AN ABSTRACT OF THE DISSERTATION OF

William Maxwell for the degree of Doctor of Philosophy in Computer Science presented on September 1, 2021.

Title: Algorithmic Problems on Simplicial Complexes

Abstract approved: ____

Amir Nayyeri

Simplicial complexes may be viewed as a high dimensional generalization of graphs; a graph is a 1-dimensional simplicial complex. Despite their similarities, generalizing graph theoretic problems to simplicial complexes can be quite challenging; many combinatorial approaches to graphs do not have natural generalizations to simplicial complexes. In order to make graph theoretic problems generalizable to higher dimensions we translate them to the language of simplicial homology which is an algebraic tool for reasoning about high dimensional cycles. In this thesis we use simplicial homology to generalize the following problems to simplicial complexes: shortest *st*-path, max-flow, min-cut, and *st*-connectivity.

The shortest path problem in graphs generalizes to the minimum bounding chain problem in sipmlicial complexes. For this problem we prove APX-hardness, and assuming the unique games conjecture that no polynomial-time constant factor approximation exists. Our hardness results hold even in the special case when the complex is a manifold embedded in $\mathbb{R}^3$. For d-dimensional simplicial complexes embedded in $\mathbb{R}^{d+1}$ we design both fixed-parameter tractable and $O(\sqrt{\log n})$-approximation algorithms. We prove nearly identical results for the closeley related minimum homologous chain problem.

We then generalize the notions of flows and cuts in graphs to simplicial complexes and show that a generalization of the max-flow/min-cut theorem holds in dimensions $d \geq 1$. We show that max-flows and min-cuts can be found with linear programming; however, unlike graphs, finding integral solutions is NP-hard. We investigate a combinatorial generalization of min-cut and show that it leads to an NP-hard problem. Finally, we provide a variant of

the Ford-Fulkerson algorithm that halts in dimensions $d > 1$ despite running in exponential time.

The graph theoretic notions of effective resistance and capacitance have natural topological definitions arising from the high dimensional definitions of flows and cuts. In graphs these quantities are associated with a pair of edges, and in simplicial complexes they are associated with a null-homologous cycle. In graphs the effective resistance and capacitance are polynomial with respect to the size of the graph. However, the quantities may be exponential in simplicial complexes. This arises from the existence of torsion in the relative homology groups of the complex. We prove upper bounds on resistance and capacitance that are polynomial in the size of the complex and the size of the torsion subgroups of the relative homology groups. We note that the size of the torsion subgroup may be exponential with respect to the size of the complex. Finally, we provide a quantum algorithm that decides if a cycle is null-homologous in a simplicial complex. This problem generalizes *st*-connectivity in graphs. The query complexity of the quantum algorithm is parameterized by the cycle's effective resistance and capacitance.

Algorithmic Problems on Simplicial Complexes

by

William Maxwell

A DISSERTATION

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented September 1, 2021
Commencement June 2022

Doctor of Philosophy dissertation of William Maxwell presented on September 1, 2021.

APPROVED:

Major Professor, representing Computer Science

Head of the School of Electrical Engineering and Computer Science

Dean of the Graduate School

I understand that my dissertation will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my dissertation to any reader upon request.

William Maxwell, Author

ACKNOWLEDGEMENTS

Firstly, I would like to thank my amazing advisors Amir and Cora. This dissertation wouldn't be possible without them. I would also like to thank my friends Jeff, Ameya, Colin, and Mitchell for the many interesting conversations that kept me motivated to pursue my research. Finally, I'd like to thank Cory, Cloud, Ryan, and the rest of the Downward Dog crew for making sure I didn't work too hard.

TABLE OF CONTENTS

Page

TABLE OF CONTENTS (Continued)

Page

LIST OF FIGURES

Chapter 1: Introduction

Simplicial complexes can be thought of as a collection of simple polytopes of varying dimension, called simplices, attached to one another in a nice way. The d-dimensional polytopes are called d-simplices; the 0-simplices are vertices, 1-simplices are edges, 2-simplices are triangles, 3-simplices are tetrahedron, and so on. Within a simplicial complex the simplices are attached to one another such that the intersection of any two simplices is a simplex of lower dimensionality. Simplicial complexes serve as a generalization of graphs; a graph may be viewed as a 1-dimensional simplicial complex.

Many graph theoretic problems have interesting generalizations to simplicial complexes. However, the typical combinatorial arguments used on graphs do not generalize to simplicial complexes and one must use tools from algebraic topology to solve these problems. As a result many polynomial-time solvable problems on graphs become NP-hard for simplicial complexes of dimension $d \geq 2$, and some problems even become undecidable.

Homology, an algebraic tool describing high dimensional cycles, may be used to translate graph theoretic problems to simplicial complexes of higher dimensions. Homology is an algebraic construction arising from the boundary operator of a simplicial complex. We will give formal definitions in Chapter 2, but for now we will provide intuition to motivate the concept. Vertices have no boundary, the boundary of an edge is a pair of vertices, the boundary of a triangle is three edges forming a cycle, the boundary of a tetrahedron is its collection of faces. In general the boundary of a d-dimensional simplex is homeomorphic to a $(d-1)$-dimensional sphere. It is important to note that once we consider formal sums of simplices the simplicial boundary operator may no longer return a geometric boundary. An interesting observation is that the edge is the only simplex whose boundary is disconnected. The consequence is that some combinatorial techniques on graphs such as divide and conquer via vertex separators do not have natural generalizations to simplicial homology.

The boundary operator extends linearly to formal sums of simplices with coefficients over any abelian group. A cycle is a collection of simplices whose sum lies in the kernel of the boundary operator. In the case of a simple cycle in a graph each vertex appears in the boundary of two edges, and as a result the second copy of the vertex cancels out the first. Every boundary is a cycle; this is the key property that makes the theory of simplicial homology work. The d-dimensional homology group of a simplicial complex is a collection of equivalence classes where each class contains cycles (elements of the kernel of the d-dimensional boundary operator) that are not boundaries (elements of the image of the $(d+1)$-dimensional boundary operator). The homology classes are the "interesting" cycles in a simplicial complex. The d-dimensional homology classes represent the d-dimensional "holes" in the simplicial complex, the cycles in the homology class can be thought to go "around" the hole. The d-dimensional homology group is the abelian group generated by the homology classes. If a graph theoretic problem can be expressed in terms of boundaries and cycles there is a good chance that it has a natural generalization to simplicial homology. For example, the shortest st-path problem asks to find the smallest collection of edges whose boundary is s and t. This generalizes to the minimum bounding chain problem, which given a $(d-1)$-dimensional cycle γ asks to find the smallest collection of d-simplices whose boundary is γ. We will discuss this problem in detail in Chapter 4.

Using homology, several interesting generalizations of graph theoretic problems can be solved in simplicial complexes. Delfinado and Edelsbrunner show that the problem of computing the dth Betti number β_d can be solved in polynomial time [27]. The Betti number β_d is the rank of the d-dimensional homology group. The rank is the number of linearly independent generators of the group. In a graph β_0 is the number of connected components and β_1 is the dimension of the cycle space. The homology groups of a simplicial complex are a generalization of the cycle space of a graph; the first homology group of a graph is its cycle space. In general computing some basis for the d-dimensional homology group of a simplicial complex can be done in polynomial time using standard matrix reduction techniques [37, Section 4.2]. However, for $d > 1$, Chen and Freedman show computing a minimum weight basis for the d-dimensional homology group is NP-hard [23]. Chen and Freedman actually prove a stronger result showing that it is NP-hard to approximate a minimum homology basis within any constant factor. Dey, Li, and Wang show that computing a minimum basis of the first homology group can be done in polynomial time [32]; Erickson and Whittlesey improve the running time when the simplicial complex is a surface [41]. These results rely on polynomial-time algorithms for solving the minimum cycle basis problems in graphs originally due to Horton [54] and improved up by Kavitha et al [62]. Their techniques were used by Borradaile et al. to compute minimum cycle and homology bases of surface embedded graphs [11]. For complexes embedded in $\mathbb{R}^3$, Dey provides an algorithm computing a basis for the first homology group in near linear time [28].

Computing a minimum spanning tree in a graph can be restated as computing a minimum subgraph whose first homology group is trivial. The generalization for a d-dimensional simplicial complex is to find a subcomplex whose dth homology group is trivial; we call such complexes acyclic. Skraba, Thoppe, and Yogeshwran show acyclic complexes can be computed in polynomial time by adapting Kruskal's algorithm [90].

Cycles can be added to one another as formal sums with coefficients coming from any abelian group. In this thesis we will primarily use coefficients from the fields $\mathbb{Z}_2$ and $\mathbb{R}$, and occasionally from the ring $\mathbb{Z}$. Choosing coefficients over a field lets us treat the homology groups as vector spaces allowing us to utilize tools from linear algebra. The different choices of coefficients result in different interpretations of the resulting homology theory. For graphs, a cycle over $\mathbb{Z}_2$ is an Eulerian subgraph, and a cycle over $\mathbb{R}$ is a circulation. Over $\mathbb{Z}_2$, a graph is viewed as undirected, however over $\mathbb{R}$ we consider directed graphs since since the sign on a real number allows us to distinguish an orientation on an edge. In Chapter 4, we consider the minimum bounding chain problem with coefficients over $\mathbb{Z}_2$, however in Chapter 5, we formulate the max-flow min-cut problem with coefficients over $\mathbb{R}$. For the minimum bounding chain problem, we compute the size of a chain to be the number of simplices contained in it. This makes $\mathbb{Z}_2$ a natural choice since chains over $\mathbb{Z}_2$ are equivalent to their underlying sets and addition over $\mathbb{Z}_2$ corresponds to the symmetric difference of the underlying sets.

For finite simplicial complexes the boundary operator can be expressed as a matrix, so we are able to reduce problems about homology to the domain of linear algebra. As a result it is fairly easy to obtain polynomial-time algorithms running in time $O(n^\omega)$ where $\omega \leq 2.3728596$ [5] is the exponent for matrix multiplication. While a superquadradic running time may still be fairly slow in practice for large simplicial complexes, assuming more structure on the complex may allow for quadratic or even faster running times. In this thesis, we obtain faster running times on complexes which admit dual graphs similar to planar duality in graphs. We can then use a high dimensional analog of cycle/cut duality to obtain faster algorithms by operating on the dual graph.

One natural class of complexes which we study in this thesis are the d-dimensional complexes which admit an embedding into $\mathbb{R}^{d+1}$. These complexes generalize planar graphs. As a consequence of the Alexander duality theorem there is a dual graph associated with every embedded complex. An embedded complex partitions Euclidean space into $\beta_d + 1$ connected components. One of these connected components is unbounded and corresponds to the unbounded face of a planar graph. The dual graph is the graph whose vertices are the connected components of the partition where two vertices are adjacent if and only if their corresponding connected components intersect along their boundary. In the case of graphs β_1 is the dimension of the cycle space of the graph, and the situation reduces to the fact that a planar graph partitions the plane into $\beta_1 + 1$ connected components. This is a consequence of the Jordan curve theorem, and Alexander duality may be viewed as a high dimensional generalization of the Jordan curve theorem.

Deciding planarity, which takes linear time [53], generalizes to deciding whether or not a d-dimensional simplical complex admits an embedding in $\mathbb{R}^{d+1}$. When $d = 1$ we have the problem of deciding planarity. Mesmay et al. show this problem is NP-hard for $d = 2$ [26]. Matoušek et al. show the problem to be NP-hard for $d = 3$ and to be undecidable for $d \geq 4$ [75]. This undecidability result comes from a result of Nabutovsky showing that for $d \geq 5$ determining whether or not a d-dimensional simplicial complex is homeomorphic to the d-dimensional sphere is undecidable [78]; Volodin, Kuznetsov, and Fomenko provide a more modern exposition of the result [95]. For $d = 1$ the sphere is a simple cycle and determining if a graph is a simple cycle is decidable in linear time. More generally, Chernavsky and Leksine show that determining if a d-dimensional simplicial complex is a manifold is undecidable [24]; 1-dimensional manifolds are simple cycles.

Another class of complexes that admit dual graphs are the weak pseudomanifolds. These are the complexes such that every $(d-1)$-simplex is incident to at most two d-simplices. In the dual graph the d-simplices are the vertices and the $(d-1)$-simplices are the edges. The dual graph of a weak pseudomanifold generalizes the dual of a surface embedded graph. Weak pseudomanifolds generalize manifolds and retain the properties needed to perform algorithms operating on the dual graph. Unlike manifolds, weak pseudomanifolds can be recognized in linear time by checking the degree of each $(d-1)$-simplex.

1.1 Contributions

In this thesis we investigate two types of problems. The first is the minimum bounding chain problem which serves as a generalization of the shortest paths problem, or more generally, the minimum T-join problem. Additionally, we study the more general problem of deciding whether or not a cycle is null-homologous. This generalizes the question of deciding whether or not a path exists between two vertices. The second type of problem is a generalization of the max-flow min-cut theorem. For both types of problems, we obtain both algorithmic and hardness result. The remainder of this section will be dedicated to summarizing the major contributions of this thesis. The definitions and assumptions necessary to state the results can be found either in the preliminaries in Chapter 2 or in the chapter containing the results.

Embeddings and Duality In this section we give a detailed description of Alexander duality and the resulting dual graph. The results in this section are not new, but it seems that they have yet to be written in this context. We show that Alexander duality leads to a higher dimensional analog of cycle/cut duality in planar graphs. Mac Lane's planarity criterion states that a graph is planar if and only if there exists a cycle basis such that every edge appears in at most two cycles [72]. Such a basis is called a 2-basis. When the graph is 2-connected and the 2-basis is expressed as a matrix; it forms the incident matrix of the dual graph. The cycles in the basis are the vertices, and the edges of the cycles are the edges of the dual graph. For embedded complexes, we show a similar result. If a d-dimensional simplicial complex admits an embedding in $\mathbb{R}^{d+1}$, then the boundaries of the connected components obtained from Alexander duality form a 2-basis for the $\mathbb{Z}_2$-homology group. We observe that the converse is not true; there exist non-embeddable complexes which admit a 2-basis for their $\mathbb{Z}_2$-homology. This implies that there exist non-embeddable complexes with a dual graph, so any algorithm operating on the dual graph can accept as input a class of complexes that is larger than the class of embeddable complexes. In general, deciding whether or not a complex is embeddable is undecidable. However, we show that there exists a polynomial-time algorithm to construct a 2-basis. Given a $\mathbb{Z}_2$-homology basis as input one can apply Tutte's algorithm for recognizing cographic matroids to obtain a dual graph if one exists. It turns out that a simplicial complex admits a dual graph if and only if its $\mathbb{Z}_2$-homology space forms a cographic matroid.

The Minimum Bounded and Homologous Chain Problems We investigate two related problems: the minimum bounding chain problem and the minimum homologous chain problem. Both problems are formulated with coefficients over $\mathbb{Z}_2$. Both problems take as input a d-dimensional simplicial complex $\mathcal{K}$. The minimum bounding chain problem asks: given a $(d-1)$-cycle γ find the smallest d-chain whose boundary is γ. The minimum homologous chain problem asks: given a $(d-1)$-chain γ find the smallest $(d-1)$-chain τ such that $\gamma \oplus \tau$ is in the image of the boundary operator; here $\oplus$ denotes addition over $\mathbb{Z}_2$. Chains over $\mathbb{Z}_2$ are equivalent to their underlying sets of simplices. A minimum chain is one that minimizes the cardinality of its underlying set of simplices. For both problems we prove APX-hardness and design both $O(\sqrt{\log n})$-approximation algorithms and fixed-parameter tractable algorithms operating on the dual graph of the complex. Our hardness results hold even when $\mathcal{K}$ is an orientable manifold embedded in $\mathbb{R}^3$. Assuming the unique games conjecture our hardness results are strengthened as no polynomial-time constant factor approximation can exist for either problem. Our algorithmic results operate on the dual graph of $\mathcal{K}$. In the case of the minimum bounding chain problem we assume that $\mathcal{K}$ is a d-complex embedded in $\mathbb{R}^{d+1}$ and in the case of the minimum homologous chain problem we assume that $\mathcal{K}$ is a $(d+1)$-dimensional weak pseudomanifold. Finally, we present a result originally due to Kirsanov and Gortler showing that under certain conditions the minimum bounding chain problem can be solved in polynomial time in embedded complexes [66]. Their algorithm reduces the problem to a minimum cut problem on the dual graph and is similar to finding the minimum st-path in a planar graph with s and t appearing on the same face. Our analysis is much simpler than the original and our presentation is combinatorial, making it more accessible to an audience of computer scientists.

