

An Incremental Span-Program-Based Algorithm and the Fine Print of Quantum Topological Data Analysis*

Mitchell Black[†]

William Maxwell[‡]

Amir Nayyeri[§]

Abstract

We introduce a new quantum algorithm for computing the Betti numbers of a simplicial complex. In contrast to previous quantum algorithms that work by estimating the eigenvalues of the combinatorial Laplacian, our algorithm is an instance of the generic Incremental Algorithm for computing Betti numbers that incrementally adds simplices to the simplicial complex and tests whether or not they create a cycle. In contrast to existing quantum algorithms for computing Betti numbers that work best when the complex has close to the maximal number of simplices, our algorithm works best for sparse complexes.

To test whether a simplex creates a cycle, we introduce a quantum span-program algorithm. We show that the query complexity of our span program is parameterized by quantities called the effective resistance and effective capacitance of the boundary of the simplex. Unfortunately, we also prove upper and lower bounds on the effective resistance and capacitance, showing both quantities can be exponentially large with respect to the size of the complex, implying that our algorithm would have to run for exponential time to exactly compute Betti numbers.

As a corollary to these bounds, we show that the spectral gap of the combinatorial Laplacian can be exponentially small. As the runtime of all previous quantum algorithms for computing Betti numbers are parameterized by the inverse of the spectral gap, our bounds show that all quantum algorithms for computing Betti numbers must run for exponentially long to exactly compute Betti numbers.

Finally, we prove some novel formulas for effective resistance and effective capacitance to give intuition for these quantities.

1 Introduction.

The past few years has seen the development of quantum algorithms with the potential to speed up computation of topological features of simplicial complexes called *Betti numbers*. Betti numbers are important topological invariants of a space; indeed, there is an entire, rapidly-growing field called *Topological Data Analysis* (TDA) that studies the application of topological invariants like Betti numbers (among other) [10, 16, 20]. Accordingly, the study of quantum algorithms for computing Betti numbers has been deemed *Quantum Topological Data Analysis* (QTDA).

Betti numbers can be both time and space inefficient for classical computers to compute. For example, a simplicial complex on vertices can be exponentially large and it can take exponential time to compute Betti numbers in arbitrary dimensions. Quantum computers offer a potential

*A preliminary version of this paper appeared at ISAAC 2021 under the title “Effective Resistance and Capacitance in Simplicial Complexes and a Quantum Algorithm.”

[†]Oregon State University: blackmit@oregonstate.edu

[‡]Oregon State University

[§]Oregon State University: nayyeria@oregonstate.edu

solution to the shortcomings of the classical algorithm. For example, quantum computers can efficiently store a simplicial complex with n vertices using only $O(\text{poly}(n))$ qubits.

However, while these quantum algorithms have certain advantages over their classical counterparts like improved space complexity, QTDA algorithms only achieve significant advantage over classical TDA algorithms under certain circumstances. QTDA algorithms only achieve significant speed up over classical algorithms when the input complex is *clique-dense*—it has close to the maximal number of simplices—and when the *spectral gap of the combinatorial Laplacian* of the complex is polynomially small. This second point is a particular problem as, before now, it was unknown how small the spectral gap of the combinatorial Laplacian could be. This makes the spectral gap of the combinatorial Laplacian an example of “fine print” [1]: an unbounded parameter in the runtime of a celebrated quantum algorithm.

1.1 Our Contributions.

- In Sections 3 and 4, we provide a novel quantum algorithm for computing Betti numbers using the framework of span programs [40, 59]. As opposed to existing QTDA algorithms that work by estimating the eigenvalues of the combinatorial Laplacian or singular values of the boundary matrices, our algorithm is more similar to classical matrix-reduction algorithms for computing Betti numbers as it works by incrementally adding simplices to the simplicial complex and testing if these simplices create or destroy a cycle. One advantage of our algorithm is that it avoids the step of creating a superposition over the k -simplices, which is a bottleneck of existing QTDA algorithms that restricts their utility to the clique-dense regime. In Section 4.3 and Section 4.4, we show that the query and time complexity of our span program algorithm for QTDA is parameterized by the maximum effective resistance and capacitance of cycles in \mathcal{K} . In Section 4.6, we compare our algorithm with existing QTDA algorithms. The culmination of this section is the following theorem.

Theorem 1.1. *Let \mathcal{K} be a simplicial complex. There is a quantum algorithm for computing the d th Betti number β_d of \mathcal{K} in time*

$$\tilde{O} \left(\left(\sqrt{\frac{\mathcal{R}_{\max} \mathcal{C}_{\max}}{\tilde{\lambda}_{\min}}} n_0 + \sqrt{\mathcal{R}_{\max} n_0} \right) (n_d + n_{d+1}) \right),$$

where

- n_i is the number of i -simplices of \mathcal{K} .
 - \mathcal{R}_{\max} is the maximum finite effective resistance $\mathcal{R}_{\partial\sigma}(\mathcal{L})$ of the boundary of any d - or $(d+1)$ -simplex $\sigma \in \mathcal{K}$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$.
 - \mathcal{C}_{\max} are the maximum finite effective capacitance $\mathcal{C}_{\partial\sigma}(\mathcal{L}, \mathcal{K})$ of the boundary of any d - or $(d+1)$ -simplex $\sigma \in \mathcal{K}$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$.
 - $\tilde{\lambda}_{\min}$ is the minimum spectral gap of the normalized up Laplacians $\tilde{L}_{d-1}^{up}[\mathcal{K}]$ and $\tilde{L}_d^{up}[\mathcal{K}]$.
- In Section 5, we provide upper bounds on the maximum effective resistance by the size of the simplicial complex and the maximal rank of the torsion subgroup of the simplicial complex, as well as looser upper bounds purely in terms of the size of the complex. These upper bounds show that the effective resistance can be at most exponentially-large with respect to the size of the complex. We also provide similar upper bounds on effective capacitance for special cases. Finally, we provide families of simplicial complexes with cycles whose effective resistance or

effective capacitance is exponentially large, thus matching the upper bound up to the base of the exponent. This implies that our algorithm for QTDA can take exponentially long in the worst case. However, in the next paragraph, we will see that our results imply all other QTDA algorithms must run for exponential time as well.

- In Section 6, we show how the upper and lower bounds for effective resistance provide lower and upper bounds for the spectral gap of the combinatorial Laplacian respectively; thus, the spectral gap is exponentially small in the worst case. Moreover, we show there are clique-dense complexes that achieve worst-case spectral gap.

Theorem 1.2. *Let K be a simplicial complex. Let n_i be the number of i -simplices of K . Let $n = \max\{\min\{n_{d-1}, n_d\}, \min\{n_d, n_{d+1}\}\}$. Then the spectral gap $\lambda_{\min}(L_d[K]) \in \Omega\left(\frac{1}{n^2 d^n}\right)$.*

Theorem 1.3. *Let $d, n \geq 1$. There are constants c_d, κ_d that depends only on d and a d -dimensional simplicial complex \mathcal{C}_d^n with $n_d = \Omega(\kappa_d \binom{n_0}{d})$ d -simplices such that the spectral gaps $\lambda_{\min}(L_{d-1}[\mathcal{C}_d^n]), \lambda_{\min}(L_d[\mathcal{C}_d^n]) \in O\left(\frac{1}{c_d^n}\right)$.*

This answers one of the most important question in QTDA: how small can the spectral gap be? As all existing QTDA algorithms are parameterized by the inverse of the spectral gap of the combinatorial Laplacian $\frac{1}{\lambda_{\min}}$, this implies that all existing QTDA algorithms need exponential time to exactly estimate Betti numbers.¹ Additionally, the space complexity of some QTDA algorithms is parameterized by $\log\left(\frac{1}{\lambda_{\min}}\right)$, so these algorithms will need space proportional to the number of $(d-1)$ -, d -, or $(d+1)$ -simplices, rather than space proportional to the number of vertices.

- We also prove some interesting formulas for effective resistance and capacitance that give intuition for these quantities. In Appendix B, we provide series and parallel formulas for effective resistance akin to the formulas for effective resistance in graphs. We also show that effective resistance satisfies a Rayleigh monotonicity property akin to effective resistance in graphs. Finally, in Appendix C, we show that effective resistance is dual to effective capacitance for embedded simplicial complexes.

1.2 Related Work.

History of QTDA. Lloyd, Garnerone, and Zanardi (LGZ) introduced the first quantum algorithm for computing Betti numbers up to a multiplicative error [48]. Their algorithm works by estimating the eigenvalues of the combinatorial Laplacian, which is inspired by Friedman’s classical algorithm for computing Betti numbers [23]. The LGZ algorithm has the advantage that its runtime is only polynomial with respect to the number of vertices, as opposed to the number of simplices like the matrix reduction algorithm. The trade-off is that this algorithm gains a dependence on the inverse of the spectral gap and the ratio of the number of simplices to the number of possible simplices. The LGZ algorithm performs best in the regime where the spectral gap of the combinatorial Laplacian is polynomially lower-bounded and the simplicial complex is clique-dense, meaning it has close to the maximal number of simplices. Subsequent works have improved the LGZ algorithm in different ways but maintain a runtime dependence on the inverse of the spectral gap and the clique density [29, 50, 66].

¹This exponential time complexity is not a result of the quantum nature of these algorithms. Some classical algorithms for computing Betti numbers are also parameterized by the inverse of the spectral gap, so these algorithms would need to run for exponential time as well [3, 23].

Another line of QTDA research has been developing algorithms for *persistent* Betti numbers. While the LGZ algorithm was initially claimed to be able to compute persistent Betti numbers, this was later disproved by Meijer [51] and Neumann and den Breeijen [55]. Hayakawa was the first to develop a quantum algorithm for computing persistent Betti numbers [33]. McArdle, Gilyén, and Berta have also developed algorithms for computing persistent Betti numbers [50].

Hardness of Computing Betti Numbers. In addition to new algorithms for computing Betti numbers, there have also been a number of works arguing computing Betti numbers is hard in general. Adamaszek and Stacho [2] show that determining if a simplicial complex has non-zero Betti number is NP-Hard when parameterized either by the number of vertices and the number of maximal simplices, or the number of vertices and number of minimal non-faces. Additionally, they show the problem is NP-Hard for clique complexes when parameterized by the number of vertices. Schmidhuber and Lloyd [61] show that computing Betti numbers of a clique complex is #P-Hard and estimating the Betti number up to a multiplicative constant is NP-Hard when parameterized by the number of vertices. Moreover, the hardness results of Schmidhuber and Lloyd hold for clique-dense clique complexes. This is an important restriction as the runtime of LGZ and other QTDA algorithms are lowest for clique-dense complexes. Here, the assumptions on the input are vital. Computing Betti numbers is in P when parameterized by the number of all simplices in the complex. This does not contradict $P \neq NP$ though, as the number of simplices can be exponentially large with respect to the number of vertices.

There have also been a number of works showing that problems related to computing Betti numbers are hard for the quantum computing complexity class DCQ-1. Crichigno and Kohler [12] showed that determining if the Betti number of a clique complex is nonzero is QMA_1 -Hard when parameterized by the number of vertices, and computing the Betti number of a clique complex is #BQP-Hard. Gyurik, Cade, and Dunjko [30] show that a generalization of Betti number estimation called *low-lying spectral density estimation (LLSD)* is DCQ1-Complete, suggesting that LLSD may be classically intractable. Cade and Crichigno [8] showed that estimating Betti numbers for general chain complexes (not just those arising from simplicial complexes) is also DCQ1-complete.

Lower Bounds on the Spectral Gap of the Combinatorial Laplacian. All existing QTDA algorithms are parameterized by the inverse of the spectral gap of the combinatorial Laplacian. While we show the spectral gap can be exponentially small, there have also been a number of exact or expected lower bounds on the spectral gap of the combinatorial Laplacian for certain families of simplicial complexes [4, 23, 28, 44, 45, 46, 63, 68]. However, these bounds place non-trivial assumptions on the simplicial complex so should not be taken to represent general simplicial complexes.

2 Preliminaries.

Algebraic Topology. A *simplicial complex* \mathcal{K} on a set of vertices V is a subset of the power set $\mathcal{K} \subseteq P(V)$ with the property that if $\sigma \in \mathcal{K}$ and $\tau \subset \sigma$ then $\tau \in \mathcal{K}$. An element of \mathcal{K} is a *simplex*. A simplex $\sigma \in \mathcal{K}$ of size $|\sigma| = d + 1$ is a *d -simplex*. The set of all d -simplices of \mathcal{K} is denoted \mathcal{K}_d , and the number of d -simplices is denoted $n_d = |\mathcal{K}_d|$. The *d -skeleton* of \mathcal{K} , denoted \mathcal{K}^d , is the simplicial complex of all simplices of \mathcal{K} of dimension at most d , i.e. $\mathcal{K}^d = \cup_{i=0}^d \mathcal{K}_i$. The *dimension* of \mathcal{K} is the largest d such that \mathcal{K} contains a d -simplex; a 1-dimensional simplicial complex is a *graph*.

The d^{th} **chain group** $C_d(\mathcal{K})$ is the vector space over \mathbb{R} with orthonormal basis \mathcal{K}_d . An element of $C_d(\mathcal{K})$ is a ***d-chain***. Unless otherwise stated, all vectors and matrices will be in the basis \mathcal{K}_d . For a chain $f \in C_d(\mathcal{K})$, we denote its σ coordinate $f(\sigma)$. Finally, the **support** of a chain f is the set of simplices given a non-zero value by f and is denoted $\text{supp}(f) = \{\sigma_i \in \mathcal{K}_d: f(\sigma_i) \neq 0\}$.

We assume there is a fixed but arbitrary order on the vertices $V = (v_1, \dots, v_n)$. Let $\sigma = \{v_{i_0}, \dots, v_{i_d}\}$ be a d -simplex in \mathcal{K} with $v_{i_j} \leq v_{i_k}$ whenever $j \leq k$. The **boundary** of σ is the $(d-1)$ -chain $\partial\sigma = \sum_{j=0}^d (-1)^j \cdot (\sigma \setminus \{v_{i_j}\})$. The **d^{th} boundary map** is the linear map $\partial_d : C_d(\mathcal{K}) \rightarrow C_{d-1}(\mathcal{K})$ defined $\partial_d f = \sum_{\sigma \in \mathcal{K}_d} f(\sigma) \partial\sigma$ where $f(\sigma)$ denotes the component of f indexed by the simplex σ . An element in $\ker \partial_d$ is a **cycle**, and an element in $\text{im } \partial_d$ is a **boundary** or a **null-homologous cycle**. See Figure 1. The boundary maps have the property that $\partial_d \circ \partial_{d+1} = 0$, so $\text{im } \partial_{d+1} \subset \ker \partial_d$. The **d^{th} homology group** is the quotient group $H_d(\mathcal{K}) = \ker(\partial_d) / \text{im}(\partial_{d+1})$. The **d^{th} Betti number** β_d is the dimension of $H_d(\mathcal{K})$. The **d^{th} coboundary map** is the map $\delta_d := \partial_{d+1}^T : C_d(\mathcal{K}) \rightarrow C_{d+1}(\mathcal{K})$. An element of $\ker \delta_d$ is a **cocycle**, and an element in $\text{im } \delta_{d-1}$ is a **coboundary**. We will use the notation $\partial[\mathcal{K}]$ and $\delta[\mathcal{K}]$ when we want to specify the complex associated with the (co)boundary operator.

While our algorithms calculate homology with real coefficients, for some of our topological results, we will need to consider homology with integer coefficients. The **integral chain group** $C_d(\mathcal{K}, \mathbb{Z})$ is the free abelian group generated by the set of d -simplices \mathcal{K}_d whose elements are formal sums $\sum_{\sigma_i \in \mathcal{K}_d} \alpha_i \sigma_i$ with coefficients $\alpha_i \in \mathbb{Z}$. The integer homology groups are constructed in the same way as the real homology groups. We define boundary maps $\partial_d : C_d(\mathcal{K}; \mathbb{Z}) \rightarrow C_{d-1}(\mathcal{K}; \mathbb{Z})$ the same way as for the real chain groups, except now the boundary maps are group homomorphisms rather than linear maps. The **integral homology groups** are the quotient groups $H_d(\mathcal{K}; \mathbb{Z}) = \ker \partial_d / \text{im } \partial_{d+1}$.

Laplacians. The **d^{th} up Laplacian** is $L_d^{\text{up}} = \partial_{d+1} \delta_d$, the **d^{th} down Laplacian** is $L_d^{\text{down}} = \delta_{d-1} \partial_d$, and the **d^{th} (combinatorial) Laplacian** is $L_d = L_d^{\text{up}} + L_d^{\text{down}}$. The Laplacians define the following orthogonal decomposition of the d^{th} chain group $C_d(\mathcal{K})$ called the **Hodge Decomposition**.

$$\begin{aligned} C_d(\mathcal{K}) &= \text{im } L_d^{\text{up}} \oplus \text{im } L_d^{\text{down}} \oplus \ker L_d \\ &= \text{im } \partial_d \oplus \text{im } \delta_{d-1} \oplus \ker L_d \end{aligned}$$

where the second equality follows from the fact that $\text{im } AA^T = \text{im } A$ for any matrix A . We call the subspaces $\text{im } \partial_d$, $\text{im } \delta_{d-1}$, and $\ker L_d$ the **boundary**, **coboundary**, and **harmonic spaces**. Arguably the fundamental theorem of the combinatorial Laplacian is the **Hodge Theorem**.

Theorem 2.1 (Hodge Theorem, Eckmann [19]). *The d^{th} harmonic space is isomorphic to the d^{th} homology group, i.e. $\ker L_d \cong H_d(\mathcal{K})$.*

Therefore, the d^{th} Betti number can equivalently be computed by computing the rank of L_d , a fact used by many existing QTDA algorithms.

The following lemma gives several properties of the spectrum of the combinatorial Laplacian.

Lemma 2.2. *Let $\text{spec}_{\mathbb{N}\mathbb{Z}}(A)$ denote the multiset of the non-zero eigenvalues of a linear operator A . Let \mathcal{K} be a simplicial complex. Let $d > 0$ be a positive integer. Then*

1. (Goldberg [24, Lemma 4.1.8]) $\text{spec}_{\mathbb{N}\mathbb{Z}}(L_d^{\text{up}}) = \text{spec}_{\mathbb{N}\mathbb{Z}}(L_{d+1}^{\text{down}})$
2. (Goldberg [24, Lemma 4.1.7]) $\text{spec}_{\mathbb{N}\mathbb{Z}}(L_d) = \text{spec}_{\mathbb{N}\mathbb{Z}}(L_d^{\text{up}}) \cup \text{spec}_{\mathbb{N}\mathbb{Z}}(L_d^{\text{down}})$

3. (Goldberg [24, Lemma 4.2.3]) If \mathcal{K} has connected components $\mathcal{K}_1, \dots, \mathcal{K}_m$, then

$$\text{spec}_{NZ}(L_d[\mathcal{K}]) = \text{spec}_{NZ}(L_d[\mathcal{K}_1]) \cup \dots \cup \text{spec}_{NZ}(L_d[\mathcal{K}_m])$$

where all unions are multiset unions.

The up, down, and combinatorial Laplacian are all **positive-semidefinite**, meaning their eigenvalues are all non-negative [24]. The **spectral gap** $\lambda_{\min}(L_d)$ is the smallest non-zero eigenvalue of the combinatorial Laplacian. Lemma 2.2 Part 2 implies the following theorem about the spectral gap of the combinatorial Laplacian.

Corollary 2.3. *Let \mathcal{K} be a simplicial complex. Let d be a positive integer. Then*

$$\lambda_{\min}(L_d[\mathcal{K}]) = \min\{\lambda_{\min}(L_d^{\text{down}}[\mathcal{K}]), \lambda_{\min}(L_d^{\text{up}}[\mathcal{K}])\}$$

In Section 6, we discuss upper and lower bounds on the spectral gap. There are also known upper and lower bounds on the *largest* eigenvalue of the combinatorial Laplacian.

Theorem 2.4. *Let \mathcal{K} be a simplicial complex with n_0 vertices. Let d be a natural number. Then the maximal eigenvalue of the combinatorial Laplacian $\lambda_{\max}(L_d) \leq n_0$.*

Proof. Let Δ_{n_0} be the complete complex on n_0 vertices. The maximum eigenvalue of the d^{th} up Laplacian is $\lambda_{\max}(L_d^{\text{up}}[\Delta_{n_0}]) = n_0$ for any dimension d [26, Lemma 2.6]. Moreover, by the interlacing theorem of eigenvalues of the up Laplacian, $\lambda_{\max}(L_d^{\text{up}}[\mathcal{K}]) \leq \lambda_{\max}(L_d^{\text{up}}[\Delta_{n_0}])$ for any subcomplex $\mathcal{K} \subset \Delta_{n_0}$ [34, Theorem 1.1]. The theorem follows as $\lambda_{\max}(L_d[\mathcal{K}]) = \max\{\lambda_{\max}(L_d^{\text{up}}[\mathcal{K}]), \lambda_{\max}(L_{d-1}^{\text{up}}[\mathcal{K}])\} \leq n_0$ \square

We also consider two variants of the Laplacian variants of the up-Laplacian: the weighted up Laplacian and the normalized up Laplacian. Let $w : \mathcal{K}_{d+1} \rightarrow \mathbb{R}^+$ be a weight function on the $(d+1)$ -simplices. Let $W : C_{d+1}(\mathcal{K}) \rightarrow C_{d+1}(\mathcal{K})$ be the diagonal matrix with $W_{\tau,\tau} = w(\tau)$. The **d^{th} weighted up Laplacian** is $L_d^{\text{up},W} = \partial_{d+1}W\delta_d$. The **degree** of a d -simplex σ is $\text{deg}(\sigma) = \sum_{\tau \in \mathcal{K}_{d+1} : \sigma \subset \tau} w(\tau)$. Let $D : C_d(\mathcal{K}) \rightarrow C_d(\mathcal{K})$ be the diagonal matrix with $D_{\sigma,\sigma} = \text{deg}(\sigma)$. The **d^{th} normalized up Laplacian** is $\tilde{L}_d^{\text{up}} = D^{-1/2}\partial_{d+1}W\delta_dD^{-1/2}$. The following theorem relates the spectral gap of the normalized and unnormalized Laplacians. A proof can be found in Appendix A

Lemma 2.5. *Let \mathcal{K} be a simplicial complex. Let d_{\min} and d_{\max} be the minimum and maximum degrees of any d -simplex in \mathcal{K} . Suppose that $d_{\min} > 0$. The normalized and unnormalized spectral gap are related as follows:*

$$\frac{1}{d_{\max}} \lambda_{\min}(L_d^{\text{up}}) \leq \lambda_{\min}(\tilde{L}_d^{\text{up}}) \leq \frac{1}{d_{\min}} \lambda_{\min}(L_d^{\text{up}})$$

Pseudoinverse of a Linear Map. Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a rank k linear operator with singular value decomposition $A = \sum_{i=1}^k \sigma_i u_i v_i^T$. The **pseudoinverse** of A is the linear operator $A^+ : \mathbb{R}^m \rightarrow \mathbb{R}^n$ defined $A^+ = \sum_{i=1}^k \sigma_i^{-1} u_i v_i^T$. While this is in the most compact definition of the pseudoinverse, it is not the most informative. Equivalently, the **pseudoinverse** of $A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the unique linear operator with the following properties: (1) A^+ maps each vector $x \in \text{im } A$ to the unique vector $y \in \text{im } A^T$ such that $Ay = x$ and (2) A^+ maps each vector in $(\text{im } A)^\perp$ to 0. The following are well-known properties of the pseudoinverse that follow from these definitions

Lemma 2.6. *Let $A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map.*

1. $(AA^T)^+ = (A^T)^+A^+$.
2. For $x \in \text{im } A$, $A^+x = \arg \min\{\|y\| : Ay = x\}$




Figure 1: Left: The 1-cycle is null-homologous as it is the boundary of the pictured 2-chain. Right: The 1-cycle is not null-homologous as it is not the boundary of any 2-chain. Coefficients on colored simplices are ± 1 . Orientations on the simplices have been omitted for simplicity.

Bra-Ket Notation. When discussing quantum algorithms, we will use *bra-ket notation* for vectors. As this paper may also be of interest to topologists who may be unfamiliar with this notation, we introduce bra-ket notation now. Assuming a fixed basis for a finite vector space, a *bra* is a row vector represented by the notation $\langle v|$. A *ket* is a column vector represented by the notation $|v\rangle$. Using bras and kets, we can represent an inner product as $\langle u|v\rangle$, an outer product as $|u\rangle\langle v|$, or a tensor product as $|u\rangle|v\rangle$.

3 The Incremental Algorithm for Computing Betti Numbers.

In this section, we review the incremental algorithm for computing Betti numbers introduced by Delfinado and Edelsbrunner [13]. The incremental algorithm is a generic framework for computing Betti numbers based around the primitive of *null-homology testing*. See Figure 1

Problem (Null-Homology Testing). *Given a simplicial complex \mathcal{K} and a cycle γ in \mathcal{K} , determine if γ is null-homologous.*

The *Incremental Algorithm for Computing Betti Numbers* computes β_d by testing if the boundary of simplices are null-homologous. Specifically, the incremental algorithm incrementally adds d -simplices to the simplicial complex and then performs a null-homology test on their boundaries to see how the dimension of the spaces $\ker \partial_d$ and $\text{im } \partial_{d+1}$ change.

To see how this works, fix an order on the d -simplices $\mathcal{K}_d = \{\sigma_1, \dots, \sigma_d\}$, and then iteratively add each simplex σ_i in increasing order of i . Adding the simplex σ_i will either increase the dimension of $\ker \partial_d$ or $\text{im } \partial_{d-1}$ by 1. If $\partial\sigma_i$ is null-homologous in $\mathcal{K}^{d-1} \cup \{\sigma_1, \dots, \sigma_{i-1}\}$, then adding σ_i will increase the dimension of $\ker \partial_d$ by 1, which also increases β_d by 1. If not, then adding σ_i will decrease $\text{im } \partial_{d-1}$ by 1, which also decreases β_{d-1} by 1.² The incremental algorithm is summarized in Algorithm 1.

In the classical matrix reduction algorithm for computing Betti numbers, testing whether $\partial\sigma_i$ is null-homologous is done by reducing the column corresponding to σ_i in the boundary matrix, which takes $O(n_{d-1}n_d)$ time in the worst case. However, there are special cases where null-homology testing can be performed much more quickly. For example, when a simplicial complex is embedded in \mathbb{R}^3 or the 3-sphere, null-homology testing can be performed in nearly-linear time using the

²In their original paper on the incremental algorithm [13], Delfinado and Edelsbrunner phrase this slightly differently as testing “whether σ_i is in [the support of] a cycle.” It is straightforward to verify that σ_i is in the support of a cycle if and only if $\partial\sigma_i$ is null-homologous.

Algorithm 1 The incremental algorithm for computing the d^{th} Betti Number β_d [13]

```

 $\beta_d \leftarrow 0$ 
for  $\sigma_i \in \mathcal{K}_d = \{\sigma_1, \dots, \sigma_{n_d}\}$  do
  if  $\partial\sigma_i$  is null-homologous in  $\mathcal{K}^{d-1} \cup \{\sigma_1, \dots, \sigma_{i-1}\}$  then
     $\beta_d \leftarrow \beta_d + 1$ 
  end if
end for
for  $\tau_j \in \mathcal{K}_{d+1} = \{\tau_1, \dots, \tau_{n_{d+1}}\}$  do
  if  $\partial\tau_j$  is not null-homologous in  $\mathcal{K}^d \cup \{\tau_1, \dots, \tau_{j-1}\}$  then
     $\beta_d \leftarrow \beta_d - 1$ 
  end if
end for
Return  $\beta_d$ 

```

union-find algorithm [13]. In the next section, we give a quantum algorithm for null-homology testing.

4 A Quantum Algorithm for Null-Homology Testing.

In this section, we provide a quantum algorithm based on the *span program* model to decide whether or not a cycle γ is null-homologous in a simplicial complex \mathcal{K} .

Our algorithm is a generalization of the quantum algorithm developed by Belovs and Reichardt to decide *st*-connectivity in a graph [5]. Their algorithm is parameterized by the effective resistance and capacitance between the vertices s and t . The query complexity of our algorithm is parameterized by higher-dimensional analogues of effective resistance and capacitance of γ that we introduce in Section 4.3.1.

Upper bounds on the effective resistance and capacitance in graphs imply a query complexity of $O(n^{3/2})$ for *st*-connectivity, where n is the number of vertices [37]. In Section 5, we provide upper bounds on the effective resistance and capacitance. Our upper bounds on effective resistance and capacitance imply that the query complexity is polynomial in both the number of d -simplices as well as the cardinality of the largest torsion subgroup of a relative homology group of \mathcal{K} . In the case that \mathcal{K} is a graph, our analysis of the witness sizes matches the $O(n^{3/2})$ upper bounds of previous analyses. Specifically, under the assumptions that \mathcal{K} is relative torsion free and that γ is the boundary of a d -simplex (which may or may not be included in the complex), we match the $O(n^{3/2})$ upper bound. These assumptions are always true for *st*-connectivity in graphs, which is why we match the query complexity for this problem. However, in Section 5.2, we provide examples of simplicial complexes where the effective resistance or capacitance of γ is exponentially large.

