1}{\epsilon^2}$ for some constant c' . Repeating phase estimation k times produces a state $|\tilde{\theta}\rangle_{P_1} \dots |\tilde{\theta}\rangle_{P_k}$ such that if we were to measure the state, we would obtain a string of estimates $\tilde{\theta}_1, \dots, \tilde{\theta}_k$. If $|\theta| \leq \delta$, each estimate would have absolute value less than $\delta + \varphi/2$ with probability at least c , and if $|\theta| \geq \delta + \varphi$, each estimate would have absolute value greater than $\delta + \varphi/2$ with probability at least c .

Instead of measuring, assume that we have extra registers $|0\rangle_{C_1} \dots |0\rangle_{C_k}$. Apply to every pair of registers $P_i C_i$ the unitary that maps $|\bar{\theta}\rangle_{P_i} |0\rangle_{C_i}$ to $|\bar{\theta}\rangle_{P_i} |0\rangle_{C_i}$ if $|\bar{\theta}| \leq \delta + \varphi/2$, and to $|\bar{\theta}\rangle_{P_i} |1\rangle_{C_i}$ otherwise. This unitary can be done with $O(\log^2 \frac{1}{\varphi})$ elementary gates. Then we unitarily do majority voting of the values of registers C_1, \dots, C_k and encode the result in a register C . This step requires only $O(\log \frac{1}{\epsilon} \log \log \frac{1}{\epsilon})$ elementary gates. By a standard Chernoff bound, for appropriately chosen c' , this majority will be 0 with amplitude at least $\sqrt{1 - \epsilon^2}$ if $|\theta| \leq \delta$, and will be 1 with amplitude at least $\sqrt{1 - \epsilon^2}$ if $|\theta| \geq \delta + \varphi$.

Grouping all registers $P_1, \dots, P_k, C_1, \dots, C_k$ into one label P we have that the state produced will be $(\beta_0 |0\rangle_C |\gamma_0\rangle_P + \beta_1 |1\rangle_C |\gamma_1\rangle_P) |\theta\rangle$ for some unit vectors $|\gamma_0\rangle$ and $|\gamma_1\rangle$, with $|\beta_1| \leq \epsilon$ whenever $|\theta| \leq \delta$, and $|\beta_0| \leq \epsilon$ whenever $|\theta| \geq \delta + \varphi$.

The number of elementary gates used in every call to standard phase estimation with precision $\varphi/2$ is $O(\log^2 \varphi)$, in addition to $O(\frac{1}{\varphi} \log \frac{1}{\epsilon})$ calls to U , from which the result follows. \square

Theorem 10 (Amplitude Estimation [BHMT02]). *Let \mathcal{A} be a quantum algorithm that, on input x , outputs $\sqrt{p(x)}|0\rangle|\Psi_0(x)\rangle + \sqrt{1-p(x)}|1\rangle|\Psi_1(x)\rangle$. Then there exists a quantum algorithm that estimates $p(x)$ to precision ϵ using $O\left(\frac{1}{\epsilon\sqrt{p(x)}}\right)$ calls to \mathcal{A} .*

We will make use of the following corollary (see [IJ16] for a proof).

Corollary 11. *Let \mathcal{A} be a quantum algorithm that outputs $\sqrt{p(x)}|0\rangle|\Psi_0(x)\rangle + \sqrt{1-p(x)}|1\rangle|\Psi_1(x)\rangle$ on input x such that either $p(x) \leq p_0$, or $p(x) \geq p_1$ for $p_1 > p_0$. Then there exists a quantum algorithm that decides if $p(x) \leq p_0$ using $O\left(\frac{\sqrt{p_1}}{p_1 - p_0}\right)$ calls to \mathcal{A} .*

2.4 Span Programs and Witness Sizes

Span programs [KW93] were first introduced to the study of quantum algorithms by Reichardt and Špalek [RŠ12]. They have since proven to be immensely important for designing quantum algorithms in the query model.

Definition 12 (Span Program). *A span program $P = (H, U, \tau, A)$ on $\{0, 1\}^N$ is made up of (I) finite-dimensional inner product spaces $H = H_1 \oplus \dots \oplus H_N$, and $\{H_{j,b} \subseteq H_j\}_{j \in [N], b \in \{0,1\}}$ such that $H_{j,0} + H_{j,1} = H_j$, (II) a vector space U , (III) a non-zero target vector $\tau \in U$, and (IV) a linear operator $A : H \rightarrow U$. For every string $x \in \{0, 1\}^N$, we associate the subspace $H(x) := H_{1,x_1} \oplus \dots \oplus H_{N,x_N}$, and an operator $A(x) := A\Pi_{H(x)}$.*

Definition 13 (Positive and Negative Witness). *Let P be a span program on $\{0, 1\}^N$ and let x be a string $x \in \{0, 1\}^N$. Then we call $|w\rangle$ a positive witness for x in P if $|w\rangle \in H(x)$, and $A|w\rangle = \tau$. We define the positive witness size of x as:*

$$w_+(x, P) = w_+(x) = \min\{\| |w\rangle \|^2 : |w\rangle \in H(x), A|w\rangle = \tau\}, \quad (12)$$

if there exists a positive witness for x , and $w_+(x) = \infty$ otherwise.

Let $\mathcal{L}(U, \mathbb{R})$ denote the set of linear maps from U to \mathbb{R} . We call a linear map $\omega \in \mathcal{L}(U, \mathbb{R})$ a negative witness for x in P if $\omega A\Pi_{H(x)} = 0$ and $\omega\tau = 1$. We define the negative witness size of x as:

$$w_-(x, P) = w_-(x) = \min\{\|\omega A\|^2 : \omega \in \mathcal{L}(U, \mathbb{R}), \omega A\Pi_{H(x)} = 0, \omega\tau = 1\}, \quad (13)$$

if there exists a negative witness, and $w_-(x) = \infty$ otherwise.

If $w_+(x)$ is finite, we say that x is positive (wrt. P), and if $w_-(x)$ is finite, we say that x is negative. We let P_1 denote the set of positive inputs, and P_0 the set of negative inputs for P .

For a function $f : X \rightarrow \{0, 1\}$, with $X \subseteq \{0, 1\}^N$, we say P decides f if $f^{-1}(0) \subseteq P_0$ and $f^{-1}(1) \subseteq P_1$. Given a span program P that decides f , one can use it to design a quantum algorithm whose output is $f(x)$ (with high probability), given access to the input $x \in X$ via queries of the form $\mathcal{O}_x : |i, b\rangle \mapsto |i, b \oplus x_i\rangle$.

The following theorem is due to [Rei09] (see [IJ16] for a version with similar notation).

Theorem 14. Let $U(P, x) = (2\Pi_{\ker A} - I)(2\Pi_{H(x)} - I)$. Fix $X \subseteq \{0, 1\}^N$ and $f : X \rightarrow \{0, 1\}$, and let P be a span program on $\{0, 1\}^N$ that decides f . Let $W_+(f, P) = \max_{x \in f^{-1}(1)} w_+(x, P)$ and $W_-(f, P) = \max_{x \in f^{-1}(0)} w_-(x, P)$. Then there is a bounded error quantum algorithm that decides f by making $O(\sqrt{W_+(f, P)W_-(f, P)})$ calls to $U(P, x)$, and elementary gates. In particular, this algorithm has quantum query complexity $O(\sqrt{W_+(f, P)W_-(f, P)})$.

Ref. [IJ16] defines the *approximate positive and negative witness sizes*, $\tilde{w}_+(x, P)$ and $\tilde{w}_-(x, P)$. These are similar to the positive and negative witness sizes, but with the conditions $|w\rangle \in H(x)$ and $\omega A\Pi_{H(x)} = 0$ relaxed.

Definition 15 (Approximate Positive Witness). For any span program P on $\{0, 1\}^N$ and $x \in \{0, 1\}^N$, we define the positive error of x in P as:

$$e_+(x) = e_+(x, P) := \min \left\{ \left\| \Pi_{H(x)^\perp} |w\rangle \right\|^2 : A|w\rangle = \tau \right\}. \quad (14)$$

We say $|w\rangle$ is an approximate positive witness for x in P if $\left\| \Pi_{H(x)^\perp} |w\rangle \right\|^2 = e_+(x)$ and $A|w\rangle = \tau$. We define the approximate positive witness size as

$$\tilde{w}_+(x) = \tilde{w}_+(x, P) := \min \left\{ \| |w\rangle \|^2 : A|w\rangle = \tau, \left\| \Pi_{H(x)^\perp} |w\rangle \right\|^2 = e_+(x) \right\}. \quad (15)$$

If $x \in P_1$, then $e_+(x) = 0$. In that case, an approximate positive witness for x is a positive witness, and $\tilde{w}_+(x) = w_+(x)$. For negative inputs, the positive error is larger than 0.

We can define a similar notion of approximate negative witnesses (see [IJ16]).

Theorem 16 ([IJ16]). Let $U(P, x) = (2\Pi_{\ker A} - I)(2\Pi_{H(x)} - I)$. Fix $X \subseteq \{0, 1\}^N$ and $f : X \rightarrow \mathbb{R}_{\geq 0}$. Let P be a span program on $\{0, 1\}^N$ such that for all $x \in X$, $f(x) = w_-(x, P)$ and define $\tilde{W}_+ = \tilde{W}_+(P) = \max_{x \in X} \tilde{w}_+(x, P)$. Then there exists a quantum algorithm that estimates f to accuracy ϵ and that uses $\tilde{O}\left(\frac{1}{\epsilon^{3/2}} \sqrt{w_-(x) \tilde{W}_+}\right)$ calls to $U(P, x)$ and elementary gates.

A span program for st -connectivity An important example of a span program is one for st -connectivity, first introduced in [KW93], and used in [BR12] to give a new quantum algorithm for st -connectivity. We state this span program below, somewhat generalized to include weighted graphs, and to allow the input to be specified as a subgraph of some parent graph G that is not necessarily the complete graph. We allow a string $x \in \{0, 1\}^N$ to specify a subgraph $G(x)$ of G in a fairly general way, as described in Section 2.2. In particular, for $i \in [N]$, let $\vec{E}_{i,1} \subseteq \vec{E}(G)$ denote the set of (directed) edges associated with the literal x_i , and $\vec{E}_{i,0}$ the set of edges associated with the literal \bar{x}_i . Note that if $(u, v, \ell) \in \vec{E}_{i,b}$ then we must also have $(v, u, \ell) \in \vec{E}_{i,b}$, since $G(x)$ is an undirected graph. We assume G has some implicit weighting function c .

Then we refer to the following span program as P_G :

$$\begin{aligned} \forall i \in [N], b \in \{0, 1\} : H_{i,b} &= \text{span}\{|e\rangle : e \in \vec{E}_{i,b}\} \\ U &= \text{span}\{|v\rangle : v \in V(G)\} \\ \tau &= |s\rangle - |t\rangle \\ \forall e = (u, v, \ell) \in \vec{E}(G) : A|u, v, \ell\rangle &= \sqrt{c(u, v, \ell)}(|u\rangle - |v\rangle) \end{aligned} \quad (16)$$

One can check that if s and t are connected, then if $|w\rangle$ represents a weighted st -path or linear combination of weighted st -paths in $G(x)$, then $|w\rangle$ is a positive witness for x . Furthermore, this is the only possibility for a positive witness, so x is a positive input for P_G if and only if $G(x)$ is st -connected, and in particular, $w_+(x, P_G) = \frac{1}{2}R_{s,t}(G(x))$ [BR12]. Since the weights $c(e)$ are positive, the set of positive inputs of P_G are independent of the choice of c , however, the witness sizes will depend on c .

3 Effective Capacitance and st -connectivity

In this section, we will prove the following theorem:

Theorem 17. *Let P_G be the span program in Eq. (16). Then for any $x \in \{0, 1\}^N$, $w_-(x, P_G) = 2C_{s,t}(G(x))$.*

Previously, the negative witness size of P_G was characterized by the size of a cut [RŠ12] or, in planar graphs, the effective resistance of a graph related to the planar dual of $G(x)$ [JK17].

We will prove Theorem 17 shortly, but first, we mention the following corollary:

Corollary 18. *Let G be a multigraph with $s, t \in V(G)$. Then for any choice of (non-negative, real-valued) implicit weight function, the bounded error quantum query complexity of evaluating $st\text{-CONN}_{G,X}$ is*

$$O\left(\sqrt{\max_{\substack{x \in X \\ st\text{-CONN}_{G,X}(x)=1}} R_{s,t}(G(x)) \times \max_{\substack{x \in X \\ st\text{-CONN}_{G,X}(x)=0}} C_{s,t}(G(x))}\right). \quad (17)$$

Proof. This follows from Theorem 17 and the fact that $w_+(x, P_G) = \frac{1}{2}R_{s,t}(G(x))$, which is proven in [BR12], and generalized to the weighted case in [JK17]. Then Theorem 14 gives the result. \square

We emphasize that Corollary 18 holds for $R_{s,t}$ and $C_{s,t}$ defined with respect to any weight function, some of which may give a significantly better complexity for solving this problem.

We are now ready to prove Theorem 17, the main result of this section.

Proof of Theorem 17. First, we prove that any unit st -potential on $G(x)$ can be transformed into a negative witness for x in P_G with witness size equal to twice the unit potential energy of that potential. This shows that $w_-(x, P_G) \leq 2C_{s,t}(G(x))$.

Given a unit st -potential $\mathcal{V} : V(G) \rightarrow \mathbb{R}$ on $G(x)$, we consider $\omega_{\mathcal{V}} = \sum_{v \in V(G)} \mathcal{V}(v) \langle v |$. Then because $\mathcal{V}(s) = 1$ and $\mathcal{V}(t) = 0$, we have $\omega_{\mathcal{V}} \tau = 1$. Secondly,

$$\begin{aligned} \omega_{\mathcal{V}} A \Pi_{H(x)} &= \sum_{u' \in V(G)} \mathcal{V}(u') \langle u' | \sum_{(u,v,\ell) \in \vec{E}(G(x))} \sqrt{c(u,v,\ell)} (|u\rangle - |v\rangle) \langle u, v, \ell | \\ &= \sum_{(u,v,\ell) \in \vec{E}(G(x))} \sqrt{c(u,v,\ell)} (\mathcal{V}(u) - \mathcal{V}(v)) \langle u, v, \ell | = 0, \end{aligned} \quad (18)$$

where we've used the definition of unit st -potential, which states that $\mathcal{V}(u) - \mathcal{V}(v) = 0$ when $(u, v, \ell) \in E(G(x))$. Thus $\omega_{\mathcal{V}}$ is a valid negative witness for input x .