Generalized Flows and Cuts We consider high-dimensional variants of the maximum flow and minimum cut problems in the setting of simplicial complexes and provide both algorithmic and hardness results. By viewing flows and cuts topologically in terms of the simplicial (co)boundary operator we can state these problems as linear programs and show that they are dual to one another. Unlike graphs, complexes with integral capacity constraints may have fractional max-flows. We show that computing a maximum integral flow is NP-hard. Moreover, we give a combinatorial definition of a simplicial cut that seems more natural in the context of optimization problems and show that computing such a cut is NP-hard. However, we provide conditions on the simplicial complex for when the cut found by the linear program is a combinatorial cut. For d-dimensional simplicial complexes embedded into $\mathbb{R}^{d+1}$ we provide algorithms operating on the dual graph: computing a maximum flow is dual to computing a shortest path and computing a minimum cut is dual to computing a minimum-cost circulation. Finally, we investigate the Ford-Fulkerson algorithm on simplicial complexes, prove its correctness for $d \geq 1$, and provide a heuristic which guarantees it to halt. In graphs with integral capacity constraints, Ford-Fulkerson is guaranteed to halt since all flows are integral. Since simplicial complexes may have fractional flows the fact that Ford-Fulkerson halts is not immediate.

Resistance and Capacitance Effective resistance and effective capacitance are quantities associated with a graph originally motivated by the study of electrical networks. Both quantities can be expressed in terms of homology which allows them to be generalized to higher dimensions. In graphs the effective resistance between two vertices s and t is the size of the smallest unit st-flow. Similarly, the effective capacitance between s and t is the size of the smallest st-cut. Here the word "size" refers to the ℓ_2-norm of the vectors representing the flow and cut. Using our high dimensional generalizations of flows and cuts we can extend the notions of effective resistance and capacitance to higher dimensions.

We provide upper bounds on the effective resistance in a simplicial complex. In graphs the effective resistance is $O(n)$ and the effective capacitance is $O(n^2)$ where n is the number of vertices. However, one can construct simplicial complexes with cycles whose effective resistance and capacitance is $\Theta(2^{2n})$. We show that the exponential upper bounds are purely a result of torsion in the $(d-1)$-dimensional relative homology groups of the complex. We provide an upper bound of $O(n^2t^2)$ on effective resistance and $O(n^3t^2)$ on effective capacitance. Here t denotes the cardinality of a torsion subgroup of a relative homology group of the complex whose number of elements is maximized. When the simplicial complex is relative torsion-free and the cycle under consideration is the boundary of a simplex we match the upper bounds known for graphs.

We provide a quantum algorithm for the following problem: given a d-dimensional simplicial complex $\mathcal{K}$, a $(d-1)$-dimensional null-homologous cycle γ, and a d-dimensional subcomplex $\mathcal{L} \subseteq \mathcal{K}$ decide whether or not γ is null-homologous in $\mathcal{L}$. Our algorithm is based on the span program model and is a generalization of the span program deciding st-connectivity devised by Belovs and Reinhardt [8]. Their algorithm is parameterized by the effective resistance and capacitance of the pair of vertices and has query complexity $O(n^{3/2})$. Our algorithm has query complexity $O\left(\sqrt{\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}\right)$ where $\mathcal{R}_{\max}(\gamma)$ and $\mathcal{C}_{\max}(\gamma)$ are the maximum effective resistance and capacitance of γ over all subcomplexes. Applying our upper bounds on effective resistance and capacitance this reduces to $O(n^{5/2}t^2)$. Under the assumptions that $\mathcal{K}$ is relative torsion-free, d is fixed, and that γ is the boundary of a d-simplex we match the query complexity of $O(n^{3/2})$. In higher dimensions the requirement that $\mathcal{K}$ is relative torsion-free is quite restrictive, however for $d = 2$ it is equivalent to forbidding a class of Möbius subcomplexes [30].

Chapter 2: Preliminaries

2.1 Simplicial Homology

We give a brief overview of some basic concepts from simplicial (co)homology that will be used throughout the paper. We recommend the book by Hatcher [52] for a more complete exposition. In this paper $\mathcal{K}$ will always be a finite d-dimensional simplicial complex which we now define. Given a finite set of vertices V we define a d-dimensional simplex, or ***d-simplex***, to be a subset of $d+1$ vertices of V. We define an ***abstract simplicial complex*** $\mathcal{K}$ to be a finite collection of simplices meeting the condition that for every $\sigma \in \mathcal{K}$ if $\tau \subset \sigma$ then $\tau \in \mathcal{K}$. Every algorithmic problem we consider will take an abstract simplicial complex as input, but sometimes we will assume the existence of a geometric realization of the complex. Hence, we will refer to 0-simplices as vertices, 1-simplices as edges, 2-simplices as triangles, and so on. We will formally define the geometric realization of a complex in the next chapter.

We call the subsets of a simplex the ***faces*** of the simplex. The dimension of $\mathcal{K}$ is dimension of its largest simplex. Further, we define an ***orientation*** on $\mathcal{K}$ by fixing a linear ordering on the vertices in V and treating simplices as ordered sets. An ***oriented simplex*** is a simplex along with a permutation of its vertices, and the orientation of the simplex is the parity of the permutation with respect to the fixed linear ordering on V. We call oriented simplices that agree with the linear ordering ***positively oriented*** otherwise we call them ***negatively oriented***.

We define a ***d-chain*** to be a formal sum of d-simplices with coefficients over a group G. In this thesis we will consider coefficients over $\mathbb{Z}, \mathbb{Z}_2$, and $\mathbb{R}$. By $\mathcal{K}_d$ we denote the ***d-skeleton*** of the simplicial complex $\mathcal{K}$ which is the set of all d-simplices in $\mathcal{K}$ and we denote its cardinality by n_d. Additionally, we may assign a weight function on the d-skeleton of $\mathcal{K}$ denoted $w \colon \mathcal{K}_d \to \mathbb{R}$. It will sometimes be convenient to view the weight function as a diagonal matrix W indexed by the d-simplices.

Given a simplicial complex $\mathcal{K}$ we define the dth ***chain group*** $C_d(\mathcal{K}, G)$ to be the free abelian group generated by $\mathcal{K}_d$ with coefficients in G. When G is a field, such as $\mathbb{Z}_2$ or $\mathbb{R}$, the chain group forms finite dimensional vector space since we assume $\mathcal{K}$ to be finite.

When the context is clear we will drop G from the notation and write the chain group as $C_d(\mathcal{K})$. Let $\sigma = [v_0, \dots, v_d]$ be a d-simplex with each v_i a vertex. We define the simplicial boundary operator $\partial_d[\mathcal{K}] \colon C_{d+1}(\mathcal{K}, G) \to C_d(\mathcal{K}, G)$ by

$$\partial_d[\mathcal{K}](\sigma) = \sum_{i=0}^{d} (-1)^i [v_0, \dots, \hat{v_i}, \dots, v_d],$$

where $[v_0, \dots, \hat{v_i}, \dots, v_d]$ denotes the $(d-1)$-simplex obtained by removing v_i from σ. Often we will shorten the notation to ∂_d or ∂ when the context is clear. The boundary operator extends linearly over G in the natural way. Since we assume $\mathcal{K}$ to be finite the chain group $C_d(\mathcal{K})$ is finite dimensional and we can represent the boundary operator as the $n_{d-1} \times n_d$ matrix such that the entry $\partial_{i,j}$ is equal to 1 or -1 (depending on the orientation) if the ith $(d-1)$-simplex is a face of the jth d-simplex, and 0 otherwise. With the boundary operator we define the ***chain complex*** on $\mathcal{K}$ to be the following sequence of groups connected by their boundary operators.

$$\dots \xrightarrow{\partial_{d+1}} C_d(\mathcal{K}) \xrightarrow{\partial_d} C_{d-1}(\mathcal{K}) \xrightarrow{\partial_{d-1}} \dots \xrightarrow{\partial_2} C_1(\mathcal{K}) \xrightarrow{\partial_1} C_0(\mathcal{K}) \to 0$$

The group of ***d-cycles*** $Z_d(\mathcal{K})$ is defined to be $\ker(\partial_d)$ and the group of ***d-boundaries*** $B_d(\mathcal{K})$ is defined to be $\mathrm{im}(\partial_{d+1})$. We call any d-cycle $\gamma \in B_d(\mathcal{K})$ a ***null-homologous*** cycle. Since $\partial_d \circ \partial_{d+1} = 0$ the quotient group $H_d(\mathcal{K}) \coloneqq Z_d(\mathcal{K})/B_d(\mathcal{K})$ is well-defined, and we call $H_d(\mathcal{K})$ the dth ***homology group*** of $\mathcal{K}$. By β_d we denote the dth Betti number of $\mathcal{K}$ which is defined to be $\beta_d = \mathrm{rank} H_d(\mathcal{K})$.

We obtain the dth ***cochain group*** by dualizing $C_d(\mathcal{K})$ in the following way. The dth cochain group $C^d(\mathcal{K})$ is defined to be the dual space of $C_d(\mathcal{K})$, that is, the space of all linear functions $f \colon C_d(\mathcal{K}) \to G$. We call f a ***cochain*** and we define the ***coboundary operator*** $\delta_d \colon C^{d-1}(\mathcal{K}) \to C^d(\mathcal{K})$ on cochains as the composition of functions $\delta_d f = f \circ \partial_d$. The ***cochain complex*** is the following sequence of cochain spaces.

$$\dots \xleftarrow{\delta_{d+1}} C^d(\mathcal{K}) \xleftarrow{\delta_d} C^{d-1}(\mathcal{K}) \xleftarrow{\delta_{d-1}} \dots \xleftarrow{\delta_2} C^1(\mathcal{K}) \xleftarrow{\delta_1} C^0(\mathcal{K}) \leftarrow 0$$

The group of ***d-cocycles*** $Z^d(\mathcal{K})$ is defined to be $\ker(\delta_{d+1})$ and the group of ***d-coboundaries*** $B^d(\mathcal{K})$ is defined to be $\mathrm{im}(\delta_d)$. Again, since $\delta_{d+1} \circ \delta_d = 0$ the quotient group $H^d(\mathcal{K}) \coloneqq Z^d(\mathcal{K})/B^d(\mathcal{K})$ is well-defined and we call $H^d(\mathcal{K})$ the dth ***cohomology group*** of $\mathcal{K}$.

When G is a field $C_d(\mathcal{K})$ and $C^d(\mathcal{K})$ are finite dimensional vector spaces we have the isomorphisms $C_d(\mathcal{K}) \cong C^d(\mathcal{K})$ and $H_d(\mathcal{K}) \cong H^d(\mathcal{K})$. The first isomorphism is the fact from linear algebra that a finite dimensional vector space is isomorphic to its dual, and the second is derived from the universal coefficient theorem. We can represent a cochain $f \in C^d(\mathcal{K}, G)$ as a row vector $f = [g_1, g_2, \dots, g_{n_d}]$ where $f(\sigma_i) = g_i$ and then the composition $f \circ \partial$ is obtained by the matrix multiplication $f\partial$ which is another row vector. Hence, for any chain Σ we have that $(f \circ \partial)(\Sigma)$ is represented by the inner product $\langle \partial^T f^T, \Sigma \rangle$. Under the isomorphism $C_d(\mathcal{K}, G) \cong C^d(\mathcal{K}, G)$ it follows that we may represent the coboundary operator δ_d as the transpose of the boundary operator: $\delta_d = \partial_d^T$.

We will view d-(co)chains as both d-dimensional vectors and as linear functions $C_d(\mathcal{K}) \to \mathbb{R}$ whenever it is convenient to do so. However, we will refer to flows as d-chains and cuts as d-cochains unless explicitly stated otherwise. We will often want to talk about the underlying set of simplices of the (co)chain and refer to this set as the ***support*** of the (co)chain; the support of a chain $\sigma = \sum \alpha_i \sigma_i$ is defined as the set $\mathrm{supp}(\sigma) = \{\sigma_i \in \mathcal{K}^d \mid \alpha_i \neq 0\}$.

We will need to define the notion of ***relative homology*** in order to cite known results about the boundary matrix of a simplicial complex. Let $\mathcal{K}_0 \subseteq \mathcal{K}$ be a subcomplex of $\mathcal{K}$. We call the quotient group $C_d(\mathcal{K}, \mathcal{K}_0) = C_d(\mathcal{K})/C_d(\mathcal{K}_0)$ the of d-dimensional ***relative chain group*** and call an element of it a ***relative chain***. There is an induced mapping $\partial_d[\mathcal{K}, \mathcal{K}_0]\colon C_d(\mathcal{K}, \mathcal{K}_0) \to C_{d-1}(\mathcal{K}, \mathcal{K}_0)$. From the induced mapping we define the groups of ***relative*** d***-cycles*** $Z_d(\mathcal{K}, \mathcal{K}_0)$, ***relative*** d***-boundaries*** $B_d(\mathcal{K}, \mathcal{K}_0)$, and ***relative*** d***-dimensional homology*** $H_d(\mathcal{K}, \mathcal{K}_0)$ as $Z_d(\mathcal{K}, \mathcal{K}_0) \coloneqq \ker \partial_d[\mathcal{K}, \mathcal{K}_0]$, $B_d(\mathcal{K}, \mathcal{K}_0) \coloneqq \mathrm{im}\, \partial_d[\mathcal{K}, \mathcal{K}_0]$, and $H_d(\mathcal{K}, \mathcal{K}_0) \coloneqq Z_d(\mathcal{K}, \mathcal{K}_0)/B_d(\mathcal{K}, \mathcal{K}_0)$ respectively. Further, let $\mathcal{L} \subseteq \mathcal{K}$ be a ***pure*** d-dimensional subcomplex; that is, every $(d-1)$-simplex in $\mathcal{L}$ is incident to some d-simplex in $\mathcal{L}$. Let $\mathcal{L}_0 \subseteq \mathcal{L}$ be a pure $(d-1)$-dimensional subcomplex of $\mathcal{L}$. The induced map on relative homology $\partial_d[\mathcal{L}, \mathcal{L}_0]\colon C_d(\mathcal{L}, \mathcal{L}_0) \to C_{d-1}(\mathcal{L}, \mathcal{L}_0)$ has a natural matrix representation. We construct the matrix $\partial_d[\mathcal{L}, \mathcal{L}_0]$ by starting with ∂_d and including the columns corresponding to d-simplices in $\mathcal{L}$ while excluding the rows corresponding to $(d-1)$-simplices in $\mathcal{L}_0$.

Throughout this paper we utilize discrete Hodge theory and recommend the survey by Lim [71] as an introduction to the topic. In particular, we use the Hodge decomposition which can be stated as a result on real valued matrices A and B satisfying $AB = 0$.

Theorem 1 (Hodge decomposition [71])**.** *Let $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$ be matrices satisfying $AB = 0$. We can decompose $\mathbb{R}^n$ into the orthogonal direct sum,*

$$\mathbb{R}^n = \mathrm{im}(A^T) \oplus \ker(A^T A + BB^T) \oplus \mathrm{im}(B).$$

Representing the (co)boundary operators as matrices and setting $A = \partial_d$ and $B = \partial_{d+1}$ yields the Hodge decomposition for simplicial complexes. The middle term of the direct sum becomes $\ker(\delta_{d+1}\partial_{d+1} + \partial_d\delta_d)$. The linear operator $\delta_{d+1}\partial_{d+1} + \partial_d\delta_d$ is known as the combinatorial Laplacian of $\mathcal{K}$ which is a generalization of the graph Laplacian. Moreover, it can be shown that $\ker(\delta_{d+1}\partial_{d+1} + \partial_d\delta_d) \cong H_d(\mathcal{K}, \mathbb{R})$. We now state the Hodge decomposition on simplicial complexes as the following isomorphism:

$$C_d(\mathcal{K}, \mathbb{R}) \cong \mathrm{im}(\delta_d) \oplus H_d(\mathcal{K}, \mathbb{R}) \oplus \mathrm{im}(\partial_{d+1}).$$

2.2 Total Unimodularity

When working with coefficients over $\mathbb{Z}$ the homology groups are not vector spaces but finitely generated abelian groups. The fundamental theorem of finitely generated abelian groups gives us the decomposition $H_d(\mathcal{K}, \mathbb{Z}) \cong \mathbb{Z}^k \oplus \mathbb{Z}_{t_1} \oplus \cdots \oplus \mathbb{Z}_{t_n}$ for some $k \in \mathbb{N}$. We call the subgroup $\mathbb{Z}_{t_1} \oplus \cdots \oplus \mathbb{Z}_{t_n}$ the ***torsion subgroup*** of $H_d(\mathcal{K}, \mathbb{Z})$ and when this subgroup is trivial we call the complex ***torsion-free***. We say $\mathcal{K}$ is relative-torsion free in dimension d if $H_d(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$ is torsion-free for all subcomplexes $\mathcal{L}$ and $\mathcal{L}_0$. There exist complexes that are torsion-free but are not relatively torsion-free; see [30] for examples.