4.1 A Brief Introduction to Span Programs.

Span programs were first defined by Karchmer and Wigderson [40] and were first used for quantum algorithms by Reichardt and Špalek [59]. Intuitively, a span program is a model of computation which encodes a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ into the geometry of two vector spaces and a linear operator between them. Encoding f into a span program implies the existence of a quantum query algorithm evaluating f (Theorem 4.1.)

Definition 1. A *span program* $\mathcal{P} = (\mathcal{H}, \mathcal{U}, |\tau\rangle, A)$ over the set of strings $\{0, 1\}^n$ is a 4-tuple consisting of:

1. A finite dimensional Hilbert space $\mathcal{H} = \mathcal{H}_1 \oplus \cdots \oplus \mathcal{H}_n$ where $\mathcal{H}_i = \mathcal{H}_{i,0} \oplus \mathcal{H}_{i,1}$,
2. a vector space \mathcal{U} ,
3. a non-zero vector $|\tau\rangle \in \mathcal{U}$, called the **target vector**
4. a linear operator $A: \mathcal{H} \rightarrow \mathcal{U}$.

For every string $x = (x_1, \dots, x_n) \in \{0, 1\}^n$ we associate the Hilbert space $\mathcal{H}(x) = \mathcal{H}_{1,x_1} \oplus \cdots \oplus \mathcal{H}_{n,x_n}$ and the linear operator $A(x) = A\Pi_{\mathcal{H}(x)}: \mathcal{H} \rightarrow \mathcal{U}$ where $\Pi_{\mathcal{H}(x)}$ is the projection of \mathcal{H} onto $\mathcal{H}(x)$. A string $x \in \{0, 1\}^n$ is a **positive instance** if $|\tau\rangle \in \text{im } A(x)$ and is a **negative witness** otherwise.

A span program \mathcal{P} **decides** the function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ if $f(x) = 1$ when x is a positive instance and $f(x) = 0$ when x is a negative instance. A span program can also evaluate a partial boolean function $g: D \rightarrow \{0, 1\}$ where $D \subset \{0, 1\}^n$ by the same criteria.

Span programs are a popular method in quantum computing because there are upper bounds on the complexity of evaluating span programs in the **query model**. The query model evaluates the complexity of a quantum algorithm by its **query complexity**, the number of times it queries an input oracle. In our case, the input oracle returns the bits of the binary string x . The **input oracle** \mathcal{O}_x takes $\mathcal{O}_x: |i\rangle|b\rangle \rightarrow |i\rangle|b \oplus x_i\rangle$ where $i \in [n]$. Observe that the states $|i\rangle$ can be stored on $\lceil \log n \rceil$ qubits. Reichardt [60] showed that the query complexity of a span program is a function of the positive and negative witness sizes of the program, which we now define.

Definition 2. Let \mathcal{P} be a span program and let $x \in \{0, 1\}^n$. A **positive witness** for x is a vector $|w\rangle \in \mathcal{H}(x)$ such that $A|w\rangle = |\tau\rangle$. The **positive witness size** of x is

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

If no positive witness exists for x , then $w_+(x, \mathcal{P}) = \infty$.

A **negative witness** for x is a linear map $\langle w|: \mathcal{U} \rightarrow \mathbb{R}$ such that $\langle w|A\Pi_{\mathcal{H}(x)} = 0$ and $\langle w|\tau\rangle = 1$. The **negative witness size** of x is

$$w_-(x, \mathcal{P}) = \min\{\|\langle w|A\|^2 : \langle w|: \mathcal{U} \rightarrow \mathbb{R}, \langle w|A\Pi_{\mathcal{H}(x)} = 0, \langle w|\tau\rangle = 1\}.$$

If no negative witness exists for x , then $w_-(x, \mathcal{P}) = \infty$.

Theorem 4.1 (Reichardt [60]). Let $D \subset \{0, 1\}^n$ and $f: D \rightarrow \{0, 1\}$. Let \mathcal{P} be a span program that decides f . Let $W_+(f, \mathcal{P}) = \max_{x \in f^{-1}(1)} w_+(x, \mathcal{P})$ and $W_-(f, \mathcal{P}) = \max_{x \in f^{-1}(0)} w_-(x, \mathcal{P})$. There is a bounded error quantum algorithm that decides f with query complexity $O\left(\sqrt{W_+(f, \mathcal{P})W_-(f, \mathcal{P})}\right)$.

A caveat to the query complexity model is that in general the time complexity of an algorithm can be much larger than its query complexity.

4.2 A Span Program for Null-Homology Testing.

In this section, we present a span program for testing if a cycle is null-homologous in a simplicial complex. This span program is a generalization of the span program for *st*-connectivity defined in [40] and used to develop quantum algorithms in [5, 9, 37, 38].

Let \mathcal{K} be a d -dimensional simplicial complex. Let $\gamma \in C_{d-1}(\mathcal{K})$ be a $(d-1)$ -cycle. Let n_d be the number of d -simplices in \mathcal{K} . Order the d -simplices $\{\sigma_1, \dots, \sigma_{n_d}\}$. Let $w: \mathcal{K}_d \rightarrow \mathbb{R}$ be a weight function on the d -simplices, and let $W: C_d(\mathcal{K}) \rightarrow C_d(\mathcal{K})$ be the diagonal weight matrix. We define a span program over the strings $\{0, 1\}^{n_d}$ as follows.

1. $\mathcal{H} = C_d(\mathcal{K})$, with $\mathcal{H}_{i,1} = \text{span}\{\sigma_i\}$ and $\mathcal{H}_{i,0} = \{0\}$.
2. $\mathcal{U} = C_{d-1}(\mathcal{K})$
3. $A = \partial_d \sqrt{W} : C_d(\mathcal{K}) \rightarrow C_{d-1}(\mathcal{K})$
4. $|\tau\rangle = \gamma$

We denote the above span program by $\mathcal{P}_{\mathcal{K}}$. Let $x \in \{0, 1\}^{n_d}$ be a binary string. We define the subcomplex $\mathcal{K}(x) := \mathcal{K}^{d-1} \cup \{\sigma_i : x_i = 1\}$; that is, $\mathcal{K}(x)$ contains the d -simplices σ_i such that $x_i = 1$. There exists a solution v to the linear system $\partial_d \sqrt{W} \Pi_{\mathcal{K}(x)} v = \gamma$ if and only if the cycle γ is null-homologous in $\mathcal{K}(x)$ if and only if x is a positive instance of $\mathcal{P}_{\mathcal{K}}$. The span program $\mathcal{P}_{\mathcal{K}}$ decides the boolean function $f : \{0, 1\}^{n_d} \rightarrow \{0, 1\}$ where $f(x) = 1$ if and only if γ is a null-homologous cycle in the subcomplex $\mathcal{K}(x)$.

Theorem 4.1 allows us to bound the query complexity of our span program by the size of positive and negative witness. In the next section, we provide bounds on the positive and negative witness size of our span program.

4.3 Witness Sizes of the Null-Homology Testing Span Program.

In this section, we bound the positive and negative witness sizes of our span program for null-homology testing. We will show that they are equal to the quantities called the effective resistance and effective capacitance of the cycle. We first introduce these quantities and show some of their properties. Then, in Section 4.3.4, we show that these quantities are the witness sizes of our span program.

4.3.1 Background: Effective Resistance and Effective Capacitance.

Let $\gamma \in C_{d-1}(\mathcal{K})$ be a cycle in a simplicial complex. We associate two quantities with γ : its *effective resistance* and *effective capacitance*. The effective resistance is finite if and only if γ is null-homologous, and the effective capacitance is finite if and only if γ is not null-homologous. We begin with the definition of effective resistance.

Definition 3. Let \mathcal{K} be a simplicial complex with weight function $w : \mathcal{K} \rightarrow \mathbb{R}^+$. Let γ be a $(d-1)$ -cycle in \mathcal{K} . The **effective resistance** of γ is

$$\mathcal{R}_{\gamma}(\mathcal{K}, W) = \begin{cases} \gamma^T \left(L_{d-1}^{up, W} \right)^+ \gamma & \text{if } \gamma \text{ is null-homologous} \\ \infty & \text{otherwise} \end{cases}$$

When obvious or when \mathcal{K} is unweighted, we drop the weights from the notation and write $\mathcal{R}_{\gamma}(\mathcal{K})$.

This definition of effective resistance is consistent with effective resistance in graphs (see [62]) and other definitions of effective resistance in simplicial complexes [42, 57, 31].³ However, this definition gives little intuition about effective resistance. We now prove there is an alternative definition of effective resistance in terms of chains with boundary γ . We begin with two definitions.

³Effective resistance can be defined even more generally using the combinatorial Laplacian. For simplicity, consider the unweighted case. For a null-homologous $(d-1)$ -cycle γ , the effective resistance can be defined $\mathcal{R}_{\gamma}(\mathcal{K}) = \gamma^T L_{d-1}^+ \gamma$. This equals the formula in Definition 3 because (1) $L_{d-1}^+ = (L_{d-1}^{up})^+ + (L_{d-1}^{down})^+$ and (2) $\gamma^T (L_{d-1}^{down})^+ \gamma = 0$ as $\gamma \in \ker L_{d-1}^{down} = \ker (L_{d-1}^{down})^+$.

Definition 4. Given a d -dimensional simplicial complex \mathcal{K} and a $(d-1)$ -dimensional null-homologous cycle γ , a **unit γ -flow** is a d -chain $f \in C_d(\mathcal{K})$ such that $\partial f = \gamma$.

In the case of graphs, a unit st -flow is a flow sending 1 unit of flow from s to t .

Definition 5. Given a d -dimensional simplicial complex \mathcal{K} with weight function $w : C_d(\mathcal{K}) \rightarrow \mathbb{R}^+$ and a unit γ -flow f , the **flow energy** of f on \mathcal{K} is

$$J(f) = \sum_{\sigma \in \mathcal{K}^{(d)}} \frac{f(\sigma)^2}{w(\sigma)} = f^T W^{-1} f$$

where W is the $n_d \times n_d$ diagonal matrix whose entries are the weights of the d -simplices.

We will now relate unit γ -flows and their energy to effective resistance. This generalizes a formula for effective resistance in graphs [6, Chapter IX Corollary 6].

Lemma 4.2. Let \mathcal{K} be a simplicial complex and let γ be a null-homologous d -cycle. The effective resistance of γ is the minimum flow energy over all unit γ -flows, i.e.

$$\mathcal{R}_\gamma(\mathcal{K}) = \min\{J(f) \mid \partial f = \gamma\}$$

Proof. Our first observation is that we can factor the weighted Laplacian as

$$\begin{aligned} L_d^{up, W} &= \partial_{d+1} W \delta_d \\ &= \partial_{d+1} W^{1/2} W^{1/2} \delta_d \\ &= (\partial_{d+1} W^{1/2}) (\partial_{d+1} W^{1/2})^T \end{aligned}$$

By Lemma 2.6 Part 1, $(L_d^{up, W})^+ = ((\partial_{d+1} W^{1/2})^T)^+ (\partial_{d+1} W^{1/2})^+$. Therefore,

$$\mathcal{R}_\gamma(\mathcal{K}) = \gamma^T ((\partial_{d+1} W^{1/2})^T)^+ (\partial_{d+1} W^{1/2})^+ \gamma = \|(\partial W^{1/2})^+ \gamma\|^2.$$

By Lemma 2.6 Part 2, $\mathcal{R}_\gamma(\mathcal{K})$ is the minimum squared-norm of a vector that $\partial_{d+1} W^{1/2}$ maps to γ . Let $f = (\partial W^{1/2})^+ \gamma$; the vector f is the unit γ -flow of minimum flow energy, which we now prove.

A vector v is mapped to γ by $\partial W^{1/2}$ iff $W^{1/2} v$ is mapped to γ by ∂ as $W^{1/2}$ is a bijection; that is all to say, $W^{1/2} v$ is a unit γ -flow. Moreover, the flow energy of $W^{1/2} v$ is

$$\begin{aligned} J(W^{1/2} v) &= (W^{1/2} v)^T W^{-1} W^{1/2} v \\ &= v^T W^{1/2} W^{-1} W^{1/2} v \\ &= v^T v \\ &= \|v\|^2 \end{aligned}$$

Therefore, the minimum flow energy of a unit γ -flow is the minimum squared-norm of a vector that $\partial W^{1/2}$ maps to γ , which we previously saw was $\mathcal{R}_\gamma(\mathcal{K})$. \square

We call $(\partial W^{1/2})^+ \gamma$ the **minimum-energy** unit γ -flow⁴.

⁴The minimum-energy unit γ -flow is unlike the optimal bounded chain [17], another minimum spanning object of a null-homologous cycle. The optimal bounded chain of γ is (often) defined with \mathbb{Z}_2 coefficients and is the smallest set of simplices with boundary γ ; it is analogous to an s, t -path in a graph. The minimum-energy γ -flow is defined with real coefficients and can have fractional values on simplices; it is analogous to an s, t -flow in a graph. While the weight of the optimal bounded chain is always minimized when fewer simplices are used, the minimum minimum-energy γ -flow will push a fraction of the flow on many chains spanning γ . Indeed, in Section B, we prove the minimum-energy γ -flow pushes a non-zero amount of flow on each chain spanning γ .




Figure 2: Left: A 1-cycle γ with ± 1 coefficients on the blue edges. Right: A unit γ -potential p with ± 1 coefficients on the red edges. If this complex is unweighted, then the potential energy of p is 1. It can be proved that p is a minimal energy unit γ -potential⁵, so $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K}) = 1$

Some of the key properties of effective resistance in graphs are the series and parallel formulas and Rayleigh Monotonicity. In Appendix B, we prove analogous results for higher-dimensional effective resistance.

While effective resistance has previously been generalized from graphs to simplicial complexes [31, 42, 57], to our knowledge, we are the first to generalize effective capacitance from graphs to simplicial complexes. Unfortunately, effective capacitance is more opaque than effective resistance, both in graphs and simplicial complexes. The definition of effective capacitance is less intuitive than the definition for effective resistance, and there are fewer results about effective capacitance in graphs than effective resistance.

Before defining effective capacitance in simplicial complexes, we review the definition of effective capacitance in graphs, which can be found in [37]. Let G be a graph such that s and t are connected in G , and let $H \subseteq G$ be a subgraph such that s and t are not connected in H . A *unit st -potential* is a function $p: V(G) \rightarrow \mathbb{R}$ such that $p(t) = 1$, $p(s) = 0$, and $p(u) = p(v)$ for any two vertices u, v in the same connected component. The *potential energy* of p is $\sum_{\{u,v\} \in E(G)} (p(u) - p(v))^2$. The *effective capacitance* of s and t is the minimum potential energy of any st -potential.

Our definition of effective capacitance in simplicial complexes will be analogous to the definition in graphs; namely, the effective capacitance of a cycle γ will be the minimum energy of a unit γ -potential.

Definition 6. Let \mathcal{L} be a simplicial complex, and let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$ -cycle that is not null-homologous in \mathcal{L} . A **unit γ -potential** in \mathcal{L} is a $(d-1)$ -chain p such that $\delta_{d-1}[\mathcal{L}]p = 0$ and $p^T \gamma = 1$.

Figure 2 shows a γ -potential in a simplicial complex.

Definition 7. Given simplicial complexes $\mathcal{L} \subset \mathcal{K}$ with weight function $w: C_d(\mathcal{K}) \rightarrow \mathbb{R}$ and a γ -potential p in \mathcal{L} , the **potential energy** of p on \mathcal{K} is

$$\mathcal{J}(p) = \sum_{\sigma \in \mathcal{K}_d} \delta[\mathcal{K}]p(\sigma)^2 w(\sigma) = (\delta[\mathcal{K}]p)^T W (\delta[\mathcal{K}]p).$$

Definition 8. Let $\mathcal{L} \subset \mathcal{K}$ be simplicial complexes, and let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$ -cycle that is not null-homologous in \mathcal{K} . If γ is not null-homologous in \mathcal{L} , the **effective capacitance** of γ in \mathcal{L} and \mathcal{K} is

⁵The proof that p is the minimal energy unit γ -potential is analogous to the proof of Theorem 5.14. The trick to proving this is to realize that because there is a single 2-simplex in $\mathcal{K} \setminus \mathcal{L}$, all unit γ -potentials have the same potential energy.

$$\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K}) = \begin{cases} \min_{p: p \text{ unit } \gamma\text{-potential}} \mathcal{J}(p) & \gamma \text{ is not null-homologous} \\ \infty & \gamma \text{ is null-homologous} \end{cases}$$

Our definition of effective capacitance in simplicial complexes matches the definition of effective capacitance in graphs; however, this may not be obvious at first glance, as our definition of st -potential is more general. A function $p : V(G) \rightarrow \mathbb{R}$ is equivalent to a 0-chain $p \in C_0(H)$, and the requirement that $p(u) = p(v)$ for any two vertices u, v in the same connected component is equivalent to saying $\delta_0[H] = 0$; however, not all chains p such that $p^T(s - t) = 1$ satisfy $p(s) = 1$ and $p(t) = 0$. (For example, it could be the case that $p(s) = \frac{1}{2}$ and $p(t) = -\frac{1}{2}$.) This difference in the definition ends up not mattering though. This is because the all-1s vectors $1 \in \ker \delta_0$ for any graph. Using this fact, we can see that for any st -potential p under our definition, there is an st -potential p' under the previous definition with the same potential energy, namely the potential $p' = p - p(t)1$.

There is one small detail left to show. It is not obvious from the definition that a unit γ -potential will even exist for γ . We prove this in the following lemma.

Lemma 4.3. *Let \mathcal{L} be a simplicial complex, and let $\gamma \in C_{d-1}(\mathcal{L})$ be a cycle. Then there exists a unit γ -potential in \mathcal{L} if and only if γ is not null-homologous in \mathcal{L} .*

Proof. Observe that $\ker \delta_{d-1}[\mathcal{L}] = (\text{im } \partial_d[\mathcal{L}])^\perp$ as $\delta_{d-1}[\mathcal{L}] = \partial_d[\mathcal{L}]^T$. Assume there is a γ -potential p in \mathcal{L} . As $\delta[\mathcal{L}]p = 0$, then $p \in \ker \delta_{d-1}[\mathcal{L}] = (\text{im } \partial_d[\mathcal{L}])^\perp$. As $\gamma^T p = 1$ we see that γ has a non-zero component in $(\text{im } \partial_d[\mathcal{L}])^\perp$, so $\gamma \notin \text{im } \partial_d[\mathcal{L}]$.

Alternatively, suppose that γ is not null-homologous in \mathcal{L} . Then γ has a non-zero component in $(\text{im } \partial_d[\mathcal{L}])^\perp = \ker \delta_{d-1}[\mathcal{L}]$. Let $q = \Pi_{\ker \delta[\mathcal{L}]} \gamma$, where $\Pi_{\ker \delta[\mathcal{L}]}$ is the projection operator onto $\ker \delta_{d-1}[\mathcal{L}]$. Then $\gamma^T q \neq 0$ and $\delta_{d-1}[\mathcal{L}]q = 0$. The vector q is not necessarily a unit γ -potential as it is not necessarily the case that $\gamma^T q = 1$, but the scaled vector $p = \frac{1}{\gamma^T q} q$ is a unit γ -potential. \square

One interesting property of effective resistance and capacitance in graphs is that, in planar graphs, the effective resistance between certain pair of nodes in the dual graph equals the effective capacitance between certain pairs of nodes in the primal graph. In Appendix C, we show that an analogous property holds for higher-dimensional embedded simplicial complexes.

4.3.2 Effective Resistance and the Spectral Gap.

In this section, we give a characterization of the spectral gap of the combinatorial Laplacian in terms of the effective resistance of a cycle. While the proof of this lemma follows from some simple linear algebra, the advantage of this theorem comes down to the fact that effective resistance is easier to work with than eigenvectors of the Laplacian (in our opinion). We first relate effective resistance to the spectral gap of the up Laplacian. We then show how this relates effective resistance to the spectral gap of the combinatorial Laplacian. We prove this relationship for unweighted simplicial complexes; however, the theorems also hold for weighted simplicial complexes.

Lemma 4.4. *The spectral gaps of the up Laplacian L_{d-1}^{up} and down Laplacian L_d^{down} are*

$$\lambda_{\min}(L_d^{down}) = \lambda_{\min}(L_{d-1}^{up}) = \min\{\mathcal{R}_\gamma^{-1}(\mathcal{K}) : \gamma \in \text{im } \partial_d, \|\gamma\| = 1\}.$$

Proof. We first prove this is the case for the spectral gap of the up Laplacian L_d^{up} . We then show the equivalence of $\lambda_{\min}(L_d^{down})$ and $\lambda_{\min}(L_{d-1}^{up})$.

The lemma follows from some standard facts about symmetric matrices. First, because L_d^{up} is symmetric, a vector x is an eigenvector of L_d^{up} with non-zero eigenvalue λ if and only if x is an eigenvector of $(L_d^{up})^+$ with non-zero eigenvalue λ^{-1} . This follows from the fact that the singular values and vectors of a symmetric matrix are also its eigenvalues and eigenvectors. Therefore, the smallest non-zero eigenvalue of L_d^{up} is the inverse of the largest non-zero eigenvalue of $(L_d^{up})^+$, or $\lambda_{\min}(L_d^{up}) = \lambda_{\max}^{-1}((L_d^{up})^+)$ for short.

Next, we can characterize the eigenvalues of the symmetric matrix $(L_d^{up})^+$ with the **Courant-Fischer Theorem**. We use a special case of the theorem, which says that $\lambda_{\max}((L_d^{up})^+) = \max\{x^T(L_d^{up})^+x : \|x\| = 1\}$ and $x_{\max} = \arg \max\{x^T(L_d^{up})^+x : \|x\| = 1\}$, where x_{\max} is an eigenvector corresponding to $\lambda_{\max}((L_d^{up})^+)$. The lemma follows from the fact that $x_{\max} \in \text{im } L_d^{up} = \text{im } \partial_d$, which is the case because x_{\max} is the eigenvector of a non-zero eigenvalue of L_d^{up} .

Finally, $\lambda_{\min}(L_d^{up}) = \lambda_{\min}(L_{d-1}^{down})$ by Lemma 2.2 Part 1. \square

4.3.3 Effective Capacitance and the Spectral Gap.

In the previous section, we saw that the effective resistance of a unit-length cycle is always bounded above by the inverse of the spectral gap of the combinatorial Laplacian. While we don't know such a bound for the effective capacitance of arbitrary cycles, we can prove such a bound for the effective capacitance for the boundaries of simplices. This is sufficient for our analysis of the incremental algorithm as the only cycles we consider are the boundaries of simplices.

Before proving our upper bound on the effective capacitance of a cycle, we need to prove an upper bound on the largest singular value of the coboundary matrix.

Lemma 4.5. *Let \mathcal{K} be a simplicial complex with n_0 vertices. For any $d \geq 1$, the largest singular value of the coboundary matrix $\delta_{d-1}[\mathcal{K}]$ is $\sigma_{\max}(\delta_{d-1}) = O(\sqrt{n_0})$.*

Proof. This follows as the squared singular values of δ_{d-1} are the eigenvalues of the up Laplacian $\delta_{d-1}^T \delta_{d-1} = L_{d-1}^{up}$. (This is true for any matrix of the form $A^T A$.) The maximum eigenvalue of L_{d-1}^{up} is known to be at most n_0 by Theorem 2.4. \square

Theorem 4.6. *Let $\mathcal{L} \subset \mathcal{K}$ be d -dimensional simplicial complexes. Let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$ -cycle that is null-homologous in \mathcal{K} but not in \mathcal{L} . Assume that $\gamma = \partial\sigma$ for a d -simplex $\sigma \notin \mathcal{L}$.⁶ The effective capacitance of γ in \mathcal{K} is bounded above by $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K}) = O(n_0 \cdot \lambda_{\min}^{-1}(L_{d-1}[\mathcal{L} \cup \{\sigma\}]))$.*

Proof. We can express the constraints of a γ -potential p in the following set of linear equations:

$$\begin{bmatrix} \delta[\mathcal{L}] \\ \gamma^T \end{bmatrix} p = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

To simplify notation, let $C = [\delta[\mathcal{L}]^T \ \gamma]^T$ and $b = [0 \ 0 \ \dots \ 1]^T$

We consider the smallest vector p which satisfies these equations, which is $p = C^+ b$. Because $\gamma = \partial\sigma$, we can see that $C = \delta[\mathcal{L} \cup \{\sigma\}]$. Therefore, $\|p\| = O(\|C^+ b\|) = O(\sigma_{\min}^{-1}(\delta[\mathcal{L} \cup \{\sigma\}]))$, where $\sigma_{\min}(\delta[\mathcal{L} \cup \{\sigma\}])$ is the smallest non-zero singular value of $\delta[\mathcal{L} \cup \{\sigma\}]$. However, we know that $\sigma_{\min} = \sqrt{\lambda_{\min}(L_{d-1}^{up}[\mathcal{L} \cup \{\sigma\}])} \in \Omega(\sqrt{\lambda_{\min}(L_{d-1}[\mathcal{L} \cup \{\sigma\}])})$. Therefore, $\|p\| \in O\left(\sqrt{\lambda_{\min}^{-1}(L_{d-1}[\mathcal{L} \cup \{\sigma\}])}\right)$.

⁶The theorem holds whether or not $\sigma \in \mathcal{K}$.

We now want to bound the potential energy of p . Using Lemma 4.5, we can bound $\|\delta[\mathcal{K}]p\|^2 \in O(n_0 \cdot \lambda_{\min}^{-1}(L_{d-1}[\mathcal{L} \cup \{\sigma\}]))$. \square

4.3.4 Connecting Effective Resistance and Capacitance to Witness Sizes.

Given a string $x \in \{0, 1\}^{n_d}$, we show in the following two lemmas that $w_+(x, \mathcal{P}_{\mathcal{K}}) = \mathcal{R}_{\gamma}(\mathcal{K}(x))$ and $w_-(x, \mathcal{P}_{\mathcal{K}}) = \mathcal{C}_{\gamma}(\mathcal{K}(x), \mathcal{K})$. The proofs are simple calculations following from the definitions of effective resistance and capacitance.

Lemma 4.7. *Let $x \in \{0, 1\}^{n_d}$ be a positive instance. There is a bijection between positive witnesses $|w\rangle$ for x and unit γ -flows f in $\mathcal{K}(x)$. Moreover, the positive witness size is equal to the effective resistance of γ in $\mathcal{K}(x)$; that is, $w_+(x, \mathcal{P}_{\mathcal{K}}) = \mathcal{R}_{\gamma}(\mathcal{K}(x))$.*

Proof. Let $|w_+\rangle \in C_d(\mathcal{K})$ be a positive witness for x , so $\partial_d \sqrt{W}|w_+\rangle = \gamma$. We construct a unit γ -flow f in $\mathcal{K}(x)$ by $f = \sqrt{W}|w_+\rangle$; f is indeed a unit γ -flow as $\partial_d f = \partial_d \sqrt{W}|w_+\rangle = \gamma$. Moreover, $|w_+\rangle = W^{-1/2}|f\rangle$. The flow energy of γ is

$$\begin{aligned} J(f) &= \langle f|W^{-1}|f\rangle \\ &= \langle W^{-1/2}f|W^{-1/2}f\rangle \\ &= \langle w_+|w_+\rangle \\ &= \| |w_+\rangle \|^2. \end{aligned}$$

Hence, the flow energy of f equals the witness size of x .