We have

$$\begin{aligned} w_-(x, P_G) &\leq \min_{\mathcal{V}} \|\omega_{\mathcal{V}} A\|^2 = \min_{\mathcal{V}} \left\| \sum_{(u,v,\ell) \in \vec{E}(G)} \sqrt{c(u,v,\ell)} (\mathcal{V}(u) - \mathcal{V}(v)) \langle u, v, \ell | \right\|^2 \\ &= 2 \min_{\mathcal{V}} \sum_{(u,v,\ell) \in E(G)} (\mathcal{V}(u) - \mathcal{V}(v))^2 c(u,v,\ell) = 2C_{s,t}(G(x)), \end{aligned} \quad (19)$$

where the minimization is over unit st -potentials on $G(x)$.

Next, we show that any negative witness ω for P_G on input x can be transformed into a unit st -potential \mathcal{V}_{ω} on $G(x)$, with negative witness size equal to twice the unit potential energy of \mathcal{V}_{ω} . This shows that $w_-(x, P_G) \geq 2C_{s,t}(G(x))$.

Given ω , a negative witness for input x , let $\mathcal{V}_{\omega}(v) = \omega(|v\rangle - |t\rangle)$ for $v \in V(G)$. Then $V_{\omega}(s) = \omega(|s\rangle - |t\rangle) = \omega \tau = 1$, and $V_{\omega}(t) = \omega(|t\rangle - |t\rangle) = 0$. Also for $(u, v, \ell) \in E(G(x))$, we have

$$\begin{aligned} V_{\omega}(u) - V_{\omega}(v) &= \omega(|u\rangle - |t\rangle) - \omega(|v\rangle - |t\rangle) = \omega(|u\rangle - |v\rangle) \\ &= \omega A |u, v, \ell\rangle = \omega A \Pi_{H(x)} |u, v, \ell\rangle = 0, \end{aligned} \quad (20)$$

because $\omega A \Pi_{H(x)} = 0$. Thus, V_ω is a st -unit potential for $G(x)$.

Then

$$\begin{aligned} w_-(x, P_G) &= \min_{\omega} \|\omega A\|^2 = \min_{\omega} \sum_{(u,v,\ell) \in \vec{E}(G)} (\omega(|u| - |v|))^2 c(u, v, \ell) \\ &= 2 \min_{\omega} \sum_{(u,v,\ell) \in E(G)} (\mathcal{V}_\omega(u) - \mathcal{V}_\omega(v))^2 c(u, v, \ell) \geq 2C_{s,t}(G(x)), \end{aligned} \quad (21)$$

where the minimization is over negative witnesses. Since $w_-(x, P_G) \geq 2C_{s,t}(G(x))$ and $w_-(x, P_G) \leq 2C_{s,t}(G(x))$, we must have $w_-(x, P_G) = 2C_{s,t}(G(x))$. \square

4 Applications

4.1 Estimating the Capacitance of a Circuit

By Theorem 17, $w_-(x, P_G) = 2C_{s,t}(G(x))$, so we can apply Theorem 16 to estimate $C_{s,t}(G(x))$. By Theorem 16, the complexity of doing this depends on $C_{s,t}(G(x))$ and $\tilde{W}_+(P_G) = \max_x \tilde{w}_+(x, P_G)$. We will prove the following theorem:

Theorem 19. *For the span program P_G , we have that $\tilde{W}_+(P_G) = O(\max_p J_{E(G)}(p))$, where the maximum runs over all st -unit flows p that are paths from s to t .*

Note that when the weights are all 1, $\max_p J_{E(G)}(p)$ is just the length of the longest self-avoiding st -path in G . Combining Theorem 16, Theorem 17, and Theorem 19, we have:

Corollary 20. *Given a network (G, c) , with $s, t \in V(G)$ and access to an oracle O_x , the bounded error quantum query complexity of estimating $C_{s,t}(G(x))$ to accuracy ϵ is $\tilde{O}(\epsilon^{-3/2} \sqrt{C_{s,t}(G(x)) \max_p J_{E(G)}(p)})$ where the maximum runs over all st -unit flows p that are paths from s to t .*

Similarly, we can show:

Corollary 21. *Let U be the cost of implementing the map*

$$|u\rangle|0\rangle \mapsto \sum_{v,\ell:(u,v,\ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell) / d_G(u)} |u, v, \ell\rangle.$$

Then the quantum time complexity of estimating $C_{s,t}(G(x))$ to accuracy ϵ is $\tilde{O}(\epsilon^{-3/2} \sqrt{C_{s,t}(G(x)) \max_p J_{E(G)}(p) U})$.

Proof. The algorithm in Theorem 16 requires $\tilde{O}(\epsilon^{-3/2} \sqrt{C_{s,t}(G(x)) \max_p J_{E(G)}(p)})$ calls to a unitary $U(P_G, x)$, and other elementary operations [IJ16]. By [JK17] (generalizing [BR12]), for any G , $U(P_G, x)$ can be implemented in cost U . \square

To prove Theorem 19, we first relate unit st -flows on G to approximate positive witnesses. Intuitively, an approximate positive witness is an st -flow on G that has energy as small as possible on edges in $E(G) \setminus E(G(x))$. Thus, we can upper bound the approximate positive witness size by the highest possible energy of any st -flow on G , which is always achieved by a flow that is an st -path.

The following claim can be proven using the technique of the proof of [JK17, Lemma 11]:

Claim 22. *Let P_G be the span program of Eq. (16). Then the positive error of x in P_G is*

$$e_+(x, P_G) = \min_{\theta} \{J_{E(G) \setminus E(G(x))}(\theta)\} \quad (22)$$

where θ runs over unit st -flows on G . The approximate positive witness size is

$$\tilde{w}_+(x, P_G) = \min_{\theta} J_{E(G)}(\theta) \quad (23)$$

where θ runs over unit st -flows on G such that $J_{E(G) \setminus E(G(x))}(\theta) = e_+(x, P_G)$.

Proof of Theorem 19. This theorem follows immediately from the observation that $\min_{\theta} J_E(\theta) \leq J_E(\theta')$ for any valid st -flow θ' on the set of edges E . Hence, we are free to choose θ' to be any st -path through E . Then, separately taking $E \mapsto E(G)$ and $E \mapsto E(G) \setminus E(G(x))$ yields the desired result. \square

4.2 Deciding Connectivity

Let $\text{CONN}_{G,X}$ be the problem of deciding, given $x \in X$, whether $G(x)$ is connected. That is:

$$\text{CONN}_{G,X} = \bigwedge_{\{u,v\}:u,v \in V(G)} uv\text{-CONN}_{G,X}. \quad (24)$$

Using the technique of converting logical AND into st -connectivity problems in series [NTS95, JK17], we note that the above problem is equivalent to $n(n-1)/2$ st -connectivity problems in series, one for each pair of distinct vertices in $V(G)$. (The approach in Ref. [Ari16] is similar, but only looks at $n-1$ instances — the pairs s and v for each $v \in V(G)$. Our approach is symmetrized over the vertices and thus makes the analysis simpler.)

More precisely, we define a graph \mathcal{G} such that:

$$\begin{aligned} V(\mathcal{G}) &= V(G) \times \{\{u, v\} : u \neq v \in V(G)\} \\ E(\mathcal{G}) &= E(G) \times \{\{u, v\} : u \neq v \in V(G)\} \end{aligned} \quad (25)$$

where \times denotes the Cartesian product, and $\{u, v\}$ is an extra label denoting that that edge or vertex is in the $\{u, v\}$ th copy of the graph G present as a subgraph in \mathcal{G} . Choose any labeling of the vertices from 1 to n (with slight abuse of notation, we use u both for the original vertex name and the label). We then label the vertex $(1, \{1, 2\})$ as s and the vertex $(n, \{n-1, n\})$ as t . Next identify vertices $(v, \{u, v\})$ and $(u, \{u, v+1\})$ if $u < v$ and $v < n$, and identify vertices $(v, \{u, v\})$ and $(u+1, \{u+1, u+2\})$ if $v = n$ and $u < n-1$. See Figure 3 for an example of this construction.




Figure 3: Example of how \mathcal{G} is formed from a graph G . We have labeled the subgraphs of \mathcal{G} according to the st -connectivity problems the subgraphs represent.

Finally, we define $\mathcal{G}(x)$ to be the subgraph of \mathcal{G} with edges

$$E(\mathcal{G}(x)) = E(G(x)) \times \{\{u, v\} : u \neq v \in V(G)\}. \quad (26)$$

We can see that any st -path in $\mathcal{G}(x)$ must go through each of the copies of $G(x)$, meaning it must include, for each $\{u, v\}$, a uv -path through the copy of $G(x)$ labeled $\{u, v\}$. Thus, there is an st -path in $\mathcal{G}(x)$ if and only if $G(x)$ is connected.

We consider the span program $P_{\mathcal{G}}$, where $c(e) = 1$ for all $e \in E(\mathcal{G})$. We will use $P_{\mathcal{G}}$ to solve st -connectivity on $\mathcal{G}(x)$. To analyze the resulting algorithm, we need to upper bound the negative and positive witness sizes $w_-(x, P_{\mathcal{G}}) = 2C_{s,t}(\mathcal{G}(x))$ and $w_+(x, P_{\mathcal{G}}) = \frac{1}{2}R_{s,t}(\mathcal{G}(x))$.

Lemma 23. *For any x such that $G(x)$ is connected, $w_+(x, P_{\mathcal{G}}) = \frac{n(n-1)}{2}R_{\text{avg}}(G(x))$, where $R_{\text{avg}}(G(x))$ is the average resistance.*

Proof. Using the rule that resistances in series add, we have:

$$R_{s,t}(\mathcal{G}(x)) = \frac{1}{2} \sum_{u,v \in V(G)} R_{u,v}(G(x)) = n(n-1)R_{\text{avg}}(G(x)). \quad (27)$$

This is equal to $2w_+(x, P_G)$. □

Now we bound $C_{s,t}(\mathcal{G}(x))$, to prove the following:

Lemma 24. *Fix $\kappa > 1$, and suppose $G(x)$ has κ connected components. Then if G is a subgraph of a complete graph (that is, G has at most one edge between any pair of vertices), we have $w_-(x, P_G) = O(1/\kappa)$. Otherwise, we have $w_-(x, P_G) = O(d_{\max}(G)/\sqrt{n\kappa})$.*

Proof. Using the rule for capacitors in series, and accounting for double counting pairs of vertices, we have

$$\frac{1}{C_{s,t}(\mathcal{G}(x))} = \frac{1}{2} \sum_{s',t' \in V(G)} \frac{1}{C_{s',t'}(G(x))}. \quad (28)$$

To put an upper bound on $C_{s,t}(\mathcal{G}(x))$, we can put an upper bound on each term $C_{s',t'}(G(x))$. To upper bound $C_{s',t'}(G(x))$, consider the following unit $s't'$ -potential for $G(x)$. Set $\mathcal{V}(v) = 1$ for all v in the same connected component as s' in $G(x)$, set $\mathcal{V}(v) = 0$ for all v in the same connected component as t' in $G(x)$, and set $\mathcal{V}(v) = v$ for all other vertices in G . We now find the minimum unit potential energy of this $s't'$ -potential, (minimizing over v). This will be an upper bound on $C_{s',t'}(G(x))$ by Definition 6, since it is not necessarily the optimal choice to set all vertices not connected to s' or t' to have the same unit potential value.

While we could calculate the unit potential energy using Definition 5, and then minimize that quantity over v , we instead use Claim 7, and the fact that our choice of unit $s't'$ -potential effectively creates a graph with three vertices: one vertex corresponds to the connected component of $G(x)$ containing s' (let $n_{s'}$ be the number of vertices in this component), one vertex corresponds to the connected component of $G(x)$ containing t' (let $n_{t'}$ be the number of vertices in this component), and one vertex corresponds to all the other vertices in the graph (let $n_a = n - n_{s'} - n_{t'}$ be the remaining number of vertices).

For any pair of vertices u and v , let D_u be the number of edges of G coming out of the component of $G(x)$ containing u , and let D_{uv} be the number of edges in G between the components of $G(x)$ containing u and v , so the number of edges between the component containing a and all components other than that containing s' is $D_{s'} - D_{s'a}$. Using the rules for calculating capacitance in series and parallel, we have

$$C_{s',t'}(G(x)) \leq D_{s't'} + \left(\frac{1}{D_{s'} - D_{s't'}} + \frac{1}{D_{t'} - D_{s't'}} \right)^{-1} = \frac{D_{s'}D_{t'} - D_{s't'}^2}{D_{s'} + D_{t'} - 2D_{s't'}}. \quad (29)$$

Using Eq. (28), we have

$$\frac{1}{C_{s,t}(\mathcal{G}(x))} \geq \frac{1}{2} \sum_{\substack{s',t' \in V(G) \\ \{s',t'\} \notin E(G(x))}} \frac{D_{s'} + D_{t'} - 2D_{s't'}}{D_{s'}D_{t'} - D_{s't'}^2}. \quad (30)$$

Now the expression on the right-hand side of Eq. (30) depends only on which connected components s' and t' are in, so instead of summing over the vertices of G , we can instead sum over the κ connected components of $G(x)$. Let n_i be the number of vertices in the i^{th} connected component. Then, continuing from Eq. (30) we have:

$$\begin{aligned} \frac{1}{C_{s,t}(\mathcal{G}(x))} &\geq \frac{1}{2} \sum_{i,j \in [\kappa]: i \neq j} n_i n_j \frac{D_i + D_j - 2D_{ij}}{D_i D_j - D_{ij}^2} = \frac{1}{2} \sum_{i,j \in [\kappa]: i \neq j} n_i n_j \frac{(D_i - D_{ij}) + (D_j - D_{ij})}{D_i D_j - D_{ij}^2} \\ &= \frac{1}{2} \left(2 \sum_{i,j \in [\kappa]: i \neq j} n_i n_j \frac{(D_i - D_{ij})}{D_i D_j - D_{ij}^2} \right) = \sum_{i,j \in [\kappa]: i \neq j} n_i n_j \frac{D_j - D_{ij}}{D_i D_j - D_{ij}^2}. \end{aligned} \quad (31)$$

First, consider the case when G is a complete graph. In that case, $D_i - D_{ij}$, the number of edges leaving component i and going to a component other than i or j , is exactly $n_i(n - n_i - n_j)$, whereas the number of edges leaving component i is $D_i = n_i(n - n_i)$. Finally, $D_{ij} = n_i n_j$ is the number of edges going from the i^{th} component to the j^{th} component. Thus, continuing from Eq. (31), we have:

$$\begin{aligned} \frac{1}{C_{s,t}(\mathcal{G}(x))} &\geq \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{n_i(n - n_i - n_j)}{n_i(n - n_i)n_j(n - n_j) - n_i^2 n_j^2} = \sum_{i,j \in [k]: i \neq j} \frac{n_i(n - n_i - n_j)}{(n - n_i)(n - n_j) - n_i n_j} \\ &= \sum_{i,j \in [k]: i \neq j} \frac{n_i(n - n_i - n_j)}{n^2 - n n_i - n n_j} = \sum_{i,j \in [k]: i \neq j} \frac{n_i}{n} = \kappa - 1. \end{aligned} \quad (32)$$

Note that this upper bound on $C_{s,t}(\mathcal{G}(x))$ applies to *any* subgraph of a complete graph, since adding edges can only increase the capacitance. Thus, we have completed the first part of the proof.