Let A be an integral matrix; we say that A is ***totally unimodular*** if every square submatrix A' of A has $\det(A') \in \{-1, 0, 1\}$. Totally unimodular matrices are important in combinatorial optimization because linear programs with totally unimodular constraint matrices are guaranteed to have integral solutions [42]. Dey, Hirani, and Krishnamoorthy have provided topological conditions on when a simplicial complex has a totally unimodular boundary matrix [30] stated below. We call a simplicial complex meeting the criteria of Theorem 2 ***relative torsion-free*** in dimension $d - 1$.

Theorem 2 (Dey et al. [30], Theorem 5.2)**.** *Let $\mathcal{K}$ be a d-dimensional simplicial complex. The boundary matrix $\partial\colon C_d(\mathcal{K}) \to C_{d-1}(\mathcal{K})$ is totally unimodular if and only if $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$ is torsion-free for all pure subcomplexes $\mathcal{L}_0, \mathcal{L}$ of $\mathcal{K}$ of dimensions $d-1$ and d where $\mathcal{L}_0 \subset \mathcal{L}$.*

We note that for a 2-dimensional simplicial complex $\mathcal{K}$ being relative torsion-free is equivalent to $\mathcal{K}$ not containing any Möbius subcomplex [30, Theorem 5.13].

Chapter 3: Embedability and Duality

We say a d-dimensional simplicial complex $\mathcal{K}$ is ***embeddable*** if it admits an ***embedding*** into $\mathbb{R}^{d+1}$, which we will now characterize. First, we recall the notion of the ***geometric realization*** $|\mathcal{K}|$ of $\mathcal{K}$. A d-dimensional simplex σ can be geometrically realized as the convex hull of $d+1$ affinely independent points in $\mathbb{R}^d$, and we denote the realization as $|\sigma|$. The faces of σ appear as the hyperplanes on the boundary of $|\sigma|$. The convex hull $|\sigma|$ is homeomorphic to the d-dimensional ball. We construct $|\mathcal{K}|$ as a quotient space by identifying the geometric realizations of the faces of each d-simplex $|\sigma|$ with the geometric realizations of the $(d-1)$-simplices on the boundary of σ. Now, we say that $\mathcal{K}$ is embeddable if there exists a continuous injective function $f\colon |\mathcal{K}| \to \mathbb{R}^{d+1}$ and we call f an embedding.

We can give an alternative combinatorial definition of an embeddable simplicial complex. We say that a d-dimensional simplicial complex $\mathcal{K}$ is embeddable if and only if there exists a triangulation of $\mathbb{R}^{d+1}$ that contains $\mathcal{K}$ as a subcomplex. This is equivalent to the existence of a triangulation of the $(d-1)$-sphere containing $\mathcal{K}$ as a subcomplex. This definition is more convenient in the context of algorithms since the triangulation contains very useful combinatorial information about the complex.

Embeddable complexes are very important as they admit a natural ***dual graph*** structure which allows us to reduce problems about simplicial complexes to graph theoretic ones. In Chapter 4 we use the dual graph structure to obtain fixed-parameter tractable and $O(\sqrt{\log n})$-approximation algorithms for the minimum bounding chain and minimum homologous chain problems. We also utilize the dual graph structure to prove hardness of approximation for both problems. In Chapter 5 we use the dual graph to solve max-flow and min-cut in simplicial complexes with a running time faster than linear programming.

In Section 3.1 we give an introduction to the Alexander duality theorem which is crucial for defining the dual graph of an embedded complex. We will show that embedded complexes and their dual graphs admit a cycle/cut duality similar to planar graphs. The ideas in Section 3.1 are not new, but the author is unaware of any text presenting them in this context. In Section 3.2 we characterize the existence of a dual graph in a way similar to Mac Lane's planarity criterion [72]. We show that, unlike graphs, there exist non-embeddable simplicial complexes that admit a dual graph. This characterization leads to an algorithm to compute the dual graph of a simplicial complex. Throughout the entire chapter we will assume all (co)homology coefficents to be coming from the field $\mathbb{Z}_2$.

3.1 Alexander duality

The Jordan curve theorem implies that a planar graph G separates the plane into $\beta_1 + 1$ connected components. Recall that β_1 is the dimension of the first homology group $H_1(G, \mathbb{Z}_2)$ which is the ***cycle space*** of the graph; this follows from the fact that graphs contain no 2-simplices. The Alexander duality theorem is a generalization of the Jordan curve theorem to simplicial complexes. For a d-dimensional simplicial complex $\mathcal{K}$ embedded in $\mathbb{R}^{d+1}$ the Alexander duality theorem implies that $\mathcal{K}$ separates $\mathbb{R}^{d+1}$ into β_d+1 connected components. Alexander duality relates the reduced d-dimensional homology $\tilde{H}_d(\mathcal{K})$ to the reduced 0-dimensional cohomology $\tilde{H}^0(\mathbb{R}^{d+1} \setminus \mathcal{K})$. We state a version of the theorem that holds in the setting of finite simplicial complexes.

Theorem 3 (Alexander Duality [4])**.** *Let $\mathcal{K}$ be a finite d-dimensional simplicial complex embedded in $\mathbb{R}^{d+1}$. We have the isomorphism $\tilde{H}_k(\mathbb{R}^{d+1} \setminus \mathcal{K}) \cong \tilde{H}^{d-k}(\mathcal{K})$ where the coefficients of the reduced (co)homology groups are taken from any abelian group.*

In the case of finite simplicial complexes with coefficients over a field $H_k(\mathcal{K})$ is a finitely generated vector space and $H^k(\mathcal{K})$ is its dual space, so we obtain $H_k(\mathcal{K}) \cong H^k(\mathcal{K})$. From Alexander duality we obtain the important isomorphism $H_d(\mathcal{K}) \cong \tilde{H}^0(\mathbb{R}^{d+1} \setminus \mathcal{K})$ which almost places the d-dimensional cycles of $\mathcal{K}$ in bijection with subsets of the connected components of $\mathbb{R}^{d+1} \setminus \mathcal{K}$ since the rank of $\tilde{H}^0(\mathbb{R}^{d+1} \setminus \mathcal{K})$ is one less than the number of connected components. The extra connected component is analogous to the outer face of a planar graph. This isomorphism will be very important for deriving one half of the cycle/cut duality between $\mathcal{K}$ and its dual graph and we will describe it in more detail later. We will now define the dual graph formally.

Let $\mathcal{K}$ be a d-dimensional simplicial complex embedded in $\mathbb{R}^{d+1}$. We define its dual graph G to be the graph whose vertices correspond to the $\beta_d + 1$ connected components of $\mathbb{R}^{d+1} \setminus \mathcal{K}$ and whose edges are in bijection with the d-simplices of $\mathcal{K}$. We call the connected components of $\mathcal{K} \setminus \mathbb{R}^{d+1}$ ***voids*** and note that each d-simplex is on the boundary of at most two voids [25, Corollary 7.1.2], so two vertices in the dual graph share an edge if and only if their corresponding voids contain a d-simplex in the intersection of their boundaries. Exactly one of the voids is unbounded. We refer to the vertex dual to the unbounded void as v_∞. It's important to note that Alexander duality actually gives us a d-dimensional simplicial complex dual to $\mathcal{K}$. Almost always we will only care about the 1-skeleton of this dual complex; which is the dual graph that we have just defined. However in Chapter 5 we will require the 2-skeleton of the dual complex for one duality theorem.

Before describing cycle/cut duality for simplicial complexes we will give a brief explanation of cycle/cut duality in planar graphs. Recall that the edge space of a graph G is actually the 1-dimensional chain group $C_1(G, \mathbb{Z}_2)$. It is well known that the cycle space and the cut space of a graph are orthogonal complements in the edge space; observe that a cycle intersects a cut an even number of times. This fact coincides with the Hodge decomposition of the edge space $C_1(G, \mathbb{Z}_2) \cong H_1(G, \mathbb{Z}_2) \oplus \operatorname{im} \delta_1[G]$. Consider the dual graph G^*, by Alexander duality we have the isomorphism $H_1(G) \cong \tilde{H}_0(\mathbb{R}^2 \setminus G)$. Since $\tilde{H}_0(\mathbb{R}^2 \setminus G)$ is equal to the number of vertices of G^* minus one we have $H_1(G) \cong C_0(G^* \setminus v_\infty)$. Recall that a cut in a graph is given by a partition of the vertices which is a function from the vertices to $\{0, 1\}$. The edges in the cut correspond to the coboundary of the vertices given the label 1. So, $\operatorname{im} \delta_1$ is generated by the coboundaries of formal sums of vertices over $\mathbb{Z}_2$. A basis for $\operatorname{im} \delta_1$ is given by the coboundaries of $|C_0(G)| - 1$ vertices, since the coboundary of the missing vertex can be obtained by summing up the coboundaries of all other vertices. Hence, the isomorphism $H_1(G) \cong C_0(G^* \setminus v_\infty)$ implies that $H_1(G) \cong \operatorname{im} \delta_1$, which says that the cycle space of G is isomorphic to the cut space of G^*. We obtain the converse $H_1(G^*) \cong C_0(G \setminus v_\infty^*)$ by the observation that for connected graphs $(G^*)^* = G$.

Now, we return to describing cycle/cut duality for simplicial complexes. For a d-dimensional simplicial complex $\mathcal{K}$ with dual graph G it is straightforward to show that $H_d(\mathcal{K}) \cong \operatorname{im} \delta_1[G]$. The reasoning is identical to the proof in planar graphs which follows directly from the construction of the dual graph. The second half of the duality, showing that $H_1(G) \cong \operatorname{im} \delta_d[\mathcal{K}]$ is more difficult. Unlike for connected graphs, we cannot use the fact that dual of the dual is the primal complex. Moreover, we need a definition a cut in a simplicial complex. We will define a ***cut*** to be an element of the image of the coboundary operator, and we will refer to the image of the d-dimensional coboundary operator as the ***d-dimensional cut space***. In Chapter 5 we will describe cuts in greater detail. For any cochain $\sigma \in \operatorname{im} \delta[\mathcal{K}]$ we have $\dim H_{d-1}(\mathcal{K}) < \dim H_{d-1}(\mathcal{K} \setminus \operatorname{supp}(\sigma))$. This follows directly from the definitions of the (co)homology groups and the isomorphism $H_{d-1}(\mathcal{K}) \cong H^{d-1}(\mathcal{K})$. In graphs this inequality is equivalent to the statement that removing a coboundary increases the number of connected components in the graph. That is, coboundaries are cuts. To prove the second half of cycle/cut duality, that the cycle space of the dual graph is isomorphic to the d-dimensional homology space of the primal complex, we will prove two lemmas. The first shows that any "minimal" collection of d-simplices whose removal increases the dimension of the $(d-1)$-homology group corresponds to a coboundary. The second shows that a cycle in the dual graph is dual to such a collection of simplices. The first lemma only requires an abstract simplicial complex, but the second lemma requires an embedding.

Lemma 1. *Let $\mathcal{K}$ be a d-dimensional simplicial complex and let C be a collection of d-simplices such that* $\dim H_{d-1}(\mathcal{K}) < \dim H_{d-1}(\mathcal{K} \setminus C)$*. Further, assume that C is minimal in the sense that for all $C' \subset C$ we have* $\dim H_{d-1}(\mathcal{K} \setminus C') < \dim H_{d-1}(\mathcal{K} \setminus C)$*. Then there exists a $(d-1)$-cochain p such that* $\operatorname{supp}(\delta(p)) = C$*.*

Proof. Recall that $H_{d-1}(\mathcal{K}) = \ker(\partial_{d-1}) / \operatorname{im}(\partial_d)$, $H^{d-1}(\mathcal{K}) = \ker(\delta_{d+1}) / \operatorname{im}(\delta_d)$, and we have the isomorphism $H_{d-1}(\mathcal{K}) \cong H^{d-1}(\mathcal{K})$. Removing the set of d-simplices C from $\mathcal{K}$ does not affect $\ker(\partial_{d-1})$ or $\operatorname{im}(\delta_{d-1})$. However, by our assumption it must decrease the dimension of $\operatorname{im}(\partial_d)$ and increase the dimension of $\ker(\delta_d)$. It follows that there must exist some $(d-1)$-cochain p such that $\delta_d[\mathcal{K}](p) \neq 0$ but $\delta_d[\mathcal{K} \setminus C](p) = 0$. Hence, $\operatorname{supp}(\delta_d(p)) \subseteq C$. Without loss of generality we choose p such that $|\operatorname{supp}(\delta_d(p))|$ is maximized. Define $C' \coloneqq C \setminus |\operatorname{supp}(\delta_d(p))|$. By minimality we have $\dim H_{d-1}(\mathcal{K} \setminus C') < \dim H_{d-1}(\mathcal{K} \setminus C)$. It follows that there exists some $(d-1)$-cochain p' with $\delta_d[\mathcal{K} \setminus C'](p') \neq 0$ but $\delta_p[\mathcal{K} \setminus C](p') = 0$. Note that by construction $\operatorname{supp}(\delta_d(p))$ and $\operatorname{supp}(\delta(p'))$ are disjoint. It follows that $\delta_d[\mathcal{K} \setminus C](p+p') = 0$ but $\delta_d[\mathcal{K}](p+p') \neq 0$ in $\mathcal{K}$, so $\operatorname{supp}(\delta_d(p+p')) \subseteq C$. Since $\operatorname{supp}(\delta_d(p))$ and $\operatorname{supp}(\delta_d(p'))$ are disjoint we have $|\operatorname{supp}(\delta_d(p))| < |\operatorname{supp}(\delta_d(p+p'))|$, contradicting our assumption that p maximizes $|\operatorname{supp}(\delta_d(p))|$. $\square$

Lemma 1 gives us a way to identify sets of d-simplices with coboundaries, which we have defined to be cuts. In Lemma 2[1] we will show that simple cycles in the dual graph are dual to sets of d-simplices meeting the assumptions of Lemma 1.

Lemma 2. *Let $\mathcal{K}$ be a d-dimensional simplicial complex embedded in $\mathbb{R}^{d+1}$ and let G be its dual graph. Let C be a simple cycle in G and let C^* be the set of d-simplices dual to C. We have the equality* $\dim H_{d-1}(\mathcal{K} \setminus C^*) = \dim H_{d-1}(\mathcal{K}) + 1$*.*

Proof. Let $\{V_1, \ldots, V_{\beta_d+1}\}$ be the closures of the voids of $\mathbb{R}^{d+1} \setminus \mathcal{K}$ viewed as subspaces of $\mathbb{R}^{d+1}$. We consider each pair V_i, V_j as if they were disjoint, and we we view them as

[1]The author would like to acknowledge an anonymous MathOverflow user for providing the proof [48].

$(d+1)$-complexes whose boundary d-cycles generate $H_d(\mathcal{K})$. Each d-simplex of $\mathcal{K}$ appears on the boundary of at most two of these subspaces. Removing a d-simplex, whose dual edge has endpoints v_i and v_j, is equivalent to glueing V_i and V_j together along that simplex. We use the Mayer-Vietoris sequence for reduced homology to compute the homology after glueing V_i and V_j together.

$$\cdots \to H_k(V_i \cap V_j) \to H_k(V_i) \oplus H_k(V_j) \to H_k(V_i \cup V_j) \to H_{k-1}(V_i \cap V_j) \to \dots$$

Since $V_i \cap V_j$ is a single d-simplex its reduced homology is trivial in all dimensions; we have $H_1(V_i \cup V_j) \cong H_1(V_i) \oplus H_1(V_j)$. It follows that for any tree T in the dual graph, the space created after glueing together its corresponding voids has first homology $H_1(T) \cong \bigoplus_{v_i \in T} H_1(V_i)$.

Now consider some simple cycle C in G. Order the vertices of C by $v_1, \dots, v_k$ which induces an order on the edges $e_1 = (v_1, v_2), \dots, e_{k-1} = (v_{k-1}, v_k), e_k = (v_k, v_1)$. We remove the d-simplices in $\mathcal{K}$ according to this order, that is at v_i we remove the d-simplex dual to the edge e_{i-1}.

Let V be space created by glueing the components in $\{V_1, \dots, V_{k-1}\}$ according to the ordering. We use the reduced Mayer-Vietoris sequence to compute the homology after removing the d-simplex corresponding to the edge e_k.

$$\cdots \to H_1(V_k \cap V) \to H_1(V_k) \oplus H_1(V) \xrightarrow{\psi} H_1(V_k \cup V) \to \tilde{H}_0(V_k \cap V) \xrightarrow{\phi} \tilde{H}_0(V_k) \oplus \tilde{H}_0(V)$$

The intersection $V_k \cap V$ consists of two disjoint d-simplices which are dual to the edges e_{k-1} and e_k. It follows that $H_1(V_k \cap V) \cong 0$ and $\tilde{H}_0(V_k \cap V) \cong \mathbb{Z}_2$, hence ψ is injective. Further, we see that ϕ is the zero map since both V_k and V consist of a single connected component. We can compute the dimension of $H_1(V_k \cup V)$ by the rank nullity theorem:

$$\begin{aligned}\dim H_1(V_k \cup V) &= \dim \mathbb{Z}_2 + \dim H_1(V_k) \oplus H_1(V) \\ &= 1 + \dim H_1(V_k) + \dim H_1(V).\end{aligned}$$

To conclude we argue that removing C^* from $\mathcal{K}$ increases the $(d-1)$-homology group of the complex by exactly one. We construct the simplicial complex $\mathcal{K} \setminus C^*$ be removing one d-simplex from $C^* = \{e_1^*, \dots, e_k^*\}$ at a time, and following the linear ordering assigned to the edges in cycle $C = \{e_1, \dots, e_k\}$. Removing the d-simplices $\{e_1^*, \dots, e_{k-1}^*\}$ does not change the $(d-1)$-homology group, since their dual edges form a tree. However, upon removing e_k^* the $(d-1)$-homology group is increased by the argument in the preceding paragraph.