Conversely, let f be a unit γ -flow in $\mathcal{K}(x)$ and define the positive witness for x as $|w_+\rangle = W^{-1/2}|f\rangle$. The same computation in the above paragraph shows that the flow energy of f equals the positive witness size of x . \square

Lemma 4.8. *Let $x \in \{0, 1\}^{n_d}$ be a negative instance. There is a bijection between negative witnesses $\langle w_-|$ for x and unit γ -potentials p in $\mathcal{K}(x)$. Moreover, the negative witness size is equal to the effective capacitance of γ in $\mathcal{K}(x)$; that is, $w_-(x, \mathcal{P}_{\mathcal{K}}) = \mathcal{C}_{\gamma}(\mathcal{K}(x))$.*

Proof. Let $\langle w_-|$ be a negative witness for x . As $\langle w|$ is a linear function from $C_{d-1}(\mathcal{K})$ to \mathbb{R} we may view it as a $(d-1)$ -chain $p^T = \langle w|$. Since $\langle w_-|\gamma\rangle = 1$, then $p^T \gamma = 1$. To show that p is a unit γ -potential we must show that the coboundary of p is zero in $\mathcal{K}(x)$. By the definition of a negative witness we have

$$\begin{aligned} 0 &= \langle w_-|\partial_d \sqrt{W}\Pi_{\mathcal{K}(x)} \\ &= \langle p|\partial_d \sqrt{W}\Pi_{\mathcal{K}(x)} \\ &= \langle \delta_d(p)|\sqrt{W}\Pi_{\mathcal{K}(x)}. \end{aligned}$$

Since \sqrt{W} is a diagonal matrix and $\Pi_{\mathcal{K}(x)}$ restricts the coboundary to the subcomplex $\mathcal{K}(x)$ we see that $\langle \delta_d(p)|\sigma\rangle = 0$ for any $\sigma \in \mathcal{K}(x)_d$. To show that the witness size of $\langle w_-|$ is equal to the potential energy of p we have

$$\begin{aligned} \| \langle w_-|\partial_d \sqrt{C}\|^2 &= \langle p\partial_d \sqrt{W}|p\partial_d \sqrt{W}\rangle \\ &= \langle \sqrt{W}\delta_d(p)|\sqrt{W}\delta_d(p)\rangle \\ &= \sum_{\sigma \in \mathcal{K}_d} \langle \delta_d(p)|\sigma\rangle^2 w(\sigma) \\ &= \mathcal{J}(p). \end{aligned}$$

Conversely, let p be a unit γ -potential for $\mathcal{K}(x)$ we construct a negative witness for x by setting $\langle w_- | := p^T$. Since the coboundary of p is zero in $\mathcal{K}(x)$ we have $\langle \delta_p(p) | \sigma \rangle = 0$ for each $\sigma \in \mathcal{K}(x)_d$ which implies $\langle w_- | \partial_d \sqrt{W} \Pi_{\mathcal{K}(x)} = 0$ by the reasoning in the previous paragraph. Also by the previous paragraph we have that the potential energy of p is equal to the negative witness size of $\langle w_- |$ which concludes the proof. \square

From these two lemmas we obtain the main theorem of the section, the quantum query complexity of γ .

Theorem 4.9. *Given a d -dimensional simplicial complex \mathcal{K} , a $(d-1)$ -dimensional cycle γ that is null-homologous in \mathcal{K} , and a d -dimensional subcomplex $\mathcal{K}(x) \subseteq \mathcal{K}$, there exists a quantum algorithm deciding whether or not γ is null-homologous in $\mathcal{K}(x)$ with quantum query complexity $O\left(\sqrt{\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}\right)$, where $\mathcal{R}_{\max}(\gamma)$ is the maximum finite effective resistance $\mathcal{R}_\gamma(\mathcal{L})$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$, and $\mathcal{C}_{\max}(\gamma)$ is the maximum finite effective capacitance $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K})$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$.*

Proof. By Theorem 4.1, the span program $\mathcal{P}_\mathcal{K}$ can be converted into a quantum algorithm whose query complexity is $O\left(\sqrt{W_+(f, \mathcal{P}_\mathcal{K})W_-(f, \mathcal{P}_\mathcal{K})}\right)$ where $W_+(f, \mathcal{P}_\mathcal{K}) = \max_{x \in f^{-1}(1)} \mathcal{R}_\gamma(\mathcal{K}(x)) = \mathcal{R}_{\max}(\gamma)$ and $W_-(f, \mathcal{P}_\mathcal{K}) = \max_{x \in f^{-1}(0)} \mathcal{C}_\gamma(\mathcal{K}(x), \mathcal{K}) = \mathcal{C}_{\max}(\gamma)$. \square

4.4 Time Efficient Implementations of the Span Program.

We have given bounds on the query complexity of null-homology testing; however, this does not imply a bound on the time complexity of evaluating this span program, as the query complexity does not account for the work outside of the oracle calls. In Appendix D, we describe the details of an implementation of this algorithm. For certain special cases, we are able to analyse the time complexity of the algorithm. We describe this special case below.

There are two obstacles to a time-efficient implementation of the span program: the weights and the input cycle γ . The weights on the d -simplices make it difficult to implement the matrix $\partial\sqrt{W}$, as the weights on the simplices can be arbitrary real numbers. The input cycle γ is difficult to create on a quantum computer as the entries of γ can also be arbitrary real numbers.

Accordingly, we can give a quantum algorithm of bounded time complexity in one particular instance: when \mathcal{K} is unweighted and γ is the boundary of a d -simplex. (We do not require the d -simplex to actually appear in the complex.) While this is only a special case of the generic null-homology testing algorithm, this is the only case we need for the incremental algorithm for computing Betti numbers (Algorithm 1). The time complexity of this case is given in the following theorem.

Theorem 4.10. *Let \mathcal{K} be an unweighted simplicial complex with n_0 vertices, let $\gamma \in C_{d-1}(\mathcal{K})$ a null-homologous cycle in \mathcal{K} , and $\mathcal{K}(x) \subset \mathcal{K}$ be a simplicial complex. Furthermore, assume that γ is the boundary of a d -simplex and the complex is unweighted. There is a quantum algorithm for deciding if γ is null-homologous in $\mathcal{K}(x)$ that runs in time*

$$\tilde{O}\left(\sqrt{\frac{\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}{\tilde{\lambda}_{\min}}}n_0 + \sqrt{\mathcal{R}_\gamma(\mathcal{K})n_0}\right),$$

where $\mathcal{R}_{\max}(\gamma)$ is the maximum finite effective resistance $\mathcal{R}_\gamma(\mathcal{L})$ of γ in any subcomplex $\mathcal{L} \subset \mathcal{K}$, $\mathcal{C}_{\max}(\gamma)$ is the maximum finite effective capacitance $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{L})$ in any subcomplex $\mathcal{K}(x)$, and $\tilde{\lambda}_{\min}$ is the spectral gap of the normalized up-Laplacian $\tilde{L}_{d-1}^{up}[\mathcal{K}]$.

4.5 Runtime of the Quantum Incremental Algorithm

In the previous section, we saw an implementation of an algorithm for testing if the boundary of a d -simplex was null-homologous. Combined with the framework of the Incremental Algorithm (Algorithm 1), this allows us to compute the d -Betti number.

Theorem 1.1. *Let \mathcal{K} be a simplicial complex. There is a quantum algorithm for computing the d th Betti number β_d of \mathcal{K} in time*

$$\tilde{O} \left(\left(\sqrt{\frac{\mathcal{R}_{\max} \mathcal{C}_{\max}}{\tilde{\lambda}_{\min}}} n_0 + \sqrt{\mathcal{R}_{\max} n_0} \right) (n_d + n_{d+1}) \right),$$

where

- n_i is the number of i -simplices of \mathcal{K} .
- \mathcal{R}_{\max} is the maximum finite effective resistance $\mathcal{R}_{\partial\sigma}(\mathcal{L})$ of the boundary of any d - or $(d+1)$ -simplex $\sigma \in \mathcal{K}$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$.
- \mathcal{C}_{\max} are the maximum finite effective capacitance $\mathcal{C}_{\partial\sigma}(\mathcal{L}, \mathcal{K})$ of the boundary of any d - or $(d+1)$ -simplex $\sigma \in \mathcal{K}$ in any subcomplex $\mathcal{L} \subset \mathcal{K}$.
- $\tilde{\lambda}_{\min}$ is the minimum spectral gap of the normalized up Laplacians $\tilde{L}_{d-1}^{up}[\mathcal{K}]$ and $\tilde{L}_d^{up}[\mathcal{K}]$.

Proof. The Incremental Algorithm (Algorithm 1) incrementally adds each d and $(d+1)$ -simplex σ to the simplicial complex and checks if the cycle $\partial\sigma$ is null-homologous. We can use the span-program algorithm of Theorem 4.10 to check if $\partial\sigma$ is null-homologous. The theorem follows by using this algorithm for each of the $(n_d + n_{d+1})$ d - and $(d+1)$ -simplices. \square

4.6 Comparison with Existing Algorithms.

In this section, we compare our algorithm to existing algorithms for QTDA. This presentation specifically compares our algorithm to the LGZ algorithm [48], but most of these ideas also hold for other existing QTDA algorithms.

Input. Our algorithm makes different assumptions about how the simplicial complex is stored compared to previous algorithms. We assume we have a list of simplices in the simplicial complex; this is required for the incremental algorithm as we must iteratively add the simplices and test if their boundaries are null-homologous. Compare this to existing quantum algorithms, which assume we have a way of checking if a simplex is included in the simplicial complex.

Our algorithm assumes we have a **list oracle** that can return the simplices in the simplicial complex:

$$\mathcal{O}_{list} : |i\rangle|0\rangle \rightarrow |i\rangle|\sigma_i\rangle,$$

where σ_i is the i th d -simplex of our simplicial complex.

Compare this to the **membership oracle** used in other QTDA algorithms that can check whether a simplex is in the simplicial complex:

$$\mathcal{O}_{memb} : |\sigma_i\rangle|j\rangle \rightarrow |\sigma_i\rangle|j \oplus b_i\rangle,$$

where b_i is a bit indicating if $\sigma_i \in \mathcal{K}_d$.

These oracles come with different trade-offs. The oracle \mathcal{O}_{memb} does not require computing the set of simplices in advanced, while \mathcal{O}_{list} does. However, algorithms that use the membership oracle \mathcal{O}_{memb} pay for this in the time it takes to compute a uniform superposition of the d -simplices, a costly operation leading to a factor of $\zeta_d = n_d / \binom{n_0}{d+1}$ in the runtime. Thus, our algorithm is better suited for *sparse* simplicial complexes—complexes where $n_d \ll \binom{n_0}{d+1}$ and where the list of simplices can be computed efficiently—a family of complexes where existing QTDA algorithms perform poorly; see the section “Runtime” below for more discussion.

Output. The LGZ algorithm estimates the d th Betti number up to an additive factor by returning a value χ_d such that $\left| \chi_d - \frac{\beta_d}{\dim C_d(\mathcal{K})} \right| \leq \epsilon$; the problem of computing χ_d has been deemed **Betti number estimation**. Our algorithm instead returns the Betti number β_k .

Runtime. To compare our algorithm to existing quantum algorithms, we bound the runtime of our algorithm with respect to the spectral gap of the combinatorial Laplacian. Note that while we can bound the runtime of our algorithm by the inverse of the spectral gap, this bound is not necessarily tight.

Corollary 4.11. *Let \mathcal{K} be a simplicial complex with n_i i -simplices. There is a quantum algorithm for computing the d th Betti number β_d in time*

$$\tilde{O}\left(\Lambda_{\min}^{-3/2} n_0^{5/2} \cdot (n_d + n_{d+1})\right)$$

where Λ_{\min} is the minimum spectral gap of $L_d[\mathcal{L}]$ over all subcomplexes $\mathcal{L} \subset \mathcal{K}$.

Proof. This follows from Theorem 1.1 by applying the bounds of Lemma 4.4, Theorem 4.6, and Lemma 2.5 to bound \mathcal{R}_{\max} , \mathcal{C}_{\max} , and $\frac{1}{\lambda_{\min}}$ respectively. The bounds on the effective resistance of Lemma 4.4 only apply to unit vectors, so one factor of n_0 is because of the fact that $\|\partial\sigma\| = \sqrt{d}$, so $\mathcal{R}_{\max}(\partial\sigma) \leq d\Lambda_{\min}^{-1}$. \square

Compare this to the LGZ algorithm, which runs in time

$$O\left(\epsilon^{-2} \lambda_{\min}^{-1} n_0^4 \sqrt{\zeta_d^{-1}}\right)$$

where n_0 is the number of vertices, ϵ is the error term, λ_{\min} is the spectral gap of the combinatorial Laplacian, and ζ_d is a density term given by

$$\zeta_d = \frac{n_d}{\binom{n_0}{d+1}}.$$

The density term is the ratio of the number k -simplices in the \mathcal{K} to number of k -simplices in the complete complex on n_0 vertices, which may be exponentially small. For example, when \mathcal{K} is sparse, meaning that the number of d -simplices is polynomial in the number of vertices, the density may be exponentially small with respect to d . Specifically, if $\zeta_d = \Omega\left(\frac{n_0^{O(1)}}{n_0^{d+1}}\right) = \Omega\left(\frac{1}{n_0^{O(d+1)}}\right)$, then the runtime of LGZ is

$$O\left(\epsilon^{-2} \lambda_{\min}^{-1} n_0^{O(d+1)}\right).$$

Compare this to our algorithm, which in this case has runtime

$$\tilde{O}\left(\Lambda_{\min}^{-3/2} n_0^{O(1)}\right).$$

In this case, our algorithm has a better asymptotic dependence on the size of the complex as it avoids the factor of $\binom{n_0}{d+1}$. This factor of $\binom{n_0}{d+1}$ shows up in many of the alternatives to the LGZ algorithm, so our algorithm has a more favorable dependence on n_0 compared to these algorithms as well. Additionally, we note that the term Λ_{\min} in our algorithm and λ_{\min} in the LGZ algorithm are similar but not directly comparable. See the following section.

Effective Resistance and Capacitance vs. Spectral Gap. Our algorithm is parameterized by the maximum effective resistance and capacitance of all subcomplexes of a simplicial complex and the square root of the inverse of spectral gap of the simplicial complex, whereas previous QTDA algorithm are only parameterized by the inverse of the spectral gap of the simplicial complex. Although for a *fixed* complex, effective resistance and capacitance are upper bounded by the inverse of the spectral gap, the fact that our algorithm is parameterized by the maximum effective resistance and capacitance over *all* subcomplexes means that the runtime of our algorithm is not entirely comparable to the runtime of existing QTDA algorithms. It is possible are complexes where the effective resistance and capacitance in subcomplexes are significantly lower than the spectral gap of the entire complex, and complexes where the effective resistance or capacitance of a cycle in a subcomplex is larger than the spectral gap of the entire complex. The complete complex is an example of the second case, as it has the maximal possible spectral gap of n_0 .

Randomized Order for the Incremental Algorithm. Building on the previous point, while our algorithm is parameterized by the maximum effective resistance and capacitance in various subcomplexes, our algorithm can also incrementally add the simplices *in any order*. This is potentially beneficial as a simplex may have smaller effective resistance or capacitance in one order than another. Of course, we likely will not know in advance whether or not a particular order of the simplices results in less or greater resistance and capacitance. However, we still may be able to use this fact to our advantage, as we could run our algorithm multiple times with different orders to gain confidence that our Betti number computations are accurate, which is not the case with previous QTDA algorithms.

5 Bounds on Effective Resistance and Capacitance.

In this section, we provide upper bounds on the resistance and capacitance of a cycle γ in an simplicial complex \mathcal{K} . Throughout this section, all simplicial complexes are **unweighted**.

Our upper bounds are polynomial in the number of d -simplices and the cardinality of the torsion subgroup of the relative homology groups. In particular, our bounds on resistance and capacitance are dependent on the maximum cardinality of the torsion subgroup of the relative homology group $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$, where $\mathcal{L} \subset \mathcal{K}$ is a d -dimensional subcomplex and $\mathcal{L}_0 \subset \mathcal{L}$ is a $(d-1)$ -dimensional subcomplex. In the worst case, our upper bounds are exponential.

In Theorem 5.9 we provide an example of a simplicial complex containing a cycle γ whose effective resistance is exponential in the number of simplices in the complex. It is important to reiterate that our bounds are in terms of the torsion of the relative homology groups, not the torsion of the (non-relative) homology groups. The simplicial complex we provide has no torsion in its homology groups, only torsion in its relative homology groups.

5.1 Upper Bounds

Our upper bounds rely on a change of basis on the boundary matrix called the Smith normal form which reveals information about the torsion subgroup of $H_{d-1}(\mathcal{K}, \mathbb{Z})$. We state the normal

form theorem below.

Theorem 5.1 (Munkres, Chapter 1 Section 11 [54]). *There are bases for $C_d(\mathcal{K})$ and $C_{d-1}(\mathcal{K})$ such that the matrix for the boundary operator $\partial_d: C_d(\mathcal{K}, \mathbb{Z}) \rightarrow C_{d-1}(\mathcal{K}, \mathbb{Z})$ is in **Smith normal form**, i.e.*

$$\tilde{\partial}_d = \begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix}$$

where D is a diagonal matrix with positive integer entries d_1, \dots, d_m such that each d_i divides d_{i+1} and each 0 is a zero matrix of appropriate dimensionality. The normal form of ∂_d satisfies the following properties:

1. The entries d_1, \dots, d_m correspond to the torsion coefficients of $H_{d-1}(\mathcal{K}, \mathbb{Z}) \cong \mathbb{Z}^{\beta_d} \oplus \mathbb{Z}_{d_1} \oplus \dots \oplus \mathbb{Z}_{d_m}$ (where $\mathbb{Z}_1 = 0$),
2. The number of zero columns is equal to the dimension of $\ker(\partial_d)$.

Moreover, the boundary matrix ∂ in the standard basis can be transformed to $\tilde{\partial}$ by elementary row and column operations. If ∂ is square, these operations multiply $\det \partial$ by ± 1 .

Using Theorem 5.1, we obtain an upper bound on the determinants of the square submatrices of the boundary matrix $\partial_d[\mathcal{K}]$ in terms of the *relative homology groups* of \mathcal{K} . Let \mathcal{L} be d -dimensional subcomplex of \mathcal{K} , and let \mathcal{L}_0 be a $(d-1)$ -dimensional subcomplex of \mathcal{K} . The **relative boundary matrix** $\partial_d[\mathcal{L}, \mathcal{L}_0]$ is the submatrix of ∂_d obtained by including the columns of the d -simplices in \mathcal{L} and excluding the rows of the $(d-1)$ -simplices in \mathcal{L}_0 . With the relative boundary matrices, one can define the **relative homology groups** as $H_d(\mathcal{L}, \mathcal{L}_0, \mathbb{Z}) = \ker \partial_d[\mathcal{L}, \mathcal{L}_0] / \text{im } \partial_{d+1}[\mathcal{L}, \mathcal{L}_0]$. More information on the relative boundary matrix can be found in [14]. We denote the cardinality of the torsion subgroup of the relative homology group $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$ by $\mathcal{T}(\mathcal{L}, \mathcal{L}_0)$. Similarly, we denote the maximum $\mathcal{T}(\mathcal{L}, \mathcal{L}_0)$ over all relative homology groups as $\mathcal{T}_{\max}(\mathcal{K})$.

Lemma 5.2. *Let $\partial_d[\mathcal{L}, \mathcal{L}_0]$ be a $k \times k$ square submatrix of ∂_d constructed by including columns for the d -simplices in \mathcal{L} and excluding rows for the $(d-1)$ -simplices in \mathcal{L}_0 . The magnitude of the determinant of $\partial_d[\mathcal{L}, \mathcal{L}_0]$ is bounded above by the cardinality of the torsion subgroup of $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$, i.e*

$$|\det(\partial_d[\mathcal{L}, \mathcal{L}_0])| \leq \mathcal{T}(\mathcal{L}, \mathcal{L}_0).$$

Proof. Without loss of generality, we assume that $\det(\partial_d[\mathcal{L}, \mathcal{L}_0]) \neq 0$; if $\det(\partial_d[\mathcal{L}, \mathcal{L}_0]) = 0$, the bound is trivial. Since $\partial_d[\mathcal{L}, \mathcal{L}_0]$ is a non-singular square matrix, its normal form $\tilde{\partial}_d[\mathcal{L}, \mathcal{L}_0]$ is a diagonal matrix $D = \text{diag}(d_1, \dots, d_k)$. The determinant is equal to $\pm \prod_{i=1}^k d_i$ and by Theorem 5.1 the torsion subgroup of $H_{d-1}(\mathcal{L}, \mathcal{L}_0)$ is $\mathbb{Z}_{d_1} \oplus \dots \oplus \mathbb{Z}_{d_k}$ which has cardinality $\mathcal{T}(\mathcal{L}, \mathcal{L}_0) = \prod_{i=1}^k d_i$. \square

5.1.1 Upper Bounds on Effective Resistance

We are now ready to upper bound the effective resistance of a cycle in a simplicial complex.

Theorem 5.3. *Let \mathcal{K} be a d -dimensional simplicial complex and γ a unit-length null-homologous $(d-1)$ -cycle in \mathcal{K} . Let $n = \min\{n_{d-1}, n_d\}$. The effective resistance of γ is bounded above by $\mathcal{R}_\gamma(\mathcal{K}) \in O(n^2 \cdot \mathcal{T}_{\max}(\mathcal{K})^2)$.*

Proof. First, we remove d -simplices from \mathcal{K} to create a new complex \mathcal{L} such that $\ker(\partial_d[\mathcal{L}]) = 0$ and $\text{im } \partial_d[\mathcal{K}] = \text{im } \partial_d[\mathcal{L}]$. Theorem B.3 proves that removing d -simplices only increases the effective resistance, so $\mathcal{R}_\gamma(\mathcal{K}) \leq \mathcal{R}_\gamma(\mathcal{L})$. As $\ker(\partial_d[\mathcal{L}]) = 0$, there is a unique unit γ -flow $f \in C_{d-1}(\mathcal{L})$ which implies $\mathcal{R}_\gamma(\mathcal{L}) = \|f\|^2$. Let $n \leq n_d$ denote the number of d -simplices in \mathcal{L} .

The matrix $\partial_d[\mathcal{L}]$ has full column rank, so we can find a non-singular $n \times n$ square submatrix of $\partial_d[\mathcal{L}]$; call this submatrix B . Let \mathcal{L}_0 be the $(d-1)$ -dimensional subcomplex that contains the $(d-1)$ -simplices corresponding to rows excluded from B ; B is the relative boundary matrix $\partial_d[\mathcal{L}, \mathcal{L}_0]$. We have that $Bf = c$, where c is the restriction of γ to the rows of B . Observe that $\|c\| \leq \|\gamma\| = 1$

We will apply Cramer's rule to upper bound the size of f . By Cramer's rule we have the equality

$$f(\sigma) = \frac{\det(B_{\sigma,c})}{\det(B)}$$

where $B_{\sigma,c}$ is the matrix obtained by replacing the column of B indexed by σ with the vector c . Since $\det(B)$ is integral, $|\det(B)| \geq 1$, so we drop the denominator and focus on the inequality $|f(\sigma)| \leq |\det(B_{\sigma,c})|$. We bound $|\det(B_{\sigma,c})|$ by its cofactor expansion,

$$\begin{aligned} |\det(B_{\sigma,c})| &= \left| \sum_{i=1}^{n_d} (-1)^i \cdot c_i \cdot \det(B_{\sigma,c}^{c,i}) \right| \\ &\leq \sum_{i=1}^{n_d} |c_i| \cdot \mathcal{T}_{\max}(\mathcal{K}) \\ &= \|c\|_1 \cdot \mathcal{T}_{\max}(\mathcal{K}) \\ &= O(\sqrt{n} \cdot \mathcal{T}_{\max}(\mathcal{K})) \end{aligned}$$

where $B_{\sigma,c}^{c,i}$ denotes the submatrix obtained by removing the column c and removing the i th row and c_i denotes the i th component of c . The first inequality comes from Lemma 5.2, as $B_{\sigma,c}^{c,i}$ is the relative boundary matrix $\partial_d[\mathcal{L} \setminus \{\sigma\}, \mathcal{L}_0 \cup \sigma_i]$, where σ_i is the $(d-1)$ -simplex corresponding to the i th row of B . The factor of \sqrt{n} comes from the fact that $\|c\|_1 \leq \sqrt{n}\|c\|_2$ and $\|c\|_2 \leq 1$. Finally, we compute the flow energy of f as

$$\begin{aligned} J(f) &= \sum_{\sigma \in \mathcal{L}_d} f(\sigma)^2 \\ &\leq \sum_{i=1}^n n \cdot \mathcal{T}_{\max}(\mathcal{K})^2 \\ &= O(n^2 \cdot \mathcal{T}_{\max}(\mathcal{K})^2). \end{aligned}$$

The effective resistance of γ is the flow energy of f , so the result follows. \square

If $\mathcal{L} \subset \mathcal{K}$, then the boundary matrix $\partial_d[\mathcal{L}]$ is a submatrix of $\partial_d[\mathcal{K}]$. In particular, $\mathcal{T}_{\max}(\mathcal{L}) \leq \mathcal{T}_{\max}(\mathcal{K})$. Therefore, the proof of Theorem 5.3 gives an upper bound on the effective resistance for any subcomplex $\mathcal{L} \subset \mathcal{K}$.

Corollary 5.4. *Let $\mathcal{L} \subset \mathcal{K}$ be a d -dimensional simplicial complex and γ a null-homologous $(d-1)$ -cycle in \mathcal{L} . Let $n = \min\{n_{d-1}[\mathcal{L}], n_d[\mathcal{L}]\}$. The effective resistance of γ in \mathcal{L} is bounded above by $\mathcal{R}_\gamma(\mathcal{L}) = O(n^2 \cdot \mathcal{T}_{\max}(\mathcal{K})^2)$.*

In Section 5.1.3, we give an upper bound on relative torsion, which implies an upper bound on the effective resistance purely in terms of the size of the complex.

5.1.2 Upper Bounds on Capacitance.

We now provide an upper bound for the effective capacitance of a cycle. While upper bounds on the effective resistance only depended on the norm of the cycle, upper bounds on the capacitance of the cycle are not. Therefore, we consider the special case where γ is the boundary of a d -simplex. This is a natural assumption as these are exactly the type of cycles considered in the incremental algorithm for computing Betti numbers (Algorithm 1). While we only prove this special case, we note that our proof could be adapted to bound the effective capacitance of a cycle whose entries have constant upper and lower bounds.

Theorem 5.5. *Let $\mathcal{L} \subset \mathcal{K}$ be d -dimensional simplicial complexes. Let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$ -cycle that is null-homologous in \mathcal{K} but not in \mathcal{L} . Let $n = \min\{n_{d-1}, n_d\}$. Assume that $\gamma = \partial\sigma$ for a d -simplex $\sigma \notin \mathcal{L}$. The effective capacitance of γ in \mathcal{K} is bounded above by $C_\gamma(\mathcal{L}, \mathcal{K}) \in O(n \cdot n_0 \cdot \mathcal{T}_{\max}(\mathcal{K})^2)$.*

Proof. Let p be a γ -potential. We upper bound the potential energy of p . By definition, $\delta[\mathcal{L}]p = 0$ and $\gamma^T p = 1$. We can express these constraints as the linear system

$$\begin{bmatrix} \delta[\mathcal{L}] \\ \gamma^T \end{bmatrix} p = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

We first remove linearly-dependent columns from this linear system until this system has full column rank. Columns of the matrix are indexed by $(d-1)$ simplices of \mathcal{L} , and rows are indexed by d -simplices of \mathcal{L} . Removing columns from $\delta[\mathcal{L}]$ changes it to the relative coboundary matrix $\delta[\mathcal{L}, \mathcal{L}_0]$ where \mathcal{L}_0 is the $(d-1)$ -subcomplex corresponding to the columns that were removed. Removing linearly-dependent columns does not change the image of the system of equation, so there is still a solution r , i.e.

$$\begin{bmatrix} \delta[\mathcal{L}, \mathcal{L}_0] \\ c^T \end{bmatrix} r = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

where c is the subvector of γ after removing the columns. The vector r is not a γ -potential as it is a vector in $C_{d-1}(\mathcal{L}, \mathcal{L}_0)$, not $C_{d-1}(\mathcal{L})$. However, we can extend r to be a γ -potential by adding zeros in the entries indexed by \mathcal{L}_0 . Adding zero-valued entries preserves the length of r .

We now want to remove rows from this matrix so that it has full row rank. Topologically, removing rows corresponds to removing d -simplices from the complex \mathcal{L} to create a new complex \mathcal{L}_1 . Note that we must always include the row c to have full row rank; otherwise, r would be a non-zero vector in the kernel of this system, meaning the system does not have full rank. Removing these rows gives the linear system

$$\begin{bmatrix} \delta[\mathcal{L}_1, \mathcal{L}_0] \\ c \end{bmatrix} r = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}.$$

Let $C = [\delta[\mathcal{L}_1, \mathcal{L}_0]^T \quad c^T]^T$ and $b = [0 \quad 0 \quad \dots \quad 1]^T$. Note that C is a square matrix of size (say) $m \times m$, where $m \leq n + 1$.

We use Cramer's rule to bound the size of $\|r\|$. By Cramer's rule, r_i , the i th entry of r , is

$$r_i = \frac{\det(C_{i,b})}{\det(C)}.$$

where $C_{i,b}$ is the matrix obtained by replacing the i th column with b .