We now continue with the more general case, where G is not necessarily a subgraph of a complete graph. Let $d = d_{\max}(G)$. Continuing from Eq. (31), and using the fact that for any component, we have:

$$\begin{aligned} \frac{1}{C_{s,t}(\mathcal{G}(x))} &\geq \frac{1}{2} \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{D_i + D_j - 2D_{ij}}{D_i D_j - D_{ij}^2} = \frac{1}{2} \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{(\sqrt{D_i} - \sqrt{D_j})^2 + 2\sqrt{D_i D_j} - 2D_{ij}}{D_i D_j - D_{ij}^2} \\ &\geq \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{\sqrt{D_i D_j} - D_{ij}}{D_i D_j - D_{ij}^2} \geq \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{\sqrt{D_i D_j - D_{ij}^2}}{D_i D_j - D_{ij}^2} \geq \sum_{i,j \in [k]: i \neq j} n_i n_j \frac{1}{\sqrt{D_i D_j}} \\ &\geq \sum_{i,j \in [k]: i \neq j} \frac{\sqrt{n_i n_j}}{d} \geq \frac{1}{d} \sqrt{\sum_{i \in [k]} \sum_{j \neq i} n_i n_j} = \frac{1}{d} \sqrt{\sum_{i \in [k]} n_i(n - n_i)}. \end{aligned} \quad (33)$$

Above we used the fact that for any component, $D_i \leq d n_i$. The sum $\sum_{i \in [k]} n_i^2$ is maximized when the n_i are as far as possible from uniform. In this case, we have $n_i \geq 1$ for all i , so $\sum_{i \in [k]} n_i^2 \leq (\kappa - 1) + (n - (\kappa - 1))^2$. Thus, continuing, we have

$$\begin{aligned} \frac{1}{C_{s,t}(\mathcal{G}(x))} &\geq \frac{1}{d} \sqrt{n \sum_{i \in [k]} n_i - \sum_{i \in [k]} n_i^2} \geq \frac{1}{d} \sqrt{n^2 - (\kappa - 1) - n^2 - (\kappa - 1)^2 + 2n(\kappa - 1)} \\ &= \frac{1}{d} \sqrt{(2n - \kappa)(\kappa - 1)} \geq \frac{1}{d} \sqrt{n\kappa}. \end{aligned} \quad (34)$$

The result follows by Theorem 17, which says that $w_-(x, P_G) = 2C_{s,t}(\mathcal{G}(x))$. \square

Combining Lemmas 23 and 24 and Theorem 14, we have the following:

Theorem 25. *For any family of graphs G such that G is a subgraph of a complete graph, and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, if $G(x)$ is connected, $R_{\text{avg}}(G(x)) \leq R$, and if $G(x)$ is not connected, it has at least κ components, the bounded error quantum query complexity of $\text{CONN}_{G,X}$ is $O(n\sqrt{R/\kappa})$.*

For any family of connected graphs G and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, if $G(x)$ is connected, $R_{\text{avg}}(G(x)) \leq R$, and if $G(x)$ is not connected, it has at least κ components, the bounded error quantum query complexity of $\text{CONN}_{G,X}$ is $O(n^{3/4} \sqrt{R d_{\max}(G)} / \kappa^{1/4})$.

Similarly, we can show:

Corollary 26. *Let U be the cost of implementing the map*

$$|u\rangle|0\rangle \mapsto \sum_{v,\ell: (u,v,\ell) \in \vec{E}(G)} \sqrt{1/d_G(u)} |u, v, \ell\rangle.$$

If G is subset of a complete graph, the quantum time complexity of $\text{CONN}_{G,X}$ is $O(n\sqrt{R/\kappa}U)$.

For any family of connected graphs G and $X \subseteq \{0,1\}^{E(G)}$ such that for all $x \in X$, if $G(x)$ is connected, $R_{\text{avg}}(G(x)) \leq R$, and if $G(x)$ is not connected, it has at least κ components, the quantum time complexity of $\text{CONN}_{G,X}$ is $O\left(n^{3/4} \sqrt{Rd_{\max}(G)}/\kappa^{1/4}U\right)$.

Proof. The algorithm of Theorem 14 makes $O\left(n^{3/4} \sqrt{Rd_{\max}(G)}/\kappa^{1/4}\right)$ calls to a unitary $U(P_G, x)$ (see [IJ16]). By [JK17] (generalizing [BR12]), for any G , $U(P_G, x)$ can be implemented in cost U . \square

5 Spectral Algorithm for Deciding Connectivity

In this section, we will give alternative quantum algorithms for deciding connectivity. We begin by presenting an algorithmic template, outlined in Algorithm 33, that requires the instantiation of a certain initial state. Since this initial state is independent of the input, we already get an upper bound on the quantum query complexity, as follows:

Corollary 27. Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices and $X \subseteq \{0,1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$ or $G(x)$ has at least κ connected components, the bounded error quantum query complexity of $\text{CONN}_{G,X}$ is $O\left(\sqrt{\frac{nd_{\max}(G)}{\kappa\lambda}}\right)$.

In Section 5.1, we describe one such initial state, and how to prepare it, leading to the following upper bound, in which U is the cost of performing one step of a quantum walk on G , and S is the cost of preparing a quantum state corresponding to the stationary distribution of a quantum walk on G :

Theorem 28. Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices and $X \subseteq \{0,1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$, or $G(x)$ has at least κ connected components, $\text{CONN}_{G,X}$ can be solved in bounded error in time

$$\tilde{O}\left(\sqrt{\frac{nd_{\text{avg}}(G)}{\kappa\lambda_2(G)}}\left(S + \sqrt{\frac{d_{\max}(G)}{\lambda}}U\right)\right).$$

In Section 5.2, we restrict our attention to the case when G is a Cayley graph, and give an alternative instantiation of the algorithm in Algorithm 33, proving the following, where Λ is the cost of computing the eigenvalues of G :

Theorem 29. Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices such that each G is a degree- d Cayley graph over an Abelian group, and $X \subseteq \{0,1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$ or $G(x)$ has at least κ connected components, $\text{CONN}_{G,X}$ can be solved in bounded error in time

$$\tilde{O}\left(\sqrt{\frac{nd}{\kappa\lambda}}U + \sqrt{\frac{nd}{\kappa\lambda_2(G)}}\Lambda\right).$$

We remark that the results in this section, in contrast to the previous connectivity algorithm, apply with respect to any weighting of the edges of G . Applying non-zero weights to the edges of G does not change which subgraphs $G(x)$ are connected, but it does impact the complexity of our algorithm. Thus, for any weight function on the edges, we get algorithms with the complexities given in Corollary 27, Theorem 28 and Theorem 29, where $d_{\max}(G)$ and $d_{\text{avg}}(G)$ are in terms of the weighted degrees, and $\lambda_2(G)$ and $\lambda_2(G(x))$ are in terms of the weighted Laplacians.

Finally, in Section 5.3, we describe how when G is a complete graph, these ideas can be used to design algorithms, not only for deciding connectivity, but also for estimating the *algebraic connectivity* of a graph, which is a measure of how connected a graph is. In particular, we show:

Theorem 30. Let G be the complete graph on n vertices. There exists a quantum algorithm that, on input x , with probability at least $2/3$, outputs an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_2(G(x))| \leq \varepsilon\lambda_2(G(x))$, where $\lambda_2(G(x))$ is the algebraic connectivity of $G(x)$, in time $\tilde{O}\left(\frac{1}{\varepsilon} \sqrt{\frac{n}{\lambda_2(G(x))}}\right)$.

Let $P_G = (H, U, A, \tau)$ be the span program for st -connectivity defined in Eq. (16). Note that only τ depends on s and t , and we will not be interested in τ here. We let $A(x) = A\Pi_{H(x)}$.

$$\begin{aligned}
A(x)A(x)^T &= \sum_{(u,v,\ell) \in \bar{E}(G(x))} \sqrt{c(u,v,\ell)}(|u\rangle - |v\rangle)\langle u, v, \ell | \sum_{(u,v,\ell) \in \bar{E}(G(x))} \sqrt{c(u,v,\ell)}|u, v, \ell\rangle(\langle u| - \langle v|) \\
&= \sum_{(u,v,\ell) \in \bar{E}(G(x))} c(u,v,\ell)(|u\rangle\langle u| - |u\rangle\langle v| - |v\rangle\langle u| + |v\rangle\langle v|) \\
&= \sum_{u \in [n]} 2d_G(u)|u\rangle\langle u| - 2\mathcal{A}_{G(x)} = 2(\mathcal{D}_{G(x)} - \mathcal{A}_{G(x)}) = 2L_{G(x)}. \tag{35}
\end{aligned}$$

Above, $L_{G(x)}$ is the Laplacian of $G(x)$ (see Section 2.2). By a similar computation to the one above, we have $AA^T = 2L_G$, where G is the parent graph, upon which A depends. Recall that for any G , the eigenvalues of L_G lie in $[0, d_{\max}]$, with $|\mu\rangle = \frac{1}{\sqrt{n}} \sum_v |v\rangle$ as a 0-eigenvalue. In our case, since G is assumed to be connected, $|\mu\rangle$ is the only 0-eigenvector of L_G , so $\text{row}(L_G)$ is the orthogonal complement of $|\mu\rangle$. For any x , $G(x)$ also has $|\mu\rangle$ as a 0-eigenvalue, but if $G(x)$ is connected, this is the only 0-eigenvalue. In general, the dimension of the 0-eigenspace of $L_{G(x)}$ is the number of components of $G(x)$. Thus, Eq. (35) implies the following.

- The multiset of nonzero eigenvalues of L_G are exactly half of the squared singular values of A , and in particular, since no eigenvalue of L_G can be larger than the maximum degree of G , $\sigma_{\max}(A) \leq \sqrt{2d_{\max}(G)}$.
- The multiset of nonzero eigenvalues of $L_{G(x)}$ are exactly half the squared singular values of $A(x)$, and in particular, if $G(x)$ is connected, then $\sigma_{\min}(A(x)) = \sqrt{2\lambda_2(G(x))}$, where $\lambda_2(G(x))$ is the second smallest eigenvalue of $L_{G(x)}$, which is non-zero if and only if $G(x)$ is connected.
- The support of L_G is $\text{col}(A)$, which is the orthogonal subspace of the uniform vector $|\mu\rangle = \frac{1}{\sqrt{n}} \sum_v |v\rangle$.

For a particular span program P , and input x , an associated unitary $U(P, x) = (2\Pi_{\ker A} - I)(2\Pi_{H(x)} - I)$ can be used to construct quantum algorithms, for example, for deciding the span program. Then by [IJ16, Theorem 3.10], which states that $\Delta(U(P, x)) \geq 2\sigma_{\min}(A(x))/\sigma_{\max}(A)$, we have the following.

Lemma 31. *Let P_G be the st -connectivity span program from Eq. (16). Then $\Delta(U(P, x)) \geq 2\sqrt{\lambda_2(G(x))/d_{\max}(G)}$.*

Our algorithm will be based on the following connection between the connectivity of $G(x)$ and the presence of a 0-phase eigenvector of $U(P, x)$ in $\text{row}(A)$.

Lemma 32. *$G(x)$ is not connected if and only if there exists $|\psi\rangle \in \text{row}(A)$ that is fixed by $U(P, x)$. Moreover, if $G(x)$ has $\kappa > 1$ components, there exists a $(\kappa - 1)$ -dimensional subspace of $\text{row}(A)$ that is fixed by $U(P, x)$.*

Proof. If $G(x)$ is not connected, and in particular, $G(x)$ has $\kappa > 1$ components, then

$$\text{rank}(A(x)) = \text{rank}(A(x)A(x)^T) = \text{rank}(L_{G(x)}) = n - \kappa < n - 1 = \text{rank}(A). \tag{36}$$

Eq. (36) implies $\ker A \subseteq \ker A(x)$, which means $\ker A(x) \cap \text{row}(A)$ has dimension $\kappa - 1 \geq 1$. Let $|\psi\rangle \in \ker A(x) \cap \text{row}(A)$. Since $|\psi\rangle \in \ker A(x)$, $A\Pi_{H(x)}|\psi\rangle = 0$, so $\Pi_{H(x)}|\psi\rangle \in \ker A$. Since $|\psi\rangle \in \text{row}(A)$, $|\psi\rangle = \Pi_{\text{row}(A)}|\psi\rangle = \Pi_{\text{row}(A)}\Pi_{H(x)^\perp}|\psi\rangle$, so $|\psi\rangle \in H(x)^\perp$. Thus since $|\psi\rangle \in H(x)^\perp \cap \text{row}(A)$, it follows that $U(P, x)|\psi\rangle = |\psi\rangle$.