□

For an illustration of the proof of Lemma 2 see Figure 3.1. Now we will wrap things up by using Lemmas 1 and 2 to prove that cycle space of the dual graph is isomorphic to the image of $\delta_d[\mathcal{K}]$.

Theorem 4. *Let $\mathcal{K}$ be a d-dimensional simplicial complex embedded in $\mathbb{R}^{d+1}$ and let G be its dual graph. The cycle space of G is isomorphic to the $(d-1)$-cut space of $\mathcal{K}$; that is, $H_1(G) \cong \operatorname{im} \delta_d[\mathcal{K}]$. Moreover, every cycle basis of G is dual to a basis for $\operatorname{im} \delta_d[\mathcal{K}]$.*

Proof. First, we note that the dimension of $H_1(G)$ is given by $\dim H_1(G) = n_d - \beta_d$ where n_d is the number of d-simplices in $\mathcal{K}$. By the Hodge decomposition we have $C_d(\mathcal{K}) \cong H_d(\mathcal{K}) \oplus \operatorname{im} \delta_d$, from which it follows that $\dim \operatorname{im} \delta_d = n_d - \beta_d$. Hence, $H_1(G) \cong \operatorname{im} \delta_d$. Moreover, by Lemma 2 every cycle in G is dual to a minimal set of d-simplices whose removal increases the dimension of $H_{d-1}(\mathcal{K})$, so by Lemma 1 is dual to a coboundary. Hence, every cycle basis of G is dual to a basis for $H_{d-1}(\mathcal{K})$. □

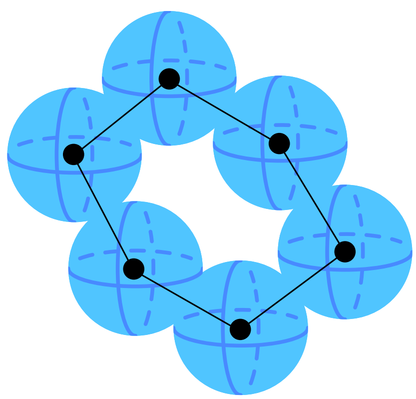

Figure 3.1: The blue spheres represent the closures of connected components in $\mathbb{R}^3 \setminus \mathcal{K}$ for some 2-complex $\mathcal{K}$. In black we have a cycle in the dual graph whose edges are the triangles in $\mathcal{K}$ contained in the intersection of two spheres. By Lemma 2 removing these triangles increases the rank of $H_1(\mathcal{K})$ by one.

3.2 Constructing the Dual Graph

Several of the algorithmic results in this thesis assume that a dual graph is given as input to the algorithm. Hence, it is very important to show that we can construct the dual graph from the embedding in polynomial time. Dey, Hou, and Mandal provide an algorithm computing the dual graph from an embedding in $O(n \log n)$ time [31]. The dual graphs that we compute in this section are different from the dual graphs arising from an embedding as defined in the previous section. However, we still refer to them as dual graphs since they retain the cycle/cut duality property. In this section we use a more general definition of a dual graph. We define a ***dual graph*** of a simplicial complex $\mathcal{K}$ to be any graph G such that $Z_d(\mathcal{K}) \cong \operatorname{im} \delta_1[G]$; that is, the d-dimensional cycles of $\mathcal{K}$ are in bijection with the cuts of G. For d-dimensional complexes $Z_d(\mathcal{K}) \cong H_d(\mathcal{K})$. Clearly, the dual graphs defined in the previous section meet this definition.

In this section we provide an alternative algorithm which computes a dual graph, if one exists, even when an embedding is not given. In most cases our approach is slower than $O(n \log n)$ time but it is interesting from a theoretical perspective. We show that there exist non-embeddable complexes that admit a dual graph; the d-cycles in the complex are in bijection with the cuts of the graph. Further, for a d-dimensional simplicial complex $\mathcal{K}$ we show that given a basis for $H_d(\mathcal{K}, \mathbb{Z}_2)$ there exists an algorithm that acts as a change of basis on $H_d(\mathcal{K}, \mathbb{Z}_2)$ whose output is a basis that can be interpreted as the incidence matrix of a dual graph for $\mathcal{K}$. By showing that the class of complexes admitting dual graphs is larger than the class of embeddable complexes we strengthen previous algorithmic results operating on dual graphs; the class of complexes that may be given as input to these algorithms is larger than previously thought. These problems include computing Betti numbers [27], the homology localization problem [23], and computing minimal persistent cycles [31].

For planar graphs computing the dual is the same as computing an embedding in the plane; every embedding yields a dual graph and every dual graph defines an embedding. Hence, computing the dual of a planar graph can be done in linear time with graph drawing algorithms. Unfortunately, this technique does not work in higher dimensions. Finding an embedding of a 2-complex in $\mathbb{R}^3$ is NP-hard [26]. Even worse, determining whether or not a $(d-1)$-complex embeds in $\mathbb{R}^d$ is undecidable for $d \geq 5$ [75]. At first it may appear that this implies computing the dual graph is undecidable, but this is not the case. Fortunately, computing the dual graph of an embedded simplicial complex is not equivalent to computing an embedding. The dual graph of an embedded d-complex tells us how the d-cycles intersect one another, but it does not characterize the d-cycles up to homeomorphism. It seems likely that constructing an embedding from the dual graph would require an algorithm solving the homeomorphism problem, which is undecidable [74]. This problem does not arise in planar graphs as every face of a planar graph is homeomorphic to a disk. Moreover, the issue can be avoided in $\mathbb{R}^3$ by using Fox's reimbedding theorem [43] which states that if a 3-manifold $\mathcal{M}$ admits an embedding in $\mathbb{R}^3$ then it admits an embedding such that the complement $\mathbb{R}^3 \setminus \mathcal{M}$ is homeomorphic to a union of handlebodies[2]. Note that this result applies to 2-complexes as in polynomial time any embeddable 2-complex can be thickened to an embeddable 3-manifold that is homotopic to the original complex [29]. Given these limitations we seek algebraic methods for computing the dual graph.

Mac Lane's planarity criterion gives us an algebraic condition equivalent to planarity [72]. It states that a graph G is planar if and only if there exists a basis B of $H_1(G, \mathbb{Z}_2)$ such that every edge of G is contained in at most two cycles of B. We call such a basis a ***2-basis***. The faces of a planar embedding clearly give rise to a 2-basis; the converse can be proved combinatorially by showing that any graph containing K_5 or $K_{3,3}$ cannot have a 2-basis [80]. When viewing B as a matrix the columns are indexed by the cycle basis and the rows are indexed by the edges. Each row of B contains at most two non-zero entries. Hence, a 2-basis for a planar graph is the incidence matrix of the dual graph minus the vertex corresponding to the outer face. We can compute the column associated with the outer face by simply taking the sum of the columns in B. We note that the dual graph obtained will be missing any loops. Bridges, edges that do not appear in any cycle, in a graph do not appear in the cycle basis, so we obtain the full dual graph if and only if G is 2-connected. However, we can always add the bridges to the dual graph in linear time.

The notion of a 2-basis easily extends to d-dimensional simplicial complexes. Let $\mathcal{K}$ be a d-complex that is not necessarily embeddable into $\mathbb{R}^{d+1}$. A 2-basis for $H_d(\mathcal{K}, \mathbb{Z}_2)$ is a homology basis B such that each d-simplex is contained in at most two basis elements. A 2-basis need not be unique, for example every embedding of a complex gives rise to a different 2-basis. As before we can view B as a matrix whose columns are d-cycles and whose rows are d-simplices. This matrix gives us the incidence matrix of a dual graph. The coboundary of every vertex in the graph is dual to a d-cycle in $\mathcal{K}$ so we immediately obtain a bijection between d-cycles in the complex and cuts in the graph.

[2]In the case of 3-manifolds a handlebody of genus g is a 3-manifold whose boundary is a surface of genus g.

When $\mathcal{K}$ is a d-complex embedded in $\mathbb{R}^{d+1}$ Alexander duality implies the existence of a 2-basis for $H_d(\mathcal{K}, \mathbb{Z}_2)$. The boundaries of the connected components of $\mathbb{R}^{d+1} \setminus \mathcal{K}$ generate $H_d(\mathcal{K}, \mathbb{Z}_2)$, and every d-simplex is contained in at most two of these boundaries. Surprisingly, there exist non-embeddable complexes that admit 2-bases. For example, a Klein bottle taken as a 2-complex does not embed into $\mathbb{R}^3$, but as a 2-basis since it only contains one 2-cycle. Using this observation we can construct non-embeddable complexes with non-trivial 2-bases. See Figure 3.2. This means that the class of simplicial complexes for which duality based algorithms work is larger than the class of embeddable simplicial complexes. Moreover, we will show that we can compute a dual graph for this class of complexes, the only input needed for these algorithms is the simplicial complex itself.

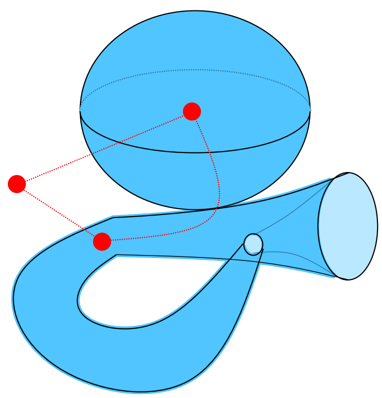

Figure 3.2: A simplicial complex consisting of a sphere and Klein bottle glued together along a common triangle. This 2-complex does not embed into $\mathbb{R}^3$, but has a dual graph.

Given a basis for $H_d(\mathcal{K}, \mathbb{Z}_2)$ we want an algorithm to compute a 2-basis if one exists. In this section we will phrase the problem in the language of matroid theory, and then use a known matroid algorithm to compute a 2-basis. The study of matroids originated with Whitney who used them as a way to generalize the notion of linear dependence [96]. There are many equivalent definitions of a matroid, however we use the circuit definition as it most naturally describes our situation. A ***matroid*** is a pair $M = (X, C)$ where X is a set and $C \subseteq P(X)$ is a collection of subsets of X called ***circuits*** such that the following two properties hold. For all $A \in C$ no proper subset of A is contained in C. If $A, B \in C$ and there exists an $x \in A \cap B$ then $(A \cup B) \setminus \{x\}$ contains a circuit.

In a graph G the cut space $\operatorname{im} \delta$ forms a vector space over $\mathbb{Z}_2$ hence they form a matroid. This fact can also be easily seen from properties of the cut space. A coboundary is a minimal set of edges whose removal disconnects the graph. By definition no subset of a minimal edge cut is an edge cut, so the first property of the circuit definition of a matroid holds. If E_S and E_T are minimal edge cuts and $e \in E_S \cap E_T$ then by properties of the symmetric difference we obtain $E_S \oplus E_T \subset (E_S \cup E_T) \setminus \{e\}$, but minimal edge cuts are coboundaries and hence are closed under symmetric difference, so the second property of the circuit definition of a matroid also holds.

The matroid of edge cuts of a graph G is called the ***bond matroid*** of G. Any matroid isomorphic to a bond matroid is called a ***cographic matroid***. Since the cut space of a graph is a vector space over $\mathbb{Z}_2$ it follows that all cographic matroids can be represented by some vector space over $\mathbb{Z}_2$. Any vector space over $\mathbb{Z}_2$ with a 2-basis is the bond matroid of some graph. The 2-basis gives rise to the incidence matrix of the graph for which it is the bond matroid of. It follows that deciding whether or not $H_d(\mathcal{K}, \mathbb{Z}_2)$ has a 2-basis is equivalent to decided if $H_d(\mathcal{K}, \mathbb{Z}_2)$ is a cographic matroid.

Tutte's algorithm for recognizing cographic[3] matroids [93] can be used to compute a dual graph of a simplicial complex with a 2-basis in polynomial time. The algorithm takes as input a basis to a vector space over $\mathbb{Z}_2$. If the vector space represents a cographic matroid, the algorithm outputs the incidence matrix of a graph whose bond matroid is isomorphic to the input vector space. Hence, when given a homology basis as input Tutte's algorithm returns a dual graph. We note that computing the initial homology basis can be performed in polynomial time [37, Section 4.2]. The runtime analysis of Tutte's algorithm given by Qu [84] gives us the following theorem.

Theorem 5. *Let $\mathcal{K}$ be a d-dimensional simplicial complex such that $H_d(\mathcal{K})$ has a 2-basis. A dual graph of $\mathcal{K}$ can be computed in $O(\beta_d^2 \cdot n_d)$ time.*

3.3 Weak pseudomanifolds

A $(d+1)$-weak pseudomanifold is a $(d+1)$-dimensional simplicial complex such that every d-simplex is incident to at most two $(d+1)$-simplices. Weak pseudomanifolds are a generalization of combinatorial manifolds. Like manifolds they admit a dual graph. However, recognizing weak pseudomanifolds can be done in linear time while recognizing manifolds is undecidable [24]. We choose the notation $(d+1)$-weak because the problems we solve on them will be formulated on the d-chains.

[3]Confusingly, Tutte used the reverse of modern terminology and used the term graphic matroid to refer to what is commonly known as a cographic matroid today.

We now formally define the dual graph G of a $(d+1)$-weak pseudomanifold. In G we have a vertex σ^* for each $(d+1)$-simplex σ and one additional vertex v_b The edges of G are in bijection with the d-simplices. If two $(d+1)$-simplices share a common d-simplex as a face then there is an edge between their dual vertices. If a d-simplex is incident to one $(d+1)$-simplex σ then there is an edge (σ^*, v_b). We will prove a weaker version of cycle/cut duality. We will show that there is a bijection between null-homologous cycles in a weak psuedomanifold and cuts in its dual graph.

Theorem 6. *Let $\mathcal{M}$ be a $(d+1)$-dimensional weak psuedomanifold with dual graph $G = (V, E)$. There is an isomorphism* $\operatorname{im} \partial_{d+1}[\mathcal{M}] \cong \operatorname{im} \delta_1[G]$.

Proof. We show a one-to-one correspondence between null-homologous d-cycles in $\mathcal{M}$ and cuts in G. Let γ be a null-homologous d-cycle in $\mathcal{M}$, then there exists some $(d+1)$-chain Γ such that $\partial\Gamma = \gamma$. This yields a partition of the $(d+1)$-skeleton $\mathcal{M}_{d+1}$ into Γ and $\mathcal{M}_{d+1} \setminus \Gamma$. Hence we have a partition of the vertices in G: Γ^* and $V \setminus \Gamma^*$. By the construction of the dual graph the coboundary of Γ^* (or equivalently of $V \setminus \Gamma^*$) is dual to the boundary of Γ which is equal to γ.

Conversely, let S and $V \setminus S$ partition V, and let E_S be the edges on the coboundary of the partition. Without loss of generality we assume that $v_b \in V \setminus S$. By duality S^* is a $(d+1)$-chain with boundary E_S^*. □

Similar to the case of embedded complexes, the dual graph construction for weak pseudomanifolds actually extends to a larger class of graphs. Any simplicial complex such that $\operatorname{im} \partial_d$ admits a 2-basis gives rise to a dual graph by the construction outlined in this section. For an example we can consider any 2-complex with the property that the 1-skeleton is a planar graph. Then $\operatorname{im} \partial$ is a subspace of the cycle space, which is planar, and admits a 2-basis. Hence, we can use Tutte's cographic matroid detection algorithm to compute a dual graph, if one exists, even if the input is not a weak pseudomanifold.

Chapter 4: The Minimum Bounded and Homologous Chain Problems

In this section we present two optimzation problems over $\mathbb{Z}_2$-homology: the minimum bounded chain problem and the minimum homologous chain problem. Both problems are minimization problems asking for a minimum sized chain meeting some constraints. By the size of a d-chain σ we mean the cardinality of its support $|\operatorname{supp}(\sigma)|$, or equivalently the number of d-simplices in the chain when viewed as a set. Throughout the entirity of this chapter we assume coefficients over $\mathbb{Z}_2$ and we will use the symbol $\oplus$ to denote addition over $\mathbb{Z}_2$; this is to emphasize the fact that addition of chains over $\mathbb{Z}_2$ is equivalent to taking the symmetric difference of the underlying sets. We begin with the formal definitions of both problems.