We first lower bound $|\det(C)|$. As C is a full-rank integral matrix, then $|\det(C)| \geq 1$. We now upper bound $|\det(C_{i,b})|$. We calculate $\det(C_{i,b})$ with the cofactor expansion on the column replaced by b . As b has 1 in its last entry and 0s elsewhere, the cofactor expansion gives $\det(C_{i,b}) = \pm 1 \cdot \det(C_{i,b}^{i,c})$ where $C_{i,b}^{i,c}$ is the matrix where we dropped the i th column and the row c from $C_{i,b}$. The matrix $C_{i,b}^{i,c}$ is a square submatrix of $\delta[\mathcal{K}]$, so we can bound $|\det(C_{i,b})| \leq \mathcal{T}_{\max}(\mathcal{K})$. Thus, $r_i = \det(C_{i,b})/\det(C) \leq \mathcal{T}_{\max}(\mathcal{K})$ and

$$\begin{aligned} \|r\| &= \sqrt{\sum_{i=1}^m r_i^2} \\ &\leq \sqrt{m \cdot \mathcal{T}_{\max}(\mathcal{K})^2} \\ &\leq \sqrt{n \cdot \mathcal{T}_{\max}(\mathcal{K})^2} \\ &= \sqrt{n} \cdot \mathcal{T}_{\max}(\mathcal{K}) \end{aligned}$$

The potential energy of r is $\|\delta[\mathcal{K}]r\|^2$. We can use Lemma 4.5 to obtain the bound $\|\delta[\mathcal{K}]r\|^2 = O(n \cdot n_0 \cdot \mathcal{T}_{\max}(\mathcal{K})^2)$. \square

5.1.3 Upper Bound on Relative Torsion.

To conclude this section, we provide an upper bound on $\mathcal{T}_{\max}(\mathcal{K})$.

Lemma 5.6. *Let \mathcal{K} be a simplicial complex. Let $n = \min\{n_{d-1}, n_d\}$. Then the maximum rank of any $(d-1)$ -dimensional relative torsion group of K is $\mathcal{T}_{\max}(\mathcal{K}) \in O((\sqrt{d+1})^n)$.*

Proof. By the proof of Lemma 5.2, the quantity $\mathcal{T}_{\max}(\mathcal{K})$ is the absolute value of the determinant of some submatrix of ∂_d . We therefore bound the determinant of such a submatrix. We prove this bound using **Hadamard's Inequality**: the determinant of an $m \times m$ matrix B is upper-bounded by the product of the norms of its columns.

Consider a square, $m \times m$ submatrix B of ∂_d . We know that $m \leq n$. Moreover, any column of B has norm bounded above by $\sqrt{d+1}$. This bound follows from the fact that each column of ∂_d has norm exactly $\sqrt{d+1}$; each column has $d+1$ nonzero entries, each of which is ± 1 . The bound of the lemma follows by Hadamard's Inequality. \square

Lemma 5.6 immediately implies the corollaries to Theorem 5.3 and Theorem 5.5.

Corollary 5.7. *Let \mathcal{K} be a d -dimensional simplicial complex and γ a unit-length null-homologous $(d-1)$ -cycle in \mathcal{K} . Let $n = \min\{n_{d-1}, n_d\}$. The effective resistance of γ is bounded above by $\mathcal{R}_\gamma(\mathcal{K}) \in O(n^2(d+1)^n)$.*

Corollary 5.8. *Let $\mathcal{L} \subset \mathcal{K}$ be d -dimensional simplicial complexes. Let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$ -cycle that is null-homologous in \mathcal{K} but not in \mathcal{L} . Let $n = \min\{n_{d-1}, n_d\}$. Assume that $\gamma = \partial\sigma$ for a d -simplex $\sigma \notin \mathcal{L}_d$. The effective capacitance of γ in \mathcal{K} is bounded above by $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K}) \in O(n \cdot n_0 \cdot (d+1)^n)$.*




Figure 3: Left: Stellar subdivision of a triangle. Middle and Right: Top and bottom view of the stellar prism of a triangle with the tetrahedra pushed apart.

5.2 Lower Bounds.

5.2.1 Lower Bounds on Effective Resistance

At the end of the previous section, we gave an exponential upper bound on the effective resistance of a $(d - 1)$ -cycle in a simplicial complex (Corollary 5.7). In this section, we describe a d -dimensional simplicial complex \mathcal{B}_d^n with a $(d - 1)$ -cycle γ with exponentially-large effective resistance with respect to the size of the complex.

Theorem 5.9. *Let d, n be positive integers. There is a constant $c_d \geq 1$ that depends only on d and a d -dimensional simplicial complex \mathcal{B}_d^n with $n_d \in \Theta((d + 1)^3 n)$ d -simplices and a unit-length null-homologous cycle $\gamma \in C_{d-1}(\mathcal{B}_d^n)$ such that $\mathcal{R}_\gamma(\mathcal{B}_d^n) \in \Theta(c_d^{n_d})$.*

The building block. Our simplicial complex will be obtained by gluing together multiple instances of the same “building block” B_d . A formal description of B_d is given in Appendix E; here, we give an intuitive description of the complex. Let Δ^d denote the closure of the d -simplex $\sigma = \{v_0, \dots, v_d\}$, and let $\partial\Delta^d$ denote the $(d - 1)$ -dimensional simplicial complex $\Delta^d \setminus \{\sigma\}$. Our construction starts with a triangulation of the space $\partial\Delta^d \times [0, 1]$ that we call the **stellar prism**. The “bottom copy” $\partial\Delta^d \times \{0\}$ is triangulated like the original complex $\partial\Delta^d$, and the “top copy” $\partial\Delta^d \times \{1\}$ is triangulated with the **stellar subdivision**, the subdivision that adds a vertex to the center of each $(d - 1)$ -simplex. See Figure 3. The relevant property of this triangulation is that the bottom copy has $d + 1$ $(d - 1)$ -simplices, and the top copy has $d \cdot (d + 1)$ $(d - 1)$ -simplices. The **building block** B_d is obtained from this triangulation by identifying the vertex in the center of each $(d - 1)$ -simplex $\tau \times \{1\}$ with the unique vertex in $\sigma \times \{1\} \setminus \tau \times \{1\}$, the unique vertex in the set $\sigma \times \{1\}$ that is not a vertex of the simplex $\tau \times \{1\}$. See Figure 4.

When we identify the vertices, each $(d - 1)$ -simplex in the top of the stellar prism is replaced by one of the $(d - 1)$ faces of $\sigma \times \{1\}$. Moreover, this replacement is d -to-1, meaning each face of $\sigma \times \{1\}$ replaces a $(d - 1)$ -simplex exactly d times, or informally, each $(d - 1)$ -simplex “appears d times” in the top copy of B_d . Of course, this is not literally true, as a simplicial complex can only contain a single copy of each simplex. However, something to this effect is true. Namely, there is a d -chain $f \in C_d(B_d)$ whose boundary assigns value ± 1 to each $(d - 1)$ -simplex in the bottom copy and value $\pm d$ to each $(d - 1)$ -simplex in the top copy. The key properties of the building block B_d are summarized in the following lemma.

Lemma 5.10. *Let $\sigma = \{v_0, \dots, v_d\}$ be a set. There is a d -dimensional simplicial complex B_d with vertices $\sigma \times \{0, 1\}$ such that*

1. B_d has $\Theta((d + 1)^3)$ d -simplices.




Figure 4: The complex \mathcal{B}_2^3 . Recall that the vertices of \mathcal{B}_2^3 are of the form (v_i, j) for $0 \leq i \leq 2$ and $0 \leq j \leq 3$; colors in the figure denote the second coordinate of the vertices. Vertices with the same label are identified.

2. there is a d -chain $f \in C_d(B_d)$ such that

- (i) $\partial f = \partial(\sigma \times \{0\}) + d \cdot \partial(\sigma \times \{1\})$
- (ii) $\|f\|^2 \in \Omega((d+1)^3)$

Note that neither of the simplices $\sigma \times \{0\}$ or $\sigma \times \{1\}$ are in B_d ; however, all of their faces are in the complex, so the boundary of these simplices are well-defined.

The total complex. The complex \mathcal{B}_d^n is obtained by gluing together n copies of B_d . We describe this gluing inductively on n . The vertices of \mathcal{B}_d^n are $\sigma \times \{0, \dots, n\}$. The base case \mathcal{B}_d^0 is the complete complex on the vertices $\sigma \times \{0\}$. Inductively, the complex \mathcal{B}_d^n is obtained from \mathcal{B}_d^{n-1} by identifying the vertices $\sigma \times \{n-1\}$ of \mathcal{B}_d^{n-1} with the vertices $\sigma \times \{1\}$ of a copy of B_d . We will denote this copy of B_d as B_d^n and the vertices $\sigma \times \{0\}$ of B_d^n as $\sigma \times \{n\}$. See Figure 4.

The key property of \mathcal{B}_d^n is that it has an exponentially-large chain with constant-sized boundary.

Lemma 5.11. *Let $d \geq 1$ and $n \geq 1$. There is a d -chain $y_n \in C_d(\mathcal{B}_d^n)$ such that*

- 1. $\|y_n\|^2 \in \Omega(d^{2n}(d+1))$
- 2. $\|\partial y_n\|^2 = d+1$

Proof. We construct this chain by induction on n . The chain y_n will have $\partial y_n = \partial(\sigma \times \{n\})$. For the base case $n=0$, the chain $y_0 = \sigma \times \{0\}$ clearly has this property.

For the inductive case, recall from Lemma 5.10 that there is a d -chain $f \in C_d(B_d)$ such that $\partial f = \partial(\sigma \times \{0\}) + d \cdot \partial(\sigma \times \{1\})$. Let f_n denote this chain in B_d^n . We define the chain $y_n := f_n - d \cdot y_{n-1}$. We now verify that y_n has boundary $\partial(\sigma \times \{n\})$.

$$\begin{aligned} \partial y_n &= \partial f_n - d \cdot \partial y_{n-1} \\ &= \partial(\sigma \times \{n\}) + d \cdot \partial(\sigma \times \{n-1\}) - d \cdot \partial(\sigma \times \{n-1\}) \\ &= \partial(\sigma \times \{n\}) \end{aligned}$$

It is clear that $\|\partial y_n\|^2 = d+1$, so we just need to lower bound $\|y_n\|^2$. We prove that $\|y_n\|^2 \in \Omega(d^{2n}(d+1))$ by induction. For the base case of $n=1$, we know that $\|y_1\|^2 \in \Omega((d+1)^3)$ by Lemma 5.10. For the inductive case, we can see that $f_n \perp y_{n-1}$ as these chains are supported on different sets of simplices. Therefore, $\|y_n\|^2 = \|f_n\|^2 + d^2 \cdot \|y_{n-1}\|^2 \in \Omega(d^{2n}(d+1))$. \square

Lemma 5.11 shows that \mathcal{B}_d^n has an exponentially-large d -chain f with constant-sized boundary. If we can show that $\ker \partial_d = 0$, this will prove that ∂f has exponentially-large effective resistance, as f will be the *only* d -chain with boundary ∂f .

Corollary 5.12. *The kernel of the boundary matrix $\partial_d[\mathcal{B}_d^n]$ is trivial, i.e. $\ker \partial_d[\mathcal{B}_d^n] = 0$.*

Proof. Lemma F.3 in Appendix F shows that \mathcal{B}_d^n collapses to a $(d - 1)$ -dimensional subcomplex; call this subcomplex L . This implies that \mathcal{B}_d^n is homotopy equivalent to L , so in particular, the d -dimensional homology $H_d(\mathcal{B}_d^n) = 0$. As \mathcal{B}_d^n is d -dimensional, then $\text{im } \partial_{d+1}[\mathcal{B}_d^n] = 0$, so the only way that $H_d(\mathcal{B}_d^n)$ can equal 0 is if $\ker \partial_d[\mathcal{B}_d^n] = 0$. \square

Proof of Theorem 5.9. Our cycle is the normalized cycle $\gamma = \partial y_n / \|\partial y_n\|$ where y_n is the d -chain from Lemma 5.11. We know that $\|y_n\|^2 / \|\partial y_n\|^2 \in \Omega(d^{2n})$. Moreover, we know that $\ker \partial_d = 0$ by Corollary 5.12. Therefore, we conclude that effective resistance $\mathcal{R}_\gamma(\mathcal{B}_d^n) \in \Theta(d^{2n})$.

Now we need to restate this bound in terms of the number of d -simplices of \mathcal{B}_d^n . Each copy of B_d has $\Theta((d + 1)^3)$ d -simplices. Therefore, the entire complex \mathcal{B}_d^n has $n_d = \Theta(n \cdot (d + 1)^3)$ d -simplices. If we substitute n_d into our bound, we find that the effective resistance is bound below by $O((d + 1)^{2n_d/c \cdot (d+1)^3})$ for some constant c . Therefore, the constant in the theorem statement is $c_d = (d + 1)^{2/c \cdot (d+1)^3}$. \square

Related Work. Variants of this construction are sometimes called the *Iterated Mapping Cylinder* and have been used as a worst-case construction for other topological properties like torsion [56], homotopy [21], or embeddability [7, 22]. However, ours is the first work showing these complexes have a cycle with exponentially-large effective resistance (or exponentially-small spectral gap, as we will see later.) We were specifically inspired by the work of Newman [56].

However, our construction is more efficient than previous constructions of the Iterated Mapping Cylinder in terms of the number of d -simplices; this is why we dedicate several pages in the appendix to this construction. As an example, the building block in Newman’s construction is the iterated suspension of the Möbius band, and the number of d -simplices in the suspension grows exponentially with the dimension d . Comparatively, the number of d -simplices in our building block only grows polynomially with the dimension. Additionally, our construction has a cycle with ± 1 coefficients that is homologous to a cycle with $\pm d$ coefficients; in contrast, previous works have a cycle with ± 1 coefficients that is homologous to a cycle with ± 2 coefficients. Both of these properties result in a larger constant c_d in Theorem 5.9. It is an open question if there is a simplicial complex with a cycle whose effective resistance exactly matches the constant of the lower bound, i.e. $c_d \in \Theta(d + 1)$.

5.2.2 Lower Bounds on Capacitance.

In this section, we describe a pair of nested simplicial complexes that have a cycle with exponentially large effective capacitance. This complex will be built from the same building block as the complex for lower bounding effective resistance, but the simplicial complex will be built by gluing together the building blocks in slightly different ways.

Recall that the building block is a simplicial complex B_d with vertices $\sigma \times \{0, 1\}$, where $\sigma = \{v_0, \dots, v_d\}$ is a d -simplex. For a natural number n , we recursively construct the simplicial complex \mathcal{Q}_d^n . The simplicial complex \mathcal{Q}_d^n is obtained by gluing together n copies of B_d . The vertices of \mathcal{Q}_d^n will be denoted $\sigma \times \{0, \dots, n\}$. The base case \mathcal{Q}_d^n is the complete complex on the vertices $\sigma \times \{0\}$. Inductively, the complex \mathcal{Q}_d^n is obtained from \mathcal{Q}_d^{n-1} by identifying the vertices $\sigma \times \{n-1\}$ of \mathcal{Q}_d^{n-1} with the vertices $\sigma \times \{0\}$ of a copy of B_d .⁷ We will denote this copy of B_d as B_d^n and the vertices $\sigma \times \{1\}$ as $\sigma \times \{n\}$. The simplicial complex \mathcal{P}_d^n is defined \mathcal{Q}_d^{n-1} minus the simplex $\sigma \times \{0\}$.

⁷It is worth comparing \mathcal{Q}_d^n to the simplicial complex \mathcal{B}_d^n from the bounds on effective resistance. Both complexes are obtained by identifying vertices of copies of the building block B_d , but the identifications are made in different ways. When a building block is added to \mathcal{B}_d^i to form \mathcal{B}_d^{i+1} , its “top” $\sigma \times \{1\}$ is identified with $\sigma \times \{i\}$. When a

We will prove that the cycle $\gamma = \partial(\sigma \times \{n\})$ has exponentially large capacitance in $\mathcal{P}_d^n \subset \mathcal{Q}_d^n$. In order for γ to have finite capacitance, the cycle γ must be null-homologous in \mathcal{Q}_d^n but not null-homologous in \mathcal{P}_d^n . We prove this in the following lemma.

Lemma 5.13. *Let \mathcal{P}_d^n and \mathcal{Q}_d^n as described above. Consider the chain $\partial(\sigma \times \{n\}) \in C_d(\mathcal{P}_d^n)$.*

1. *The $(d-1)$ -chain $\partial(\sigma \times \{n\}) + (-1)^{n-1} \frac{1}{d^n} \partial(\sigma \times \{0\})$ is null-homologous in \mathcal{P}_d^n .*
2. *The $(d-1)$ -chain $\partial(\sigma \times \{n\})$ is null-homologous in \mathcal{Q}_d^n .*
3. *The $(d-1)$ -chain $\partial(\sigma \times \{n\})$ is not null-homologous in \mathcal{P}_d^n .*

Proof. Proof of Part (1) We will prove that the cycle $\partial(\sigma \times \{i\}) + \frac{1}{d^i} \partial(\sigma \times \{0\})$ is null-homologous in \mathcal{Q}_d^i by induction on i . Specifically, we will find a d -chain y_i such that $\partial y_i = \partial(\sigma \times \{i\}) + \frac{1}{d^i} \partial(\sigma \times \{0\})$. For each copy of the building block B_d^i , let $f_i \in C_d(B_d^i)$ be the d -chain guaranteed by Lemma 5.10 such that $\partial f_i = \partial(\sigma \times \{i-1\}) + d \cdot \partial(\sigma \times \{i\})$. For the base case of $i=1$, the d -chain $y_1 = \frac{1}{d} f_1$. Inductively, we define the chain $y_i = \frac{1}{d} f_i - \frac{1}{d} y_{i-1}$. We can verify that y_i has the claimed boundary as

$$\begin{aligned} \partial y_i &= \frac{1}{d} \partial f_i - \frac{1}{d} \partial y_{i-1} \\ &= \partial(\sigma \times \{i\}) + \frac{1}{d} \partial(\sigma \times \{i-1\}) - \frac{1}{d} \partial(\sigma \times \{i-1\}) + (-1)^{i-1} \frac{1}{d^i} \partial(\sigma \times \{0\}) \\ &= \partial(\sigma \times \{i\}) + (-1)^{i-1} \frac{1}{d^i} \partial(\sigma \times \{0\}) \end{aligned}$$

Proof of Part (2) As the simplex $(\sigma \times \{0\}) \in \mathcal{Q}_d^n$, then the chain $y_n + (-1)^n \frac{1}{d^n} \cdot (\sigma \times \{0\})$ obviously has boundary $\partial(\sigma \times \{n\})$.

Proof of Part (3) To prove that $\partial(\sigma \times \{n\})$ is not null-homologous in \mathcal{P}_d^n , we will use two lemmas we prove in the appendix. Lemma F.4 shows that \mathcal{P}_d^n collapses to a $(d-1)$ -dimensional subcomplex—call it L —containing the support of $\partial(\sigma \times \{n\})$. Lemma F.2 shows that for nested complexes $L \subset K$ such that K collapses to L , a cycle is null-homologous in L if and only if it is null-homologous in K . As $\partial(\sigma \times \{n\})$ is not null-homologous in L (L is $(d-1)$ -dimensional, so $\text{im } \partial_d[L] = 0$), then $\partial(\sigma \times \{n\})$ is not null-homologous in \mathcal{P}_d^n either. \square

Theorem 5.14. *Let d, n be positive integers. There is a pair of nested d -dimensional simplicial complexes $\mathcal{P}_d^n \subset \mathcal{Q}_d^n$ with $n_d \in \Theta((d+1)^3 n)$ d -simplices, a unit-length null-homologous cycle $\gamma \in C_{d-1}(\mathcal{Q}_d^n)$, and a constant $c_d \geq 1$ that depends only on d such that $\mathcal{C}_\gamma(\mathcal{P}_d^n, \mathcal{Q}_d^n) \in \Theta(c_d^{n_d})$.*

Proof. The complexes \mathcal{P}_d^n and \mathcal{Q}_d^n are the complexes described in the preceding paragraphs. The cycle $\gamma = \partial(\sigma \times \{n\}) / \sqrt{d+1}$, where $\sqrt{d+1}$ is a normalization factor. We must show that any unit γ -potential p has exponentially-large potential energy.

Let p be any unit γ -potential in \mathcal{P}_d^n . We know that $\gamma^T p = 1$ and $\delta_{d-1}[\mathcal{P}_d^n] p = 0$. As $\delta_{d-1}[\mathcal{P}_d^n] = \partial_d^T[\mathcal{P}_d^n]$, the second condition is equivalent to saying that $p^T b = 0$ for any vector $b \in \text{im } \partial_d[\mathcal{P}_d^n]$. Lemma 5.13 Part 1 proves that $\partial(\sigma \times \{n\}) + (-1)^{n-1} d^{-n} \partial(\sigma \times \{0\}) \in \text{im } \partial_d[\mathcal{P}_d^n]$, so the previous two facts imply $p^T (\partial(\sigma \times \{n\}) + (-1)^{n-1} d^{-n} \partial(\sigma \times \{0\})) = 0$. As $p^T \partial(\sigma \times \{n\}) = \sqrt{d+1}$ and $p^T (\partial(\sigma \times$

building block is added to \mathcal{Q}_d^i to form \mathcal{Q}_d^{i+1} , its “bottom” $\sigma \times \{0\}$ is identified with $\sigma \times \{i\}$. It is easiest to see the difference between \mathcal{B}_d^n and \mathcal{Q}_d^n by considering the cycle $\partial(\sigma \times \{n\})$ in both complexes. The unique d -chain $f \in C_d(\mathcal{B}_d^n)$ with boundary $\partial(\sigma \times \{n\})$ assigns exponentially-large coefficients to some d -simplices. In contrast, the unique d -chain $f \in C_d(\mathcal{Q}_d^n)$ assigns exponentially-small coefficients to some d -simplices. (This is a corollary to the proof of Lemma 5.13.)

$\{n\}) + (-1)^{n-1}d^{-n}(\partial(\sigma \times \{0\}))) = 0$, then we conclude that $p^T \partial(\sigma \times \{0\}) = (-1)^{n-1}d^n/\sqrt{d+1}$. Moreover, as the only d -simplex in \mathcal{Q}_d^n that is not in \mathcal{P}_d^n is $\sigma \times \{0\}$, then the fact that $\delta_{d-1}[\mathcal{P}_d^n]p = 0$ implies that the potential energy $\|\delta_{d-1}[\mathcal{Q}_d^n]p\|^2 = (p^T \partial(\sigma \times \{0\}))^2 = \Omega(d^{2n})$.

This shows that p has exponentially-large potential energy with respect to n . By the same argument as in the proof of Theorem 5.9, we can show that p also has exponentially-large potential energy with respect to n_d , but for a different (but constant) base of the exponent c_d . \square

6 Bounds on the Spectral Gap

Our lower and upper bounds on effective resistance imply lower and upper bounds on the spectral gap of the combinatorial Laplacian. This is because *the spectral gap is the inverse of the maximum effective resistance of all unit-length, null-homologous cycles*, a fact we proved in Lemma 4.4. Therefore, a corollary of Theorem 5.9 is that the spectral gap of the combinatorial Laplacian can be exponentially-small in the worst-case. This resolves one of the most important open questions in the field of Quantum Topological Data Analysis and shows that Betti Number Estimation algorithms must run for an exponentially-long time to exactly compute Betti numbers.

6.1 Exponentially-small spectral gap.

Based on this connection between the spectral gap and effective resistance of Lemmas 4.4 and 2.3, we can derive lower and upper bounds on the spectral gap of the d -combinatorial Laplacian as corollaries of the upper and lower bounds on the effective resistance (Corollary 5.7 and Theorem 5.9). While lower bounds on the spectral gap were previously known (see for example [22, Proof of Theorem 1.2]), one advantage of our proof is that it provides a necessary condition for large spectral gap, namely the existence of a subcomplex with exponentially-large relative torsion.

Theorem 1.2. *Let K be a simplicial complex. Let n_i be the number of i -simplices of K . Let $n = \max\{\min\{n_{d-1}, n_d\}, \min\{n_d, n_{d+1}\}\}$. Then the spectral gap $\lambda_{\min}(L_d[K]) \in \Omega\left(\frac{1}{n^{2d^n}}\right)$.*

Theorem 6.1. *Let $d, n \geq 1$. There is a d -dimensional simplicial complex \mathcal{B}_d^n with $n_d \in \Theta((d+1)^3 \cdot n)$ d -simplices and a constant $c_d \geq 1$ that depends only on d such that the spectral gaps of $L_{d-1}[\mathcal{B}_d^n]$ and $L_d[\mathcal{B}_d^n]$ are $O\left(\frac{1}{c_d^{n_d}}\right)$.*

We remark that we can also derive lower bound of the spectral gap in terms of relative torsion as a corollary to Theorem 5.3. This implies bounded spectral gap for simplicial complexes with no relative torsion, as remarked upon by Friedman [23, Theorem 7.2] in the case of orientable d -manifolds.

6.2 Many Small Eigenvalues.

Corollary 6.1 shows there is a simplicial complex with a single small eigenvalue. It is natural to ask whether there is a bound on the number of very small eigenvalues a simplicial complex can have. This is relevant to QTDA algorithms that work by counting the number of eigenvalues smaller than a given threshold. Here, we provide a complex \mathcal{M}_d^n with a polynomial number of exponentially-small eigenvalues.

Corollary 6.2. *Let $n, d \geq 1$. There exists a simplicial complex \mathcal{M}_d^n with $n_d \in \Theta((d+1)^3 \cdot n)$ d -simplices and a constant $c_d > 1$ that depends only on d such that both $L_{d-1}[\mathcal{M}_d^n]$ and $L_d[\mathcal{M}_d^n]$ have $\Omega(\sqrt{n_d})$ eigenvalues of size $O\left(\frac{1}{c_d^{\sqrt{n_d}}}\right)$.*

Proof of Theorem 6.2. This follows from Part 2 of Lemma 2.2 relating the spectrum of a complex to the spectra of its connected components. We define the complex \mathcal{M}_d^n to be n_d disjoint copies of the \mathcal{B}_d^n of Corollary 6.1. The complex \mathcal{M}_d^n has n_d^2 d -simplices and n_d eigenvalues of size $O(\frac{1}{c_d^{n_d}})$. \square

6.3 Clique-Dense Complexes.

Existing QTDA algorithms perform best when the simplicial complex is clique-dense, meaning that the simplicial complex has close to the maximal number of d -simplices, i.e $n_d \sim \binom{n_0}{d+1}$; However, the simplicial complex we constructed in Theorem 6.1 is sparse: it only has $n_d \in \Theta((d+1)^2 n_0)$ d -simplices. Therefore, Theorem 6.1 does not rule out the possibility that clique-dense complexes avoid worst-case spectral gap.

However, in this section, we show that we can extend the construction of Theorem 6.1 to clique-dense complexes (at the expense of making the constant c_d of the exponent smaller.) To do this, we use a probabilistic coloring argument of Newman [56] that reduces the number of vertices of a simplicial complex while preserving the number of d -simplices and the Laplacian.

We begin with definitions. A **coloring** of a simplicial complex is a map on its vertices $c : \mathcal{K}_0 \rightarrow \mathbb{N}$. The **pattern complex** of a simplicial complex K with coloring c is the simplicial complex $\mathcal{K}_c = \{\{c(v) : v \in \sigma\} : \sigma \in \mathcal{K}\}$; intuitively, the pattern complex of \mathcal{K} is the simplicial complex obtained by identify all vertices of \mathcal{K} of the same color and identifying all simplices whose vertices have the same set of colors. While a simplex in \mathcal{K} may be mapped to a lower-dimensional simplex in the pattern complex \mathcal{K}_c if two of its vertices are the same color, we are only considered with colorings where this does not happen. A **proper coloring** of a d -dimensional simplicial complex \mathcal{K} is a coloring $c : \mathcal{K}_0 \rightarrow \mathbb{N}$ such that (1) the endpoints of each edge in \mathcal{K} are different colors and (2) $\{c(v) : v \in \sigma_1\} \neq \{c(v) : v \in \sigma_2\}$ for any distinct $(d-1)$ -simplices $\sigma_1, \sigma_2 \in K$. Note that for a proper coloring, condition (1) guarantees that each simplex in \mathcal{K} corresponds to a simplex of the same dimension in \mathcal{K}_c . Additionally, \mathcal{K} and \mathcal{K}_c have the same set of $(d-1)$ and d -simplices up to recoloring. Proper colorings are relevant to our paper as they preserve the spectral gap.

Lemma 6.3. *Let \mathcal{K} be a d -dimensional simplicial complex and let c be a proper coloring of \mathcal{K} . Then $\lambda_{\min}(L_{d-1}^{up}[\mathcal{K}]) = \lambda_{\min}(L_{d-1}^{up}[\mathcal{K}_c])$*

Proof. This follows as \mathcal{K} and \mathcal{K}_c have the same set of $(d-1)$ and d -simplices up to recoloring, so the boundary maps $\partial_d[\mathcal{K}]$ and $\partial_d[\mathcal{K}_c]$ are the same up to the signs on the simplices; however, the spectrum of the up Laplacians $L_{d-1}^{up}[\mathcal{K}]$ and $L_{d-1}^{up}[\mathcal{K}_c]$ are unaffected by different orientations of the simplices ([24, Theorem 4.1.1]), so the lemma follows. \square

Newman's method also requires bounds on a generalized notion of degree. For a d -dimensional simplicial complex \mathcal{K} and natural numbers $0 \leq i < j \leq d$, define $\Delta_{i,j}(\mathcal{K}) = \max_{\sigma \in \mathcal{K}_i} |\{\tau \in \mathcal{K}_j : \sigma \subset \tau\}|$ and $\Delta(\mathcal{K}) = \max_{1 \leq i < j \leq d} \Delta_{i,j}(\mathcal{K})$. Newman used a probabilistic argument to show that for a d -dimensional simplicial complex \mathcal{K} there was always a proper coloring with a bounded number of colors.