On the other hand, if $G(x)$ is connected, then both $A(x)$ and A have rank $n - 1$, so $\ker A = \ker A(x)$ because $\ker A \subseteq \ker A(x)$. If $|\psi\rangle \in \text{row}(A)$ is fixed by $U(P, x)$, then $|\psi\rangle \in H(x)^\perp$, meaning $A(x)|\psi\rangle = 0$, hence $|\psi\rangle \in \ker A(x) = \ker A$, which is a contradiction. \square

Thus, to determine if $G(x)$ is connected, it is sufficient to detect the presence of *any* 0-phase eigenvector of $U(P, x)$ on $\text{row}(A)$. Let $\{|\psi_i\rangle\}_{i=1}^{n-1}$ be any basis for $\text{row}(A)$, not necessarily orthogonal, and suppose we have access to an operation that generates

$$|\psi_{\text{init}}\rangle = \sum_{i=1}^{n-1} |i\rangle |\psi_i\rangle.$$

Such a basis is independent of the input, so we can certainly perform such a map with 0 queries. We will later discuss cases in which we can implement such a map time efficiently.

Algorithm 33. Assume there is a known constant λ such that if $G(x)$ is connected, then $\lambda_2(G(x)) \geq \lambda$. Let $\{|\psi_i\rangle\}_i$ be some states that span the row space of A , whose choice determines the cost of the amplitude estimation step.

1. Prepare $|\psi_{\text{init}}\rangle = \sum_{i=1}^{n-1} \frac{1}{\sqrt{n-1}} |i\rangle |\psi_i\rangle$.
2. Perform the phase estimation of $U(P, x)$ (see Theorem 8) on the second register, to precision $\sqrt{\lambda/d_{\max}(G)}$, and accuracy ϵ .
3. Use amplitude estimation (see Theorem 10) to determine if the amplitude on $|0\rangle$ in the phase register is 0, in which case, output “connected”, or > 0 , in which case, output “not connected.”

The algorithm proceeds by first preparing the initial state $|\psi_{\text{init}}\rangle = \sum_{i=1}^{n-1} \frac{1}{\sqrt{n-1}} |i\rangle |\psi_i\rangle$. Next, the algorithm performs phase estimation on the second register, as described in Theorem 8, with precision $\sqrt{\lambda/d_{\max}(G)}$ and accuracy ϵ . First, suppose that there are $\kappa - 1 > 0$ orthonormal 0-phase eigenvectors of $U(P, x)$ in $\text{row}(A)$, and let Π be the orthonormal projector onto their span. By Theorem 8, for each i , the phase estimation step will map $|i\rangle (\Pi|\psi_i\rangle)$ to $|i\rangle |0\rangle (\Pi|\psi_i\rangle)$. Thus, the squared amplitude on $|0\rangle$ in the phase register will be at least:

$$\epsilon := \frac{\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2}{\|\psi_{\text{init}}\rangle\|^2} = \frac{1}{\|\psi_{\text{init}}\rangle\|^2} \sum_{i=1}^{n-1} \|\Pi|\psi_i\rangle\|^2 > 0, \quad (37)$$

since the $|\psi_i\rangle$ span $\text{row}(A)$.

On the other hand, suppose $G(x)$ is connected, so there is no 0-phase eigenvector in $\text{row}(A)$. Then all phases will be at least $\Delta(U(P, x)) \geq \sqrt{\lambda/d_{\max}(G)}$, by Lemma 31, so the phase register will have squared overlap at most ϵ with $|0\rangle$.

Setting $\epsilon = \epsilon/2$, we just need to distinguish between an amplitude of $\geq \epsilon$ and an amplitude of $\leq \epsilon/2$ on $|0\rangle$. Using Corollary 11, we can distinguish these two cases in $\frac{1}{\sqrt{\epsilon}}$ calls to steps 1 and 2. By Theorem 8,

Step 2 can be implemented using $\sqrt{\frac{d_{\max}(G)}{\lambda}} \log \frac{1}{\epsilon}$ calls to $U(P, x)$. By [JK17, Theorem 13], if U is the cost of implementing, for any $u \in V$, the map

$$|u, 0\rangle \mapsto \sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell)/d_G(u)} |u, v, \ell\rangle, \quad (38)$$

which corresponds to one step of a quantum walk on G , then $U(P, x)$ can be implemented in time $O(U)$. We thus get the following:

Theorem 34. Fix $\lambda > 0$. Let Init denote the cost of generating the initial state $|\psi_{\text{init}}\rangle$, and U the cost of the quantum walk step in Eq. (38). Let ϵ be as in Eq. (37). Then for any family of connected graphs G and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$ or $G(x)$ is not connected, $\text{CONN}_{G, X}$ can be decided by a quantum algorithm with cost $O\left(\frac{1}{\sqrt{\epsilon}} \left(\text{Init} + \sqrt{\frac{d_{\max}(G)}{\lambda}} U \log \frac{1}{\epsilon}\right)\right)$.

In Section 5.1 and Section 5.2, we will discuss particular implementations of this algorithm, but if we only care about query complexity, we already have the following.

Corollary 27. Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$ or $G(x)$ has at least κ connected components, the bounded error quantum query complexity of $\text{CONN}_{G, X}$ is $O\left(\sqrt{\frac{nd_{\max}(G)}{\kappa\lambda}}\right)$.

Proof. First, observe that $U(P, x)$ can be implemented with 2 queries.

Next, let $\{|\psi_i\rangle\}_{i=1}^{n-1}$ be any orthonormal basis for $\text{row}(A)$. Then in $\text{Init} = 0$ queries, we can generate the state

$$|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{n-1}} \sum_{i=1}^{n-1} |i\rangle |\psi_i\rangle.$$

Then if there are $\kappa - 1$ orthonormal 0-phase vectors of $U(P, x)$ in $\text{row}(A)$, $|\phi_1\rangle, \dots, |\phi_{\kappa-1}\rangle$, setting $\Pi = \sum_{j=1}^{\kappa-1} |\phi_j\rangle\langle\phi_j|$, we have

$$\varepsilon = \|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 = \frac{1}{n-1} \sum_{j=1}^{\kappa-1} \sum_{i=1}^{n-1} |\langle\phi_j|\psi_i\rangle|^2 = \frac{\kappa-1}{n-1}.$$

Then the result follows from Theorem 34. \square

5.1 A Connectivity Algorithm for any G

Let $|\mu\rangle = \frac{1}{\sqrt{n}} \sum_{u \in V(G)} |u\rangle$ and for $j \in \{1, \dots, n-1\}$ let $|\hat{j}\rangle = \frac{1}{\sqrt{n}} \sum_{u \in [n]} e^{2\pi i j u/n} |u\rangle$. Then it is easily checked that $|\hat{1}\rangle, \dots, |\widehat{n-1}\rangle$ are an orthonormal basis for the column space of A for a connected graph G . Let $|\psi_j\rangle = A^T |\hat{j}\rangle$. Then $\{|\psi_1\rangle, \dots, |\psi_{n-1}\rangle\}$ is a basis for $\text{row}(A)$, but it is not necessarily orthogonal unless the $|\hat{j}\rangle$ form an eigenbasis of L_G , (as in the case of Cayley graphs, discussed in Section 5.2), and in general, the $|\psi_j\rangle$ are not normalized. We have

$$\begin{aligned} |\psi_j\rangle = A^T |\hat{j}\rangle &= \sum_{(u,v,\ell) \in \vec{E}(G)} \sqrt{c(u,v,\ell)} |u,v,\ell\rangle (\langle u|\hat{j}\rangle - \langle v|\hat{j}\rangle) \\ &= \sum_{(u,v,\ell) \in \vec{E}(G)} \sqrt{c(u,v,\ell)} \left(\frac{1}{\sqrt{n}} e^{2\pi i j u/n} - \frac{1}{\sqrt{n}} e^{2\pi i j v/n} \right) |u,v,\ell\rangle, \end{aligned} \quad (39)$$

from which we can compute

$$\begin{aligned} \|\psi_j\|^2 = \|A^T |\hat{j}\rangle\|^2 &= \frac{1}{n} \sum_{(u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) \left| e^{2\pi i j u/n} - e^{2\pi i j v/n} \right|^2 \\ &= \frac{1}{n} \sum_{(u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) 2 \left(1 - \cos \frac{2\pi j(v-u)}{n} \right). \end{aligned} \quad (40)$$

For any graph, we can use as initial state a normalization of $\sum_{j=1}^{n-1} |j\rangle |\psi_j\rangle$. From Eq. (40), we have, using Lagrange's identity:

$$\begin{aligned} \sum_{j=1}^{n-1} \|\psi_j\|^2 &= \frac{2}{n} \sum_{(u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) \sum_{j=1}^{n-1} \left(1 - \cos \frac{2\pi j(v-u)}{n} \right) \\ &= \frac{2}{n} \sum_{(u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) \left(n-1 - \left(\frac{1}{2} \frac{\sin \left((n-1/2) \frac{2\pi(v-u)}{n} \right)}{\sin \left(\frac{\pi(v-u)}{n} \right)} - 1 \right) \right) \\ &= \frac{2}{n} \sum_{(u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) \left(n - \frac{\sin \left(-\frac{\pi(v-u)}{n} \right)}{2 \sin \left(\frac{\pi(v-u)}{n} \right)} \right) \\ &= \frac{2}{n} \left(n + \frac{1}{2} \right) \sum_{u \in V} \sum_{v,\ell: (u,v,\ell) \in \vec{E}(G)} c(u,v,\ell) \\ &= 2 \left(1 + \frac{1}{2n} \right) \sum_{u \in V} d_G(u) = 2 \left(1 + \frac{1}{2n} \right) n d_{\text{avg}}(G), \end{aligned} \quad (41)$$

where $d_{\text{avg}} = d_{\text{avg}}(G)$ is the average weighted degree in G .

Define the initial state as the unit vector:

$$|\psi_{\text{init}}\rangle = \frac{\sum_{j=1}^{n-1} |j\rangle |\psi_j\rangle}{\sqrt{2(1 + \frac{1}{2n})nd_{\text{avg}}}}. \quad (42)$$

Then we can lower bound the overlap with the 0-phase space of $U(P, x)$ in the case where $G(x)$ is not connected, as follows:

Lemma 35. *Let $|\psi_{\text{init}}\rangle$ be as in Eq. (42), and suppose $G(x)$ has $\kappa > 1$ connected components. Let Π be the projector onto a $(\kappa - 1)$ -dimensional subspace of $\text{row}(A)$ that is in the 0-phase space of $U(P, x)$. Then*

$$\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 \geq \frac{(\kappa - 1)\lambda_2(G)}{(1 + \frac{1}{2n})nd_{\text{avg}}}.$$

Proof. By Lemma 32, there exist orthonormal vectors $|\phi_1\rangle, \dots, |\phi_{\kappa-1}\rangle \in \text{row}(A)$ that are fixed by $U(P, x)$. Let $\Pi = \sum_{i=1}^{\kappa-1} |\phi_i\rangle\langle\phi_i|$. We have:

$$\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 = \frac{\sum_{i=1}^{\kappa-1} \sum_{j=1}^{n-1} |\langle\phi_i|\psi_j\rangle|^2}{2(1 + \frac{1}{2n})nd_{\text{avg}}} = \frac{\sum_{i=1}^{\kappa-1} \sum_{j=1}^{n-1} |\langle\phi_i|A^T|\hat{j}\rangle|^2}{2(1 + \frac{1}{2n})nd_{\text{avg}}} = \frac{\sum_{i=1}^{\kappa-1} \|A|\phi_i\rangle\|^2}{2(1 + \frac{1}{2n})nd_{\text{avg}}}.$$

Since $|\phi_i\rangle \in \text{row}(A)$ for each i , $\|A|\phi_i\rangle\|^2 \geq \sigma_{\min}(A)^2$. Thus

$$\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 \geq \frac{(\kappa - 1)\sigma_{\min}(A)^2}{2(1 + \frac{1}{2n})nd_{\text{avg}}} = \frac{(\kappa - 1)2\lambda_2(G)}{2(1 + \frac{1}{2n})nd_{\text{avg}}}. \quad (43)$$

Here we used the fact that G is assumed to be connected, the second smallest eigenvalue of $2L_G$ is the smallest non-zero eigenvalue, so $\lambda_2(2L_G) = \lambda_2(AA^T) = \sigma_{\min}(A)^2$. \square

We remark that although this initial state lends itself well to analysis, it might be a particularly bad choice of an initial state, because it ensures lower weight on lower eigenvalue eigenstates of G , which may have high overlap with the 0-eigenvalue eigenstates of $G(x)$.

We next describe how we can construct the initial state.

Lemma 36. *Let S be the cost of generating the stationary state of the graph G*

$$\sum_{u \in V} \sqrt{\frac{d_G(u)}{d_{\text{avg}}n}} |u\rangle.$$

Let U be the cost of implementing a step of the quantum walk on G , that is, generating, for any $u \in V$, a state of the form

$$\sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell) / d_G(u)} |v, \ell\rangle.$$

Then the map $|0\rangle \mapsto |\psi_{\text{init}}\rangle$ can be implemented with error probability at most ϵ in time complexity $O((S + U + \log n) \log \frac{1}{\epsilon})$.

Proof. We have

$$\begin{aligned} |\psi_j\rangle &= A^T|\hat{j}\rangle = \sum_{(u, v, \ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell)} \frac{1}{\sqrt{n}} (e^{2\pi iju/n} - e^{2\pi ijv/n}) |u, v, \ell\rangle \\ &= \frac{1}{\sqrt{n}} \sum_{u \in V} e^{2\pi iju/n} |u\rangle \sum_{\substack{v, \ell: \\ (u, v, \ell) \in \vec{E}(G)}} \sqrt{c(u, v, \ell)} |v, \ell\rangle - \frac{1}{\sqrt{n}} \sum_{u \in V} |u\rangle \sum_{\substack{v, \ell: \\ (u, v, \ell) \in \vec{E}(G)}} e^{2\pi ijv/n} \sqrt{c(u, v, \ell)} |v, \ell\rangle. \end{aligned} \quad (44)$$

We first note that we can generate the state

$$\sum_{u \in V} \sqrt{\frac{d_G(u)}{nd_{\text{avg}}}} |u\rangle \quad (45)$$

in cost S from which we can generate, for any $j \in [n]$,

$$\sum_{u \in V} e^{2\pi i j u / n} \sqrt{\frac{d_G(u)}{nd_{\text{avg}}}} |u\rangle \quad (46)$$

using a generalized Z_n^j gate, which performs the map $|u\rangle \mapsto e^{2\pi i u j / n} |u\rangle$ with complexity $O(\log n)$. From this, with one step of the quantum walk, we can get

$$|\alpha_j\rangle = \frac{1}{\sqrt{nd_{\text{avg}}}} \sum_{u \in V} e^{2\pi i j u / n} |u\rangle \sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell)} |v, \ell\rangle \quad (47)$$

in cost U . The total cost of constructing $|\alpha_j\rangle$ is $O(S + U + \log n)$.