Definition 1 (Minimum Bounded Chain Problem)**.** *Given a d-dimensional simplicial complex $\mathcal{K}$ and a null-homologous $(d-1)$-cycle γ, the* ***minimum bounded chain problem*** *asks to find a d-chain Γ such that $\partial\Gamma = \gamma$ and $|\sup(\Gamma)|$ is minimized.*

The minimum bounded chain problem asks to find the smallest d-chain whose boundary is γ. The minimum homologous chain problem is similar which asks to find the smallest $(d-1)$-chain that is homologous with a given $(d-1)$-chain.

Definition 2 (Minimum Homologous Chain Problem)**.** *Given a d-dimensional simplicial complex $\mathcal{K}$ and a $(d-1)$-chain τ, the* ***minimum homologous chain problem*** *asks to find a $(d-1)$-chain σ such that there exists a d-chain Γ with $\partial\Gamma = \tau \oplus \sigma$ and that $|\operatorname{supp}(\sigma)|$ is minimized.*

The minimum bounded chain problem has been previously studied with homology coefficients over $\mathbb{Z}$ and $\mathbb{R}$ where linear programming techniques can be utilized. Sullivan described the problem as a discretation of the minimal spanning surface problem [92] on the closely related cellular complexes, but with the restriction that the complex admits an embedding into Euclidean space. Kirsanov and Gortler reduce the problem to a minimum cut problem in the dual graph under the assumption that $\mathcal{K}$ embeds into $\mathbb{R}^{d+1}$ and that there exists a bounding chain contained in the boundary of the unbounded void V_∞ of $\mathbb{R}^{d+1} \setminus \mathcal{K}$ [66]. This assumption makes the problem an analog of the shortest st-path problem in a planar graph when s and t appear on the same face. In Section 4.4 we provide a combinatorial proof of their result.

The linear programming approach was then applied to the minimum bounded chain problem (over $\mathbb{Z}$) by Dunfield and Hirani [34]. Moreover, they show the minimum bounded chain problem is NP-complete via a reduction from 1-in-3 SAT. The gadget they use was originally used by Agol, Hass and Thurston to show that the minimal spanning area problem is NP-complete [2]. Linear programming techniques have also been used by Chambers and Vejdemo-Johansson to solve the minimum bounded chain problem in the context of $\mathbb{R}$-homology [20]. In $\mathbb{R}$-homology Carvalho et al provide an algorithm finding a (not necessarily minimum) bounded chain in a manifold by searching the dual graph of the manifold [16].

Research on minimum homologous chain has largely worked in $\mathbb{Z}$-homology. Dey, Hirani and Krishnamoorthy formulate the minimum homologous chain problem over $\mathbb{Z}$ as an integer linear program and describe topological conditions for the linear program to be totally unimodular, and so polynomial-time solvable [30]. Of course, integer linear programming approaches do not extend to $\mathbb{Z}_2$-homology.

Special cases of the minimum homologous chain problem have been studied in $\mathbb{Z}_2$-homology. The homology localization problem is the case when the input chain is a cycle. The homology localization problem over $\mathbb{Z}_2$ in surface-embedded graphs is known to be NP-hard via a reduction from maximum cut by Chambers et al. [17]; our reduction is from the complement problem minimum uncut. On the algorithmic side, Erickson and Nayyeri provide a $2^{O(g)}n \log n$ time algorithm where g is the genus of the surface [39]. Using the idea of annotated simplices, Busaryev et al. generalize this algorithm for homology localization of 1-cycles in simplicial complexes; the algorithm runs in $O(n^\omega) + 2^{O(g)}n^2 \log n$ time where ω is the exponent of matrix multiplication, and g is the dimension of the first homology group of the complex [14].

Using a reduction from the nearest codeword problem Chen and Freedman showed that homology localization with coefficients over $\mathbb{Z}_2$ is not only NP-hard, but it cannot be approximated within any constant factor in polynomial time [23]. These hardness results hold for a 2-dimensional simplicial complex, but not necessarily for 2-dimensional complexes embedded in $\mathbb{R}^3$. They also give a polynomial-time algorithm for the special case of d-dimensional simplicial complex that is embedded in $\mathbb{R}^d$. This is different from our setting of a d-dimensional simplicial complex that is embedded in $\mathbb{R}^{d+1}$; however the algorithm also reduces to a minimum cut problem in a dual graph, much like that of

Kirsanov and Gortler.

The minimum bounded chain problem over $\mathbb{Z}_2$ can be stated as a linear algebra problem, but this has little algorithmic use since the resulting problems are intractable. The algebraic formulation is to find a vector x of minimum Hamming weight that solves an appropriately defined linear system $Ax = b$. It is possible to reduce in the reverse direction, but the resulting complex is not embeddable in general, and so provides no new results.

In coding theory this algebraic problem is a well-studied decoding problem known as maximum likelihood decoding, and it was shown to be NP-hard by Berlekamp, McEliece and van Tilborg [9, 94]. Downey, Fellows, Vardy and Whittle show that maximum likelihood decoding is W[1]-hard [33]. Further, Austrin and Khot show that maximum likelihood decoding is hard to approximate within a factor of $2^{(\log n)^{1-\epsilon}}$ under the assumption that $\mathsf{NP} \not\subseteq \mathsf{DTIME}\left(2^{(\log n)^{O(1)}}\right)$ [7]. This work was continued by Bhattacharyya, Gadekar, Ghosal and Saket who showed that maximum likelihood decoding is still W[1]-hard when the problem is restricted to $O(k \log n) \times O(k \log n)$ sized matrices for some constant k [10].

4.1 Summary of main results

We prove similar results for both problems. Algorithmically we show the existence of both approximation and exact fixed-parameter algorithms which are summarized in the following four theorems.

Theorem 7. *There exists an $O(\sqrt{\log \beta_d})$-approximation algorithm for the minimum bounded chain problem for a simplicial complex $\mathcal{K}$ embedded in $\mathbb{R}^{d+1}$, with dth Betti number β_d.*

Theorem 8. *There exists an $O(15^k \cdot k \cdot n_d^3)$ time exact algorithm for the minimum bounded chain problem for simplicial complexes embedded in $\mathbb{R}^{d+1}$, where k is the number of d-simplices in the optimal solution.*

Theorem 9. *There exists an $O(\sqrt{\log n_{d+1}})$-approximation algorithm for the minimum homologous chain problem for d-chains in $(d+1)$-manifolds.*

Theorem 10. *There exists an $O(15^k \cdot k \cdot n_d^3)$ time exact algorithm for the minimum homologous chain problem for d-chains in $(d+1)$-manifolds, where k is the size of the optimal solution.*

The running times for the first two theorems is computed assuming that the dual graph of the complex in $\mathbb{R}^{d+1}$ is available. The last two theorems hold, more generally, for weak pseudomanifolds studied by Dey et al. in [31].

Our hardness results show that both problems are APX-hard and that constant factor approximations are unlikely. The hardness of constant factor approximation is dependent upon the unique games conjecture which is an assumption necessary is many interesting results in computational topology [47].

Theorem 11. *The minimum bounded chain problem is*

(i) hard to approximate within a $(1+\varepsilon)$ factor for some $\varepsilon > 0$ assuming $\mathsf{P} \neq \mathsf{NP}$, and

(ii) hard to approximate within any constant factor assuming the unique games conjecture,

even if $\mathcal{K}$ is a 2-dimensional simplicial complex embedded in $\mathbb{R}^3$ with input cycle γ embedded on the boundary of the unbounded volume in $\mathbb{R}^3 \setminus \mathcal{K}$.

The condition that γ be on the unbounded volume of $\mathbb{R}^3 \setminus \mathcal{K}$ is important. In Section 4.4 we will see that the problem becomes polynomial-time solvable when γ has a bounded chain on the unbounded volume (note that this implies γ is on the unbounded volume as well).

Theorem 12. *The minimum homologous chain problem is*

(i) hard to approximate within a $(1+\varepsilon)$ factor for some $\varepsilon > 0$ assuming $\mathsf{P} \neq \mathsf{NP}$, and

(ii) hard to approximate within any constant factor assuming the unique games conjecture,

even when the input chain is a 1-cycle on an orientable 2-manifold.

For the sake of completeness, we also give a more general presentation of the result of Kirsanov and Gortler [66], that minimum bounded chain is polynomial-time solvable for a d-dimensional simplicial complex $\mathcal{K}$ embedded in $\mathbb{R}^{d+1}$ and input chain γ null-homologous on the boundary of the unbounded region in $\mathbb{R}^{d+1} \setminus \mathcal{K}$. This can be found in Section 4.4. This algorithmic result is likely to be the most general possible, given Theorem 11.

4.2 Approximation algorithm and fixed-parameter tractability

In this section, we describe approximation algorithms and parameterized algorithms for both minimum bounded chain and minimum homologous chain problems. Our algorithms work with the dual graph of the input space. In order to simplify our presentation we assume that the dual graph of the input complex contains no loops. The following lemma shows that we can make this assumption without any loss of generality.

Lemma 3. *In polynomial time we can preprocess an instance of the minimum bounded chain problem $(\mathcal{K}, \gamma)$ into a new instance $(\mathcal{K}', \gamma')$ such that (i) the dual graph $(\mathcal{K}')^*$ contains no loops and (ii) an α-approximation algorithm for $(\mathcal{K}', \gamma')$ implies an α-approximation algorithm for $(\mathcal{K}', \gamma')$.*

Proof. Let F be the set of d-simplices corresponding to the loops in $\mathcal{K}^*$. By cycle/cut duality we see that no d-simplex $f \in F$ can be on a d-cycle; as a loop cannot be on any cut in the dual graph. Therefore, for any X, Y with boundary γ, f is either on both of them, or on neither of them. This follows because $X \oplus Y$ is a d-cycle. Thus, each $f \in F$ is either on all d-chains with boundary γ, or none of them. Let $F_{all} \subseteq F$ be the d-simplices that are on all d-chains X with boundary γ, and let $F_{none} \subseteq F$ be the d-simplices that are on no X with boundary γ, we have $F_{all} \cup F_{none} = F$.

Now, we compute a feasible solution Y with $\partial Y = \gamma$ by solving the linear system using standard methods [56]. Using Y we can partition F into F_{all} and F_{none}: a d-simplex $f \in F$ is in F_{all} if it is in Y, and in F_{none} otherwise. We can remove F_{none} from $\mathcal{K}$ without changing the optimal solution. Further, any chain X with boundary γ contains F_{all}. That is, we can write $X = X' \oplus F_{all}$. It follows that

$$\partial X = \partial X' \oplus \partial F_{all} = \gamma \Rightarrow \partial X' = \gamma \oplus \partial F_{all} = \gamma'.$$

Hence, we can find the minimum chain X'_{opt} in $\mathcal{K}' = \mathcal{K} \setminus \partial F_{all}$ with boundary γ'. Then, $X_{opt} = X'_{opt} \oplus F_{all}$ is the minimum bounding chain for γ in $\mathcal{K}$. Furthermore, any approximation algorithm for $(\mathcal{K}', \gamma')$ implies an approximation algorithm with the same ratio for $(\mathcal{K}, \gamma)$. To see that, let X'_{apx} be an approximation of X'_{opt} in $(\mathcal{K}', \gamma)$, and let $X_{apx} = X'_{apx} \oplus F_{all}$. So, we have:

$$\frac{|X_{apx}|}{|X_{opt}|} = \frac{|X'_{apx} \oplus F_{all}|}{|X'_{opt} \oplus F_{all}|} = \frac{|X'_{apx} \cup F_{all}|}{|X'_{opt} \cup F_{all}|} = \frac{|X'_{apx}| + |F_{all}|}{|X'_{opt}| + |F_{all}|} \leq \frac{|X'_{apx}|}{|X'_{opt}|}.$$

The second equality holds as X'_{apx} and X'_{opt} are disjoint from F_{all}; as they are solutions in $\mathcal{K}'$ that does not contain F. The last equality holds as $|X'_{apx}|$, $|X'_{apx}|$, and $|F_{all}|$ are non-negative and $|X'_{apx}| \geq |X'_{opt}|$. $\square$

4.2.1 Reductions to the minimum cut completion problem

Given a graph $G = (V, E)$ and edge set $E' \subseteq E$, the ***minimum cut completion*** problem asks for a cut $(S, \overline{S})$ with edge set E_S that minimizes the symmetric difference $|E_S \oplus E'|$. Essentially, the minimum cut completion problem asks for the cheapest $(S, \overline{S})$ cut such that the edges in E' are free. In the next lemma we show that the minimum cut completion problem generalizes the minimum bounded chain problem for embedded complexes.

Lemma 4. *For any d-dimensional instance of the minimum bounded chain problem, $(\mathcal{K}, \gamma)$, there exists an instance of the minimum cut completion problem $(G = (V, E), E')$ that can be computed in polynomial time, and a one-to-one correspondence between cuts in G and d-chains with boundary γ in $\mathcal{K}$. Moreover, if the cut $(S, \overline{S})$ with edge set E_S in G corresponds to the d-chain Q in $\mathcal{K}$ then $|E_S \oplus E'| = |Q|$.*

Proof. Let F be a feasible solution to the linear system $\partial F = \gamma$, such a solution can be computed in polynomial time. We construct the minimum cut completion instance by setting $G = \mathcal{K}^*$ and $E' = F^*$.

Let Q be any other d-chain with $\partial Q = \gamma$, so $\partial(Q \oplus F) = 0$. By Alexander duality $Q \oplus F$ partitions $\mathbb{R}^{d+1}$ into an interior and exterior. Let $(S, \overline{S})$ be the corresponding dual cut in $\mathcal{K}^*$ and let E_S be the edge set of this cut. We have $|E_S \oplus E'| = |E_S \oplus F^*| = |Q^*| = |Q|$.

Conversely, let $(S, \overline{S})$ be a cut in $\mathcal{K}^*$ with edge set E_S. By cycle/cut duality $\partial E_S^* = 0$. Now, let $Q = E_S^* \oplus F$. Hence, $\partial Q = \gamma$. Finally, we have $|Q| = |E_S^* \oplus F| = |E_S \oplus F^*| = |E_S \oplus E'|$. $\square$

To make the bijection in the above proof explicit we summarize it as the following composition of bijections

$$d\text{-chain in } \mathcal{K} \quad \xleftrightarrow{-\oplus F} \quad d\text{-cycle in } \mathcal{K} \quad \xleftrightarrow{\text{Duality}} \quad \text{edge cut in } \mathcal{K}^*.$$

Next, we show via a similar argument that the cut completion problem also generalizes the minimum homologous chain problem when the input complex is a weak pseudomanifold.

Lemma 5. *For any d-dimensional instance of the minimum homologous chain problem $(\mathcal{M}, D)$, where $\mathcal{M}$ is a weak pseudomanifold, there exists an instance of the minimum cut completion problem $(G = (V, E), E')$ that can be computed in polynomial time, and a one-to-one correspondence between cuts in G and d-chains in $\mathcal{M}$ that are homologous to D.*

Moreover, if the cut $(S, \overline{S})$ *with edge set* E_S *in* G *corresponds to the* d*-chain* Q *in* $\mathcal{K}$ *then* $|E_S \oplus E'| = |Q|$.

Proof. We construct an instance (G, E') of minimum cut completion as follows. For each $(d+1)$-simplex v^* in $\mathcal{M}$ we have a dual vertex v in G. $V(G)$ contains one additional vertex, v_B. For each d-simplex e^* if it is a face of two $(d+1)$-simplices u^* and v^* we add the edge (u, v) to $E(G)$. Otherwise, if e^* is a face of only one $(d+1)$-simplex u^* we add the edge (u, v_B) to $E(G)$. Finally, let E' be the set of edges dual to γ.

Let Q be any d-chain homologous with γ, so $Q \oplus \gamma$ is null-homologous. That is, there exists a $(d+1)$-chain S^* with $\partial S^* = Q \oplus \gamma$. The cut $(S, \overline{S})$ in G has cost $|E_S \oplus E'| = |Q|$ since E' is dual to γ and E_S is dual to $Q \oplus \gamma$. Conversely, let $(S, \overline{S})$ be a cut in G with edge set E_S. It follows that E_S^* is null-homologous, so $Q = E_S^* \oplus \gamma$ is homologous with γ and its cost is $|E_S \oplus E'|$. □

Again, we summarize the bijection in the proof of the preceding lemma. We start with a d-chain homologous to γ, solve the linear system to obtain a $(d+1)$-chain, and then use duality to obtain an edge cut. This corresponds to the following diagram

$$
\begin{aligned}
d\text{-chain in } \mathcal{M} \quad &\xleftrightarrow{-\oplus Q} \quad d\text{-cycle in } \mathcal{M} \quad \xleftrightarrow{\partial^{-1}} \quad (d+1)\text{-chain in } \mathcal{M} \\
&\xleftrightarrow{\text{Duality}} \quad \text{edge cut in } \mathcal{M}^*.
\end{aligned}
$$

The operator ∂^{-1} is not a well-defined function as there exist up to two $(d+1)$-chains bounding every null-homologous cycle γ in a weak pseudomanifold. However, these $(d+1)$-chains uniquely determine γ and we may view ∂^{-1} as a function mapping a d-cycle to a partition of the $(d+1)$-simplices.