Lemma 6.4 (Lemma 3, Newman [56]). *Let \mathcal{K} be a d -dimensional simplicial complex such that $\Delta(\mathcal{K}) \geq 4$. Then there is a proper coloring c of \mathcal{K} with at most $18(\Delta(\mathcal{K}) + 1)^6 d^6 \sqrt[d]{n_0}$ colors.*

This implies the following corollary of our bound on the spectral gap.

Theorem 1.3. *Let $d, n \geq 1$. There are constants c_d, κ_d that depends only on d and a d -dimensional simplicial complex \mathcal{C}_d^n with $n_d = \Omega(\kappa_d \binom{n_0}{d})$ d -simplices such that the spectral gaps $\lambda_{\min}(L_{d-1}[\mathcal{C}_d^n]), \lambda_{\min}(L_d[\mathcal{C}_d^n]) \in O(\frac{1}{c_d^{n_d}})$.*

Proof. This is a corollary to Theorem 6.1, Lemma 6.3, and Lemma 6.4. The simplicial complex \mathcal{C}_d^n is a pattern complex of a coloring c of the complex \mathcal{B}_d^n from Theorem 6.1. This coloring c is the coloring guaranteed by Lemma 6.4. This coloring has $f(d) \sqrt[d]{n_0}$ colors for some function f ; we can see this by bounding $\Delta(\mathcal{B}_d^n)$ by a function of d . Examining its construction in Section 5.2, each simplex in \mathcal{B}_d^n is in at most two different building blocks. Each building block has $2(d+1)$ vertices, so for any $1 \leq i < j \leq d$, an i -simplex in \mathcal{B}_d^n is incident to at most $\binom{4d-i-1}{j-i}$ j -simplices. Thus, $\Delta(\mathcal{B}_d^n)$ is at most some function of d . Therefore, there is a simplicial complex with $\tilde{n}_0 = f(d) \sqrt[d]{n_0}$ vertices and $\theta\left(\frac{(d+1)^3}{f(d)^d} \tilde{n}_0^d\right)$ d -simplices. As $\binom{n_0}{d} \leq \left(\frac{en_0}{d}\right)^d$, then there are $\Omega(\kappa_d(\tilde{n}_0))$ d -simplices in \mathcal{C}_d^n for some appropriate function $f(d) = \Omega\left(\frac{(d+1)^3}{f(d)^d e^d}\right)$. \square

6.4 Variants of the Laplacian.

We finish this section by showing how our results imply upper and lower bounds on the spectral gap of several variants of the Laplacian.

6.4.1 Boundary Matrix.

The QTDA algorithm of McArdle, Gilyén, and Berta [50] is not parameterized by the spectral gap of the combinatorial Laplacian; rather, it is parameterized by the spectral gap of the boundary matrices. However, the non-zero singular values of the d^{th} boundary matrix ∂_d are the square roots of the eigenvalues of the $(d-1)^{\text{st}}$ up Laplacian L_{d-1}^{up} as $L_{d-1}^{\text{up}} = \partial_d \partial_d^T$. Therefore, Theorem 1.2 and Theorem 6.1 imply exponential upper and lower bounds on the spectral gap of the boundary matrix.

6.4.2 Normalized Laplacian.

We now show that the normalized up Laplacian \tilde{L}_d^{up} can also have exponentially-small spectral gap. While the eigenvalues of the unnormalized d^{th} up Laplacian L_d^{up} are in the range $[0, n_0]$, the eigenvalues of the normalized d^{th} up Laplacian are in the range $[0, d+2]$ [35, Theorem 3.2.i]. As the normalized up Laplacian has a constant upper bound on its eigenvalues, it is reasonable to suspect the normalized up Laplacian also has a constant lower bound on its eigenvalues. Theorem 6.5 shows this is not the case.

Corollary 6.5. *Let $d, n \geq 1$. There is a d -dimensional simplicial complex \mathcal{B}_d^n with $n_d \in \Theta(\text{poly}(d) \cdot n)$ d -simplices and a constant $c_d \geq 1$ that depends only on d such that the spectral gap the normalized up Laplacian is $\lambda_{\min}(\tilde{L}_d^{\text{up}}) \in O\left(\frac{1}{c_d^d}\right)$.*

Proof. This follows from Theorem 1.2 and Lemma 2.5. The statement follows as $d_{\min}[\mathcal{B}_d^n] = 1$. \square

6.4.3 Persistent Laplacian.

Recently, Wang, Nguyen, and Wei [67] introduced the **persistent Laplacian** of simplicial filtrations as a generalization of the combinatorial Laplacian. The spectral gap of the persistent Laplacian has since appeared as a parameter of quantum algorithms for computing persistent Betti number [33], so lower bounding it is also of interest to QTDA. Mémoli, Wan, and Wang [52, Theorem SM5.8] prove that persistent Laplacians preserve effective resistance of cycles; therefore, the bounds on the spectral gap of the combinatorial Laplacian also apply to the persistent Laplacian by Lemma 4.4.

7 Conclusion and Open Questions.

In this paper, we propose a new span-program-based quantum algorithm for computing Betti numbers. This algorithm is a novel approach to QTDA that is more similar to classical incremental algorithms for computing Betti numbers than previous QTDA algorithms. Unfortunately, we show that, in the worst case, the span-program based algorithm takes exponential time due to cycles with exponentially-large effective resistance or effective capacitance. However, as a corollary to exponentially-large effective resistance, we prove that the spectral gap of the combinatorial Laplacian can be exponentially small. This proves that all known QTDA algorithms also require exponential time in the worst case. Below we discuss some of the questions left open by our work.

Incremental Quantum Algorithm for Persistent Betti numbers. Our algorithm incrementally computes the Betti number of a simplicial complex. While the classical algorithm for computing *persistent* Betti numbers is incremental [69], our algorithm is unable to perform persistent pairing. In other words, our algorithm can identify when a homology class dies, but it cannot identify when that homology class was born. It is an open question whether our algorithm can be adapted to compute the persistent Betti numbers of a simplicial complex. There are quantum algorithms for computing persistent Betti numbers [33, 50], but these algorithms are not incremental.

Lower Bounds or Expectation of the Spectral Gap. Theorem 6.1 shows that the spectral gap of the combinatorial Laplacian can be exponentially small. However, it is an open question how common these sorts of worst-case complexes are. While there are exact or expected lower bounds for certain families of simplicial complexes [4, 23, 28, 44, 45, 46, 63, 68], it is still unknown what the expected spectral gap is, or if there are lower bounds on the spectral gap, for all simplicial complexes, or for families of simplicial complexes of interest like Vietoris-Rips complexes.

Cheeger Inequalities and Implications of Exponentially-Small Spectral Gaps. The existence of simplicial complexes with exponentially-small spectral gap implies that existing QTDA algorithms cannot exactly compute Betti numbers without running for an exponentially long time; however, they can solve the related problem of *Approximate Betti Number Estimation* [30] (ABNE) of counting the number of eigenvalues of the Laplacian smaller than a given threshold. However, it remains an open question how useful approximate Betti number estimation is in practice.

A potential interpretation for approximate Betti number estimation could come in the form of a higher-dimensional Cheeger inequality. The *Cheeger inequality* in graphs relates the smallest eigenvalue(s) of the graph Laplacian to a value called the *Cheeger constant* that measures the existence of (multi-way) sparse graph cuts [11, 43]. Intuitively, if the Hodge Theorem (Theorem 2.1) says that the graph Laplacian has more than one zero eigenvalue if and only if the graph is disconnected, then the Cheeger inequality says that it has small non-zero eigenvalues if and only if it is “almost” disconnected.

However, higher-dimensional generalizations of the Cheeger inequality remain elusive. Ideally (for the purpose of ABNE), a higher-dimensional Cheeger inequality would say something similar: a simplicial complex “almost has non-trivial d -homology” or “has a sparse cut” if and only if the d^{th} combinatorial Laplacian has small non-zero eigenvalues. One hurdle is that it is not clear how to generalize the notion of “sparse cut” to higher dimensions. While there have been several definitions proposed for a Cheeger constant for higher-dimensional Laplacians [25, 47, 53, 58], one

or both sides of a Cheeger inequality have failed for these constants [27, 28, 58, 64]. Our work provides another counterexample to these Cheeger inequalities; the spectral gap of our worst-case complexes is exponentially small, but the proposed notions of Cheeger constant cannot be.

A recent paper presents a two-sided Cheeger inequality [39] that connects the spectral gap of the combinatorial Laplacian to a Cheeger constant based on the 1-norm of chains. This Cheeger inequality does not have the same interpretation as the graph Cheeger inequality of implying that simplicial complexes with small eigenvalues “almost have non-trivial homology” though. It remains an open question if such a higher-dimensional Cheeger inequality exists.

Acknowledgements.

This work was supported by NSF grants CCF-1816442 and CCF-1617951

References

- [1] Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, Apr 2015.
- [2] Michał Adamaszek and Juraj Stacho. Complexity of simplicial homology and independence complexes of chordal graphs. *Computational Geometry*, 57:8–18, 2016.
- [3] Simon Apers, Sayantan Sen, and Dániel Szabó. A (simple) classical algorithm for estimating betti numbers. *arXiv preprint arXiv:2211.09618*, 2022.
- [4] Orr Beit-Aharon and Roy Meshulam. Spectral expansion of random sum complexes. *Journal of Topology and Analysis*, 12(04):989–1002, 2020.
- [5] Aleksandrs Belovs and Ben W. Reichardt. Span programs and quantum algorithms for st -connectivity and claw detection. In *Algorithms – ESA 2012*, pages 193–204. Springer Berlin Heidelberg, 2012.
- [6] Béla Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 1998.
- [7] Ulrich Brehm and Karanbir S Sarkaria. Linear vs. piecewise-linear embeddability of simplicial complexes. 1992.
- [8] Chris Cade and P. Marcos Crichigno. Complexity of supersymmetric systems and the cohomology problem, 2021.
- [9] Chris Cade, Ashley Montanaro, and Aleksandrs Belovs. Time and space efficient quantum algorithms for detecting cycles and testing bipartiteness. *Quantum Info. Comput.*, 18(1–2):18–50, February 2018.
- [10] Gunnar Carlsson and Mikael Vejdemo-Johansson. *Topological data analysis with applications*. Cambridge University Press, 2021.
- [11] Fan RK Chung. Laplacians of graphs and cheeger’s inequalities. *Combinatorics, Paul Erdos is Eighty*, 2(157-172):13–2, 1996.
- [12] Marcos Crichigno and Tamara Kohler. Clique homology is qma1-hard, 2022.

- [13] Cecil Jose A. Delfinado and Herbert Edelsbrunner. An incremental algorithm for betti numbers of simplicial complexes. In *Proceedings of the Ninth Annual Symposium on Computational Geometry*, SCG '93, page 232–239, New York, NY, USA, 1993. Association for Computing Machinery.
- [14] Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM J. Comput.*, 40(4):1026–1044, July 2011.
- [15] Tamal K. Dey, Tao Hou, and Sayan Mandal. Computing minimal persistent cycles: Polynomial and hard cases. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, page 2587–2606, USA, 2020. Society for Industrial and Applied Mathematics.
- [16] Tamal Krishna Dey and Yusu Wang. *Computational topology for data analysis*. Cambridge University Press, 2022.
- [17] Nathan M. Dunfield and Anil N. Hirani. The least spanning area of a knot and the optimal bounding chain problem. In *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry*, SoCG '11, page 135–144, New York, NY, USA, 2011. Association for Computing Machinery.
- [18] Christoph Dürr, Mark Heiligman, Peter HOyer, and Mehdi Mhalla. Quantum query complexity of some graph problems. *SIAM Journal on Computing*, 35(6):1310–1328, January 2006.
- [19] Beno Eckmann. Harmonische funktionen und randwertaufgaben in einem komplex. *Commentarii Mathematici Helvetici*, 17(1):240–255, Dec 1944.
- [20] Herbert Edelsbrunner and John Harer. *Computational Topology - an Introduction*. American Mathematical Society, 2010.
- [21] Marek Filakovský, Peter Franek, Uli Wagner, and Stephan Zhechev. Computing simplicial representatives of homotopy group elements. *Journal of Applied and Computational Topology*, 2(3):177–231, Dec 2018.
- [22] Michael Freedman and Vyacheslav Krushkal. Geometric complexity of embeddings in \mathbb{R}^d . *Geometric and Functional Analysis*, 24(5):1406–1430, Oct 2014.
- [23] J. Friedman. Computing betti numbers via combinatorial laplacians. *Algorithmica*, 21(4):331–346, Aug 1998.
- [24] Timothy Goldberg. Combinatorial laplacians of simplicial complexes. B.S. Thesis, Bard College, Annandale-on-Hudson, NY, 2002.
- [25] Mikhail Gromov. Singularities, expanders and topology of maps. part 2: from combinatorics to topology via algebraic isoperimetry. *Geometric and Functional Analysis*, 20(2):416–526, Aug 2010.
- [26] Anna Gundert. *On expansion and spectral properties of simplicial complexes*. PhD thesis, ETH Zurich, Zürich, 2013.
- [27] Anna Gundert and May Szedlák. Higher dimensional discrete cheeger inequalities. *J. Comput. Geom.*, 6(2):54–71, 2015.

- [28] Anna Gundert and Uli Wagner. On eigenvalues of random complexes. *Israel Journal of Mathematics*, 216:545–582, 2016.
- [29] Sam Gunn and Niels Kornerup. Review of a quantum algorithm for betti numbers, 2019.
- [30] Casper Gyurik, Chris Cade, and Vedran Dunjko. Towards quantum advantage via topological data analysis. *Quantum*, 6:855, November 2022.
- [31] Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4):315–358, Aug 2019.
- [32] Allen Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2000.
- [33] Ryu Hayakawa. Quantum algorithm for persistent Betti numbers and topological data analysis. *Quantum*, 6:873, December 2022.
- [34] Danijela Horak and Jürgen Jost. Interlacing inequalities for eigenvalues of discrete laplace operators. *Annals of Global Analysis and Geometry*, 43:177–207, 2013.
- [35] Danijela Horak and Jürgen Jost. Spectra of combinatorial laplace operators on simplicial complexes. *Advances in Mathematics*, 244:303–336, 2013.
- [36] Tsuyoshi Ito and Stacey Jeffery. Approximate span programs. *Algorithmica*, 81(6):2158–2195, November 2018.
- [37] Michael Jarret, Stacey Jeffery, Shelby Kimmel, and Alvaro Piedrafita. Quantum algorithms for connectivity and related problems. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:13, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [38] Stacey Jeffery and Shelby Kimmel. Quantum algorithms for graph connectivity and formula evaluation. *Quantum*, 1:26, August 2017.
- [39] Jürgen Jost and Dong Zhang. Cheeger inequalities on simplicial complexes, 2023.
- [40] M. Karchmer and A. Wigderson. On span programs. In *[1993] Proceedings of the Eighth Annual Structure in Complexity Theory Conference*. IEEE Comput. Soc. Press, 1993.
- [41] A. Yu. Kitaev. Quantum measurements and the Abelian Stabilizer Problem. *arXiv e-prints*, pages quant-ph/9511026, November 1995.
- [42] Woong Kook and Kang-Ju Lee. Simplicial networks and effective resistance. *Advances in Applied Mathematics*, 100:71–86, 2018.
- [43] James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):1–30, 2014.
- [44] Alan Lew. Spectral gaps, missing faces and minimal degrees. *Journal of Combinatorial Theory, Series A*, 169:105127, 2020.
- [45] Alan Lew. The spectral gaps of generalized flag complexes and a geometric hall-type theorem. *International Mathematics Research Notices*, 2020(11):3364–3395, 2020.

- [46] Alan Lew. Garland’s method for token graphs. *arXiv preprint arXiv:2305.02406*, 2023.
- [47] Nathan Linial and Roy Meshulam. Homological connectivity of random 2-complexes. *Combinatorica*, 26(4):475–487, aug 2006.
- [48] Seth Lloyd, Silvano Garnerone, and Paolo Zanardi. Quantum algorithms for topological and geometric analysis of data. *Nature Communications*, 7(1):10138, Jan 2016.
- [49] Frédéric Magniez, Ashwin Nayak, Jérémie Roland, and Miklos Santha. Search via quantum walk. *SIAM Journal on Computing*, 40(1):142–164, Jan 2011.
- [50] Sam McArdle, András Gilyén, and Mario Berta. A streamlined quantum algorithm for topological data analysis with exponentially fewer qubits, 2022.
- [51] RHAJ Meijer. Clustering using quantum persistent homology. M.S. Thesis, Utrecht University, Utrecht, NL, 2019.
- [52] Facundo Mémoli, Zhengchao Wan, and Yusu Wang. Persistent laplacians: Properties, algorithms and implications. *SIAM Journal on Mathematics of Data Science*, 4(2):858–884, 2022.
- [53] R. Meshulam and N. Wallach. Homological connectivity of random k-dimensional complexes. *Random Struct. Algorithms*, 34(3):408–417, may 2009.
- [54] James R. Munkres. *Elements of Algebraic Topology*. CRC Press, March 1984.
- [55] Niels Neumann and Sterre den Breeijen. Limitations of clustering using quantum persistent homology, 2019.
- [56] Andrew Newman. Small simplicial complexes with prescribed torsion in homology. *Discrete & Computational Geometry*, 62(2):433–460, March 2018.
- [57] Braxton Osting, Sourabh Palande, and Bei Wang. Spectral sparsification of simplicial complexes for clustering and label propagation. *Journal of Computational Geometry*, 11(1):176–211, 2020.
- [58] Ori Parzanchevski, Ron Rosenthal, and Ran J. Tessler. Isoperimetric inequalities in simplicial complexes. *Combinatorica*, 36(2):195–227, apr 2016.
- [59] Ben Reichardt and Robert Spalek. Span-program-based quantum algorithm for evaluating formulas. *Theory of Computing*, 8(1):291–319, 2012.
- [60] Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, October 2009.
- [61] Alexander Schmidhuber and Seth Lloyd. Complexity-theoretic limitations on quantum algorithms for topological data analysis, 2022.
- [62] Daniel Spielman. *Spectral and Algebraic Graph Theory*. 2019.
- [63] John Steenbergen, Caroline Klivans, and Sayan Mukherjee. A cheeger-type inequality on simplicial complexes. *Advances in Applied Mathematics*, 56:56–77, 2014.
- [64] John Steenbergen, Caroline Klivans, and Sayan Mukherjee. A cheeger-type inequality on simplicial complexes. *Advances in Applied Mathematics*, 56:56–77, 2014.

- [65] Mario Szegedy. Quantum speed-up of markov chain based algorithms. In *Proceedings of the Forty-Fifth Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 32–41, 11 2004.
- [66] Shashanka Ubaru, Ismail Yunus Akhalwaya, Mark S. Squillante, Kenneth L. Clarkson, and Lior Horesh. Quantum topological data analysis with linear depth and exponential speedup, 2021.
- [67] Rui Wang, Duc Duy Nguyen, and Guo-Wei Wei. Persistent spectral graph. *International journal for numerical methods in biomedical engineering*, 36(9):e3376, 2020.
- [68] Taiki Yamada. Spectrum of the laplacian on simplicial complexes by the ricci curvature. *arXiv preprint arXiv:1906.07404*, 2019.
- [69] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. In *Proceedings of the twentieth annual symposium on Computational geometry*, pages 347–356, 2004.

A Proof of Lemma 2.5

Lemma 2.5. *Let \mathcal{K} be a simplicial complex. Let d_{\min} and d_{\max} be the minimum and maximum degrees of any d -simplex in \mathcal{K} . Suppose that $d_{\min} > 0$. The normalized and unnormalized spectral gap are related as follows:*

$$\frac{1}{d_{\max}} \lambda_{\min}(L_d^{up}) \leq \lambda_{\min}(\tilde{L}_d^{up}) \leq \frac{1}{d_{\min}} \lambda_{\min}(L_d^{up})$$

Proof. Following the same steps as in the proof of Lemma 4.4, we can see that $\lambda_{\min}(\tilde{L}_d^{up}) = \lambda_{\max}^{-1}((\tilde{L}_d^{up})^+)$. Therefore, we will prove the follow equivalent statement:

$$d_{\min} \lambda_{\max}((L_d^{up})^+) \leq \lambda_{\max}((\tilde{L}_d^{up})^+) \leq d_{\max} \lambda_{\max}((L_d^{up})^+)$$

Also from the proof of Lemma 4.4, we see that $\lambda_{\max}((\tilde{L}_d^{up})^+) = \max\{x^T (\tilde{L}_d^{up})^+ x : \|x\| = 1, x \in \text{im } D^{-1/2} \partial_{d+1}\}$. For a vector $x \in \text{im } \tilde{L}_d^{up}$, following the steps of Lemma 4.2, we conclude that

$$x^T (\tilde{L}_d^{up})^+ x = \min\{\|y\|^2 : D^{-1/2} \partial_{d+1} y = x\}$$

As $D^{1/2}$ is a bijection, then we further conclude that

$$\begin{aligned} x^T (\tilde{L}_d^{up})^+ x &= \min\{\|y\|^2 : D^{-1/2} \partial_{d+1} y = x\} \\ &= \min\{\|y\|^2 : \partial_{d+1} y = D^{1/2} x\} \\ &= R_{D^{1/2} x}(\mathcal{K}) \end{aligned}$$

We have the bound $R_{D^{1/2} x}(\mathcal{K}) = \|D^{1/2} x\|^2 \mathcal{R}_{\bar{x}}(\mathcal{K}) \geq d_{\min} \mathcal{R}_{\bar{x}}(\mathcal{K})$. As the chain $D^{1/2} x \in \text{im } \partial_{d+1}$ if and only if $x \in \text{im } D^{-1/2} \partial_{d+1}$, we have the general bound:

$$\begin{aligned} \lambda_{\max}((\tilde{L}_d^{up})^+) &= \max\{x^T (\tilde{L}_d^{up})^+ x : \|x\| = 1, x \in \text{im } D^{-1/2} \partial_{d+1}\} \\ &= \max\{R_{D^{1/2} x}(\mathcal{K}) : \|x\| = 1, x \in \text{im } D^{-1/2} \partial_{d+1}\} \\ &\geq d_{\min} \max\{R_x(\mathcal{K}) : \|\bar{x}\| = 1, \bar{x} \in \text{im } \partial_{d+1}\} \\ &= d_{\min} \lambda_{\max}((L_d^{up})^+) \end{aligned} \tag{Lemma 4.4}$$

To prove $\lambda_{\max}((\tilde{L}_d^{up})^+) \leq d_{\max} \lambda_{\max}(L_d^{up})$, we can similarly prove that

$$\lambda_{\max}((L_d^{up})^+) = \{(D^{-1/2} x)^T (\tilde{L}_d^{up})^+ (D^{-1/2} x) : \|x\| = 1, x \in \text{im } \partial_{d+1}\},$$

and the rest of the proof follows similarly. \square

B Properties of Effective Resistance: Parallel, Series, and Monotonicity Formulas.

We now prove there are formulas for effective resistance in simplicial complexes analogous to the series and parallel formulas for effective resistance in graphs. These formulas not only are useful for calculating effective resistance, but they also provide intuition for effective resistance. In particular, they provide justification for the claim that effective resistance measures “how null-homologous” a cycle is in a complex.

Theorem B.1 (Series Formula). *Let \mathcal{K}_1 and \mathcal{K}_2 be simplicial complexes with $\gamma \in C_{d-1}(\mathcal{K}_1) \cap C_{d-1}(\mathcal{K}_2)$, $C_d(\mathcal{K}_1) \cap C_d(\mathcal{K}_2) = \emptyset$, and γ null-homologous in \mathcal{K}_1 and \mathcal{K}_2 . Let $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2$. Then*

$$\mathcal{R}_\gamma(\mathcal{K}) \leq \mathcal{R}_{\gamma_1}(\mathcal{K}_1) + \mathcal{R}_{\gamma_2}(\mathcal{K}_2).$$

Equality is achieved when γ_1 and γ_2 are the unique chains in $C_d(\mathcal{K}_1)$ and $C_d(\mathcal{K}_2)$ that sum to γ .

Proof. Let γ_1 and γ_2 be null-homologous cycles in \mathcal{K}_1 and \mathcal{K}_2 respectively that sum to γ , and let f_1 and f_2 be the minimum-energy unit γ_1 - and γ_2 -flows, respectively. Then $f = f_1 + f_2$ is a unit γ -flow, and we can bound $\mathcal{R}_\gamma(\mathcal{K}) \leq J(f) = J(f_1) + J(f_2) = \mathcal{R}_{\gamma_1}(\mathcal{K}_1) + \mathcal{R}_{\gamma_2}(\mathcal{K}_2)$; the equality $J(f) = J(f_1) + J(f_2)$ follows from the fact that \mathcal{K}_1 and \mathcal{K}_2 have disjoint sets of d -simplices.

To prove the other direction, observe that γ can always be written as the sum of two null-homologous chains $\gamma_1 \in C_{d-1}(\mathcal{K}_1)$ and $\gamma_2 \in C_{d-1}(\mathcal{K}_2)$. Any unit γ -flow g defines null-homologous $(d-1)$ -cycles γ_1 and γ_2 that sum to γ ; namely, if g_1 and g_2 are the restriction of g to \mathcal{K}_1 and \mathcal{K}_2 respectively, then $\gamma_1 = \partial g_1$ and $\gamma_2 = \partial g_2$.

If γ can be uniquely decomposed as $\gamma = \gamma_1 + \gamma_2$, then any unit γ -flow f can be decomposed as a unit γ_1 -flow f_1 and a unit γ_2 -flow f_2 . It follows that the energy of f is minimized when the energy of f_1 and f_2 are both minimized. Hence, $\mathcal{R}_\gamma(\mathcal{K}) = \mathcal{R}_{\gamma_1}(\mathcal{K}_1) + \mathcal{R}_{\gamma_2}(\mathcal{K}_2)$. \square

Theorem B.2 (Parallel Formula). *Let \mathcal{K}_1 and \mathcal{K}_2 be simplicial complexes with $\gamma \in C_{d-1}(\mathcal{K}_1) \cap C_{d-1}(\mathcal{K}_2)$, $C_d(\mathcal{K}_1) \cap C_d(\mathcal{K}_2) = \emptyset$, and γ null-homologous in \mathcal{K}_1 and \mathcal{K}_2 . Let $\mathcal{K} = \mathcal{K}_1 \cup \mathcal{K}_2$. Then*

$$\mathcal{R}_\gamma(\mathcal{K}) \leq \left(\frac{1}{\mathcal{R}_\gamma(\mathcal{K}_1)} + \frac{1}{\mathcal{R}_\gamma(\mathcal{K}_2)} \right)^{-1}$$

Equality is achieved when $\text{im } \partial_d[\mathcal{K}_1] \cap \text{im } \partial_d[\mathcal{K}_2] = \text{span}\{\gamma\}$.

Proof. Let f_1 and f_2 be the minimum energy unit γ -flows in \mathcal{K}_1 and \mathcal{K}_2 resp. For any $t \in \mathbb{R}$, the chain $g_t = t f_1 + (1-t) f_2$ is a unit γ -flow in \mathcal{K} . We can therefore bound the effective resistance over the minimum of these combinations as $\mathcal{R}_\gamma(\mathcal{K}) \leq \min_{t \in \mathbb{R}} J(g_t)$.

To get the tightest bound of $\mathcal{R}_\gamma(\mathcal{K})$, we now derive $t_{\text{opt}} := \arg \min_{t \in \mathbb{R}} J(g_t)$. Observe that $J(g_t) = t^2 J(f_1) + (1-t)^2 J(f_2) = t^2 \mathcal{R}_\gamma(\mathcal{K}_1) + (1-t)^2 \mathcal{R}_\gamma(\mathcal{K}_2)$; this follows from the fact that \mathcal{K}_1 and \mathcal{K}_2 have disjoint sets of d -simplices. The quantity $J(g_t)$ is a positive quadratic with respect to t , so t_{opt} is the value of t where the derivative of $J(g_t)$ is 0. Taking the derivative, we find that



Figure 5: Left: The cycle γ is the blue boundary with ± 1 coefficients, and the unique unit γ -flow is the 6 triangles in series. If the complex is unweighted, then the effective resistance of γ is 6. Right: The cycle γ is the equator of the sphere, and the two hemispheres are two unit γ -flows in parallel. If each hemisphere has flow energy c , then the effective resistance of γ is $\frac{c}{2}$.