Next, we can construct the state

$$\sum_{u \in V} \sqrt{\frac{d_G(u)}{nd_{\text{avg}}}} |u\rangle$$

in cost S , from which we can construct

$$\frac{1}{\sqrt{nd_{\text{avg}}}} \sum_{u \in V} |u\rangle \sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} \sqrt{c(u, v, \ell)} |v, \ell\rangle$$

in cost U . Finally, for any $j \in [n]$, applying a generalized Z_n^j gate on the second register, we get

$$|\beta_j\rangle = \frac{1}{\sqrt{nd_{\text{avg}}}} \sum_{u \in V} |u\rangle \sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} e^{2\pi i j v / n} \sqrt{c(u, v, \ell)} |v, \ell\rangle,$$

for total cost $O(U + S + \log n)$. One can now see that $|\alpha_j\rangle - |\beta_j\rangle = \frac{1}{\sqrt{d_{\text{avg}}}} |\psi_j\rangle$.

To construct $|\psi_{\text{init}}\rangle$, generate the state:

$$\frac{1}{\sqrt{2(n-1)}} |0\rangle \sum_{j=1}^{n-1} |j\rangle |\alpha_j\rangle - \frac{1}{\sqrt{2(n-1)}} |1\rangle \sum_{j=1}^{n-1} |j\rangle |\beta_j\rangle.$$

This costs $O(S + U + \log n)$. Next, apply a Hadamard gate to the first register to get:

$$\begin{aligned} & \frac{1}{2\sqrt{n-1}} |0\rangle \sum_{j=1}^{n-1} |j\rangle (|\alpha_j\rangle - |\beta_j\rangle) + \frac{1}{2\sqrt{n-1}} |1\rangle \sum_{j=1}^{n-1} |j\rangle (|\alpha_j\rangle + |\beta_j\rangle) \\ &= \frac{1}{2\sqrt{(n-1)d_{\text{avg}}}} |0\rangle \sum_{j=1}^{n-1} |j\rangle |\psi_j\rangle + \frac{1}{2\sqrt{n-1}} |1\rangle \sum_{j=1}^{n-1} |j\rangle (|\alpha_j\rangle + |\beta_j\rangle). \end{aligned} \quad (48)$$

By Eq. (41), we have $\left\| \sum_{j=1}^{n-1} |j\rangle |\psi_j\rangle \right\| = \sqrt{2(1 + 1/(2n))nd_{\text{avg}}}$, so the amplitude on the $|0\rangle$ part of the state is at least $\frac{1}{\sqrt{2}}$. Thus, we can measure the first register, and post select on measuring $|0\rangle$ to obtain $|\psi_{\text{init}}\rangle$. With $\log \frac{1}{\epsilon}$ repetitions, we succeed with probability $1 - \epsilon$. \square

We can now give an upper bound on the complexity of deciding connectivity for any family of parent graphs G , in terms of U and S , in Theorem 28, below. We first note that it is reasonable to assume that these costs should be low in many natural cases. The cost U is the cost of implementing a step of a quantum walk on G , and note that G is input-independent, so as long as it is sufficiently structured, this shouldn't be a particularly large cost. For example, if for any vertex in G , we can efficiently query its degree, and its i^{th} neighbour for any i , then $U = O(\log n)$. Note that this is *not* the same as assuming we can efficiently query the i^{th} neighbour of a vertex in $G(x)$, which is not an operation that we can easily implement in the edge-query input model. Similarly, we might hope that S is also $O(\log n)$ in many cases of interest. Indeed, whenever G is d -regular, it's simply the cost of generating the uniform superposition over all vertices.

Theorem 28. *Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$, or $G(x)$ has at least κ connected components, $\text{CONN}_{G,X}$ can be solved in bounded error in time*

$$\tilde{O} \left(\sqrt{\frac{nd_{\text{avg}}(G)}{\kappa\lambda_2(G)}} \left(S + \sqrt{\frac{d_{\text{max}}(G)}{\lambda}} U \right) \right).$$

Proof. By Lemma 36, the complexity of generating $|\psi_{\text{init}}\rangle$ is $\text{Init} = O(S + U + \log n)$, and by Lemma 35, the initial state has overlap at least $\varepsilon = \Omega\left(\frac{\kappa\lambda_2(G)}{nd_{\text{avg}}}\right)$ with any unit vector in $\ker A(x) \cap \text{row}(A)$. Plugging these values into the expression in Theorem 34 gives (neglecting polylogarithmic factors)

$$O \left(\frac{1}{\sqrt{\varepsilon}} \left(\text{Init} + \sqrt{\frac{d_{\text{max}}(G)}{\lambda}} U \right) \right) = \tilde{O} \left(\sqrt{\frac{nd_{\text{avg}}}{\kappa\lambda_2(G)}} \left(S + \sqrt{\frac{d_{\text{max}}(G)}{\lambda}} U \right) \right). \quad \square$$

5.2 An Algorithm for Cayley Graphs

When the parent graph G is a Cayley graph for a finite Abelian group, we can use the extra structure to construct an orthonormal basis of $\text{row}(A)$. We first define a Cayley graph. Let Γ be a finite Abelian group, and S a symmetric subset of Γ , meaning that if $g \in S$, then $-g \in S$. The Cayley graph $\text{Cay}(\Gamma, S)$ is the graph that has Γ as its vertex set, and edge set $\{\{a, b\} : b - a \in S\}$.

For a positive integer m , let $\omega_m = e^{2\pi i/m}$. For an Abelian group $\Gamma = \mathbb{Z}/m_1\mathbb{Z} \times \cdots \times \mathbb{Z}/m_k\mathbb{Z}$ and an element $g \in \Gamma$, we define the character $\chi_g : G \rightarrow \mathbb{C}$ as the function $\chi_g(s) = \omega_{m_1}^{g_1 s_1} \cdots \omega_{m_k}^{g_k s_k}$.

When $G = \text{Cay}(\Gamma, S)$ is a Cayley graph, with $n = |\Gamma|$ and $d = |S|$, it is easily verified that the eigenvectors of L_G are exactly the Fourier vectors:

$$|\hat{g}\rangle = \frac{1}{\sqrt{n}} \sum_{h \in \Gamma} \chi_g(h) |h\rangle.$$

For $g \neq 0$, these are also the left-singular vectors of A . Most importantly, the vectors $A^T |\hat{g}\rangle$ for $g \neq 0$ are an orthogonal basis of $\text{row}(A)$ since they are proportional to the right-singular vectors of A . We define

$$\begin{aligned} |\psi_g\rangle &= A^T |\hat{g}\rangle = \sum_{u \in \Gamma} \sum_{v: v-u \in S} \sum_{h \in \Gamma} \frac{\chi_g(h)}{\sqrt{n}} |u, v\rangle (\langle u| - \langle v|) |h\rangle = \sum_{u \in \Gamma} \sum_{v: v-u \in S} \frac{1}{\sqrt{n}} (\chi_g(u) - \chi_g(v)) |u, v\rangle \\ &= \sum_{u \in \Gamma} \frac{\chi_g(u)}{\sqrt{n}} \sum_{s \in S} (1 - \chi_g(s)) |u, u+s\rangle. \end{aligned} \quad (49)$$

We have $\| |\psi_g\rangle \|^2 = \| A^T |\hat{g}\rangle \|^2 = \lambda_g$, where λ_g is the eigenvalue of L_G associated with $|\hat{g}\rangle$. In particular, $\lambda_g = d - \sum_{s \in S} \chi_g(s)$ (See for example [Bol13]).

We define

$$|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{n-1}} \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle |\psi_g\rangle. \quad (50)$$

We first lower bound the overlap of the initial state with the 0-phase space of $U(P, x)$ when $G(x)$ is not connected.

Lemma 37. Suppose $G(x)$ has at least $\kappa > 1$ components. Let $|\psi_{\text{init}}\rangle$ be as in Eq. (50), and let $\{|\phi_i\rangle\}_{i=1}^{\kappa-1}$ be orthonormal 0-phase vectors of $U(P, x)$ in $\text{row}(A)$. Let $\Pi = \sum_{i=1}^{\kappa-1} |\phi_i\rangle\langle\phi_i|$. Then

$$\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 \geq \frac{\kappa - 1}{n - 1}.$$

Proof. Let $|\bar{\psi}_g\rangle = |\psi_g\rangle / \sqrt{\lambda_g}$. Then $\{|\bar{\psi}_g\rangle\}_{g \neq 0}$ is an orthonormal basis for $\text{row}(A)$. We have

$$\|(I \otimes \Pi)|\psi_{\text{init}}\rangle\|^2 = \frac{1}{n-1} \sum_{i=1}^{\kappa-1} \sum_{g \in \Gamma \setminus \{0\}} |\langle\phi_i|\bar{\psi}_g\rangle|^2 = \frac{1}{n-1} \sum_{i=1}^{\kappa-1} 1 = \frac{\kappa - 1}{n - 1}. \quad \square \quad (51)$$

Next, we give an upper bound on the time complexity of constructing the initial state. We will use the following fact:

Claim 38 ((See, for example, [Bol13])). Let G be any connected graph, with non-zero eigenvalues $\lambda_2, \dots, \lambda_n$. Then $R_{\text{avg}}(G) = \frac{1}{n-1} \sum_{i=2}^n \frac{1}{\lambda_i}$.

Lemma 39. Let U be the cost of generating the state $\frac{1}{\sqrt{d}} \sum_{s \in S} |s\rangle$. Let Λ be the cost of implementing, for $g \in \Gamma$, $|g\rangle|0\rangle \mapsto |g\rangle|\lambda_g\rangle$. Then the cost of generating the state $|\psi_{\text{init}}\rangle$ as in Eq. (50) with success probability $1 - \epsilon$ is

$$O\left(\left(\sqrt{R_{\text{avg}}(G)d}(U + \log n) + \Lambda\sqrt{\frac{d}{\lambda_2(G)}}\right) \log \frac{1}{\epsilon}\right).$$

Proof. The proof is similar to that of Lemma 36. Using a Fourier transform, we can generate the state

$$|g\rangle \mapsto \frac{1}{\sqrt{n}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \quad (52)$$

for any $g \in \Gamma$ in time $O(\log n)$. We can then generate $\sum_{s \in S} \frac{1}{\sqrt{d}} |s\rangle$ in time U , and perform the map

$$\frac{1}{\sqrt{n}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \sum_{s \in S} \frac{1}{\sqrt{d}} |s\rangle \mapsto \frac{1}{\sqrt{dn}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \sum_{s \in S} |u+s\rangle =: |\alpha_g\rangle \quad (53)$$

for a total complexity of $O(\log n + U)$ to generate $|\alpha_g\rangle$.

Alternatively, we can use the generalized Z_Γ^g gate, which maps $|s\rangle$ to $\chi_g(s)|s\rangle$ in time $O(\log n)$ to get

$$\frac{1}{\sqrt{nd}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \sum_{s \in S} |s\rangle \mapsto \frac{1}{\sqrt{nd}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \sum_{s \in S} \chi_g(s) |s\rangle \mapsto \frac{1}{\sqrt{nd}} \sum_{u \in \Gamma} \chi_g(u) |u\rangle \sum_{s \in S} \chi_g(s) |u+s\rangle =: |\beta_g\rangle \quad (54)$$

for a total complexity of $O(\log n + U)$ to generate $|\beta_g\rangle$. Observe that $|\alpha_g\rangle - |\beta_g\rangle = \frac{1}{\sqrt{d}} |\psi_g\rangle$.

Now to construct $|\psi_{\text{init}}\rangle$, we first construct $\sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle$, as follows. We first generate:

$$\sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{n-1}} |g\rangle |\lambda_g\rangle, \quad (55)$$

in cost Λ . Next, we map this to:

$$\sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{n-1}} |g\rangle |\lambda_g\rangle \left(\sqrt{\frac{\lambda_2(G)}{\lambda_g}} |0\rangle + \sqrt{1 - \frac{\lambda_2(G)}{\lambda_g}} |1\rangle \right). \quad (56)$$

We can perform this map because for all $g \in \Gamma \setminus \{0\}$, $\lambda_2(G)/\lambda_g \leq 1$. We uncompute $|\lambda_g\rangle$, and then do amplitude amplification on $|0\rangle$ in the last register to get the desired state. The squared amplitude on $|0\rangle$ is:

$$\left\| \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{n-1}} \sqrt{\frac{\lambda_2(G)}{\lambda_g}} |g\rangle \right\|^2 = \frac{\lambda_2(G)}{n-1} \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\lambda_g} = \lambda_2(G) R_{\text{avg}}(G). \quad (57)$$

So we can generate the normalized state

$$\frac{1}{\sqrt{(n-1)R_{\text{avg}}(G)}} \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle \quad (58)$$

with constant success probability in time complexity $O(\Lambda/\sqrt{\lambda_2(G)R_{\text{avg}}(G)})$.

Next, we map this state to:

$$\mapsto \frac{1}{\sqrt{2(n-1)R_{\text{avg}}}} |0\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle |\alpha_g\rangle - \frac{1}{\sqrt{2(n-1)R_{\text{avg}}}} |1\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle |\beta_g\rangle \quad (59)$$

at an additional cost of $O(U + \log n)$. Next, we apply a Hadamard gate to the first qubit to get

$$\begin{aligned} & \frac{1}{2\sqrt{(n-1)R_{\text{avg}}}} |0\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle (|\alpha_g\rangle - |\beta_g\rangle) + \frac{1}{2\sqrt{(n-1)R_{\text{avg}}}} |1\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle (|\alpha_g\rangle + |\beta_g\rangle) \\ = & \frac{1}{2\sqrt{(n-1)R_{\text{avg}}d}} |0\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle |\psi_g\rangle + \frac{1}{2\sqrt{(n-1)R_{\text{avg}}}} |1\rangle \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle (|\alpha_g\rangle + |\beta_g\rangle). \end{aligned} \quad (60)$$

The total cost to make one copy of this state with constant success probability is $O(U + \log n + \Lambda/\sqrt{\lambda_2(G)R_{\text{avg}}})$.