4.2.2 Algorithms for the minimum cut completion problem

In this section we obtain our approximation and fixed-parameter tractable algorithms by reducing minimum cut completion to 2CNF deletion. The 2CNF deletion problem asks: given an instance of 2SAT find the minimum number of clauses whose deletion makes the instance satisfiable. We apply the following two lemmas to obtain our results. In in the statements of the following lemmas n is the number of variables and m is the number of clauses.

Lemma 6 (Agarwal et al. [1])**.** *There is a polynomial time algorithm for finding an* $O(\sqrt{\log n})$*-approximation for the 2CNF deletion problem.*

Lemma 7 (Razgon and O'Sullivan [85])**.** *For an instance of 2CNF deletion that admits a solution of size* k *there is an* $O(15^k \cdot k \cdot m^3)$ *time algorithm solving the instance.*

It is worth noting that the algorithm in Lemma 6 is randomized as it relies on a randomized semidefinite relaxation presented by Arora, Rao, and Vazirani [6]. We are now ready to show our reduction.

Lemma 8. *For the cut completion problem* $(G = (V, E), E')$*,*

1. *there is a polynomial time* $O(\sqrt{\log |V|})$*-approximation algorithm, and*
2. *there is an* $O(15^k \cdot k \cdot |E|^3)$ *time exact algorithm, where* k *is the size of the optimal solution.*

Proof. Given a cut completion instance $(G = (V, E), E')$ we construct a 2CNF deletion instance B_G in polynomial time. Moreover for any cut $(S, \overline{S})$ with edge set E_S the number of unsatisfied clauses in B_G is $|E_S \oplus E'|$. The result follows from Lemmas 6 and 7.

We construct our 2CNF deletion instance B_G from (G, E') as follows. For each vertex $v \in V$ we have a variable $b(v)$. For each edge $(u, v) \in E$ we consider two cases. If $(u, v) \in E'$ we add the clauses $b(u) \vee b(v)$ and $\neg b(u) \vee \neg b(v)$ to B_G. If $(u, v) \notin E'$ we add the clauses $b(u) \vee \neg b(v)$ and $\neg b(u) \vee b(v)$ to B_G. In both cases any truth assignment satisfies at least one of the clauses, moreover there exists truth assignments of $b(u)$ and $b(v)$ that satisfy both clauses.

Let $(S, \overline{S})$ be a cut with edge set E_S. Let b_S be the indicator vector, indexed by $v \in V$ where $b_S(v) = 1$ if and only if $v \in S$. Note that B_S is a truth assignment for B_G. We will show that $|E_S \oplus E'|$ is equal to the number of clauses that are not satisfied by b_S in B_G. Specifically we show that for each edge $(u, v) \in E_S \oplus E'$ exactly one of its corresponding clauses is satisfied and for each $(u, v) \notin E_S \oplus E'$ both of its corresponding clauses are satisfied.

If $(u, v) \in E_S \oplus E'$ there are two cases to consider. In the first case we have $(u, v) \in E_S$ and $(u, v) \notin E'$, so we have the clauses $b(u) \vee \neg b(v)$ and $\neg b(u) \vee b(v)$ in B_G. Given any truth assignment for $b(u)$ and $b(v)$ exactly one of these clauses is satisfied. In the second case we have $(u, v) \notin E_S$ and $(u, v) \in E'$, so we have the clauses $b(u) \vee b(v)$ and $\neg b(u) \vee \neg b(v)$.

Since $(u, v) \in E_S$ both $b(u)$ and $b(v)$ have the same truth assignment, hence only one of the clauses is satisfied.

When $(u, v) \notin E_S \oplus E'$ there are also two cases to consider. In the first case we have $(u, v) \in E_S$ and $(u, v) \in E'$, so we have clauses $b(u) \vee b(v)$ and $(\neg b(u) \vee \neg b(v)$ in B_G. Since $(u, v) \in E_S$ we have that $b(u)$ and $b(v)$ are given opposite assignments, so both clauses are satisfied. Finally, when $(u, v) \notin E_S$ and $(u, v) \notin E'$ we have clauses $b(u) \vee \neg b(v)$ and $\neg b(u) \vee b(v)$. In this case $b(u)$ and $b(v)$ are given the same truth assignment, hence both clauses are satisfied. □

Theorems 7, 8, 9, and 10 follow from the observation that the dual graph of an embedded d-complex has $\beta_d + 1$ vertices and n_d edges, and the dual graph of a $(d + 1)$-weak pseudomanifold has $n_{d+1} + 1$ vertices and n_d edges.

4.3 Hardness of approximation

In this section we show hardness of approximation results for both minimum bounded chain and minimum homologous chain. Our hardness results hold for low dimensions; $d = 2$ for minimum bounded chain and $d = 1$ for minimum homologous chain. In our result for minimum bounded chain we point out that the problem is still hard even when the input cycle in contained in the boundary of V_∞. This restriction on the cycle is notable as it closely resembles the assumption made for the polynomial time special case discussed in Section 4.4. Our result for the minimum homologous chain problem shows hardness of approximation for the homology localization problem on an orientable manifold.

In Sections 4.3.1 and 4.3.2 we reduce the minimum bounded chain and minimum homologous chain problems to the minimum cut completion problem defined in 4.2.1. We show hardness of approximation results for the minimum cut completion problem in Section 4.3.3.

4.3.1 Minimum bounded chain to minimum cut completion

We show that the minimum cut completion problem reduces to a 2-dimensional instance of the minimum bounded chain problem $(\mathcal{K}, \gamma)$, where the boundary of V_∞ is in fact a manifold and γ is a (possibly not connected) cycle on the boundary of V_∞. Our hardness of approximation result for the minimum bounded chain problem is based on this reduction.

The high level idea motivating our reduction is rather simple. Given a graph $G = (V, E)$ we construct a 2-complex as follows. For each each vertex $v \in V$ we have a topological sphere $\hat{v}$ punctured with $\deg(v)$ boundary components, and for each $e \in E$ we have a topological sphere $\hat{e}$ punctured with two boundary components. Finally, we glue the boundary components of each $\hat{e}$ to the boundary components corresponding to the endpoints of e. The resulting topological space is a surface with genus equal to the dimension of $H_1(G)$. In the reduction we will geometrically construct each $\hat{v}$ as a cube and each $\hat{e}$ as a tube connecting two cubes. Finally, inside each tube $\hat{e}$ we insert a wall dividing the tube into two halves and call the resulting complex $\mathcal{K}$. We have that $\dim(H_2(\mathcal{K})) = |V|$, and each generator of $H_2(\mathcal{K})$ is a cube along with its attached half-tubes. Moreover, the dual graph $\mathcal{K}^*$ is isomorphic to G along with the extra vertex v_∞. Our reduction naturally follows by identifying a subset $E' \subseteq E$ with the cycle consisting of all boundaries of dividers on tubes corresponding with E'. See Figure 4.1 for an illustration (up to homeomorphism) of our construction. We are now ready to give our formal construction.

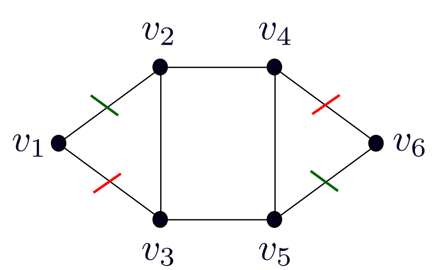

(a) The edges marked in green correspond to the set E' in the cut completion problem. The edges marked in red are chosen to complete the cut.

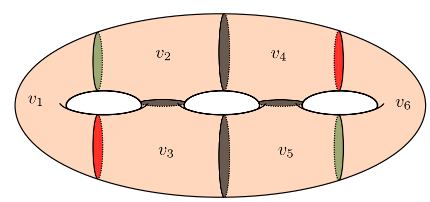

(b) The graph transformed into a surface. The cut is represented by the green and red cycles whose symmetric difference is null-homologous.

Figure 4.1

Lemma 9. *Let $(G = (V, E), E')$ be any instance of the minimum cut completion problem. There exists an instance of the 2-dimensional minimum bounded chain problem $(\mathcal{K}, \gamma)$ with γ on the boundary of V_∞ that can be computed in polynomial time, and a one-to-one correspondence between cuts in G and 2-chains with boundary γ in $\mathcal{K}$. Moreover, if the cut $(S, \overline{S})$ with edge set E_S in G corresponds to the 2-chain Q in $\mathcal{K}$ then*

$$\frac{|Q|}{\tau} - 1 \leq |E_S \oplus E'| \leq \frac{|Q|}{\tau},$$

where $\tau = 58m + 2$ *and* m *is the number of edges in* G.

Proof. Our construction is simple in high-level. We start from any embedding of G in $\mathbb{R}^3$, and we thicken it to obtain a space, in which each edge corresponds to a tube. We insert a disk in the middle of each tube; we call these disks *edge disks*. Then we triangulate all of the 2-dimensional pieces. The dual graph of the complex that we build is almost G, except for one extra vertex corresponding to V_∞, and a set of extra edges, all incident to the extra vertex v_∞. We give our detailed construction below.

We consider the following piecewise linear embedding of G in $\mathbb{R}^3$; let n and m be the number of vertices and edges of G, respectively. First, map the vertices of G into $\{1, 2, \ldots, n\}$ on the x-axis. Now, consider $m + 2$ planes $h_0, h_1, \ldots, h_{m+1}$ all containing the x-axis with normals being evenly spaced vectors ranging from $(0, 1, 1)$ to $(0, 1, -1)$. We use $h_1, \ldots, h_m$ for drawing the edges G. We arbitrarily assign edges of G to these plane, so each plane will contain exactly one edge. Each edge is drawn on its plane as a three-segment curve; the first and the last segment are orthogonal to x-axis and the middle one is parallel. All edges are drawn in the upper half-space of $\mathbb{R}^3$. See Figure 4.5, left[1].

Next, we place an axis parallel cube around each vertex. The size of the cubes must be so that they do not intersect, fix the width of each cube to be $1/10$. We refer to these cubes as *vertex cubes* Then, we replace the part of each edge outside the cubes with a cubical tube, called *edge tube*. We choose the thickness of these tubes sufficiently small so that they are disjoint. We also puncture the cubes so that the union of all vertex cubes and edges tubes form a surface; see Figure 4.5, left. (This surface will have genus $m - n + 1$ by Euler's formula, which is the dimension of the cycle space of G)

Next, we subdivide each tube by placing a square in its middle; see Figure 4.5, right. We refer to these squares as *edge squares*. Edge squares partition the inside of the surface into n volumes. We observe that each of these volumes contains exactly one vertex of the drawing of G, thus, we call them *vertex volumes*.

For our reduction to work, we need the weight of each 2-cycle to be dominated by the weight of its edge squares. To achieve that we finely triangulate each edge square. For an edge tube, we first subdivide its surface to 16 quadrangles as shown in Figure 4.3, left. Then, we obtain a triangulation with 32 triangles by splitting each quadrangle into two triangles. For a vertex cube, note that all the punctures are on the top face by our construction.

[1]Thanks to Amir Nayyeri for Figures 4.5, 4.3, and 4.4 in this chapter, which appeared in the published version of this work [13].

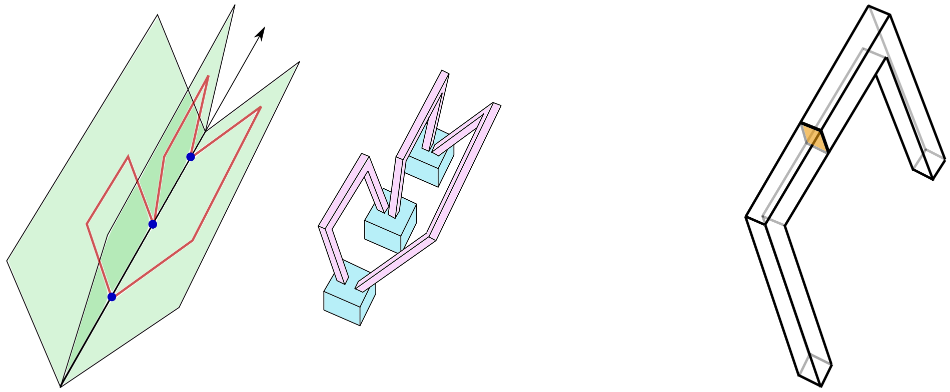

Figure 4.2: Left: an embedding of K_3 in $\mathbb{R}^3$, and the thickened surface composed of blue vertex cubes and pink edge tubes, right: an edge tube subdivided by an edge square.

We split all the other faces by dividing each of them into two triangles. For the top face, we can obtain a triangulation in polynomial time; this triangulation will have $4\deg(v)+8$ triangles by Euler's formula, where $\deg(v)$ is the degree of the vertex corresponding to the cube. Therefore, the triangulation of each vertex cube will have $4\deg(v)+18$ triangles, see Figure 4.3, right. Therefore, there are $\left(\sum_{v\in V} 4\deg(v)+18\right)+32m \leq 58m$ triangles that are not part of edge squares. Finally, we triangulate each edge square into $58m+2$ triangles so that the cost of one edge square is greater than the sum of all triangles not contained in edge squares. This triangulation can be done efficiently by subdividing triangles. The subdivision is performed by inserting a vertex into the interior of the triangle and connecting it with an edge to each vertex on the boundary of the triangle. The result is a new complex, homeomorphic to the original, with two additional triangles. Overall, our complex $\mathcal{K}$ has $O(m^2)$ triangles.

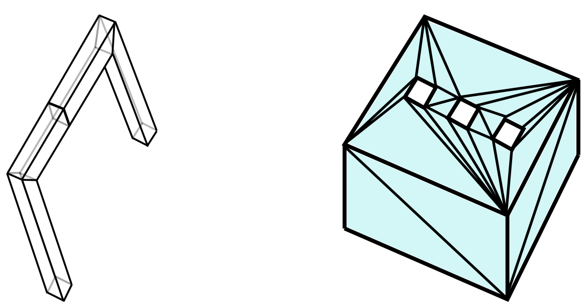

Figure 4.3: Left: subdividing the surface of an edge-tube to quadrangles, right: triangulating the surface of a vertex cube.

We are now done with the construction of $\mathcal{K}$. Let B be the set of all triangles in edge squares that correspond to edges in E'. Then, let $\gamma = \partial B$. We show an (almost) cost preserving one-to-one correspondence between cuts in the cut completion problem in G and chains with boundary γ in $\mathcal{K}$.

Let $(S, \overline{S})$ be a cut with edge set E_S, note that the cost of this cut is $|E_S \oplus E'|$ in the cut completion problem (G, E'). In $\mathcal{K}$, let $\mathcal{V}_S$ be the symmetric difference of the vertex volumes that correspond to vertices of S. The total weight of $\mathcal{V}_S$ is between $|E_S| \cdot (58m+2)$ and $|E_S| \cdot (58m+2) + 58m$. Similarly, the total weight of $\mathcal{V}_S \oplus B$ is between $|E_S \oplus E'| \cdot (58m+2)$ and $|E_S \oplus E'| \cdot (58m+2) + 58m$. Since we cannot get an exact count on the number of edges in the subgraph induced by S we have a range of values for the weight of $\mathcal{V}_S$ instead of an exact weight. However, if E_S and $E_{S'}$ are two cuts with $|E_S| < |E_{S'}|$ then the weight of $\mathcal{V}_S$ is strictly less than the weight of $\mathcal{V}_{S'}$ by the construction of the edge squares.

On the other hand, let Q be a 2-chain with boundary γ in $\mathcal{K}$. As γ does not intersect the interior of any edge square, for each edge square either Q contains all of its triangles or none of them. Also, $Q \oplus B$ has no boundary, thus its complement $\mathbb{R}^3 \setminus (Q \oplus B)$ is disconnected. The interior of each vertex volume is completely inside one of the connected components of $\mathbb{R}^3 \setminus (Q \oplus B)$, as by the construction $Q \oplus B$ must either contain the entire vertex volume or none of it. Now, let S be the set of all vertices whose corresponding vertex volumes are in the unbounded connected component of $\mathbb{R}^3 \setminus (Q \oplus B)$. The edges of the cut $(S, \overline{S})$ correspond to edge squares in $Q_s \oplus B$, where Q_s is the set of edge square triangles of Q. As B is in one-to-one correspondence to E', it follows that the cut completion cost of $(S, \overline{S})$ is $\frac{|Q_s|}{58m+2}$. We have $|Q| = |Q_s| + |Q_r|$ where Q_r is the set of triangles in Q not contained in edge squares. The size of $|Q_s|$ is $58m+2$ per edge square, and $|Q_r| \leq 58m$ by construction. It follows that we have our desired inequality,

$$\frac{Q}{58m+2} - 1 \leq |E_S \oplus E'| \leq \frac{Q}{58m+2}.$$

□

The next lemma shows that an approximation algorithm for the minimum bounded chain problem implies an approximation algorithm with almost the same quality for the minimum cut completion problem.