$t_{\text{opt}} = \mathcal{R}_\gamma(\mathcal{K}_2)/(\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2))$. Plugging t_{opt} into $J(g_{t_{\text{opt}}})$, we find that

$$\begin{aligned}
J(g_{t_{\text{opt}}}) &= \left(\frac{\mathcal{R}_\gamma(\mathcal{K}_2)}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \right)^2 \mathcal{R}_\gamma(\mathcal{K}_1) + \left(1 - \frac{\mathcal{R}_\gamma(\mathcal{K}_2)}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \right)^2 \mathcal{R}_\gamma(\mathcal{K}_2) \\
&= \left(\frac{\mathcal{R}_\gamma(\mathcal{K}_2)}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \right)^2 \mathcal{R}_\gamma(\mathcal{K}_1) + \left(\frac{\mathcal{R}_\gamma(\mathcal{K}_1)}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \right)^2 \mathcal{R}_\gamma(\mathcal{K}_2) \\
&= \left(\frac{1}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \right)^2 (\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)) \mathcal{R}_\gamma(\mathcal{K}_1) \mathcal{R}_\gamma(\mathcal{K}_2) \\
&= \frac{\mathcal{R}_\gamma(\mathcal{K}_1) \mathcal{R}_\gamma(\mathcal{K}_2)}{\mathcal{R}_\gamma(\mathcal{K}_1) + \mathcal{R}_\gamma(\mathcal{K}_2)} \\
&= \left(\frac{1}{\mathcal{R}_\gamma(\mathcal{K}_1)} + \frac{1}{\mathcal{R}_\gamma(\mathcal{K}_2)} \right)^{-1}
\end{aligned}$$

This implies the upper bound on $\mathcal{R}_\gamma(\mathcal{K})$ in the theorem statement. To get the lower bound, observe that any unit γ -flow g in \mathcal{K} can be orthogonally decomposed into chains $g_1 \in C_d(\mathcal{K}_1)$ and $g_2 \in C_d(\mathcal{K}_2)$, so $\partial g_1 + \partial g_2 = \gamma$. We claim that $\partial g_1 = t\gamma$ and $\partial g_2 = (1-t)\gamma$ for some value of t ; if not, then $\partial g_1 = t\gamma + \eta$ and $\partial g_2 = (1-t)\gamma - \eta$ for some non-zero chain $\eta \notin \text{span}\{\gamma\}$, which cannot be the case as $\text{im } \partial_{\mathcal{K}_1} \cap \text{im } \partial_{\mathcal{K}_2} = \text{span}\{\gamma\}$. This proves the chain g is a linear combination of a unit γ -flow in \mathcal{K}_1 and a unit γ -flow in \mathcal{K}_2 . The chain g_t is the lowest energy such linear combination. \square

Figure 5 shows examples of unit γ -flows in series and parallel. These formulas justify the claim that the effective resistance of a null-homologous cycle γ is a measure of *how* null-homologous γ is. The more chains with boundary γ , the smaller the effective resistance of γ by the parallel formula. The smaller the chains bounding γ , the lower the effective resistance by the series formula.

Another important property of effective resistance in graphs is **Rayleigh monotonicity**. Rayleigh monotonicity says that adding edges to the graph can only decrease the effective resistance between any pair of vertices; this reinforces the notion that effective resistance measures how well-connected a pair of vertices are, as adding an edge can only make a pair of vertices better connected. We prove a similar result for simplicial complexes.

Theorem B.3 (Rayleigh Monotonicity). *Let $\mathcal{L} \subset \mathcal{K}$ be simplicial complexes. Let $\gamma \in C_{d-1}(\mathcal{K}) \cap C_{d-1}(\mathcal{L})$ be a cycle that is null-homologous in both complexes. Then $\mathcal{R}_\gamma(\mathcal{K}) \leq \mathcal{R}_\gamma(\mathcal{L})$.*

Proof. As $C_d(\mathcal{L}) \subset C_d(\mathcal{K})$, then any unit γ -flow in \mathcal{L} is also a unit γ -flow in \mathcal{K} . As the effective resistance is the minimum energy of a unit γ -flow, then clearly $\mathcal{R}_\gamma(\mathcal{K}) \leq \mathcal{R}_\gamma(\mathcal{L})$. \square

C Duality of Resistance and Capacitance in Embedded Complexes.

In this section, we consider the effective capacitance of a cycle in the special case when \mathcal{K} is a d -dimensional simplicial complex with a given embedding into \mathbb{R}^{d+1} . In this case, \mathcal{K} is called an **embedded complex**. Embedded complexes serve as a high-dimensional generalization of planar graphs and naturally admit a dual graph. We will show that the effective capacitance of certain $(d-1)$ -cycles γ in \mathcal{K} are equal to the effective resistance between a pair of vertices in the dual graph that are “dual” to γ . This theorem generalizes the analysis of capacitance in planar graphs given by Jeffery and Kimmel [38]. Hence, we can parameterize the quantum algorithm deciding if γ is null-homologous (Theorem 4.9) in terms of the effective resistance of γ in \mathcal{K} and the effective resistance between the pair of vertices in the dual graph. Specifically, we will generalize the special case of planar graphs for which the vertices s and t appear on the boundary of the same face. Throughout this section we assume we are given the embedding as input. Computing the dual graph from an embedding can be done in polynomial time [15].

C.1 Duality in Embedded Complexes.

The **Alexander Duality theorem** [32, Corollary 3.45] states that for a d -dimensional simplicial complex \mathcal{K} with an embedding into \mathbb{R}^{d+1} , the complement $\mathbb{R}^{d+1} \setminus \mathcal{K}$ consists of $\beta_d[\mathcal{K}] + 1$ connected components. We call these connected components **voids**. Exactly one of these voids is unbounded. We denote the bounded voids as V_i for $1 \leq i \leq \beta_d$ and the unbounded void as V_∞ . Moreover, the boundaries of the bounded voids generate the homology group $H_d(\mathcal{K})$. The embedding implies that each d -simplex is contained on the boundary of at most two voids, and we make the assumption that the d -simplices are oriented consistently with respect to the voids. That is, if a d -simplex is on the boundary of two voids it is oriented positively on one void, and negatively on the other. We have a boundary matrix ∂_{d+1} whose columns are the voids and whose rows are the d -simplices. From the embedding and the consistent orientation we see that ∂_{d+1} is the edge-vertex incident matrix of the **directed dual graph**: the directed graph whose vertices are in bijection with the voids and whose edges are in bijection with the d -simplices of \mathcal{K} . The direction of the edges are inherited from the orientations of the d -simplices. For a d -simplex σ on the boundary of voids V_1 and V_2 we denote the dual edge by $\sigma^* = \{v_1^*, v_2^*\}$ and we define the dual weight function by $w^*(\sigma^*) = 1/w(\sigma)$.

We construct an additional chain group $C_{d+1}(\mathcal{K})$ with the bounded voids as basis elements. This is a purely algebraic construction, meaning basis elements of $C_{d+1}(\mathcal{K})$ do *not* correspond to $(d+1)$ -simplices. We can also define a boundary map $\partial_{d+1} : C_{d+1}(\mathcal{K}) \rightarrow C_d(\mathcal{K})$, where ∂V_i is the boundary of the void V_i as described above. This gives rise to a new chain complex

$$\cdots 0 \rightarrow C_{d+1}(\mathcal{K}) \xrightarrow{\partial_{d+1}} C_d(\mathcal{K}) \xrightarrow{\partial_d} \cdots \xrightarrow{\partial_1} C_0(\mathcal{K}).$$

Since the boundaries of the voids generate the d^{th} homology group of \mathcal{K} and $C_{d+1}(\mathcal{K})$ is generated by these voids we obtain a valid chain complex. Moreover, we have that $\dim H_d(\mathcal{K}) = 0$ in our new chain complex.

In addition to a dual graph, we define the **dual complex** of \mathcal{K} , denoted \mathcal{K}^* , by defining the dual chain groups $C_k(\mathcal{K}^*)$ via the isomorphism $C_{d-k+1}(\mathcal{K}) \cong C_k(\mathcal{K})$. Importantly, note that only $C_0(\mathcal{K}^*)$ and $C_1(\mathcal{K}^*)$ correspond to the chain groups of a simplicial complex (the dual graph),

while the higher chain groups $C_i(\mathcal{K})$ for $i \geq 2$ are a purely algebraic construction. Moreover, we define the dual boundary operator $\partial_k^*: C_k(\mathcal{K}^*) \rightarrow C_{k-1}(\mathcal{K}^*)$ to be the coboundary operator $\delta_{d-k+1}: C_{d-k+1}(\mathcal{K}) \rightarrow C_{d-k}(\mathcal{K})$ of \mathcal{K} , and the dual coboundary operator $\delta_k^*: C_{k-1}(\mathcal{K}^*) \rightarrow C_k(\mathcal{K}^*)$ to be the boundary operator $\partial_{d-k+1}: C_{d-k+1}(\mathcal{K}) \rightarrow C_{d-k}(\mathcal{K})$ of \mathcal{K} . In other words the (co)boundary operators commute with the duality isomorphism. We summarize the construction in the following commutative diagram.

$$\begin{array}{ccccccc}
C_{d+1}(\mathcal{K}) & \xleftarrow{\frac{\partial_{d+1}}{\delta_d}} & C_d(\mathcal{K}) & \xleftarrow{\frac{\partial_d}{\delta_{d-1}}} & \cdots & \xleftarrow{\frac{\partial_1}{\delta_1}} & C_0(\mathcal{K}) \\
\uparrow \cong & & \uparrow \cong & & & & \uparrow \cong \\
C_0(\mathcal{K}^*) & \xleftarrow{\frac{\delta_0^*}{\partial_1^*}} & C_1(\mathcal{K}^*) & \xleftarrow{\frac{\delta_1^*}{\partial_2^*}} & \cdots & \xleftarrow{\frac{\delta_d^*}{\partial_{d+1}^*}} & C_{d+1}(\mathcal{K}^*)
\end{array}$$

We need to make one additional assumption on the location of the input $(d-1)$ -dimensional cycle γ which makes our setup a generalization of a planar graph with two vertices s and t appearing on the same face. We assume that there exists a void V_i with two unit γ -flows Γ_1 and Γ_2 such that $\text{supp}(\Gamma_1) \cap \text{supp}(\Gamma_2) = \emptyset$ and $\text{supp}(\Gamma_1) \cup \text{supp}(\Gamma_2) = \text{supp}(\partial_{d+1}V_i)$. That is, there exist two unit γ -flows whose supports partition the boundary of the void V_i . For example, the support of the cycle γ could be the equator of a sphere, and the support of the cycles Γ_1 and Γ_2 could be the north and south hemispheres. This generalizes the fact in planar graphs that when s and t are on the same face we can find two st -paths which partition the boundary of the face. In planar graphs we are guaranteed to find two such paths, however for an arbitrary $(d-1)$ -cycle γ we are not guaranteed to find two unit γ -flows partitioning the boundary of some void. More specifically, we take Γ_2 to be a unit $(-\gamma)$ -flow so that $\partial_d\Gamma_2 = -\gamma$. In the planar graph analogy this is equivalent as viewing Γ_1 as a path from s to t and viewing Γ_2 as a path from t to s . We add an additional basis element Σ to $C_d(\mathcal{K})$ such that $\partial_d\Sigma = -\gamma$. In planar graphs this is equivalent to adding an edge directed from t to s . The addition of this edge splits the face containing s and t into two. In higher dimensions, the geometry of adding another d -simplex to fill γ is more complicated, but the addition of Σ to $C_d(\mathcal{K})$ allows us to perform a purely algebraic operation to our chain complex that behaves as if V_i has been split into two; namely, we remove V_i from $C_{d+1}(\mathcal{K})$ and replace it with two new basis elements V_s and V_t . Next, we extend the boundary operator to V_s and V_t in the following way: $\partial_{d+1}V_s = \Gamma_1 - \Sigma$ and $\partial_{d+1}V_t = \Gamma_2 + \Sigma$. In the dual complex, the vertices dual to V_s and V_t are denoted s^* and t^* , and the edge dual to Σ is denoted $\Sigma^* = \{t^*, s^*\}$.

C.2 Effective capacitance is Dual to Effective Resistance.

In this next section, we will show that the effective capacitance of γ in a subcomplex $\mathcal{L} \subset \mathcal{K}$ is equal to the effective resistance between s^* and t^* in the dual complex $\mathcal{L}^* \subset \mathcal{K}^*$, where \mathcal{L}^* is defined as being the subgraph of \mathcal{K}^* with all of the vertices of \mathcal{K}^* but only the edges dual to the d -simplices **not** in \mathcal{L} .

The effective resistance between s^* and t^* in \mathcal{L}^* is determined by the unit s^*t^* -flows in \mathcal{L}^* . However, it will be convenient to work with *circulations* instead of flows. A unit s^*t^* -circulation f is a 1-cycle such that $f(\Sigma^*) = 1$. Recall that Σ^* is the edge directed from t^* to s^* , so a unit s^*t^* -circulation is just a unit s^*t^* -flow with additional flow on the edge Σ^* to the cycle. Clearly, there is a bijection between unit s^*t^* -flows and unit s^*t^* -circulations. We define the flow energy of a circulation to be equal to the flow energy of its corresponding flow.

Theorem C.1. *Let \mathcal{K} be a d -dimensional simplicial complex embedded into \mathbb{R}^{d+1} , and let $\mathcal{L} \subset \mathcal{K}$ be a subcomplex. Let $\gamma \in C_{d-1}(\mathcal{K})$ be a $(d-1)$ -cycle such that there exist two unit γ -flows $\Gamma_1, \Gamma_2 \in$*

$C_d(\mathcal{K})$ with supports that partition the boundary of some void V_i . Then the effective capacitance $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{K})$ is equal to the effective resistance $\mathcal{R}_{s^*t^*}(\mathcal{L}^*)$.

Proof. Let p be a unit γ -potential in \mathcal{L} , and define $f \in C_1(\mathcal{L}^*)$ to be the image of $\delta_{d-1}p$ under the duality isomorphism; that is, $f = \partial_2^*p^*$. The 1-chain f is a circulation in the 1-skeleton of \mathcal{L}^* . Further, since $\Sigma^* = \{t^*, s^*\}$ the circulation f corresponds to a unit s^*t^* -flow by the following calculation:

$$f^T \Sigma^* = p^T \partial_d \Sigma = p^T \gamma = 1.$$

Next, we calculate the flow energy of f and show it is equal to the potential energy of p .

$$\begin{aligned} J(f) &= \sum_{e^* \in \mathcal{L}_1^*} \frac{f(e^*)^2}{w^*(e^*)} \\ &= \sum_{e^* \in \mathcal{L}_1^*} f(\sigma^*)^2 w(e) \\ &= \sum_{\sigma \in \mathcal{K}_d \setminus \mathcal{L}_d} \delta p(\sigma)^2 w(\sigma) \\ &= \mathcal{J}(p) \end{aligned}$$

Conversely, let f^* be the minimal-energy unit s^*t^* -circulation in \mathcal{L}^* . By the assumptions outlined in the previous section, we have $\dim H_d(\mathcal{K}) = 0$, which in turn gives us $\dim H_1(\mathcal{K}^*) = 0$. Hence, $f^* \in \text{im } \partial_2^*$. Let $p^* \in C_2(\mathcal{L}^*)$ with $\partial_2^*p^* = f^*$; we will show that p is the unit γ -potential in \mathcal{L} in bijection with f^* . To see that p is a unit γ -potential we compute its value on γ :

$$p^T \gamma = p^T \partial_d \Sigma = (\delta_{d-1}p)^T \Sigma = (\partial_2^*p^*)^T \Sigma^* = (f^*)^T \Sigma^* = 1.$$

Moreover, as $\partial_2^*p^* = f^*$ and f^* has its support on \mathcal{L}_1^* (which is defined to only have edges dual to the d -simplices in $\mathcal{K} \setminus \mathcal{L}$), then $\partial_{d-1}[\mathcal{L}]p = 0$, so p is a unit γ -potential in \mathcal{L} .

It remains to show that the potential energy of p is equal to the flow energy of f^* . We have the following calculation:

$$\begin{aligned} \mathcal{J}(p) &= \sum_{\sigma \in \mathcal{K}_d \setminus \mathcal{L}_d} \delta_{d-1}p(\sigma)^2 w(\sigma) \\ &= \sum_{\sigma \in \mathcal{K}_d \setminus \mathcal{L}_d} \frac{\delta_{d-1}p(\sigma)^2}{w^*(\sigma^*)} \\ &= \sum_{\sigma^* \in \mathcal{L}_1^*} \frac{\partial_2^*p^*(\sigma^*)^2}{w^*(\sigma^*)} \\ &= \sum_{\sigma^* \in \mathcal{L}_1^*} \frac{f^*(\sigma^*)^2}{w^*(\sigma^*)} \\ &= J(f^*). \end{aligned}$$

□

D Evaluating the span program for null-homology testing.

In this section, we give a quantum algorithm for evaluating the null-homology span program. Our algorithm is inspired by and generalizes the quantum algorithm for evaluating st -connectivity

span program in graphs. The first quantum algorithm for evaluating the st -connectivity span program was given by Belovsz and Reichardt [5]; however, we follow the slightly different algorithm introduced by Ito and Jeffery [36]. We are also greatly indebted to the presentation of this algorithm given by Jeffery and Kimmel [38], from which our algorithm is adapted.

The algorithm for evaluating a general span program $\mathcal{P} = (\mathcal{H}, \mathcal{U}, |\tau\rangle, A)$ is to perform phase estimation of the vector $|w_0\rangle := A^+|\tau\rangle$ on the unitary operator $U = R_{\mathcal{H}(x)}R_{\ker A}$ where the notation R_S denotes the reflection about the subspace S . (The unitary $R_S = 2\Pi_S - I$, where Π_S is the projection onto S .) Intuitively, if x is a positive instance, then $|w_0\rangle$ will be close to an eigenvector of U with phase 0. If x is a negative instance, then $|w_0\rangle$ will be far from any eigenvector of U of phase 0. If we want to evaluate the function $f : D \rightarrow \{0, 1\}$, we need to perform phase estimation to precision $O\left(1/\sqrt{W_-(f, \mathcal{P})W_+(f, \mathcal{P})}\right)$. The algorithm for phase estimation of a unitary U to precision $O(\delta)$ performs $O(1/\delta)$ implementations of the unitary U [41], so the algorithm for evaluating the span program $\mathcal{P} = (\mathcal{H}, \mathcal{U}, |\tau\rangle, A)$ requires $O\left(\sqrt{W_-(f, \mathcal{P})W_+(f, \mathcal{P})}\right)$ implementations of U . We now analyze the time complexity of implementing the unitary U .

The reflection $R_{\mathcal{H}(x)}$ can be implemented with two queries to \mathcal{O}_x . This reflection is the same as the reflection across the good states in Grover's Algorithm. The rest of this section is devoted to an implementation of $R_{\ker \partial}$.

Recall that $\ker \partial_d \subset C_d(\mathcal{K})$. The idea behind the implementation of $R_{\ker \partial}$ is that instead of reflecting across $\ker \partial_d$ directly, we can embed $C_d(\mathcal{K})$ into $C_{d-1}(\mathcal{K}) \otimes C_d(\mathcal{K})$ by sending $|\tau\rangle \rightarrow c|\partial\tau\rangle|\tau\rangle$ (where c is a normalization constant). We can then implement the reflection $R_{\ker \partial}$ by implementing a series of "local reflections" on the basis $|\partial\tau\rangle|\tau\rangle$.

We consider two subspaces B and C of $C_{d-1}(\mathcal{K}) \otimes C_d(\mathcal{K})$. The spaces B and C are defined:

$$B = \text{span} \left\{ |b_\tau\rangle := \frac{1}{\sqrt{d+1}}|\partial\tau\rangle|\tau\rangle : \tau \in \mathcal{K}_d \right\}$$

and

$$C = \text{span} \left\{ |c_\sigma\rangle := \sum_{\sigma \subset \tau} \sqrt{\frac{w(\sigma)}{\deg(\sigma)}}|\sigma\rangle|\tau\rangle : \sigma \in \mathcal{K}_{d-1} \right\}.$$

The space $C_{d-1}(\mathcal{K}) \otimes C_d(\mathcal{K})$ has basis $\{|\sigma\rangle|\tau\rangle \mid \sigma \in \mathcal{K}_{d-1}, \tau \in \mathcal{K}_d\}$. The vector $|b_\tau\rangle$ is non-zero on a basis element $|\sigma\rangle|\tau\rangle$ if and only if σ is in the support of the boundary of τ . Similarly, a component of $|c_\sigma\rangle$ is non-zero on $|\sigma\rangle|\tau\rangle$ if and only if τ is in the support of the coboundary of σ . The vector $|b_\tau\rangle$ can be thought of as being *like* the boundary of τ , with the additional property that the set $\{|b_\tau\rangle \mid \tau \in C_d(\mathcal{K})\}$ is orthonormal. Similarly, the vector $|c_\sigma\rangle$ is *like* the coboundary of σ but orthonormal.

We also define operators that embed $C_d(\mathcal{K})$ and $C_{d-1}(\mathcal{K})$ into B and C respectively. We define linear operators $M_B : C_d(\mathcal{K}) \rightarrow B$ and $M_C : C_{d-1}(\mathcal{K}) \rightarrow C$ as follows:

$$M_B := \sum_{\tau \in \mathcal{K}_d} |b_\tau\rangle\langle\tau|,$$

and

$$M_C := \sum_{\sigma \in \mathcal{K}_{d-1}} |c_\sigma\rangle\langle\sigma|.$$

As the columns of M_B and M_C are orthonormal, both operators are isometries.

We introduce the matrices M_C and M_B as they have the property that $\ker M_C^\dagger M_B = \ker \partial$, which we prove in the follow lemma. This fact will give us a way to implement $R_{\ker \partial}$.

Lemma D.1. $\ker M_C^\dagger M_B = \ker \partial$.

Proof. We first calculate the matrix $M_C^\dagger M_B$. We then argue that $\ker M_C^\dagger M_B = \ker \partial$. For a $(d-1)$ -simplex σ and a d -simplex τ , we have that

$$\langle c_\sigma | b_\tau \rangle = \sum_{\tau' \in \mathcal{K}_d: \sigma \subset \tau'} \frac{w(\tau)}{\sqrt{\deg(\sigma)}} \langle \sigma | \partial \tau \rangle \langle \tau' | \tau \rangle = \begin{cases} \sqrt{\frac{w(\tau)}{(d+1)\deg(\sigma)}} \langle \sigma | \partial \tau \rangle & \text{if } \sigma \subset \tau \\ 0 & \text{otherwise} \end{cases}.$$

So $\langle c_\sigma | b_\tau \rangle$ is non-zero if and only if σ is in the boundary of τ . We use this to calculate the product $M_C^\dagger M_B$:

$$\begin{aligned} M_C^\dagger M_B &= \sum_{\sigma \in \mathcal{K}_{d-1}} \sum_{\tau \in \mathcal{K}_d} |\sigma\rangle \langle c_\sigma | b_\tau \rangle \langle \tau| \\ &= \frac{1}{\sqrt{(d+1)}} \sum_{\sigma \subset \tau} \frac{\sqrt{w(\tau)} \langle \sigma | \partial \tau \rangle}{\sqrt{\deg(\sigma)}} |\sigma\rangle \langle \tau| \\ &= \frac{1}{\sqrt{(d+1)}} \left(\sum_{\sigma \in \mathcal{K}_{d-1}(\mathcal{K})} \frac{|\sigma\rangle \langle \sigma|}{\sqrt{\deg(\sigma)}} \right) \sum_{\tau \in \mathcal{K}_d} \sqrt{w(\tau)} |\partial \tau\rangle \langle \tau| \\ &= \frac{1}{\sqrt{(d+1)}} \left(\sum_{\sigma \in \mathcal{K}_{d-1}(\mathcal{K})} \frac{|\sigma\rangle \langle \sigma|}{\sqrt{\deg(\sigma)}} \right) \partial \sqrt{W} =: \hat{\partial}. \end{aligned}$$

The term $\frac{|\sigma\rangle \langle \sigma|}{\sqrt{\deg(\sigma)}}$ is all-zeros matrix except for the (σ, σ) -entry, which is $\frac{1}{\sqrt{\deg(\sigma)}}$. The sum $\sum_{\sigma \in \mathcal{K}_{d-1}} \frac{|\sigma\rangle \langle \sigma|}{\sqrt{\deg(\sigma)}}$ is a diagonal matrix. Accordingly, the matrix $\hat{\partial}$ is $\partial \sqrt{W}$ with each row scaled.

Scaling the rows of a matrix does not change its row space or kernel, so $\ker M_C^\dagger M_B = \ker \partial$. \square

The spaces B and C and the matrices M_B and M_C are inspired by the follow lemma of Szegedy which is necessary for implementing $R_{\ker \partial}$.

Lemma D.2 (Szegedy [65], Theorem 1). *Let M_B and M_C be matrices with the same number of rows and orthonormal columns, and let $B = \text{span } M_B$ and $C = \text{span } M_C$. The matrix $M_C^\dagger M_B$ has singular values at most 1. Let $\cos \theta_1, \dots, \cos \theta_k$ be the singular values of $M_C^\dagger M_B$ in the range $(0, 1)$. Let $U = R_C R_B$. We can decompose the eigenspaces of U as*

- The $(+1)$ -eigenspace of U is $(B \cap C) \oplus (B^\perp \cap C^\perp)$.
- The (-1) -eigenspace of U is $(B \cap C^\perp) \oplus (B^\perp \cap C)$.
- The remaining eigenvalues of U are $e^{\pm 2i\theta_j}$ for $1 \leq j \leq k$.

The following lemma gives us a way to implment the $R_{\ker \partial}$. Let R_{U-} be the rotation about (-1) -eigenspace of U , and let $V = M_B^\dagger R_{U-} M_B$. The matrix V embeds $C_d(\mathcal{K})$ into B with M_B , performs a reflection on B about the (-1) -eigenspace of U , and unembeds with M_B^\dagger . The following lemma proves that $V = R_{\ker \partial}$.

Lemma D.3. *The matrix $V = M_B^\dagger R_{U-} M_B$ satisfies the equality $V = R_{\ker \partial}$.*

Proof. We first verify that V is a reflection; that is, we show the eigenvalues of V are 1 and -1. The matrices M_B and M_C have orthonormal columns, so we can use Lemma D.2 to characterize the eigenspaces of U . The (-1) -eigenspace of U is $(B \cap C^\perp) \oplus (B^\perp \cap C)$ and the $(+1)$ -eigenspace of U is $(B \cap C) \cap (B^\perp \cap C^\perp)$. As the spaces $(B \cap C)$ and $(B \cap C^\perp)$ span B , then R_{U^-} restricted to B has eigenvalues 1 and -1. As $B = \text{im } M_B$ and $V = M_B^\dagger R_{U^-} M_B$, then we conclude that V has eigenvalues 1 and -1 as well.

Now that we have determined that V is a reflection, we need to determine which subspace V reflects across. A corollary of the previous paragraph is that a vector $|\psi\rangle \in C_d(\mathcal{K})$ is in the $(+1)$ -eigenspace of V if and only if $M_B|\psi\rangle$ is in the (-1) -eigenspace of U . Specifically, a vector $|\psi\rangle$ is in the $(+1)$ -eigenspace of V if and only if $M_B|\psi\rangle \in C^\perp$. As $C^\perp = \ker M_C^\dagger$, the vector $|\psi\rangle$ is in the $(+1)$ -eigenspace of V if and only if $|\psi\rangle \in \ker M_C^\dagger M_B$. We proved in Lemma D.1 that $\ker M_C^\dagger M_B = \ker \partial$, so we conclude that $V = R_{\ker \partial}$ \square

We have a matrix V that implements $R_{\ker \partial}$; next, we analyze the complexity of implementing V . We start by analyzing the complexity of implementing R_{U^-} , the reflection across the (-1) -eigenspace of U .

We implement the reflection around the (-1) -eigenspace of U using phase estimation, an algorithm introduced by Magniez et. al. [49]. The algorithm is as follows. We first estimate the phase of U to some degree of accuracy to be specified shortly. Intuitively, we need to estimate the phase of U to high enough accuracy to distinguish between -1 eigenvalues of U and eigenvalues of U close to -1 . We then perform a reflection controlled on the estimated phase.

The **phase gap** of a unitary U with eigenvalues $\{e^{i\theta_1}, \dots, e^{i\theta_k}\}$ is $\min\{|\theta_i| : \theta_i \neq 0\}$. The following lemma shows that the phase gap determines the complexity of reflecting across the 1-eigenspace of U .