We can then get $|\psi_{\text{init}}\rangle$ by doing amplitude amplification on the $|0\rangle$ part of this state. The amplitude on the $|0\rangle$ part of the state is given by

$$\left\| \frac{1}{2\sqrt{(n-1)R_{\text{avg}}d}} \sum_{g \in \Gamma \setminus \{0\}} \frac{1}{\sqrt{\lambda_g}} |g\rangle |\psi_g\rangle \right\| = \frac{1}{2\sqrt{R_{\text{avg}}d}},$$

so using $O(\sqrt{R_{\text{avg}}(G)d})$ rounds of amplitude amplification is sufficient to generate $|\psi_{\text{init}}\rangle$ with constant probability, for a total cost of (neglecting constants):

$$\sqrt{R_{\text{avg}}(G)d} \left(U + \log n + \frac{\Lambda}{\sqrt{\lambda_2(G)R_{\text{avg}}(G)}} \right) = \sqrt{R_{\text{avg}}(G)d} (U + \log n) + \Lambda \sqrt{\frac{d}{\lambda_2(G)}}. \quad (61)$$

We can amplify this to success probability $1 - \epsilon$ at the cost of a $\log(1/\epsilon)$ multiplicative factor. \square

Theorem 29. Fix any $\lambda > 0$ and $\kappa > 1$. For any family of connected graphs G on n vertices such that each G is a degree- d Cayley graph over an Abelian group, and $X \subseteq \{0, 1\}^{E(G)}$ such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$ or $G(x)$ has at least κ connected components, $\text{CONN}_{G,X}$ can be solved in bounded error in time

$$\tilde{O} \left(\sqrt{\frac{nd}{\kappa\lambda}} U + \sqrt{\frac{nd}{\kappa\lambda_2(G)}} \Lambda \right).$$

Proof. Combining Lemma 39 and Lemma 37 with Theorem 34, we get complexity:

$$\tilde{O} \left(\sqrt{ndR_{\text{avg}}(G)/\kappa} U + \sqrt{\frac{nd}{\lambda\kappa}} U + \sqrt{\frac{nd}{\lambda_2(G)\kappa}} \Lambda \right). \quad (62)$$

First, note that $\lambda_2(G(x)) \leq \lambda_2(G)$. Thus if $\lambda > \lambda_2(G)$, then there is no x such that $\lambda_2(G(x)) \leq \lambda$, so X is empty. Thus, we can assume $\lambda \leq \lambda_2(G)$. By Claim 38, we can see that $R_{\text{avg}}(G) \leq \frac{1}{\lambda_2(G)} \leq \frac{1}{\lambda}$. The claim follows. \square

We now look at specific examples where it is particularly efficient to compute λ_g , as well as prepare a step of the walk $\sum_{s \in S} |s\rangle$. We first consider the complete graph on n vertices, in which $\Gamma = \mathbb{Z}_n$, and $S = \Gamma \setminus \{0\}$.

Corollary 40. *Fix any $\lambda > 0$, and integer $\kappa > 1$ and let G be the complete graph. Let $X \subseteq E(G)$ be such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$, or $G(x)$ has at least κ components. Then $\text{CONN}_{G,X}$ can be solved in bounded error in time*

$$\tilde{O}\left(\frac{n}{\sqrt{\kappa\lambda}}\right).$$

Proof. It is easily verified that for G a complete graph, $L_G = (n-1)I - (J-I)$, where J is the all-ones matrix, so the eigenvalues consist of a single 0, and n with multiplicity $n-1$ — that is, all non-zero eigenvalues are n . Thus, the mapping $|g\rangle \mapsto |g\rangle|\lambda_g\rangle = |g\rangle|n\rangle$ can be implemented trivially in $O(\log n)$ complexity.

Next, we can generate the state $\sum_{s \in S} \frac{1}{\sqrt{d}} |s\rangle = \sum_{s=1}^{n-1} \frac{1}{\sqrt{n-1}} |s\rangle$ in complexity $S = O(\log n)$. Then the result follows from Theorem 29. \square

Next, we consider the Boolean hypercube, in which $\Gamma = \mathbb{Z}_2^d$, so $n = 2^d$, and $S = \{e_i\}_{i=1}^d$, where e_i is 0 everywhere except the i^{th} entry, which is 1.

Corollary 41. *Fix any $\lambda > 0$, and integer $\kappa > 1$ and let G be the Boolean hypercube on $n = 2^d$ vertices. Let $X \subseteq E(G)$ be such that for all $x \in X$, either $\lambda_2(G(x)) \geq \lambda$, or $G(x)$ has at least κ components. Then $\text{CONN}_{G,X}$ can be solved in bounded error in time*

$$\tilde{O}\left(\sqrt{\frac{n}{\kappa\lambda}}\right).$$

Proof. The eigenvalues of the Boolean hypercube are well known to be $\lambda_g = 2|g|$, for $g \in \mathbb{Z}_2^d$, where $|g|$ denotes the Hamming weight of g . Thus, the map $|g\rangle \mapsto |g\rangle|\lambda_g\rangle$ can be implemented in cost $O(\log n)$. Finally, the state $\sum_{s \in S} |s\rangle = \sum_{i=1}^d |e_i\rangle$ can be generated in time $O(\log n)$. Then by Theorem 29, the time complexity is (neglecting polylog factors):

$$\sqrt{\frac{nd}{\kappa\lambda}} = \tilde{O}\left(\sqrt{\frac{n}{\kappa\lambda}}\right). \quad (63)$$

\square

5.3 Estimating the connectivity when G is a complete graph

For the remainder of this section, let G be the complete graph on n vertices, K_n . In that case, we can not only decide if $G(x)$ is connected, but estimate $\lambda_2(G(x))$. The idea is to relate the smallest phase of $U(P, x)$ on $\text{row}(A)$ to $\lambda_2(G(x))$, and estimate this value using quantum phase estimation.

Let $\Delta(U(P, x))$ denote the smallest nonzero phase of $U(P, x)$, which we want to estimate. We will shortly show that there is a vector $|u\rangle$ in $\text{row}(A)$ in the $\pm\Delta(U(P, x))$ -phase space of $U(P, x)$. Then, if $\{|\psi_i\rangle\}_{i=1}^{n-1}$ is an orthonormal basis for $\text{row}(A)$, applying quantum phase estimation on the second register of

$$|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{n-1}} \sum_{i=1}^{n-1} |i\rangle |\psi_i\rangle \quad (64)$$

to some constant precision, and then applying amplitude estimation to determine if there is amplitude at least $\frac{1}{\sqrt{2(n-1)}}$ on phases less than $1/2$, we can distinguish between the case in which, say, $\Delta(U(P, x)) \leq 1/3$, and $\Delta(U(P, x)) \geq 2/3$. To get an accurate estimate of $\Delta(U(P, x))$, we will make repeated calls to such a phase-estimation-followed-by-amplitude-estimation subroutine, reducing the size of the interval where $\Delta(U(P, x))$ sits at every iteration. We will first describe the connection between $\Delta(U(P, x))$ and $\lambda_2(G(x))$, and then formally present the algorithm for estimating $\Delta(U(P, x))$.

Connection between $\lambda_2(G(x))$ and $\Delta(U(P, x))$ We use the following theorem, relating the phases of the product of two reflections $U = (2\Pi_A - I)(2\Pi_B - I)$ to the singular values of its discriminant, defined $\Pi_A \Pi_B$.

Theorem 42 ([Sze04]). Let $\Pi_A = \sum_{i=1}^a |\alpha_i\rangle\langle\alpha_i|$ and $\Pi_B = \sum_{i=1}^b |\beta_i\rangle\langle\beta_i|$ be orthogonal projectors into some subspaces of the same inner product space, and define $U = (2\Pi_A - I)(2\Pi_B - I)$. Let $D = \Pi_A \Pi_B$ be the discriminant of U , and suppose it has singular value decomposition $D = \sum_{j=1}^r \cos \theta_j |u_j\rangle\langle v_j|$ with $\theta_j \in [0, \pi/2)$. Then U has 1-eigenspace $(A \cap B) \oplus (A^\perp \cap B^\perp)$ and (-1) -eigenspace $(A \cap B^\perp) \oplus (A^\perp \cap B)$. The only other eigenvalues of U are exactly $\{e^{\pm 2i\theta_j}\}_{j=1}^r$, and for each j , the $e^{2i\theta_j}$ - and $e^{-2i\theta_j}$ -eigenvectors are, respectively, $|\theta_j^+\rangle = |v_j\rangle - e^{i\theta_j}|u_j\rangle$ and $|\theta_j^-\rangle = |v_j\rangle - e^{-i\theta_j}|u_j\rangle$.

We derive several consequences of this theorem, and specialize them to our particular setting.

Lemma 43. Let $U = (2\Pi_A - I)(2\Pi_B - I)$ and $D = \Pi_A \Pi_B$ be its discriminant. Then $\Delta(-U) = 2 \sin^{-1}(\sigma_{\min}(D))$. Moreover, when G is a complete graph on n vertices, we have for any x , $\lambda_2(G(x)) = n \sin^2(\Delta(U(P, x))/2)$.

Proof. Assume without loss of generality that $\sigma_{\min}(D) = \cos \theta_1$. Since \cos is a decreasing function in the interval $[0, \frac{\pi}{2}]$, it follows that $\theta_1 \geq \theta_j \forall j > 1$. By Theorem 42, the spectrum of U outside of its (± 1) -eigenspace is $\{e^{\pm 2i\theta_j}\}_{j=1}^r$, hence the biggest eigenphase smaller than π in absolute value must be $2\theta_1$. In other words, $\pi - 2\theta_1$ is the phase gap of $-U$.

Let α be the complementary angle to θ_1 , i.e. $\pi/2 = \alpha + \theta_1$. Then $2\alpha = \pi - 2\theta_1 = \Delta(-U)$. but $\cos(\theta_1) = \sin(\alpha)$, which means $\sigma_{\min}(D) = \sin(\alpha) = \sin(\Delta(-U)/2)$. The first result follows from applying the arcsine function on both sides of the last equality.

In particular, if we take $U = -U(P, x) = (2\Pi_{\text{row}(A)} - I)(2\Pi_{H(x)} - I)$, then we have

$$D = \Pi_{\text{row}(A)} \Pi_{H(x)} = A^+ A \Pi_{H(x)} = A^+ A(x),$$

with $\sin(\Delta(U(P, x))/2) = \sigma_{\min}(D)$.

We have $AA^T = 2L_G$, for G the complete graph. It is known that when G is a complete graph K_n , L_G has 0-eigenspace spanned by the uniform vector $|\mu\rangle = \frac{1}{\sqrt{n}} \sum_{u \in [n]} |u\rangle$, and moreover, we have

$$L_G = (n-1)I - (J - I) = nI - J = nI - n|\mu\rangle\langle\mu| = n \sum_{i=1}^{n-1} |b_i\rangle\langle b_i|,$$

where $\{|b_i\rangle\}_{i=1}^{n-1}$ is any orthonormal basis for $\text{span}\{|\mu\rangle\}^\perp$, which is $\text{col}(A)$. This implies that for some orthonormal basis for $\text{row}(A)$, $\{|\psi_i\rangle\}_{i=1}^{n-1}$:

$$A = \sum_{i=1}^{n-1} \sqrt{2n} |b_i\rangle\langle\psi_i|.$$

Thus

$$A^+ = \sum_{i=1}^{n-1} \frac{1}{\sqrt{2n}} |\psi_i\rangle\langle b_i|.$$

Next, note that since $L_G|\mu\rangle = 0$ for any G , and $2L_G = A(x)A(x)^T$, we have $A(x)^T|\mu\rangle = 0$, so the columnspace of $A(x)$ is in $\text{span}\{|\mu\rangle\}^\perp$, and in particular, if $G(x)$ is connected, it's exactly $\text{span}\{|\mu\rangle\}^\perp$. The basis $\{|b_i\rangle\}_{i=1}^{n-1}$ can be chosen to be any basis of $|\mu\rangle^\perp$, so let's choose it to be the right singular basis of $A(x)$. That is, there exist $|\phi_i\rangle$ and σ_i such that

$$A(x) = \sum_{i=1}^{n-1} \sigma_i |b_i\rangle\langle\phi_i|$$

is a singular value decomposition for $A(x)$. Then we have:

$$D = A^+ A(x) = \sum_{i=1}^{n-1} \frac{\sigma_i}{\sqrt{2n}} |\psi_i\rangle\langle\phi_i|.$$

Since $2L_{G(x)} = A(x)A(x)^T$, the σ_i are just the square roots of twice the nonzero eigenvalues $\lambda_2, \dots, \lambda_n$ of $L_{G(x)}$, so the singular values of D are

$$\left\{ \sqrt{\frac{2\lambda_2}{2n}}, \dots, \sqrt{\frac{2\lambda_n}{2n}} \right\}.$$

We conclude that $\sigma_{\min}(D) = \sqrt{\frac{\lambda_2(G(x))}{n}}$, which, combined with $\sigma_{\min}(D) = \sin(\Delta(U(P, x))/2)$, gives $\lambda_2(G(x)) = n \sin^2(\Delta(U(P, x))/2)$. \square

Another consequence of Theorem 42 is the following, which allows us to restrict our attention to $\text{row}(A)$ in searching for the smallest phase of $U(P, x)$:

Lemma 44. *Let $U = (2\Pi_A - I)(2\Pi_B - I)$, and let $|\Delta_+\rangle$ be a $\Delta(U)$ -phase eigenvector of U , and $|\Delta_-\rangle$ a $(-\Delta(U))$ -phase eigenvector of U . Then there exists a vector $|u\rangle$ in the support of A such that $|u\rangle \in \text{span}\{|\Delta_+\rangle, |\Delta_-\rangle\}$. In particular, if $|\Delta_\pm\rangle$ are $\pm\Delta(U(P, x))$ -phase eigenvectors of $U(P, x)$, then there exists a vector $|u\rangle$ in $\text{row}(A)$ such that $|u\rangle \in \text{span}\{|\Delta_+\rangle, |\Delta_-\rangle\}$.*

Proof. Let $\theta_j, |u_j\rangle, |\theta_j^+\rangle$ and $|\theta_j^-\rangle$ be as in Theorem 42, so in particular, $|u_j\rangle$ is in the support of Π_A . Note that for any j , we have

$$|u_j\rangle = \frac{1}{2i \sin \theta_j} (|\theta_j^+\rangle - |\theta_j^-\rangle). \quad (65)$$

In particular, this is true for the j such that θ_j is minimized, i.e. such that $\theta_j = \Delta(U)$. The statement follows. \square

Algorithm for estimating $\Delta(U(P, x))$ We will actually estimate the value $\tau = \Delta(U(P, x))/\pi$, getting an estimate in $[0, 1]$, which we will then transform into an estimate of $\lambda_2(G(x))$. At every iteration, c will denote a lower bound for τ and C will denote the current upper bound. At the beginning of the algorithm we have $c = 0, C = 1$, and every iteration will result in updating either C or c in such a manner that the new interval for τ is reduced by a fraction of $2/3$. The algorithm is described in Algorithm 45.