Lemma 10. *Let $(G = (V, E), E')$ be any instance of the minimum cut completion problem. For any $\alpha \geq 1$ and any $\varepsilon > 0$, there exists an instance of the 2-dimensional min-*

imum bounded chain problem $(\mathcal{K}, \gamma)$ that can be computed in polynomial time, such that an α-approximation algorithm for $(\mathcal{K}, \gamma)$ implies a $((1+\varepsilon)\alpha)$-approximation algorithm for (G, E'), and γ is on the boundary of V_∞.

Proof. Let $\varepsilon > 0$. Given an α-approximation algorithm for the minimum bounded chain problem, we describe an $((1+\varepsilon)\alpha)$-approximation algorithm for the cut completion problem. Let $G = (V, E)$, and $E' \subseteq E$ be any instance of the cut completion problem, and let $(S_{opt}, \overline{S_{opt}})$ with edge set be an optimal solution for this instance. Our algorithm considers two cases, based on whether $|E_{S_{opt}} \oplus E'| < 1/\varepsilon$ or not. It solves the problem under each assumption and outputs the best solution it obtains in the end.

If $|E_{S_{opt}} \oplus E'| < 1/\varepsilon$, then our algorithm finds the optimal solution in $O(n^{1/\varepsilon+O(1)})$ time by considering all subsets of edges E'' of size at most $1/\varepsilon$ as candidates for $E_{S_{opt}} \oplus E'$. From all candidates, we return the minimum E'' such that $E'' \oplus E'$ is a cut. Note this is an exact algorithm, so in this case we find the optimal solution.

Otherwise, if $|E_{S_{opt}} \oplus E'| \geq 1/\varepsilon$, we use the given α-approximation algorithm for the minimum bounded chain problem for a simplicial complex $\mathcal{K}$, and chain γ that corresponds to (G, E') by Lemma 9. Note that $\mathcal{K}$ is an unweighted simplicial complex piecewise linearly embedded in $\mathbb{R}^3$ and γ is a cycle contained in the boundary of V_∞.

Let Q_{opt} be the corresponding 2-chain to $(S_{opt}, \overline{S_{opt}})$ in $\mathcal{K}$. Thus, $\frac{Q_{opt}}{\tau} - 1 \leq |E_{S_{opt}} \oplus E'| \leq \frac{|Q_{opt}|}{\tau}$. In addition, let Q be the surface with boundary γ that the α-approximation algorithm finds, so $|Q| \leq \alpha \cdot |Q_{opt}|$. Finally, let $(S, \overline{S})$ be the cut corresponding to Q in G via the one-to-one correspondence of Lemma 9. Therefore, $\frac{Q}{\tau} - 1 \leq |E_S \oplus E'| \leq \frac{|Q|}{\tau}$. Putting everything together,

$$|E_S \oplus E'| \leq \frac{|Q|}{\tau} \leq \alpha \cdot \frac{|Q_{opt}|}{\tau} \leq \alpha \cdot \big(|E_{S_{opt}} \oplus E'| + 1\big). \tag{4.1}$$

Since $|E_{S_{opt}} \oplus E'| \geq 1/\varepsilon$, we have: $|E_{S_{opt}} \oplus E'| + 1 \leq (1+\varepsilon) \cdot |E_{S_{opt}} \oplus E'|$. Therefore, together with (4.1), we have a $((1+\varepsilon)\alpha)$-approximation algorithm, as desired. $\square$

4.3.2 Minimum homologous cycle to minimum cut completion

We show a similar reduction from the cut completion problem to the minimum homologous cycle problem for 1-dimensional cycles on orientable 2-manifolds. The minimum homologous cycle problem is the special case of the minimum homologous chain problem when the input chain is required to be a cycle, so showing hardness of approximation for it implies hardness of approximation for the more general minimum homologous chain problem.

Lemma 11. *Let $(G = (V, E), E')$ be any instance of the minimum cut completion problem. For any $\alpha \geq 1$, there exists an instance of the* 1*-dimensional minimum homologous cycle problem $(\mathcal{M}, \gamma)$ that can be computed in polynomial time such that an α-approximation for $(\mathcal{M}, \gamma)$ implies an α-approximation for (G, E').*

Proof. We construct a 2-manifold $\mathcal{M}$ as in the proof of Lemma 9, but we omit the edge squares. Each edge of G corresponds to a cycle with 4 edges in $\mathcal{M}$; these cycles are the boundaries of the omitted edge squares. We call these cycles *edge rings*. The connected components of $\mathcal{M}$ after removing the edge rings correspond to the vertices of G, we call these connected components *vertex regions*. We set γ to be equal to the set of edge rings corresponding to E'. Intuitively, if τ is the minimum cycle homologous to γ we do not want $\tau \oplus \gamma$ to intersect the interior of any vertex region. That is, $\tau \oplus \gamma$ is a collection of edge rings and corresponds to a cut in G. To achieve this, we subdivide each edge not contained in an edge ring into a long path. The result is an embedded graph with non-triangular faces, which is not a simplicial complex. To fix this, we triangulate the inside of each non-triangular face such that the shortest path between any two vertices on the face remains the shortest path after the triangulation. Given any α-approximation of the new complex we can obtain a smaller solution using only the edge rings, which corresponds to a cut in G. Our formal construction follows.

Let $\tau = 4\lceil\alpha\rceil|E| + 1$; we subdivide each edge not contained in an edge ring τ times. For each face of length $\ell > 3$ we triangulate by adding $\ell + 1$ concentric cycles, each with ℓ vertices, labeled $\gamma_0, \ldots, \gamma_\ell$, where γ_0 is the original face from the subdivided version of $\mathcal{M}$. By $v_{i,j}$ we denote the jth vertex in γ_i. We add the edges $(v_{i,j}, v_{i+1,j}$ and $(v_{i,j}, v_{i,j+1 \mod \ell})$. To complete the triangulation we add one additional vertex $\overline{v}$ at the center of γ_ℓ and add an edge between it and each vertex on γ_ℓ. We call the new simplicial complex $\mathcal{M}'$. See Figure 4.4 for an example.

Let $(S_{opt}, \overline{S_{opt}})$ be an optimal solution to the minimum cut completion instance (G, E'). Suppose we can compute an α-approximation τ_{apx} of the minimum homologous cycle instance $(\mathcal{M}', \gamma)$, hence $|\tau_{apx}| \leq \alpha|\tau_{opt}|$. By our construction an optimal solution to $(\mathcal{M}', \gamma)$ has the same size as an optimal solution to $(\mathcal{M}, \gamma)$. As τ_{apx} is a cycle, if τ_{apx} crosses a cycle γ_0 it must cross it an even number of times. For any two consecutive vertices $u, v \in \gamma_0$ in τ_{apx} we replace the path between them with the shortest path contained in γ_0. We call the

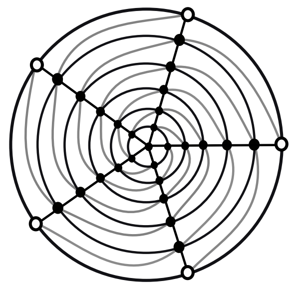

Figure 4.4: Subdividing a face of length five; the outer face with white vertices is the original face.

new cycle τ'_{apx}, since $\tau'_{apx} \leq \tau_{apx}$ we have that τ'_{apx} is also an α-approximation for $(\mathcal{M}', \gamma)$. Note that τ'_{apx} is a union of edge rings, otherwise $|\tau'_{apx}| > \alpha|\tau_{opt}|$. It follows that τ'_{apx} corresponds to a cut $E_{S'}$ with $|\tau'_{apx}| = 4|E_{S'} \oplus E'|$. Hence, we have $|E_{S'} \oplus E'| \leq \alpha|E_{S_{opt}} \oplus E'|$. Thus, $E_{S'}$ is an α-approximation for (G, E'). □

4.3.3 Hardness of minimum cut completion

It remains to show that the cut completion problem is hard to approximate. We show this via a straightforward reduction from the *minimum uncut problem*: given a graph $G = (V, E)$, find a cut with minimum number of uncut edges. Note that the optimal cuts for the minimum uncut problem and the maximum cut problem coincide, yet, approximation algorithms for one problem do not necessarily imply approximation algorithm for the other one.

Lemma 12. *The minimum uncut problem is a special case of the minimum cut completion problem.*

Proof. Consider the cut completion problem for $G = (V, E)$, and let $E' = E$. Let $(S, \overline{S})$ be any cut with edge set E_S. The cut completion cost of this cut is

$$|E_S \oplus E'| = |E_S \oplus E| = |E \setminus E_S|,$$

which is the number of uncut edges by $(S, \overline{S})$. □

Now, we are ready to prove our hardness results.

Proof of Theorem 11 and 12. The minimum uncut problem is hard to approximate within a factor of $(1+\varepsilon)$ for some $\varepsilon > 0$ [83]. In addition, it is hard to approximate within any constant factor assuming the unique games conjecture [73, 63, 22, 64]. By Lemma 12, the cut completion problem generalizes the minimum uncut problem. Finally, by Lemma 11 and 10, for any $\alpha > 1$ and $\varepsilon > 0$, an α-approximation algorithm for the minimum bounded chain problem or the minimum homologous cycle problem implies a $((1+\varepsilon)\alpha)$-approximation algorithm for the cut completion problem. $\square$

4.4 A polynomial time special case

We have shown hardness results, approximation algorithms and parameterized algorithms for the minimum bounded chain problem. We showed that the problem is hard to approximate even when the input cycle γ is on the boundary of V_∞ of an unweighted 2-manifold embedded in $\mathbb{R}^3$. If γ is null-homologous on the boundary of V_∞ there is an exact polynomial time algorithm to find the minimum chain bound by γ. The assumption that γ is null-homologous on the boundary of V_∞ allows us to treat the problem as a generalization of the shortest st-path problem in planar graphs when s and t are contained on the boundary of the unbounded face. Hence, we can generalize the duality between shortest paths and minimum cuts in planar graphs to d-complexes embedded in $\mathbb{R}^{d+1}$. The algorithm was first found by Kirsanov and Gortler in the context of continuous variational problems [66]. We provide a combinatorial proof of their theorem which is more useful in the context of algorithms. Before describing the algorithm we prove the following lemma about graph cuts, which will be useful in the proof of correctness of the algorithm.

Lemma 13. *Let $(S, \overline{S})$ and $(S', \overline{S'})$ be two (s,t)-cuts of a graph G with edge sets E_S and $E_{S'}$, respectively. The symmetric difference $E_S \oplus E_{S'}$ is the set of edges of a cut that has s and t on the same side.*

Proof. We show the edge set of the cut $(S \oplus S', \overline{S \oplus S'})$ is $E_S \oplus E_{S'}$. The statement follows as $s, t \in \overline{S \oplus S'}$.

Let $e = (u,v) \in E_S \oplus E_{S'}$. Either, $e \in E_S$ and $e \notin E_{S'}$ or $e \notin E_S$ and $e \in E_{S'}$.

In the first case, there are two possibilities up to symmetry of (u,v). Either $u \in S \cap S'$ and $v \in \overline{S} \cap S'$, which implies $u \in \overline{S \oplus S'}$ and $v \in S \oplus S'$, or $u \in S \cap \overline{S'}$ and $v \in \overline{S} \cap \overline{S'}$, which implies $u \in S \oplus S'$ and $v \in \overline{S \oplus S'}$.

In the second case, there are again two possibilities up to symmetry of (u,v). Either

$u \in S \cap S'$ and $v \in S \cap \overline{S'}$, which implies $u \in \overline{S \oplus S'}$ and $v \in S \oplus S'$, or $u \in \overline{S} \cap S'$ and $v \in \overline{S} \cap \overline{S'}$, which implies $u \in S \oplus S'$ and $v \in \overline{S \oplus S'}$. □

By $\mathrm{Bd}(V_\infty)$ we denote the set of d-simplices on the boundary of V_∞. Let $F \subseteq \mathrm{Bd}(V_\infty)$ be a d-chain such that $\partial F = \gamma$. Such an F exists by the assumption of this section. We define a cut problem based on F. Let $\mathcal{K}^*$ be the dual graph of the complex $\mathcal{K}$. By the definition of $\mathrm{Bd}(V_\infty)$, each edge of $\mathrm{Bd}(V_\infty)^*$ is adjacent to v_∞. We build the graph H from $\mathcal{K}^*$ by splitting v_∞ as follows. We replace v_∞ with two vertices v_∞^+ and v_∞^-. We replace the incident edges to v_∞ as follows:

(i) A loop that is dual to a face in F is replaced by a (v_∞^-, v_∞^+) edge.

(ii) A loop that is dual to a face *not* in F is replaced by v_∞^+-loops.

(iii) A non-loop edge (v_∞, u) that is dual to a face in F is replaced by a (v_∞^-, u)-edge.

(iv) A non-loop edge (v_∞, u) that is dual to a face *not* in F is replaced by a (v_∞^+, u)-edge.

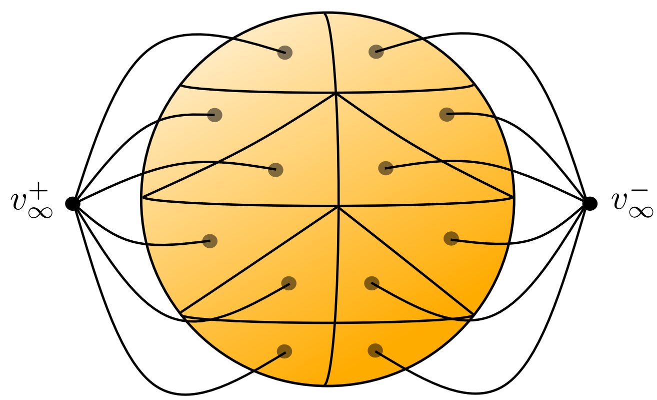

Figure 4.5: The modified dual graph of a simplicial complex whose outer shell is a triangulated sphere. The vertical line represents the boundary input boundary γ which partitions $\mathrm{Bd}(V_\infty)$ into two regions.

Note that all of the faces of F correspond to edges that are incident to v_∞^-. We are now ready to prove the main theorem of the section.

Theorem 13. *Let $\mathcal{K}$ be a simplicial complex embedded in $\mathbb{R}^{d+1}$ and γ be a null-homologous $(d-1)$-cycle in $\mathrm{Bd}(V_\infty)$. A d-chain Γ is a minimum d-chain bounded by γ if and only if Γ^* is a minimum (v_∞^+, v_∞^-)-cut in H.*

Proof. We show a one-to-one correspondence between d-chains with boundary γ in $\mathcal{K}$ and (v_∞^-, v_∞^+)-cuts in H that preserves the cost. Let Γ be a d-chain with $\partial\Gamma = \gamma$. Since γ is null-homologous in $\mathrm{Bd}(V_\infty)$, there exists $F \subseteq \mathrm{Bd}(V_\infty)$ such that $\partial F = \gamma$. It follows that $\partial(\Gamma \oplus F) = \partial\Gamma \oplus \partial F = 0$, that is $D \oplus F$ is a d-cycle. Thus, by cycle/cut duality $\Gamma^* \oplus F^*$ is a cut in $\mathcal{K}^*$ that partitions the vertices into two sets X and Y. Assume, without loss of generality, that $v_\infty \in X$, and note that by splitting v_∞ we obtain a $(X\backslash\{v_\infty\}\cup\{v_\infty^-, v_\infty^+\}, Y)$-cut. Hence, $\Gamma^* \oplus F^*$ is a cut in both $\mathcal{K}^*$ and H.

We show that any simple $v_\infty^- v_\infty^+$-path of H crosses Γ^*, therefore Γ^* is a $v_\infty^- v_\infty^+$-cut. Let $\beta = (v_\infty^- = v_0, v_1, \ldots, v_k = v_\infty^+)$ be a simple $v_\infty^- v_\infty^+$-path in H. Let α be the closed simple cycle in $\mathcal{K}^*$ obtained by identifying v_0 and v_k in β. Since α is a closed cycle and $\Gamma^* \oplus F^*$ is a cut in $\mathcal{K}^*$, α crosses $\Gamma^* \oplus F^*$ an even number of times. Therefore, β crosses $\Gamma^* \oplus F^*$ in H an even number of times; as each edge of β is in $\Gamma^* \oplus F^*$ in H if and only if the corresponding edge of it in α is in $\Gamma^* \oplus F^*$ in $\mathcal{K}^*$. On the other hand, $v_0 = v_\infty^-$ is only incident to edges from F^*. In particular, $(v_0, v_1) \in F^*$. If $(v_0, v_1) \in \Gamma^*$, then β crosses Γ^* and so the statement holds. Otherwise, if $(v_0, v_1) \notin \Gamma^*$, then the path $(v_1, \ldots, v_k)$ must cross $\Gamma^* \oplus F^*$ at least once. Since all F-edges are incident to v_0 and β is simple we have $v_i \neq v_0$ for any $0 < i \leq k$. Therefore, $(v_1, \ldots, v_k)$ must cross Γ^* and so the statement holds.