Lemma D.4 (Magniez et. al. [49], Paraphrase of Theorem 6). *Let U be a unitary with phase gap θ . A reflection around the 1-eigenspace of U can be performed to constant precision with $O\left(\frac{1}{\theta}\right)$ applications of U .*

The phase gap measures gap between the 1-eigenspace of a unitary and all other eigenvalues. We are interested in the gap in phase between the (-1) -eigenspace of U and the other eigenvalues of U . This is precisely the phase gap of $-U$. The following lemma analyzes the phase gap of $-U$ and gives the complexity of reflecting about the (-1) -eigenspace of U .

Lemma D.5. *We can implement R_{U^-} with $O\left(\sqrt{\frac{d+1}{\lambda_{\min}}}\right)$ calls to U , where $\tilde{\lambda}_{\min}$ is the smallest non-zero eigenvalue of the normalized up-Laplacian of \mathcal{K} .*

Proof. We need to calculate the phase gap of $-U$ to determine the precision to which we need to estimate the phase of U . Observe that if θ is the phase of an eigenvalue of U , then $\theta + \pi$ is the phase of an eigenvalue of $-U$. We can bound the phase gap of $-U$ using Lemma D.2. The non-zero eigenvalues of U are $\{e^{\pm i2\theta_j}\}_j$, where $\{\cos \theta_j\}_j$ were the singular values of $M_C^\dagger M_B$. Therefore, the phases of $-U$ are $\{\pm|\pi - 2\theta_j|\}_j$. Using the inequality that $\pi/2 - \theta_j \geq \cos \theta_j$ for $\theta_j \in [0, \pi/2]$, then the phase gap of $-U$ is bounded below by

$$|\pi - 2\theta_j| \geq 2 \cos \theta_j \geq 2 \cdot \sigma_{\min}(M_C^\dagger M_B)$$

where $\sigma_{\min}(M_C^\dagger M_B)$ is the smallest singular value of $M_C^\dagger M_B$.

We can actually relate the smallest singular value of $M_C^\dagger M_B$ to something more meaningful. By the proof of Lemma D.1, the matrix $M_C^\dagger M_B = \frac{1}{\sqrt{d+1}} D^{-1/2} \partial \sqrt{W}$, where D is the diagonal

matrix with the degrees of the $(d-1)$ -simplices on the diagonal. Thus, $(M_C^\dagger M_B)(M_C^\dagger M_B)^\dagger = \frac{1}{d+1} D^{-1/2} \partial W \delta D^{-1/2} = \frac{1}{d+1} D^{-1/2} L D^{-1/2}$. Recall from Section 2 that the matrix $D^{-1/2} L D^{-1/2}$ is the normalised up-Laplacian. The singular values of a matrix A are the square roots of the eigenvalues of AA^T . Thus, the smallest singular value of $M_C^\dagger M_B$, and the phase gap of $-U$, is $\Omega\left(\sqrt{\frac{\tilde{\lambda}_{\min}}{d+1}}\right)$, where $\tilde{\lambda}_{\min}$ is the smallest eigenvalue of $D^{-1/2} L D^{-1/2}$. Therefore, by Lemma D.4, we can implement R_{U^-} with $O\left(\sqrt{\frac{d+1}{\tilde{\lambda}_{\min}}}\right)$ calls to U . \square

We are almost ready to give the running time for $V = R_{\ker \partial}$, but first, we need to make a delicate distinction. The matrices M_B and M_C have orthonormal columns, but they are *not* unitary. We can see this as $\ker M_B^\dagger \neq 0$ and $\ker M_C^\dagger \neq 0$. As M_B and M_C are not unitary, they cannot be implemented on a quantum computer. Fortunately, it suffices to implement unitaries U_B and U_C such that $U_B|_{C_d(\mathcal{K})} = M_B$ and $U_C|_{C_{d-1}(\mathcal{K})} = M_C$. With this in mind, we can give the running time for $V = R_{\ker \partial}$.

Lemma D.6. *There is an algorithm to perform $R_{\ker \partial}$ in time $\tilde{O}\left(\sqrt{\frac{d+1}{\tilde{\lambda}_{\min}}}(T_B + T_C)\right)$, where T_B and T_C are the times to perform the unitaries U_B and U_C respectively.*

Proof. Lemma D.3 shows that $R_{\ker \partial} = V = M_B^\dagger R_{U^-} M_B$. We can equivalently run $U_B^\dagger R_{U^-} U_B$. As U_B takes T_B by definition, we only need to show we can implement R_{U^-} in $\tilde{O}\left(\sqrt{(d+1)/\tilde{\lambda}_{\min}}(T_B + T_C)\right)$ time. Lemma D.5 shows we can implement R_{U^-} with $O\left(\sqrt{(d+1)/\tilde{\lambda}_{\min}}\right)$ calls to U , so we need to show we can implement U in $\tilde{O}(T_B + T_C)$. The unitary $U = R_C R_B$, and we claim we can implement R_C and R_B in $\tilde{O}(T_B)$ and $\tilde{O}(T_C)$ respectively. We can implement R_B as $U_B R_{\mathcal{K}_d} U_B^\dagger$, where $R_{\mathcal{K}_d}$ reflects across the basis states $\{|0\rangle|\sigma\rangle \mid \sigma \in \mathcal{K}_d\}$. We can check if a quantum state is of the form $|0\rangle|\sigma\rangle$ in $O(\log n_d)$ gates (specifically, by checking if the basis state is within a certain range), so the unitary $R_{\mathcal{K}_d}$ takes $O(\log n_d)$ gates, and R_B takes $\tilde{O}(T_B)$ time. The unitary R_C takes $\tilde{O}(T_C)$ time by the same argument. \square

The running time T_B is dependent on how the boundary maps are loaded into the quantum algorithm. We propose a method of storing the boundary maps in a quantum computer called the **incidence array**. The incidence array is adapted from the *adjacency array* introduced by Durr et al. [18] to store the adjacency between pairs of vertices in a graph.

For a d -simplex $\tau = \{v_0, \dots, v_d\}$, the **down-incidence array** is the function $g : |\tau\rangle|j\rangle|0\rangle \rightarrow |\tau\rangle|j\rangle|\tau \setminus \{v_j\}\rangle$ for $0 \leq j \leq d$. The simplices in the boundary of τ have alternating sign. To address this, we also perform a negation conditioned on the parity of $|j\rangle$ to compute $(-1)^j |\tau\rangle|j\rangle|\tau \setminus \{v_j\}\rangle$.

Durr et al. [18] claim that queries to the incidence array can be performed in logarithmic time. As the down-incidence array is identical to the adjacency array⁸, queries to the down-incidence also take logarithmic time. We can compute the state $|\partial\tau\rangle|\tau\rangle$ with the down-incidence array and the following lemma.

Lemma D.7 (Cade, Montanaro, Belovs [9], Implicit in the proof of Lemma 2). *Let $f : [m] \rightarrow [k]$ be a function, and let \mathcal{O}_f be an oracle that computes $\mathcal{O}_f : |i\rangle|0\rangle \rightarrow |i\rangle|f(i)\rangle$. The state $\frac{1}{\sqrt{m}} \sum_{i=1}^m |f(i)\rangle$ can be computed with $O(\sqrt{m})$ queries to \mathcal{O}_f and $O(\text{polylog}(m))$ additional gates.*

Corollary D.8. *The unitary U_B can be implemented in $O(\sqrt{d})$ queries to the up-incidence array and $\tilde{O}(\sqrt{d})$ time.*

⁸The down-incidence array is actually an adjacency array of a graph related to simplicial complexes, namely, the incidence graph between the $(d-1)$ - and d -simplices.

It is harder to produce a generic implementation of U_C than U_B . The d -simplices can have arbitrary weights, so constructing the states $|c_\sigma\rangle$ in general requires constructing arbitrary quantum states with real coefficients. However, the weights on the simplices do not affect whether or not a cycle is null-homologous. Therefore, we can always run our null-homology test on the unweighted complex; the trade-off is that the effective resistance or effective capacitance might be higher in the unweighted complex. We analyze the running time of the unweighted case in Section D.1.

We now analyze the complexity of constructing the initial state $|w_0\rangle/\|w_0\|$. To construct the dummy state, we start by adding an additional “ d -cell” $|\emptyset\rangle$ to the complex with boundary $|\gamma\rangle$ (really, we just add $|\gamma\rangle$ as a column to ∂ .) The new cell will have non-trivial overlap with $|w_0\rangle$, so we can construct $|w_0\rangle$ by amplifying this component of $|\emptyset\rangle$. We outline this method in the proof of Theorem D.10, but first, we state Lemma D.9 which is a generalization of the parallel formula for effective resistance; its proof is nearly identical to the proof of Theorem B.2.

Lemma D.9. *Let $V = V_1 \oplus V_2$ be a vector space. Let $A : V \rightarrow U$ be a linear map, and let $A_1 : U_1 \rightarrow V$ and $A_2 : U_2 \rightarrow V$ be the restriction of A to U_1 and U_2 . Let $|t\rangle \in \text{im } A_1 \cap \text{im } A_2 \subset U$. If $|s\rangle = A^+|t\rangle$, $|s_1\rangle = A_1^+|t\rangle$, and $|s_2\rangle = A_2^+|t\rangle$, then*

$$\|s\|^2 \leq \left(\frac{1}{\|s_1\|^2} + \frac{1}{\|s_2\|^2} \right)^{-1}$$

Equality is achieved when $\text{im } A_1 \cap \text{im } A_2 = \text{span}\{|t\rangle\}$. In this case, $|s\rangle = t|s_1\rangle + (1-t)|s_2\rangle$ where $t = \|s_2\|^2/(\|s_1\|^2 + \|s_2\|^2)$.

Theorem D.10. *Let \mathcal{O}_γ be the oracle that takes $\mathcal{O}_\gamma : |0\rangle \rightarrow |\gamma\rangle$. Let T_γ be the time it takes to implement \mathcal{O}_γ . The state $|w_0\rangle = \partial^+|\gamma\rangle$ can be created in $\tilde{O}((\sqrt{1/\mathcal{R}_\gamma(\mathcal{K})} + \sqrt{\mathcal{R}_\gamma(\mathcal{K})})(T_B + T_C + T_\gamma))$ time.*

Proof. We append $|\gamma\rangle$ as a column to ∂ to create a new matrix $\hat{\partial}$. Let $|\emptyset\rangle$ be index of the new column, so $\hat{\partial} = \partial + |\gamma\rangle\langle\emptyset|$. Let $|w'_0\rangle = \hat{\partial}^+|\gamma\rangle$. We conclude that $|\emptyset\rangle = |w'_0\rangle + |w'_0{}^\perp\rangle$ where $|w'_0{}^\perp\rangle \in \ker \hat{\partial}$, as the projection $\Pi_{\ker \hat{\partial}}|\emptyset\rangle = \hat{\partial}^+\hat{\partial}|\emptyset\rangle = \hat{\partial}^+|\gamma\rangle = |w'_0\rangle$.

We construct $|w_0\rangle/\|w_0\|$ in two steps. First, we use amplitude amplification to amplify the $|w'_0\rangle$ component of $|\emptyset\rangle$. We then use a second amplitude amplification to amplify the $|w_0\rangle$ component of $|w'_0\rangle$. These amplitude amplifications are nested, as we need to perform the first to create the initial state for the second.

If we perform constant time phase estimation of $|\emptyset\rangle$ on the unitary $R_{\ker \hat{\partial}}$, then we can map $|\emptyset\rangle$ to $|0\rangle|w'_0\rangle + |1\rangle|w'_0{}^\perp\rangle$. We can then amplify the amplitude of $|0\rangle|w_0\rangle$ part arbitrarily close to $|w_0\rangle/\|w_0\|$ using $O(\|w_0\|^{-1})$ calls to $R_{\ker \hat{\partial}}$.

We calculate $\|w_0\|$ using the formula from the lemma. The vector $|\emptyset\rangle$ has length 1, so Lemma D.9 shows that

$$\|w'_0\|^2 = \left(\frac{1}{\mathcal{R}_\gamma(\mathcal{K})} + 1 \right)^{-1} = \frac{\mathcal{R}_\gamma(\mathcal{K})}{\mathcal{R}_\gamma(\mathcal{K}) + 1}.$$

Thus, we need to perform the reflection $R_{\ker \hat{\partial}}$ a total of $O(\|w_0\|^{-1}) = O(\sqrt{(\mathcal{R}_\gamma(\mathcal{K}) + 1)/\mathcal{R}_\gamma(\mathcal{K})})$ times to create $|w_0\rangle/\|w_0\|$.

The next step in our algorithm is to amplify the $|w_0\rangle/\|w_0\|$ component of $|w'_0\rangle/\|w'_0\|$. By Lemma D.9, the state $\|w'_0\| = t\|w_0\| + (1-t)\|\emptyset\|$ for $t = 1/(\mathcal{R}_\gamma(\mathcal{K}) + 1)$. Therefore, the

$|w_0\rangle/\|w_0\|$ component of $|w'_0\rangle/\|w'_0\|$ has norm

$$\begin{aligned} t\|w_0\|/\|w'_0\| &= \frac{1}{\mathcal{R}_\gamma(\mathcal{K}) + 1} \sqrt{\mathcal{R}_\gamma(\mathcal{K})} \sqrt{\frac{\mathcal{R}_\gamma(\mathcal{K}) + 1}{\mathcal{R}_\gamma(\mathcal{K})}} \\ &= \sqrt{\frac{1}{\mathcal{R}_\gamma(\mathcal{K}) + 1}} \end{aligned}$$

To return the state $|w_0\rangle/\|w_0\|$, we need to perform amplitude amplification again. We can create the state $|w'_0\rangle$ using the amplitude amplification from the previous two paragraphs with $O(\sqrt{(\mathcal{R}_\gamma(\mathcal{K}) + 1)/\mathcal{R}_\gamma(\mathcal{K})})$ applications of $R_{\ker \partial'}$, and we can reflect across $|\emptyset\rangle$ in constant time as it is a basis state. To create $|w_0\rangle/\|w_0\|$, we need

$$O\left(\sqrt{(\mathcal{R}_\gamma(\mathcal{K}) + 1)/\mathcal{R}_\gamma(\mathcal{K})} \sqrt{(\mathcal{R}_\gamma(\mathcal{K}) + 1)}\right) = O\left(\sqrt{\mathcal{R}_\gamma(\mathcal{K})} + \sqrt{\frac{1}{\mathcal{R}_\gamma(\mathcal{K})}}\right)$$

applications of $R_{\ker \hat{\partial}}$.

We now argue that we can compute $R_{\ker \hat{\partial}}$ in $O(T_C + T_B + T_\gamma)$ time. As was the case with $R_{\ker \partial}$, we decompose $R_{\ker \hat{\partial}} = M_{\hat{B}}^\dagger R_{\hat{U}^-} M_{\hat{B}}$ for space \hat{B} and \hat{C} defined

$$\hat{B} = B \cup \{b_\emptyset := |\gamma\rangle\}$$

$$\hat{C} = \text{span} \left\{ |c_\sigma\rangle := \frac{1}{\sqrt{\deg(\sigma) + 1}} |\sigma\rangle |\emptyset\rangle + \sum_{\sigma \subset \tau} \frac{w(\sigma)}{\sqrt{\deg(\sigma) + 1}} |\sigma\rangle |\tau\rangle : \sigma \in \mathcal{K}_{d-1} \right\}.$$

The unitaries $M_{\hat{B}}$, $M_{\hat{C}}$, and $\mathbb{R}_{\hat{U}^-}$ are defined analogously to M_B and M_C . We can implement the unitary version of these matrices $U_{\hat{B}}$ in $O(T_B + T_\gamma)$ and $U_{\hat{C}}$ in $O(T_C)$. \square

We now summarize this section in the following theorem.

Theorem D.11. *Let \mathcal{K} be a simplicial complex, $\gamma \in C_{d-1}(\mathcal{K})$ a null-homologous cycle, and $\mathcal{K}(x) \subset \mathcal{K}$ be a simplicial complex. There is a quantum algorithm for deciding if γ is null-homologous in $\mathcal{K}(x)$ that runs in time*

$$\tilde{O} \left(\sqrt{\frac{(d+1)\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}{\tilde{\lambda}_{\min}}} (\sqrt{d} + T_C)} + \left(\sqrt{\frac{1}{\mathcal{R}_\gamma(\mathcal{K})}} + \sqrt{\mathcal{R}_\gamma(\mathcal{K})} \right) (\sqrt{d} + T_C + T_\gamma) \right)$$

where \mathcal{R}_{\max} is the maximum effective resistance of γ in any subcomplex $\mathcal{K}(y)$ and $\mathcal{C}_{\max}(\gamma)$ is the maximum effective capacitance γ in any subcomplex $\mathcal{K}(y)$ and $\tilde{\lambda}_{\min}$ is the smallest eigenvalue of the normalized d up-Laplacian.

D.1 Special cases.

We now consider a few special cases of the null-homology span program. These special cases will allow us to replace the terms T_B and T_γ in Theorem D.11 with concrete running times.

Unweighted simplicial complexes We now consider the case where there are no weights on the d -simplices, or equivalently, when $w(\tau) = 1$ for each d -simplex τ . While computing M_C is hard in general, in the unweighted case, we can implement the unitary M_C using a straightforward oracle. For a $(d-1)$ -simplex σ that is incident to the d -simplices $\{\tau_1, \dots, \tau_m\}$, the **up-incidence array** is the oracle is the function that maps $h : |\tau\rangle|j\rangle|0\rangle \rightarrow |\tau\rangle|j\rangle|0\rangle$. By Lemma D.7, the up-incidence array can be used to compute $|c_\tau\rangle = \frac{1}{\sqrt{m}}|\tau\rangle|\delta\tau\rangle$ in $O(\sqrt{m})$ time. The unitary M_C computes the state $|c_\sigma\rangle$ in parallel, so computing M_C will take $\sqrt{d_{\max}}$ queries, where $d_{\max} = \max_{\tau \in C_{d-1}(\mathcal{K})} \deg(\tau)$ time. This is summarized in the following lemma.

Lemma D.12. *If \mathcal{K} is an unweighted simplicial complex, the unitary U_C can be implemented in $O(\sqrt{d_{\max}})$ queries to the up incidence array and $T_C = \tilde{O}(\sqrt{d_{\max}}) = \tilde{O}(\sqrt{n_0})$ time.*

Additionally, if \mathcal{K} is an unweighted complex, we can upper bound the quantity $\frac{1}{\mathcal{R}_\gamma(\mathcal{K})}$.

Lemma D.13. *Let \mathcal{K} be an unweighted simplicial complex with n_0 vertices, and let γ be a unit-length null-homologous $(d-1)$ -cycle in \mathcal{K} . Then $\frac{1}{\mathcal{R}_\gamma(\mathcal{K})} \leq n_0$.*

Proof. Recall that in an unweighted simplicial complex that $\mathcal{R}_\gamma(\mathcal{K}) = \gamma^T(L_{d-1}^{up})^+\gamma$. As the eigenvalues of L_{d-1}^{up} are bounded above by n_0 (Theorem 2.4), then the non-zero eigenvalues of $(L_{d-1}^{up})^+$ are bounded below by $\frac{1}{n_0}$. As $\gamma \in \text{im } L_{d-1}^{up}$, then $\gamma^T(L_{d-1}^{up})^+\gamma$ is bounded below by the smallest non-zero eigenvalue of $(L_{d-1}^{up})^+$, i.e $\frac{1}{n_0}$. Therefore, $\frac{1}{\mathcal{R}_\gamma(\mathcal{K})} \leq n_0$ \square

Cycle is the boundary of a d -simplex. We now consider the case that the input cycle γ is the boundary of a d -simplex. In this case, we can implement the oracle \mathcal{O}_γ with the down incidence array used to implement M_B . We get the same running time for T_γ as T_B .

Lemma D.14. *If γ is the boundary of a d -simplex, there is a quantum algorithm implementing \mathcal{O}_γ in $O(\sqrt{d})$ queries to the down incidence array and $T_\gamma \in \tilde{O}(\sqrt{d}) \subset \tilde{O}(\sqrt{n_0})$ time.*

Summing Up. If we combine our bounds for the case where our complex is unweighted (Lemma D.12 and Lemma D.13) and the cycle is the boundary of d -simplex (Lemma D.14) with the bound of the time complexity (Theorem D.11), we get the following bound on the time complexity.

Theorem 4.10. *Let \mathcal{K} be an unweighted simplicial complex with n_0 vertices, let $\gamma \in C_{d-1}(\mathcal{K})$ a null-homologous cycle in \mathcal{K} , and $\mathcal{K}(x) \subset \mathcal{K}$ be a simplicial complex. Furthermore, assume that γ is the boundary of a d -simplex and the complex is unweighted. There is a quantum algorithm for deciding if γ is null-homologous in $\mathcal{K}(x)$ that runs in time*

$$\tilde{O} \left(\sqrt{\frac{\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}{\tilde{\lambda}_{\min}}}n_0} + \sqrt{\mathcal{R}_\gamma(\mathcal{K})n_0} \right),$$

where $\mathcal{R}_{\max}(\gamma)$ is the maximum finite effective resistance $\mathcal{R}_\gamma(\mathcal{L})$ of γ in any subcomplex $\mathcal{L} \subset \mathcal{K}$, $\mathcal{C}_{\max}(\gamma)$ is the maximum finite effective capacitance $\mathcal{C}_\gamma(\mathcal{L}, \mathcal{L})$ in any subcomplex $\mathcal{K}(x)$, and $\tilde{\lambda}_{\min}$ is the spectral gap of the normalized up-Laplacian $\tilde{L}_{d-1}^{up}[\mathcal{K}]$.

D.2 Space Complexity.

We now comment on the space complexity of our algorithm. The inner product space $C_{d-1}(\mathcal{K}) \oplus C_d(\mathcal{K})$ has dimension $n_{d-1}n_d \in O(\binom{n_0}{d}\binom{n_0}{d+1})$, so vectors in this space can be represented with

$O(\log(\binom{n_0}{d} \binom{n_0}{d+1})) = O(dn)$ qubits. Additionally, the phase estimation step from Lemma D.5 takes $O(\log(1/\tilde{\lambda}_{\min}))$ ancillary qubits, which is $O(dn_d)$ qubits by the bounds on the normalized spectral gap (Lemma 2.5 and Theorem 1.2). Finally, the outer phase estimation of $R_{\mathcal{H}(x)}R_{\ker A}$ takes $O(\log(\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)))$ qubits. Other gates require polylogarithmic ancillary qubits. The space complexity of our algorithm is comparable to some other QTDA algorithms, as other QTDA algorithms require $O(\log(1/\lambda_{\min}))$ to perform phase estimation of the combinatorial Laplacian [29].

E Construction of the Building Block.




Figure 6: Left to Right: Stellar subdivision, prism, and stellar prism of a triangle.

In this section, we formally describe the “building block” B_d that we use to construct the simplicial complexes \mathcal{B}_d^n , \mathcal{P}_d^n , and \mathcal{Q}_d^n with exponentially-large effective resistance and capacitance in Section 5.2. A more intuitive, but informal, description of the building block can be found in Section 5.2

The building block B_d is constructed using two constructions from algebraic topology: the stellar subdivision and the prism. We will combine these two constructions to make a new construction we call the stellar prism. The building block B_d is then a quotient of the stellar prism. This section will first describe the stellar subdivision, prism, and stellar prism. We then present the construction of B_d and prove some of its relevant properties.

E.1 Stellar Subdivision.

Let K be a d -dimensional simplicial complex. The **stellar subdivision** of K is the d -dimensional simplicial complex SK that is the union of the $(d-1)$ -skeleton of K and, for each d -simplex $\sigma \in K_d$, the set of simplices $\{\tau \cup \{v_\sigma\} : \tau \subsetneq \sigma\}$, where v_σ is a new vertex. See Figure 6. The property of the stellar subdivision that is key to our construction is that it increases the number of d -simplices in the simplicial complex.

Observation E.1. *Let K be a simplicial complex with n_d d -simplices. Then SK has $d \cdot n_d$ d -simplices.*

We can map chains from our original complex K to chains in its stellar subdivision. For each d -simplex $\sigma \in K_d$ and any integer k , we define a map $b_\sigma : C_k(K) \rightarrow C_{k+1}(SK)$. Let $\tau \subsetneq \sigma$, and suppose that v_σ is the $(i+1)$ st element of $\tau \cup \{v_\sigma\}$ with respect to the ordering on the vertices⁹; that is, if $\tau = \{v_0, \dots, v_k\}$, then $\tau \cup \{v_\sigma\} = \{v_0, \dots, v_{i-1}, v_\sigma, v_i, \dots, v_k\}$. We define $b_\sigma(\tau) = (-1)^i \cdot (\tau \cup \{v_\sigma\})$ for all k -simplices of the form $\tau \subsetneq \sigma$, and $b_\sigma(\varphi) = 0$ for any other k -simplex φ .

⁹It is worth noting that it is arbitrary where v_σ is in the ordering of the vertices, so i can be any number in the

Lemma E.2. Consider $b_\sigma : C_k(K) \rightarrow C_{k+1}(SK)$ as defined above. The map b_σ satisfies $\partial \circ b_\sigma = \mathbb{I} - b_\sigma \circ \partial$.

Proof. As b_σ is a linear map, it suffices to prove this for the basis k -chains $\tau \in C_k(K)$. The boundary of $b_\sigma(\tau)$ is

$$\begin{aligned} \partial(b_\sigma(\tau)) &= (-1)^i \cdot \partial(\tau \cup \{v_\sigma\}) \\ &= (-1)^i \left[\sum_{j=0}^{i-1} (-1)^j \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) + (-1)^i \cdot \tau + \sum_{j=i}^k (-1)^{j+1} \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) \right] \\ &= \tau + (-1)^i \left[\sum_{j=0}^{i-1} (-1)^j \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) + \sum_{j=i}^k (-1)^{j+1} \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) \right] \end{aligned}$$

We claim that the second term in this sum equals $(-1) \cdot b_\sigma(\partial\tau)$. Observe that for $j < i$, v_σ will be the i th vertex of $\tau \cup \{v_\sigma\} \setminus \{v_j\}$. Likewise, for $i \geq j$, v_σ will be the $(i+1)$ st vertex of $\tau \cup \{v_\sigma\} \setminus \{v_j\}$. Therefore, we have that

$$\begin{aligned} & (-1)^i \left[\sum_{j=0}^{i-1} (-1)^j \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) + \sum_{j=i}^k (-1)^{j+1} \cdot (\tau \cup \{v_\sigma\} \setminus \{v_j\}) \right] \\ &= (-1)^i \left[\sum_{j=0}^{i-1} (-1)^j (-1)^{i-1} b_\sigma(\tau \setminus \{v_j\}) + \sum_{j=i}^k (-1)^{j+1} (-1)^i b_\sigma(\tau \setminus \{v_j\}) \right] \\ &= (-1) \left[\sum_{j=0}^k (-1)^j b_\sigma(\tau \setminus \{v_j\}) \right] = (-1) \cdot b_\sigma(\partial\tau) \end{aligned}$$

□

Using the maps b_σ , we now define another map $S_* : C_k(K) \rightarrow C_k(SK)$. For $k < d$, we define the map $S_* : C_k(K) \rightarrow C_k(SK)$ just to be the inclusion map. For $k = d$, define $S : C_d(K) \rightarrow C_d(SK)$ on each d -simplex $\sigma = \{v_0, \dots, v_d\}$ as $S_*(\sigma) = b_\sigma(\partial\sigma)$.

Lemma E.3. Consider $S_* : C_k(K) \rightarrow C_k(SK)$ as defined above. The map S_* satisfies $\partial \circ S_* = S_* \circ \partial$.

Proof. For $i < d$, this is obvious as S_* is the inclusion map. For $i = d$, this follows from Lemma E.2 as $\partial S_*(\sigma) = \partial b_\sigma(\partial\sigma) = \partial\sigma - b_\sigma \partial\partial(\sigma) = \partial\sigma = S_* \partial\sigma$. □

E.2 Prisms.

The *prism* of a simplicial complex K is a triangulation of the space $K \times [0, 1]$. To define the prism of K , we will define the prism of a single simplex $\sigma \in K$; the prism of the entire complex K is then the union of the prism of its simplices, i.e. $PK = \cup_{\sigma \in K} P\sigma$. See Figure 6.

Assume that $\sigma = \{v_0, \dots, v_d\}$. The *prism* of σ is a $(d+1)$ -dimensional simplicial complex $P\sigma$ with vertices $\sigma \times \{0, 1\}$. The prism contains all simplices of the form $\tau \times \{0\}$ and

range $0 \leq i \leq k$. Lemmas E.2 and E.3 summarize the relevant properties of the stellar subdivision for our paper, and as we will see, these lemmas will hold wherever v_σ is in the vertex order, so long as the map b_σ is defined appropriately.