Algorithm 45. To begin, let $c = 0$ and $C = 1$.

1. Set $\varphi = \frac{C-c}{3}, \epsilon = \frac{1}{\sqrt{2n}}, \delta = c + \varphi$.
2. For $j = 1, \dots, 4 \log(n/\epsilon)$:
 - (a) Prepare $\sum_{i=1}^{n-1} \frac{1}{\sqrt{n-1}} |i\rangle |\psi_i\rangle |0\rangle_C |0\rangle_P$.
 - (b) Perform the gapped phase estimation algorithm $GPE(\varphi, \epsilon, \delta)$ of Theorem 9 applying $U(P, x)$ on the second register.
 - (c) Use amplitude estimation (see Theorem 10) to distinguish between the case when the amplitude on $|0\rangle_C$ is $\geq \frac{1}{\sqrt{n}}$, in which case output " $a_j = 0$ ", and the case where the amplitude is $\leq \frac{1}{\sqrt{2n}}$, in which case, output " $a_j = 1$ ".
3. Compute $\tilde{a} = \text{Maj}(a_1, \dots, a_{4 \log(n/\epsilon)})$. If the result is 0, set $C = \delta + \varphi$. If the result is 1, set $c = \delta$. If $C - c \leq 2\epsilon c$, then output $n \sin^2\left(\frac{\pi(C+c)}{4}\right)$. Otherwise, return to Step 1.

Analysis of the algorithm We say an iteration of the algorithm *succeeds* if $\tilde{a} = \text{Maj}(a_1, \dots, a_{4 \log(n/\epsilon)})$ correctly indicates whether the amplitude on $|0\rangle_C$ is $\geq \frac{1}{\sqrt{n}}$ or $\leq \frac{1}{\sqrt{2n}}$. This happens with probability $\Omega(1 - (\epsilon/n)^4)$. Since we will shortly see that the algorithm runs for at most $\tilde{O}\left(\frac{n}{\epsilon \sqrt{\lambda_2(G(x))}}\right) \leq \tilde{O}\left(\frac{n^2}{\epsilon}\right)$ steps, the probability that every iteration succeeds is at least

$$\Omega\left(\left(1 - (\epsilon/n)^4\right)^{(n/\epsilon)^2}\right) = \Omega\left(1 - (\epsilon/n)^4(n/\epsilon)^2\right) = \Omega\left(1 - (\epsilon/n)^2\right). \quad (66)$$

It is therefore reasonable to assume that every iteration succeeds, since this happens with high probability. We first prove that if every iteration succeeds, throughout the algorithm we have $\tau = \Delta(U(P, x))/\pi \in [c, C]$.

Lemma 46. *Let $\tau = \Delta(U(P, x))/\pi$. For any φ and δ , if $\tau \geq \delta + \varphi$, applying $GPE(\varphi, \epsilon, \delta)$ to $|\psi_{\text{init}}\rangle$ results in a state with amplitude at most $\frac{1}{\sqrt{2n}}$ on $|0\rangle_C$ in register C; and if $\tau \leq \delta$, this results in a state with amplitude at least $\frac{1}{\sqrt{n}}$ on $|0\rangle_C$ in register C. Thus, if every iteration succeeds, at every iteration, we have $\tau \in [c, C]$.*

Proof. First, suppose $\tau \geq \delta + \varphi$. By Lemma 32, when $G(x)$ is connected there is no vector in $\text{row}(A)$ in the 1-eigenspace of $U(P, x)$, so the 1-eigenspace of $U(P, x)$ is contained in $\ker(A) \cap H(x) \subseteq \ker A$. Thus each $|u_j\rangle$ is in the span of $e^{i\pi\theta}$ -eigenvectors of $U(P, x)$ with $|\theta| \geq \delta + \varphi$. Thus, applying $GPE(\varphi, \epsilon, \delta)$ will map each $|u_j\rangle_R|0\rangle_C|0\rangle_P$ to a state $\beta_0|0\rangle_C|\gamma_0\rangle_{PR} + \beta_1|1\rangle_C|\gamma_1\rangle_{PR}$ such that $|\beta_0| \leq \epsilon$. Then, by linearity, the total amplitude on $|0\rangle_C$ in register C will be at most $\epsilon = \frac{1}{\sqrt{2n}}$.

On the other hand, suppose $\tau \leq \delta$. By Lemma 44, there exists a vector $|u_1\rangle \in \text{row}(A)$ such that $|u_1\rangle$ is in the span of the $e^{\pm i\pi\tau}$ -eigenvectors of $U(P, x)$. Applying $GPE(\varphi, \epsilon, \delta)$ will map $|u_1\rangle_R|0\rangle_C|0\rangle_P$ to a state $\beta_0|0\rangle_C|\gamma_0\rangle_{PR} + \beta_1|1\rangle_C|\gamma_1\rangle_{PR}$ such that $|\beta_1| \leq \epsilon$, so $|\beta_0| \geq \sqrt{1 - \epsilon^2}$. Let $|u_2\rangle, \dots, |u_{n-1}\rangle$ be any orthonormal set such that $|u_1\rangle, \dots, |u_{n-1}\rangle$ is an orthonormal basis for $\text{row}(A)$. Then there exists some (unknown) orthonormal set $\{|\tilde{j}\rangle\}_{j=1}^{n-1}$ such that

$$|\psi_{\text{init}}\rangle = \frac{1}{\sqrt{n-1}} \sum_{j=1}^{n-1} |\tilde{j}\rangle |u_j\rangle |0\rangle_C |0\rangle_P. \quad (67)$$

So after applying $GPE(\varphi, \epsilon, \delta)$ to $|\psi_{\text{init}}\rangle$, the amplitude on $|0\rangle_C$ in register C will be at least $\sqrt{\frac{1-\epsilon^2}{n-1}} \geq \frac{1}{\sqrt{n}}$. This proves the first part of the statement.

By Corollary 11, we can distinguish the case when the amplitude on $|0\rangle_C$ is at least $\frac{1}{\sqrt{n}}$ or at most $\frac{1}{\sqrt{2n}}$ with bounded error using $O(\frac{\sqrt{p_0}}{p_0-p_1}) = O(\sqrt{n})$ calls to $GPE(\varphi, \frac{1}{\sqrt{2n}}, \delta)$ where $p_0 := \frac{1}{n}$ and $p_1 := \frac{1}{2n}$. Thus, by repeating the procedure $4 \log(n/\epsilon)$ times and taking the majority, we succeed at every iteration with high probability.

We now prove by induction that we always have $\tau \in [c, C]$, as long as every iteration succeeds. At the beginning of the first iteration, we have $[c, C] = [0, 1]$. Since $\Delta(U(P, x)) \in [0, \pi]$, $\tau \in [0, 1]$. Next, suppose in some arbitrary iteration, we have $\tau \in [c, C]$. If $\tau \geq \delta + \varphi$, then there will be amplitude at most $\frac{1}{\sqrt{2n}}$ on $|0\rangle_C$, and assuming the iteration succeeds, we will have $\tilde{a} = 1$. In that case, we will set $c = \delta \leq \delta + \varphi \leq \tau$, so we will still have $\tau \in [c, C]$. If $\tau \leq \delta$, then there will be amplitude at least $\frac{1}{\sqrt{n}}$ on $|0\rangle_C$, and assuming the iteration succeeds, we will have $\tilde{a} = 0$. In that case, we will set $C = \delta + \varphi \geq \delta \geq \tau$, so we will still have $\tau \in [c, C]$.

We finally consider what happens if $\delta \leq \tau \leq \delta + \varphi$. In that case, there is no guarantee on the output of amplitude estimation; it can either output 0 or 1. However, we can still use the result to update our bounds for τ . If we get $\tilde{a} = 1$, and set $c = \delta$, we have $\delta \leq \tau$, so $\tau \in [c, C]$. If we get $\tilde{a} = 0$, and set $C = \delta + \varphi$, we have $\delta + \varphi \geq \tau$, so $\tau \in [c, C]$. \square

Next, we analyze the running time of Algorithm 45.

Theorem 47. *With probability $\Omega(1 - (\epsilon/n)^2)$, Algorithm 45 will terminate after time $\tilde{O}\left(\frac{n}{\epsilon\sqrt{\lambda_2(G(x))}}\right)$.*

Proof. With probability $\Omega(1 - (\epsilon/n)^2)$, each of the first $(n/\epsilon)^2 \geq \tilde{O}\left(\frac{n}{\epsilon\sqrt{\lambda_2(G(x))}}\right)$ iterations of the algorithm will succeed, so we assume this to be the case. We first bound the number of (successful) iterations before the algorithm terminates. Figure 4 shows the interval $[c, C]$, which represents the algorithm's current state




Figure 4: The interval $[c, C]$

of knowledge of where $\tau = \Delta(U(P, x))/\pi$ lies. The values δ and $\delta + \varphi$ are at $1/3$ and $2/3$ of the interval, respectively. It is not difficult to convince ourselves that at every iteration, the interval $[c, C]$ will be $2/3$ the size it had in the previous iteration. Hence, since the interval initially has length 1, after k iterations, we will have an interval of size $(\frac{2}{3})^k$. The execution terminates when the interval becomes sufficiently small. Specifically, let T be the smallest integer such that $T \geq \frac{\log \frac{2}{\tau\epsilon}}{\log \frac{2}{3}}$, and let $[c, C]$ be the interval after T steps, so $C - c = (2/3)^T \leq \tau\epsilon/2$. Suppose $(2/3)^T = C - c \geq 2\epsilon c$, so $\tau \geq 4c$. This implies that $C \geq 4c$, so $C/c \geq 4$. We will argue that this is a contradiction.

First, suppose $c = 0$. That means that

$$\tau \leq C = (2/3)^T \leq \tau\epsilon/2 \leq \tau/2,$$

which is a contradiction, since $\tau > 0$. Thus, we must have $c > 0$. Consider the first setting of c and C such that $c \neq 0$. Since the previous value of c was 0, we set the new value as $c = \delta = 0 + \varphi = (C - 0)/3 = C/3$, so the ratio C/c satisfies $C/c = 3$. This ratio can only decrease, because subsequent steps either decrease C , or increase c . Thus, after T steps, $C/c \leq 3$. Thus, $C - c \geq 2\epsilon c$ leads to a contradiction, so we can conclude that after T steps, $C - c \leq 2\epsilon c$, so the algorithm terminates in at most T steps.

We can now analyze the total running time by adding up the cost of all iterations. Step 1 of the algorithm is defining the variables $\varphi = \frac{C-c}{4}$, $\delta = c + \varphi$ and $\epsilon = \frac{1}{\sqrt{2n}}$, which will contribute negligibly to the complexity.

Step 2(a) begins by constructing the initial state $|\psi_{\text{init}}\rangle = \sum_{i=1}^{n-1} \frac{1}{\sqrt{n-1}} |i\rangle |\psi_i\rangle |0\rangle_P |0\rangle_C$ where $\{|\psi_i\rangle\}_{i=1}^{n-1}$ is a basis of $\text{row}(A)$. Because $G = K_n$ can be seen as a particularly simple kind of Cayley graph with group $\Gamma = \mathbb{Z}/n\mathbb{Z}$ and $S = \Gamma \setminus \{0\}$, we can use the construction of Section 5.2 to generate $|\psi_{\text{init}}\rangle$. In fact, combining the remarks in the proof of Corollary 40 with Lemma 39 it follows that this state can be constructed in time $O(\log n \log \frac{n}{\epsilon})$ with success probability $1 - (\epsilon/n)^4$.

Step 2(b) consists of applying the unitary procedure $GPE(\varphi, \epsilon, \delta)$ described in Theorem 9 on the last three registers with $\varphi, \epsilon = \frac{1}{\sqrt{2n}}, \delta$ defined in Step 1. By Theorem 9, this makes $O(\varphi^{-1} \log \epsilon^{-1}) = O(\varphi^{-1} \log n)$ calls to $U(P, x)$, for a total query complexity of $O(\varphi^{-1} \log^2 n)$.

Step 2(c) then uses amplitude estimation, repeating Steps 2(a) and 2(b) $O(\sqrt{n})$ times, by Corollary 11. Let $\varphi^{(i)}$ denote the value of φ at the i^{th} iteration of the algorithm. Neglecting polylog(n/ϵ) factors, the running time of the i^{th} iteration is

$$Q_i := \frac{\sqrt{n}}{\varphi^{(i)}}. \quad (68)$$

During the i^{th} iteration, we begin with $C - c = (2/3)^{i-1}$, and so $\varphi^{(i)} = \frac{1}{3}(2/3)^{i-1}$. Thus, we can compute the total complexity of the algorithm as (neglecting polylogarithmic factors):

$$\sum_{i=1}^T Q_i = \sqrt{n} \sum_{i=1}^T 3(3/2)^{i-1} = 3\sqrt{n} \frac{(3/2)^T - 1}{3/2 - 1} = \tilde{O} \left(\sqrt{n} (3/2)^{\frac{\log(2/(\tau\epsilon))}{\log(3/2)}} \right) = \tilde{O} \left(\frac{\sqrt{n}}{\tau\epsilon} \right). \quad (69)$$

By Lemma 43, we have $\lambda_2(G(x))/n = \sin^2(\Delta(U(P, x))/2) \leq \Delta(U(P, x))^2/4$, so Filling in $\tau = \Delta(U(P, x))/\pi \geq \sqrt{\lambda_2(G(x))/(2n)}$, we get a total query complexity of $\tilde{O} \left(\frac{n}{\epsilon \sqrt{\lambda_2(G(x))}} \right)$. \square

Finally, we prove that the algorithm outputs an estimate that is within ϵ multiplicative error of $\lambda_2(G(x))$.