$\tau \times \{1\}$ for $\tau \subset \sigma$. The remaining simplices of $P\sigma$ are the closure of the $(d+1)$ -simplices $\{(v_0, 0), \dots, (v_j, 0), (v_j, 1), \dots, (v_d, 1)\}$ for $0 \leq j \leq d$. We denote such a simplex

$$\sigma^j := \{(v_0, 0), \dots, (v_j, 0), (v_j, 1), \dots, (v_d, 1)\}.$$

Note that for a paired of nested simplices $\tau \subset \sigma$, their prisms are also nested, i.e. $P\tau \subset P\sigma$.

As was the case for stellar subdivision, we will define several maps between chains in K and chains in PK . We abuse notation and define $K \times \{1\} := \{\sigma \times \{1\} : \sigma \in K\}$. The first map $I_1 : C_k(K) \rightarrow C_k(PK)$ maps chains in K to chains in $K \times \{1\}$. Specifically, for a d -simplex $\sigma \in K_d$, we define $I_1(\sigma) = \sigma \times \{1\}$. We define a map I_0 analogously. The following lemma is obvious.

Lemma E.4. *Consider $I_i : C_k(K) \rightarrow C_k(PK)$ as defined above for $i = 0, 1$. The map I_i satisfy $\partial \circ I_i = I_i \circ \partial$ for $i = 0, 1$.*

We now define a map $P_* : C_k(K) \rightarrow C_{k+1}(PK)$. Specifically, for a k -simplex $\sigma \in K$ with $\sigma = \{v_0, \dots, v_k\}$, the corresponding $(k+1)$ -chain is defined $P_*(\sigma) = \sum_{i=0}^k (-1)^i \sigma^i$.

Lemma E.5. *Consider $P_* : C_k(K) \rightarrow C_{k+1}(PK)$ as defined above. The map P_* satisfies $P_* \circ \partial + \partial \circ P_* = I_1 - I_0$.*

Proof. The key points of this proof are essentially identical to those provided by Hatcher [32] in the proof of Theorem 2.10. \square

E.3 Stellar Prisms.

Now we propose a way of combining stellar subdivisions and prisms that we call the **stellar prism**. Intuitively, the stellar prism is a triangulation of the space $K \times [0, 1]$ where the bottom copy $K \times \{0\}$ is triangulated the same way as K and the top copy $K \times \{1\}$ is triangulated using the stellar subdivision. See Figure 6.

We first define the stellar prism of a d -simplex; the stellar prism of a d -dimensional simplicial complex K is the union of the prism of the $(d-1)$ -skeleton and the stellar prisms of the d -simplices, i.e. $SPK = PK^{d-1} \cup_{\sigma \in K_d} SP\sigma$.

Let σ be a d -simplex. The stellar prism of σ is the $(d+1)$ -dimensional simplicial complex $SP\sigma$ described as follows. The vertices of $SP\sigma$ are $(\sigma \times \{0, 1\}) \cup \{v_\sigma\}$, where v_σ is a new vertex. For any simplex $\tau \subsetneq \sigma$, $SP\sigma$ contains both the complex $P\tau$ and all simplices of the form $\{\varphi \cup \{v_\sigma\} : \varphi \in P\tau\}$. Additionally, $SP\sigma$ contains the d -simplex $\sigma \times \{0\}$ and the $(d+1)$ simplex $(\sigma \times \{0\}) \cup \{v_\sigma\}$. If we let $v_{\sigma \times \{1\}} = v_\sigma$, note that $SP\sigma$ contains the subdivision $S(\sigma \times \{1\})$.

We now define a linear map between $SP_* : C_i(K) \rightarrow C_{i+1}(SPK)$. For $i < d$, we simply define $SP_* = P_*$. For $i = d$, we define $SP_*(\sigma) = -b_\sigma I_0(\sigma) - b_\sigma P_* \partial(\sigma)$

Lemma E.6. *Consider $SP_* : C_i(K) \rightarrow C_{i+1}(SPK)$ as defined above. The map SP_* satisfies $\partial SP_* + SP_* \partial = S_* I_1 - I_0$*

Proof. For $i < d$, this follows from Lemma E.5 and the fact that $SP_* = P_*$ and $S_* = \mathbb{I}$. We now verify this for $i = d$.

We will analyse the boundary of $\partial SP_*(\sigma)$. We find that

$$\partial SP_*(\sigma) = -\partial b_\sigma I_0(\sigma) - \partial b_\sigma P_* \partial(\sigma). \quad (1)$$

We analyse the two terms in this sum. Using Lemma E.2 and E.4, the first term of Equation (1) evaluates to

$$\begin{aligned} \partial b_\sigma I_0(\sigma) &= I_0(\sigma) - b_\sigma \partial I_0(\sigma) \\ &= I_0(\sigma) - b_\sigma I_0 \partial(\sigma) \end{aligned}$$

Using Lemma E.5, the second term of Equation (1) evaluates to

$$\begin{aligned}
\partial b_\sigma P_* \partial(\sigma) &= P_* \partial(\sigma) - b_\sigma \partial P_* \partial(\sigma) \\
&= P_* \partial(\sigma) - b_\sigma (I_1 - I_0 - P_* \partial) \partial(\sigma) \\
&= P_* \partial(\sigma) - b_\sigma I_1 \partial(\sigma) + b_\sigma I_0 \partial(\sigma) + b_\sigma P_* \partial \partial(\sigma) \\
&= P_* \partial(\sigma) - b_\sigma I_1 \partial(\sigma) + b_\sigma I_0 \partial(\sigma)
\end{aligned}$$

The first term in this sum is $P_* \partial(\sigma) = SP_* \partial(\sigma)$ as $P_* = SP_*$ for all dimensions less than d . The term $b_\sigma I_1 \partial(\sigma) = b_\sigma \partial I_1(\sigma) = S_* I_1(\sigma)$. Combining the two terms of Equation (1), we find that

$$\begin{aligned}
\partial SP_*(\sigma) &= -\partial b_\sigma I_0(\sigma) - \partial b_\sigma P_* \partial(\sigma) \\
&= -(I_0(\sigma) - b_\sigma I_0 \partial(\sigma)) - (P_* \partial(\sigma) - S_* I_1(\sigma) + b_\sigma I_0 \partial(\sigma)) \\
&= -I_0(\sigma) - P_* \partial(\sigma) + S_* I_1(\sigma)
\end{aligned}$$

□

E.4 Building Block.

We now describe the building block B_d . Let Δ^d be the closure of the d -simplex $\sigma = \{v_0, \dots, v_d\}$, and let $\partial\Delta^d$ denote the $(d-1)$ -dimensional simplicial complex $\Delta^d \setminus \{\sigma\}$. The building block B_d is derived from the stellar prism $SP(\partial\Delta^d)$. Denote the $(d-1)$ -simplex $\sigma \times \{1\} \setminus \{(v_i, 1)\}$ as σ_i . Observe that the vertices of $SP(\partial\Delta^d)$ are $(\sigma \times \{0, 1\}) \cup \{v_{\sigma_i} : 0 \leq i \leq d\}$. The building block B_d is the simplicial complex obtained by replacing each vertex v_{σ_i} in $SP(\partial\Delta^d)$ with the vertex $(v_i, 1)$. See the main body of the text for a more intuitive description (but less formal) of B_d .

We now prove the relevant properties of B_d .

Lemma 5.10. *Let $\sigma = \{v_0, \dots, v_d\}$ be a set. There is a d -dimensional simplicial complex B_d with vertices $\sigma \times \{0, 1\}$ such that*

1. B_d has $\Theta((d+1)^3)$ d -simplices.
2. there is a d -chain $f \in C_d(B_d)$ such that

$$\begin{aligned}
(i) \quad \partial f &= \partial(\sigma \times \{0\}) + d \cdot \partial(\sigma \times \{1\}) \\
(ii) \quad \|f\|^2 &\in \Omega((d+1)^3)
\end{aligned}$$

Proof of Lemma 5.10, Part 1. To keep track of the simplices, we introduce some notation. Recall that $\sigma_i = \sigma \setminus \{v_i\}$. Additionally, we will denote $\sigma_{ij} = \sigma \setminus \{v_i, v_j\}$. For $0 \leq k \leq d$, denote the set

$$\sigma^k := \{(v_0, 0), \dots, (v_k, 0), (v_k, 1), \dots, (v_d, 1)\}.$$

Note that σ^k is not a simplex in any of the complexes we consider, but some of its subsets will be. For distinct values $i, j \neq k$, denote the simplex $\sigma_{ij}^k := \sigma_k \setminus \{(v_i, 0), (v_i, 1), (v_j, 0), (v_j, 1)\}$. Note that the subtraction in this definition is redundant, as only one of the vertices $(v_i, 0)$ or $(v_i, 1)$ will be contained in σ^k .

First, we count the number of d -simplices in $SP(\partial\Delta^d)$. Fix a $(d-1)$ -simplex σ_i , and consider a $(d-2)$ -simplex σ_{ij} in its boundary. The $(d-1)$ -simplices in $P\sigma_{ij}$ are those of the form σ_{ij}^k for $0 \leq k \leq d$, $k \neq i, j$. The d -simplices in $SP\sigma_i$ are those simplices of the form $\sigma_{ij}^k \cup \{v_{\sigma_i}\}$. Additionally, $SP(\partial\Delta^d)$ contains those simplices $(\sigma_i \times \{0\}) \cup \{v_{\sigma_i}\}$. Therefore, the number of d -simplices in $SP(\partial\Delta^d)$ equals the number of choices of i, j , and k , plus $d+1$ for the d -simplices of

the form $(\sigma_i \times \{0\}) \cup \{v_{\sigma_i}\}$. In total, there are $(d+1) \cdot d \cdot (d-1) + (d+1) \in \Theta((d+1)^3)$ d -simplices in $SP(\partial\Delta^d)$

Now we count the number of d -simplices in B_d . As we will see, replacing the vertex v_{σ_i} will reduce the number of d -simplices, but only by a constant factor. In the case that $i > k$, then $\sigma_{ij}^k \cup \{(v_i, 1)\} = \sigma_j^k$, but in the case that $i < k$, no such simplification is possible. Therefore, the $(d+1)$ -simplices in B_d fall into two sets: $\{\sigma_j^k : 0 \leq k < d, j \neq k\}$ and $\{\sigma_{ij}^k \cup \{(v_i, 1)\} : 0 \leq k \leq d, j \neq k, i < k\}$. Note that the simplices σ_j^k may arise from different stellar prisms $SP\sigma_i$ for different values of i , whereas the simplices of the form $\sigma_{ij}^k \cup \{v_i\}$ only arise from the stellar prism $SP\sigma_i$. However, as there are $O((d+1)^2)$ simplices of the form σ_j^k , $O((d+1)^3)$ simplices of the form $\sigma_{ij}^k \cup \{(v_i, 1)\}$, and $d+1$ simplices of the form $(\sigma_i \times \{0\}) \cup \{(v_i, 1)\}$, we conclude that there are $\Theta((d+1)^3)$ d -simplices in B_d . \square

Proof of Lemma 5.10, Part 2. The simplicial complex B_d is the one described in the paragraphs preceding this proof, which is derived from the simplicial complex $SP(\partial\Delta^d)$ by replacing some vertices. We first describe the chain f in the complex $SP(\partial\Delta^d)$. We then analyse this chain after we replace the vertices.

Let $\sigma = \{v_0, \dots, v_d\}$. The chain f is defined $f := SP_*(\partial\sigma)$. It is important to note that while the simplex σ is not contained in $\partial\Delta^d$, its boundary $\partial\sigma$ is still a well-defined $(d-1)$ -chain in $C_{d-1}(\partial\Delta^d)$. By Lemma E.6, we know that

$$\begin{aligned} \partial f &= \partial SP_*(\partial\sigma) = S_*I_1\partial(\sigma) - I_0\partial(\sigma) - SP_*\partial\partial(\sigma) \\ &= S_*I_1\partial(\sigma) - I_0\partial(\sigma) \end{aligned}$$

We know $I_0(\partial\sigma) = \partial(\sigma \times \{0\})$ by the definition of I_0 , so to finish the proof, we only need to verify that $S_*I_1\partial(\sigma) = -d \cdot \partial(\sigma \times \{1\})$ after we replace the vertices.

For this, we separately consider the value of f on each d -simplex in its support. Recall the notation $\sigma_i := \sigma \setminus \{v_i\}$. We can expand this chain using the definition of the boundary map as

$$S_*I_1(\partial\sigma) = \sum_{i=0}^d (-1)^i S_*I_1(\sigma_i).$$

Let us now investigate the terms $S_*I_1(\sigma_i)$.

Recall that the operator b_{σ_i} assigns different signs to a simplex depending on where v_{σ_i} is in the order of the simplex's vertices, so assume the vertex v_{σ_i} is ordered between $(v_{i-1}, 1)$ and $(v_{i+1}, 1)$ in the ordering of the simplices. In this case, v_{σ_i} is the i th vertex in the simplex $I_1(\sigma_{ij}) \cup \{v_{\sigma_i}\}$ for $j < i$ and in the $(i+1)$ st position for $j > i$. Therefore

$$\begin{aligned} S_*I_1(\sigma_i) &= b_{\sigma_i} \partial I_1(\sigma_i) \\ &= b_{\sigma_i} \left(\sum_{j=0}^{i-1} (-1)^j I_1(\sigma_{ij}) + \sum_{j=i+1}^d (-1)^{j-1} I_1(\sigma_{ij}) \right) \\ &= \sum_{j=0}^{i-1} (-1)^j (-1)^{i-1} I_1(\sigma_{ij}) \cup \{v_{\sigma_i}\} + \sum_{j=i+1}^d (-1)^{j-1} (-1)^i I_1(\sigma_{ij}) \cup \{v_{\sigma_i}\} \\ &= (-1)^{i-1} \left[\sum_{j=0}^{i-1} (-1)^j I_1(\sigma_{ij}) \cup \{v_{\sigma_i}\} + \sum_{j=i+1}^d (-1)^j I_1(\sigma_{ij}) \cup \{v_{\sigma_i}\} \right] \end{aligned}$$

Therefore, when we replace the vertex v_{σ_i} with $(v_i, 1)$, we find that

$$S_*I_1(\partial\{v_0, \dots, v_d\}) = \sum_{i=0}^d (-1) \left[\sum_{j=0}^{i-1} (-1)^j I_1(\sigma_j) + \sum_{j=i+1}^d (-1)^j I_1(\sigma_j) \right]$$

Note that each term $I_1(\sigma_j)$ appears d times in this sum, once for each $i \neq j$. Moreover, each time the term appears, it has sign $(-1)^j$. We conclude that $S_*I_1(\partial\{v_0, \dots, v_d\}) = -d \cdot \partial(\sigma \times \{1\})$, as claimed

We now analyze the size of f . We know that $f = SP_*(\partial\sigma)$. We can analyze this sum linearly as

$$\begin{aligned} f &= SP_*(\partial\sigma) = \sum_{i=0}^d (-1)^i SP_*(\sigma_i) \\ &= \sum_{i=0}^d (-1)^i (-b_{\sigma_i} I_0(\sigma_i) - b_{\sigma_i} P_* \partial(\sigma_i)) \end{aligned}$$

By the first term in the sum, the chain f assigns value ± 1 to each d -simplex $(\sigma_i \times \{0\}) \cup \{v_i\}$. Likewise, were we to expand the second sum, we would find that f assigns value ± 1 to all simplices of the form $\sigma_{ij}^k \cup \{(v_i, 1)\}$. By Part 1 of the lemma, we know there are $\Theta((d+1)^3)$ such simplices. We conclude that the squared norm $\|f\|^2$ is $\Omega((d+1)^3)$. \square

F Collapsibility of the complex.

In this section, we will prove that the complexes \mathcal{B}_d^n and \mathcal{P}_d^n from Section 5.2 collapse to $(d-1)$ -dimensional subcomplexes. This will imply that $\ker \partial_d[\mathcal{B}_d^n] = 0$ and $\ker \partial_d[\mathcal{P}_d^n] = 0$.

F.1 Preliminaries.

In this section, we introduce the necessary background on collapsibility.

A $(d-1)$ -simplex τ is a **face** of a d -simplex σ iff $\tau \subset \sigma$. The pair (σ, τ) is a **collapse pair** in K iff (i) τ is a face of σ , and (ii) τ is not the face of any other simplex in K . The complex K **collapses** to the complex $K \setminus \{\sigma, \tau\}$ if (σ, τ) is a collapse pair in K . More generally, a complex K collapses into a complex K'' if there exists a complex K' such that K collapses to K' and K' collapses to K'' . A complex K is **collapsible** if it collapses into a single vertex.

By the inductive definition of collapsibility, whenever a complex K collapse into a complex L , there exists a sequence of complexes $K = K_0 \supset K_1 \supset \dots \supset K_k = L$ such that for any $0 < i \leq k$, $K_i = K_{i-1} \setminus \{\sigma_i, \tau_i\}$ where (σ_i, τ_i) is a collapse pair in K_{i-1} . This sequence is called a **collapsing sequence**.

We now give two facts about consequences of collapsibility we use in this paper. Lemma F.1 is a standard fact about collapsibility we state without proof. Lemma F.2 is a non-standard fact about collapsibility, but one that will come as no surprise to those familiar with collapsibility.

Lemma F.1. *Let $L \subset K$ be simplicial complexes such that K collapses to L . Then L is a deformation retract of K . Consequently, L and K have isomorphic homology groups in all dimensions.*

Lemma F.2. *Let $L \subset K$ be simplicial complexes such that K collapses to L . Let γ be a $(d-1)$ -cycle in L . Then γ is null-homologous in L if and only if γ is null-homologous in K .*

Proof of Lemma F.2. Suppose that γ is null-homologous in L , and let f_L be a d -chain in $C_d(L)$ such that $\partial f_L = \gamma$. We can extend f to a d -chain f_K in K by setting $f_K(\sigma) = 0$ for any d -simplex $\sigma \in K_d \setminus L_d$. It is straightforward to see that $\partial f_K = \gamma$ by the linearity of the boundary map, so γ is null-homologous in K .

Conversely, suppose that γ is null-homologous in K , and let f_K be a d -chain in $C_d(K)$ such that $\partial f_K = \gamma$. Consider a collapsing sequence $K = K_0 \supset K_1 \supset \dots \supset K_k = L$ with collapses pairs (σ_i, τ_i) . We will prove by induction on i that there is a chain $f_{K_i} \in C_d(K_i)$ for $0 \leq i \leq k$ such that $\partial f_{K_i} = \gamma$. This will imply that γ is null-homologous in L .

For the base case of $i = 0$, the chain $f_{K_0} = f_K$ satisfies the claim by assumption. Now suppose there is a chain $f_{K_i} \in C_d(K_i)$ such that $\partial f_{K_i} = \gamma$. Consider the collapse pair (σ_i, τ_i) . In the case that neither σ_i nor τ_i are d -simplices, then K_i and K_{i+1} have the same set of d -simplices, so f_{K_i} is a valid d -chain in $C_d(K_{i+1})$ and we set $f_{K_{i+1}} = f_{K_i}$. In the case that σ_i is a d -simplex, then τ_i is a $(d-1)$ -simplex. As $\tau_i \notin L_{d-1}$, then $\gamma(\tau_i) = 0$. This implies that it must be the case that $f_{K_i}(\sigma_i) = 0$. If not, then $\partial f_{K_i}(\sigma_i) = \pm f_{K_i}(\sigma_i)$ (a contradiction) as the only d -simplex incident to τ_i is σ_i . Therefore, we can set $f_{K_{i+1}} = f_{K_i}$. Finally, in the case that τ_i is a d -simplex, then σ_i is a $(d+1)$ -simplex. In this case, we set $f_{K_{i+1}} = f_{K_i} - (\partial\sigma_i(\tau_i)) \cdot (f_{K_i}(\tau_i)) \cdot \partial\sigma_i$. (Note that in this expression, $\partial\sigma_i(\tau_i) = \pm 1$, $f_{K_i}(\tau_i)$ is a scalar, and $\partial\sigma_i$ is a d -chain.) We need to verify two things about $f_{K_{i+1}}$: (1) $\partial f_{K_{i+1}} = \gamma$ and (2) $f_{K_{i+1}}(\tau_i) = 0$. Condition (1) is easy to verify as

$$\begin{aligned} \partial f_{K_{i+1}} &= \partial f_{K_i} - (\partial\sigma_i(\tau_i)) \cdot (f_{K_i}(\tau_i)) \cdot \partial\sigma_i \\ &= \partial(f_{K_i}) && \text{(as } \partial\sigma_i = 0) \\ &= \gamma && \text{(Induction Hypothesis)} \end{aligned}$$

Condition (2) is also straightforward to verify.

$$\begin{aligned} f_{K_{i+1}}(\tau_i) &= f_{K_i}(\tau_i) - (\partial\sigma_i(\tau_i)) \cdot (f_{K_i}(\tau_i)) \cdot \partial\sigma_i(\tau_i) \\ &= f_{K_i}(\tau_i) - f_{K_i}(\tau_i) && \text{(as } (\partial\sigma_i(\tau_i))^2 = 1) \\ &= 0 \end{aligned}$$

Therefore, $f_{K_{i+1}}$ is a d -chain in $C_d(K_{i+1})$ with boundary γ .

In all three cases, we can find a d -chain $f_{K_{i+1}}$ such that $\partial f_{K_{i+1}} = \gamma$. This proves that γ is null-homologous in L . \square

F.2 The main lemma.

In this section, we will prove the following lemmas about the complexes \mathcal{B}_d^n and \mathcal{P}_d^n from Section 5.2.

Lemma F.3. *The simplicial complex \mathcal{B}_d^n collapses to a $(d-1)$ -dimensional subcomplex.*

Lemma F.4. *The simplicial complex \mathcal{P}_d^n collapses to a $(d-1)$ -dimensional subcomplex.*

For this, we need to prove an auxiliary lemma about the building block from Section E.

Lemma F.5. *There is a collapsing sequence that collapses all d -simplices of B_d . Furthermore, no $(d-1)$ -simplices that are subsets of $\sigma \times \{1\}$ are involved in the collapsing sequence.*

Proof. This proof relies on notation introduced in the proof of Lemma 5.10 Part 1. Recall that the d -simplices of B_d are of the form $\{\sigma_j^k : 0 \leq k < d, j \neq k\}$, $\{\sigma_{ij}^k \cup \{(v_i, 1)\} : 0 \leq k \leq d, j \neq k, i < k\}$,

and $\{(\sigma_i \times \{0\}) \cup \{(v_i, 1)\} : 0 \leq i \leq d\}$. We will collapse these d -simplices using the $(d-1)$ -simplices of the form $\sigma_j^k \setminus \{(v_k, 1)\}$, $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$, and $\sigma_i \times \{0\}$. To define an appropriate collapsing sequence, we need to know which d -simplices are incident to which $(d-1)$ -simplices.

The $(d-1)$ -simplices $\sigma_i \times \{0\}$ are incident to exactly one d -simplex: $(\sigma_i \times \{0\}) \cup \{(v_i, 1)\}$. We can see this as any other d -simplex will have two vertices with second coordinate 1.

The simplices $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$ are incident to exactly two d -simplices. It is straightforward to verify that $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$ is incident to $\sigma_{ij}^k \cup \{(v_i, 1)\}$ and $\sigma_{ij}^{k+1} \cup \{(v_i, 1)\}$ for $j \neq k, k+1$ and $k < d$; we can see this as $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\} = \sigma_{ij}^{k+1} \cup \{(v_i, 1)\} \setminus \{(v_{k+1}, 0)\}$. For the case of special case of $j = k+1$, we can see that $\sigma_{i(k+1)}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\} = \sigma_{i(k+1)}^{k+2} \cup \{(v_i, 1)\} \setminus \{(v_{k+2}, 0)\}$, so $\sigma_{i(k+1)}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$ is instead incident to $\sigma_{i(k+1)}^k \cup \{(v_i, 1)\}$ and $\sigma_{i(k+1)}^{k+2} \cup \{(v_i, 1)\}$. Finally, when $k = d$, the simplices $\sigma_{ij}^d \cup \{(v_i, 1)\} \setminus \{(v_d, 1)\}$ are incident to $\sigma_{ij}^d \cup \{(v_i, 1)\}$ and $(\sigma_i \times \{0\}) \cup \{(v_i, 1)\}$.

Lastly, the simplex $\sigma_j^k \setminus \{(v_k, 1)\}$ is incident to exactly three d -simplices. We see that $\sigma_j^k \setminus \{(v_k, 1)\} = \sigma_j^{k+1} \setminus \{(v_{k+1}, 0)\} = \sigma_{(k+1)j}^{k+2} \cup \{(v_{k+1}, 1)\} \setminus \{(v_{k+2}, 0)\}$ when $k \neq j+1$ and $k < d-1$, and $\sigma_{k+1}^k \setminus \{(v_k, 1)\} = \sigma_k^{k+2} \setminus \{(v_{k+2}, 0)\} = \sigma_{(k+2)j}^{k+3} \cup \{(v_{k+2}, 1)\} \setminus \{(v_{k+3}, 0)\}$ when $k = j+1$ and $k < d-1$. Finally, we consider the case of $k = d-1$. By the construction of B_d , we know that $\sigma_j^{d-1} = \sigma_{ij}^{d-1}$ for some i . We know that $i > d-1$, so we conclude that $i = d$. This implies $\sigma_j^{d-1} \setminus \{(v_{d-1}, 1)\} = \sigma_{dj}^{d-1} \cup \{(v_d, 1)\} \setminus \{(v_{d-1}, 1)\}$ is incident to the two d -simplices: σ_j^{d-1} and $(\sigma_d \times \{0\}) \cup \{(v_d, 1)\}$.

We now describe a collapsing sequence for B . We know that $\sigma_i \times \{0\}$ is only incident to $\sigma_i \times \{0\} \cup \{(v_i, 1)\}$. We therefore collapse each simplex $\sigma_i \times \{0\}$ onto $\sigma_i \times \{0\} \cup \{(v_i, 1)\}$.

Now fix $0 \leq i \neq j \leq d$. For k starting at d and iterating backwards to i , we can collapse the simplices $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$ into the simplex $\sigma_{ij}^k \cup \{(v_i, 1)\}$; this collapse is valid as the other simplex incident to $\sigma_{ij}^k \cup \{(v_i, 1)\} \setminus \{(v_k, 1)\}$ was collapsed in a previous iteration. These collapses remove all simplices of the form $\sigma_{ij}^k \cup \{(v_i, 1)\}$.

Now fix an integer j . For k starting at $d-1$ and iterating backwards to 0 (and skipping j), we will collapse $\sigma_j^k \setminus \{(v_j, 1)\}$ into σ_j^k . Initially, this is a valid collapse as the simplex $\sigma_j^{d-1} \setminus \{(v_{d-1}, 1)\}$ is only incident to σ_j^{d-1} and $(\sigma_d \times \{0\}) \cup \{(v_d, 1)\}$, and the second simplex has already been collapsed. For the other values of k , the only simplices incident to $\sigma_j^k \setminus \{(v_k, 1)\}$ were either removed in the previous iteration or in the series of collapses from the previous paragraph.

Finally, no $(d-1)$ -simplex in $K \times \{1\}$ was collapsed, as all of the simplices that were collapsed have at least one vertex with second coordinate 0. \square

We can now prove Lemmas F.3 and F.4.

Proof of Lemma F.3. Recall that \mathcal{B}_d^n is constructed by gluing together a d -simplex, denoted B_d^0 , and n copies of B_d , denoted B_d^i for $1 \leq i \leq n$. We will therefore prove that we can collapse each of the d -simplices of B_d^i for $n \geq i \geq 1$ by induction on i in reverse order.

For the base case of $i = n$, we know we can collapse each of the d -simplices of B_d^n . By Lemma F.5, we know we can collapse all d -simplices B_d , and the collapsing sequence does not collapse any $(d-1)$ -simplex on the vertices $\sigma \times \{1\}$. As the only $(d-1)$ -simplices in B_d^n that are incident to other d -simplices are simplices in $\sigma \times \{n-1\}$, then we know we can collapse all d -simplices of B_d^n .

Inductively, we know that none of the $(d-1)$ -simplices needed to collapse B_n^i have previously been collapsed, as the only $(d-1)$ -simplices of B_n^i that are $(d-1)$ -simplices of B_n^j for $j > i$ are subsets of $\sigma \times \{1\}$ in B_n^{i+1} , which we know have not been collapsed. Finally, we can collapse the unique d -simplex in B_d^0 , as none of the $(d-1)$ -simplices of B_d^0 have previously been collapsed. \square

Proof of Lemma F.4. The proof of this theorem is nearly identical to the proof of Lemma F.3 above, so we exclude details and instead sketch the proof. Recall that \mathcal{P}_d^n is constructed by gluing together n copies of B_d , denoted B_d^i for $1 \leq i \leq n$. As in the previous proof, we can prove that we collapse each of the d -simplices of B_d^i for $1 \leq i \leq n$ by induction. This works because, inductively, the only $(d-1)$ -simplices in B_d^i that are incident to d -simplices outside of B_d^i are subsets of $\sigma \times \{1\}$. As in the previous proof, this is sufficient to prove that we can collapse the d -simplices of B_d^i . \square