Theorem 48 (Correctness). *With probability at least $\Omega(1 - (\epsilon/n)^2)$, Algorithm 45 outputs an estimate $\tilde{\lambda}$ such that $|\lambda_2(G(x)) - \tilde{\lambda}| \leq \frac{\pi^2 3}{4} \epsilon \lambda_2(G(x))$.*

Proof. We will assume that all iterations succeed, which happens with probability at least $\Omega(1 - (\epsilon/n)^2)$. Then the algorithm outputs $\tilde{\lambda} := n \sin^2 \left(\frac{\pi(C+c)}{4} \right)$ for some c and C such that $c \leq \tau \leq C$, and $C - c \leq 2\epsilon c \leq 2\epsilon\tau$. Using $\tau = \Delta(U(P, x))/\pi$ and $\lambda_2(G(x)) = n \sin^2(\Delta(U(P, x))/2)$, we have:

$$|\lambda_2(G(x)) - \tilde{\lambda}| = \left| n \sin^2(\pi\tau/2) - n \sin^2(\pi(C+c)/4) \right|. \quad (70)$$

From $c \leq \tau \leq C$ and $C - c \leq 2\varepsilon\tau$, we have

$$\left| \frac{\pi(C+c)}{4} - \frac{\pi\tau}{2} \right| \leq \frac{\pi\varepsilon\tau}{2}. \quad (71)$$

Let $\delta = \pi(C+c)/4 - \pi\tau/2$, so $\pi(C+c)/4 = \pi\tau/2 + \delta$. Then we have:

$$\begin{aligned} \left| \sin^2(\pi\tau/2) - \sin^2(\pi\tau/2 + \delta) \right| &= \left| \frac{1 - \cos(\pi\tau)}{2} - \frac{1 - \cos(\pi\tau + 2\delta)}{2} \right| \\ &= \frac{1}{2} |\cos(\pi\tau + 2\delta) - \cos(\pi\tau)| \\ &= |\sin(\pi\tau + \delta) \sin(-\delta)| \\ &\leq |\delta(\pi\tau + \delta)| \leq \pi^2\tau^2 \frac{\varepsilon}{2} \left(1 + \frac{\varepsilon}{2}\right) \leq \frac{3\varepsilon}{4} \pi^2\tau^2 \end{aligned} \quad (72)$$

where we used $|\delta| \leq \pi\varepsilon\tau/2$. Then, plugging this into Eq. (70), we have:

$$\begin{aligned} |\lambda_2(G(x)) - \tilde{\lambda}| &\leq \frac{3\varepsilon}{4} n\pi^2\tau^2 \\ &= \frac{3\varepsilon}{4} n\pi^2 \frac{\Delta(U(P, x))^2}{\pi^2} \leq \frac{3\varepsilon}{4} n\pi^2 \sin^2\left(\frac{\Delta(U(P, x))}{2}\right) = \pi^2 \frac{3\varepsilon}{4} \lambda_2(G(x)), \end{aligned} \quad (73)$$

using the fact that $\frac{x^2}{\pi^2} \leq \sin^2(x/2)$ when $x \in [-\pi, \pi]$. \square

Theorem 30 now follows, restated below for convenience.

Theorem 30. *Let G be the complete graph on n vertices. There exists a quantum algorithm that, on input x , with probability at least $2/3$, outputs an estimate $\tilde{\lambda}$ such that $|\tilde{\lambda} - \lambda_2(G(x))| \leq \varepsilon\lambda_2(G(x))$, where $\lambda_2(G(x))$ is the algebraic connectivity of $G(x)$, in time $\tilde{O}\left(\frac{1}{\varepsilon} \frac{n}{\sqrt{\lambda_2(G(x))}}\right)$.*

6 Acknowledgments

SJ is supported by an NWO WISE Grant and NWO Veni Innovational Research Grant under project number 639.021.752. SK completed some of this work while at the Joint Center for Quantum Information and Computer Science (QuICS) at the University of Maryland. This research was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development and Innovation.

References

- [Åri16] A. Āriņš. *Span-Program-Based Quantum Algorithms for Graph Bipartiteness and Connectivity*, pages 35–41. Springer International Publishing, Cham, 2016.
- [Bel12a] A. Belovs. Learning-graph-based quantum algorithm for k -distinctness. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, pages 207–216, 2012.
- [Bel12b] A. Belovs. Span programs for functions with constant-sized 1-certificates. In *Proceedings of the 44th Symposium on Theory of Computing (STOC 2012)*, pages 77–84, 2012.
- [BHMT02] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.
- [Bol13] B. Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 2013.

- [BR12] A. Belovs and B. W. Reichardt. Span programs and quantum algorithms for st -connectivity and claw detection. In *Proceedings of the 20th European Symposium on Algorithms (ESA 2012)*, pages 193–204, 2012.
- [CEMM98] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 454(1969):339–354, 1998.
- [CKS17] A. M. Childs, R. Kothari, and R. D. Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.
- [CMB16] C. Cade, A. Montanaro, and A. Belovs. Time and space efficient quantum algorithms for detecting cycles and testing bipartiteness, 2016. [arXiv:1610.00581](https://arxiv.org/abs/1610.00581).
- [CRR⁺96] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1996.
- [DHHM06] C. Dürr, M. Heiligman, P. Høyer, and M. Mhalla. Quantum query complexity of some graph problems. *SIAM Journal on Computing*, 35(6):1310–1328, 2006.
- [DS84] P. G. Doyle and J. L. Snell. *Random Walks and Electrical Networks*, volume 22 of *The Carus Mathematical Monographs*. The Mathematical Association of America, 1984.
- [IJ16] T. Ito and S. Jeffery. Approximate span programs. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, pages 12:1–12:14, 2016. [arXiv:1507.00432](https://arxiv.org/abs/1507.00432).
- [JK17] S. Jeffery and S. Kimmel. Quantum algorithms for graph connectivity and formula evaluation. *Quantum*, 1:26, August 2017.
- [Kit95] A. Kitaev. Quantum measurements and the Abelian stabilizer problem, 1995. [arXiv:quant-ph/9511026](https://arxiv.org/abs/quant-ph/9511026).
- [KW93] M. Karchmer and A. Wigderson. On span programs. In *Proceedings of the 8th Annual IEEE Conference on Structure in Complexity Theory*, pages 102–111, 1993.
- [NTS95] N. Nisan and A. Ta-Shma. Symmetric logspace is closed under complement. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing (STOC 1995)*, pages 140–146, New York, NY, USA, 1995. ACM.
- [Rei09] B. W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every Boolean function. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 544–551, 2009. [arXiv:quant-ph/0904.2759](https://arxiv.org/abs/quant-ph/0904.2759).
- [Rei10] B. W. Reichardt. Span programs and quantum query algorithms. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:110, 2010.
- [Rei11] B. W. Reichardt. Reflections for quantum query algorithms. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)*, pages 560–569. SIAM, 2011.
- [RŠ12] B. W. Reichardt and R. Špalek. Span-program-based quantum algorithm for evaluating formulas. *Theory of Computing*, 8(13):291–319, 2012.
- [Sze04] M. Szegedy. Quantum speed-up of Markov chain based algorithms. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2004)*, pages 32–41, Washington, DC, USA, 2004. IEEE Computer Society.

A Effective Capacitance and Effective Conductance

In this section, we provide more intuition for Definition 6 based on the definition of effective capacitance and effective conductance. This section roughly follows the explication of [Bol13, IX.1].

Let G be a graph with implicit weights c . As in Section 2.2, we associate a subgraph $G(x)$ of G with an electrical circuit in the following way: we put a 0-resistance wire at every edge $e \in E(G(x))$, and at every edge $e \in E(G) \setminus E(G(x))$, we put a capacitor with capacitance $c(e)$. Then we consider the effective capacitance of this circuit when a voltage source is connected between s and t .

When a voltage is applied between s and t , some amount of charge flows from the initially uncharged part of the circuit connected to t to the initially uncharged part of the circuit connected to s . Eventually, some steady-state accumulation of charge on the capacitors is reached. Then the effective capacitance is given by Q/\mathcal{E} , where Q is the total amount of charge that moves from the t component to the s component, and \mathcal{E} is the voltage applied to the battery.

We will show that Q/\mathcal{E} is equal to the definition of effective capacitance given in Definition 6. Without loss of generality, we set $\mathcal{E} = 1$. To determine Q , we first use the fact that charge is conserved. We define a function $q : \vec{E}(G) \rightarrow \mathbb{R}$, that tracks how charge moves in the circuit. For $(u, v, \ell) \in \vec{E}(G(x))$, we define $q(u, v, \ell)$ to be the amount of charge that is shifted through the edge from u to v from the time that the battery is connected until a steady state is reached. For $(u, v, \ell) \in E(G) \setminus E(G(x))$, we define $q(u, v, \ell)$ to be the amount of charge that accumulates on the capacitor across $(\{u, v\}, \ell)$, specifically on the side of the capacitor closest to vertex u .

Then by conservation of charge, we have

1. $\forall (u, v, \ell) \in \vec{E}(G), q(u, v, \ell) = -q(v, u, \ell)$;
2. $\sum_{v, \ell: (s, v, \ell) \in \vec{E}(G)} q(s, v, \ell) = \sum_{v, \ell: (v, t, \ell) \in \vec{E}(G)} q(v, t, \ell) = Q$; and
3. $\forall u \in V(G) \setminus \{s, t\}, \sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} q(u, v, \ell) = 0$.

We define $\bar{\mathcal{V}} : V(G) \rightarrow \mathbb{R}$ to be the voltage at each vertex in the circuit at steady state, where without loss of generality, we set $\bar{\mathcal{V}}(t) = 0$. Then $\bar{\mathcal{V}}$ is a unit st -potential. To see this, note first that $\bar{\mathcal{V}}(s) = 1$ because the voltage difference between s and t must be 1, and the voltage difference across edges in $E(G(x))$ must be zero because these vertices are connected by 0-resistance wires.

However, $\bar{\mathcal{V}}$ must also satisfy the capacitance ratio across each individual edge with a capacitor. That is, for each edge $(u, v, \ell) \in \vec{E}(G) \setminus \vec{E}(G(x))$,

$$c(u, v, \ell) = \frac{q(u, v, \ell)}{\bar{\mathcal{V}}(u) - \bar{\mathcal{V}}(v)}. \quad (74)$$

Rearranging terms, applying the conservation of charge condition, and using the fact that for $(u, v, \ell) \in \vec{E}(G(x))$, we have $\bar{\mathcal{V}}(u) - \bar{\mathcal{V}}(v) = 0$, we find that for each $u \in V(G) \setminus \{s, t\}$,

$$\sum_{v, \ell: (u, v, \ell) \in \vec{E}(G)} (\bar{\mathcal{V}}(u) - \bar{\mathcal{V}}(v)) c(u, v, \ell) = 0. \quad (75)$$

Looking at Eq. (75), we see that

$$\bar{\mathcal{V}} = \arg \min_{\mathcal{V}} \frac{1}{2} \sum_{(u, v, \ell) \in \vec{E}(G)} (\mathcal{V}(u) - \mathcal{V}(v))^2 c(u, v, \ell), \quad (76)$$

where the minimization is over unit st -potentials on $G(x)$. To see this, note that the minimum occurs when the derivative with respect to $\mathcal{V}(v)$ is zero.

Therefore

$$\begin{aligned} \min_{\mathcal{V}} \frac{1}{2} \sum_{(u, v, \ell) \in \vec{E}(G)} (\mathcal{V}(u) - \mathcal{V}(v))^2 c(u, v, \ell) &= \frac{1}{2} \sum_{(u, v, \ell) \in \vec{E}(G)} (\bar{\mathcal{V}}(u) - \bar{\mathcal{V}}(v))^2 c(u, v, \ell) \\ &= \frac{1}{2} \sum_{(u, v, \ell) \in \vec{E}(G)} (\bar{\mathcal{V}}(u) - \bar{\mathcal{V}}(v)) q(u, v, \ell). \end{aligned} \quad (77)$$

Using conservation of charge, Eq. (77) becomes

$$\mathcal{V}(s) \sum_{v,\ell:(s,v,\ell) \in \vec{E}(G)} q(s,v,\ell) - \mathcal{V}(t) \sum_{v,\ell:(v,t,\ell) \in \vec{E}(G)} q(v,t,\ell) = Q. \quad (78)$$

Thus

$$\min_{\mathcal{V}} \frac{1}{2} \sum_{(u,v,\ell) \in \vec{E}(G)} (\mathcal{V}(u) - \mathcal{V}(v))^2 c(u,v,\ell) = Q/\mathcal{E}, \quad (79)$$

so the two notions of capacitance (from electrical circuits and Definition 6) coincide. (Recall we have set $\mathcal{E} = 1$.)




Figure 5: Edges in $G(x)$ are shown using solid lines, while edges in $G \setminus G(x)$ are shown using dashed lines. In $G(x)^G$, each of the four connected components of $G(x)$ becomes a vertex, and the number of edges between vertices depends on the number of edges in $G \setminus G(x)$ connecting one component to another.

Other readers may recognize Definition 6 as the formula for effective conductance of a network (see e.g. [Bol13]), where the effective conductance is the inverse of the effective resistance of a graph. However, the network for which Definition 6 is the effective conductance is a bit strange. Given a network $\mathcal{N} = (G, c)$ and a subgraph $G(x)$, create a network $(G(x)^G, c^G)$ where each connected component w in $G(x)$ corresponds to a vertex v_w in $G(x)^G$. Then for every edge $(u, v, \ell) \in E(G) \setminus E(G(x))$ such that u is in one connected component w in $G(x)$, and v is in another connected component y in $G(x)$, create an edge in $G(x)^G$ between v_w and v_y with weight $c^G(\{v_w, v_y\}, \ell) = c(\{u, v\}, \ell)$. Let v_s be the connected component containing s and v_t be the connected component containing t . Then the effective conductance of $G(x)^G$ between v_s and v_t is given by Definition 6. Figure 5 shows the correspondence between a graph $G(x)$ and $G(x)^G$.