Conversely, let Γ^* be a $v_\infty^- v_\infty^+$-cut in H. Since F^* is composed of all edges incident to v_∞^-, it is a $v_\infty^- v_\infty^+$-cut as well. It follows that $\Gamma^* \oplus F^*$ is a cut in H that has v_∞^- and v_∞^+ on the same side by Lemma 13. Therefore, $\Gamma^* \oplus F^*$ is a cut in $\mathcal{K}^*$; obtained after identifying v_∞^- and v_∞^+. Now, by cycle/cut duality, $\Gamma \oplus F$ is a cycle in $\mathcal{K}$, that is $\partial(\Gamma \oplus F) = 0$. As $\partial F = \gamma$, we have $\partial\Gamma = \gamma$ and the proof is complete. $\square$

Now we compute the runtime of the presented algorithm. The time required to perform the minimum cut computation dominates the preprocessing we perform on the dual graph. A minimum st-cut in a graph with n vertices and m edges can be computed in $O(nm)$ time via the maximum flow algorithm of Orlin [81]. If $\mathcal{K}$ has m facets then the dual graph $\mathcal{K}^*$ will have m edges. The number of vertices in H is equal to $\beta_d + 2$. Hence, we can compute the cut in $O(\beta_d \cdot m)$ time.

Chapter 5: Generalized Flows and Cuts

Computing flows and cuts are fundamental algorithmic problems in graphs. In this chapter we explore generalizations of these algorithmic problems in higher dimensional simplicial complexes. Flows and cuts in simplicial complexes have natural algebraic definitions arising from the theory of simplicial (co)homology. A flow is an element of the kernel of the simplicial boundary operator, and a cut is an element of the image of the simplicial coboundary operator. Note that when working with flows we assume coefficients over the field $\mathbb{R}$. These subspaces serve as generalizations of the cycle and cut spaces of a graph. Note that when working with flows we assume coefficients over the field $\mathbb{R}$. This generalization has been studied by Duval, Klivans, and Martin in the setting of CW complexes [35]. We formulate the algorithmic problems of computing max-flows and min-cuts algebraically. By forgetting about the underlying graph structure and focusing on the (co)boundary operators, we obtain methods that naturally generalize to high dimensions. The topological study of max-flows and min-cuts has been done by Ghrist and Krishnan [45]. They prove a max-flow min-cut theorem for graphs using the directed homology theory; which is a way to formalize the intuitive idea of "homology over the natural numbers". Of course, the natural numbers do not form a group and the intuitive idea is not well-defined. In contrast, we work with the standard notion of homology and provide our max-flow min-cut theorem in the context of linear programming.

In a graph an st-flow is an assignment of real values to the edges satisfying the conservation of flow constraints: the net flow out of any vertex other than s and t is zero, and, thus, the net flow that leaves s is equal to the net flow that enters t. Therefore, each st-flow can be viewed as a circulation in another graph with an extra edge that connects t to s. Circulations are elements of the cycle space of the graph with coefficients taken over $\mathbb{R}$. In a d-dimensional simplicial complex $\mathcal{K}$ the d-dimensional cycles are the formal sums (over $\mathbb{R}$) of d-dimensional simplices whose boundary is zero. Because there are no $(d+1)$-simplices flows are the elements of the dth homology group $H_d(\mathcal{K}, \mathbb{R})$. The maximum flow problem in a simplicial complex asks to find an optimal element of $H_d(\mathcal{K}, \mathbb{R})$ subject to capacity constraints.

The max-flow min-cut theorem states that in a graph the value of a maximum st-flow is equal to the value of a minimum st-cut. This result is a special case of linear programming duality. By rewriting the linear program in terms of the (co)boundary operator we obtain a similar result for simplicial complexes. The question of whether or not a similar max-flow min-cut theorem holds for simplicial complexes was asked, and left open, in a paper by Latorre [70]. We give a positive answer to this question, but with a caveat. When viewing flows and cuts from a topological point of view their linear programs are dual to one another. However, we also provide a combinatorial definition of a cut which feels more natural for a minimization problem. Topological and combinatorial cuts are equivalent for graphs, but they become different in dimensions $d > 1$. Flows in higher dimension, are dual to topological cuts, but not combinatorial cuts in general. From a computational complexity viewpoint the two notions of cuts are very different. We show that computing a minimum topological cut can be solved via linear programming, but that computing a minimum combinatorial cut is NP-hard.

A closely related problem is the problem of computing a max-flow in a graph which admits an embedding into some topological space. The most well-studied cases are planar graphs and the more general case when the graph embeds into a surface [12, 18, 19, 40, 44, 50, 51, 57, 58, 76, 88]. Max-flows and min-cuts are computationally easier to solve in surface embedded graphs, especially planar graphs. We consider this problem generalized to simplicial complexes. Planar graphs are 1-dimensional complexes embedded in $\mathbb{R}^2$, in Section 5.3 we consider the special case when a d-dimensional simplicial complex admits an embedding into $\mathbb{R}^{d+1}$. These complexes naturally admit a dual graph which we use to compute maximum flows and minimum cuts (both topological and combinatorial). We show that a maximum flow in a simplicial complex can be found by solving a shortest paths problem in its dual graph. This idea was used by Hassin to solve the maximum flow problem in planar graphs [50]. Further, we show that finding a minimum topological cut can be done by finding a minimum cost circulation in its dual graph. By setting the demand equal to one in the minimum cost circulation problem we obtain an algorithm computing a minimum combinatorial cut.

Maximum flows in graphs can be computed using the Ford-Fulkerson algorithm. Moreover, the fact that the Ford-Fulkerson algorithm halts serves as a proof that there exists a maximum integral flow when the graph has integral capacity constraints. In dimensions $d > 1$ the maximum flow may be fractional, even with integral capacity constraints. The problem arises due to the existence of torsion in simplicial complexes of dimension $d > 1$. We show that despite the maximum flow being fractional the Ford-Fulkerson algorithm halts on simplicial complexes. However, in order for it to halt a special heuristic on picking the high dimensional analog of an augmenting path must be implemented. Despite the algorithm halting it could we could not prove a polynomial upper bound on the number of iterations it takes.

5.1 Flows and cuts

In this section we give an overview of our generalizations of flows and cuts from graphs to simplicial complexes. Flows and cuts in higher dimensional settings have been studied previously. Duval, Klivans, and Martin have generalized cuts and flows to the setting of CW complexes [35]. Their definitions are algebraic; defining cuts to be elements of $\mathrm{im}(\delta)$ and flows to be elements of $\ker(\partial)$. Our definitions are closely related, but are motivated by the algorithmic problems of computing maximum flows and minimum cuts. In Section 5.1.1 we give definitions of flows and cuts from from the perspective of algebraic topology, and in Section 5.1.2 we give a combinatorial definition of a cut in a simplicial complex. The distinction between the two types of cuts will be important when formulating the minimum cut problem on simplicial complexes.

5.1.1 Topological flows and cuts

First we briefly recall the definition of an st-flow in a directed graph $G = (V, E)$. An st-flow f is a function $f\colon E \to \mathbb{R}$ satisfying the conservation of flow constraint: for all $v \in V \setminus \{s, t\}$ we have $\sum_{(u,v)\in E} f(u, v) = \sum_{(v,u)\in E} f(v, u)$. That is, the amount of flow entering the vertex equals the amount of flow leaving the vertex. The value of f is equal to the amount of flow leaving s (or equivalently, entering t). Alternatively, we may view f as a 1-chain and we have $\partial f = k(t - s)$ where k is the value of f. Note that $t - s$ is a null-homologous 0-cycle. More generally, for any null-homologous $(d - 1)$-cycle γ we call a d-chain f with $\partial f = k\gamma$ a γ-flow of value k. Note that under our naming convention an "st-flow" in a graph would be called a $(t - s)$-flow. However, in the case of graphs we use the traditional naming convention and call a flow from s to t an st-flow.

Definition 3 (γ-flow)**.** *Let $\mathcal{K}$ be a d-dimensional simplicial complex and γ be a null-homologous $(d-1)$-cycle in $\mathcal{K}$. A γ-**flow** is a d-chain f with $\partial f = k\gamma$ where $k \in \mathbb{R}$. We call k the value of the flow f and denote the value of f by $\|f\|$. We say that f is feasible with respect to a capacity function $c\colon \mathcal{K}_d \to \mathbb{R}^+$ if $0 \le f(\sigma) \le c(\sigma)$ for all $\sigma \in \mathcal{K}^d$.*

Our definition of a γ-flow is very close to the algebraic definition which is element of $\ker(\partial)$. Given a simplicial complex $\mathcal{K}$ and a γ-flow f of value k we convert f into a *circulation*, where a circulation is defined to be an element of $\ker(\partial)$. To convert f into a circulation we add an additional basis element to $C_d(\mathcal{K})$, call it Σ, whose boundary is $\partial\Sigma = -\gamma$. This operation is purely algebraic; we should think of it as operating on the chain complex rather than the underlying topological space. Now we construct the circulation $f' = f + k\Sigma$. We call any circulation built from a γ-flow a γ-circulation. Clearly, $f' \in \ker(\partial)$ in the new chain complex. Moreover, there is a clear bijection between γ-flows and γ-circulations. The value of the circulation is the value of $f'(\Sigma)$, so this bijection preserves the value.

We now shift our focus to the generalization of cuts to a simplicial complex. The algebraic definition, elements of $\mathrm{im}(\delta)$, is natural. The cut space of a graph is commonly defined to be the space spanned by the coboundaries of each vertex. In a simplicial complex $\mathcal{K}$, removing the support of a d-chain in $\mathrm{im}(\delta)$ increases $\dim H_{d-1}(\mathcal{K})$. In a graph G, removing the support of any 1-chain in $\mathrm{im}(\delta)$ increases $\dim H_0(G)$ which is equivalent to increasing the number of connected components of G.

The above definition implies that a cut is a d-chain in a d-dimensional simplicial complex. However, for our purposes we will define a cut to be a $(d-1)$-cochain. To motivate our definition we recall the notion of an st-cut in a graph. An st-cut in a graph is a partition of the vertices into sets S and T such that $s \in S$ and $t \in T$. Define $p\colon V(G) \to \{0, 1\}$ such that $p(v) = 1$ if $v \in S$ and $p(v) = 0$ if $v \in T$. The support of the coboundary of p is a set of edges whose removal destroys all st-paths. That is, upon removing the support, the 0-cycle $t - s$ is no longer null-homologous. Moreover, p is a 0-cochain with $p(t - s) = -1$. The sign of $p(t - s)$ will be important when we consider directed cuts. With this in mind we define our notion of a γ-cut.

Definition 4 (γ-cut)**.** *Let $\mathcal{K}$ be a d-dimensional simplicial complex with weight function $c\colon \mathcal{K}_d \to \mathbb{R}^+$ and γ be a null-homologous $(d-1)$-cycle in $\mathcal{K}$. A γ-**cut** is a $(d-1)$-cochain p such that $p(\gamma) = -1$. Denote the coboundary of p as the formal sum $\delta(p) = \sum \alpha_i \sigma_i$, we define the size of a γ-cut to be $\|p\| = \sum |\alpha_i c(\sigma_i)|$.*

Because of the requirement that $p(\gamma) = -1$ we call p a ***unit*** γ***-cut***. By relaxing this requirement to $p(\gamma) < 0$ the cochain p still behaves as a γ-cut, but its size can become arbitrarily small by multiplying by some small value $\epsilon > 0$. We justify our definition with the following proposition which shows that removing the support of the coboundary of a

γ-cut prevents γ from being null-homologous.

Proposition 1. *Let $\mathcal{K}$ be d-dimensional simplicial complex and p be a γ-cut. The cycle γ is not null-homologous in the subcomplex $\mathcal{K} \setminus \operatorname{supp}(\delta(p))$.*

Proof. By way of contradiction let Γ be a d-chain in $\mathcal{K} \setminus \operatorname{supp}(\delta(p))$ such that $\partial\Gamma = \gamma$. Since $\delta(p) = 0$ in $\mathcal{K} \setminus \operatorname{supp}(\delta(p))$, we have that $\langle \Gamma, \delta(p) \rangle = 0$ in $\mathcal{K} \setminus \operatorname{supp}(\delta(p))$. However, this implies that $0 = \langle \Gamma, \delta(p) \rangle = \langle \partial\Gamma, p \rangle = \langle \gamma, p \rangle = p(\gamma) = 0$, a contradiction as p is a γ-cut and $p(\gamma) = -1$. $\square$

5.1.2 Combinatorial cuts

Alternatively, we can view a γ-cut as a discrete set of d-simplices rather than a d-chain. In the case of graphs a combinatorial st-cut is just a set of edges whose removal disconnects s from t. This distinction will become important when we consider the minimization problem of finding a minimum cost set of d-simplices whose removal prevents γ from being null-homologous.

Definition 5 (Combinatorial γ-cut)**.** *Let $\mathcal{K}$ be a d-dimensional simplicial complex with weight function $c\colon \mathcal{K}_d \to \mathbb{R}^+$ and γ be a null-homologous $(d-1)$-cycle in $\mathcal{K}$. A **combinatorial γ-cut** is a set of d-simplices $C \subseteq \mathcal{K}^d$ such that γ is not null-homologous in $\mathcal{K} \setminus \operatorname{supp}(C)$. The size of a combinatorial γ-cut is defined by the sum of the weights of the d-simplices $\|C\| = \sum_{\sigma \in C} c(\sigma)$.*

Lemma 1 in Chapter 3 shows a relationship between γ-cuts and combinatorial γ-cuts. Removing a combinatorial γ-cut C from $\mathcal{K}$ increases $\dim H_{d-1}(\mathcal{K})$. This is because removing C must decrease the rank of ∂_d and by duality this also decreases the rank of δ_d which increases the dimension of $H^{d-1}(\mathcal{K}) \cong H_{d-1}(\mathcal{K})$. It follows that C must contain the support of some coboundary. Given an additional minimality condition on C we show that C is equal to the support of some coboundary.

In graphs the linear program solving the minimum st-cut problem takes as input a directed graph and returns a set of directed edges whose removal destroys all directed st-paths. This is called a directed cut. After removing the directed cut the 0-cycle $t - s$ may still be null-homologous; we can find a 1-chain with boundary $t - s$ using negative coefficients to traverse an edge in the backwards direction. In order to generalize the minimum cut linear program to simplicial complexes we will need to define a directed combinatorial γ-cut, which requires the additional assumption that the d-simplices of $\mathcal{K}$ are oriented.

Definition 6 (Directed combinatorial γ-cut)**.** *Let $\mathcal{K}$ be an oriented d-dimensional simplicial complex with weight function $c\colon \mathcal{K}_d \to \mathbb{R}$ and γ be a null-homologous $(d-1)$-cycle in $\mathcal{K}$. A **directed combinatorial γ-cut** is a set of d-simplices $C \subset \mathcal{K}^d$ such that in $\mathcal{K}\setminus\operatorname{supp}(C)$ there exists no d-chain Γ with non-negative coefficients such that $\partial\Gamma = \gamma$. The size of a directed combinatorial γ-cut is defined by the sum of the weights of the d-simplices $\|C\| = \sum_{\sigma\in C} c(\sigma)$.*

Given a directed graph consider an st-cut given by the cochain definition. That is, a 0-cochain $p\colon V(G) \to \{0,1\}$ with $p(s) = 1$ and $p(t) = 0$ partitioning V into S and T. The support of $\delta(p)$ consists of two types of edges: edges leaving S and entering T, and edges leaving T and entering S. If $e \in E$ leaves S and enters T we have $(p \circ \partial)(e) = -1$ and if e leaves T and enters S we have $(p \circ \partial)(e) = 1$. To construct a directed st-cut we simply take all of the edges mapped to -1. The following proposition shows that we can build a directed combinatorial γ-cut from a coboundary just like in the case of directed graphs.

Proposition 2. *Let p be a γ-cut with coboundary $\delta(p) = \sum \alpha_i\sigma_i$. The set of d-simplices $C = \{\sigma_i \mid \alpha_i < 0\}$ is a directed combinatorial γ-cut.*

Proof. By way of contradiction let Γ be a non-negative d-chain in $\mathcal{K} \setminus C$ with $\partial\Gamma = \gamma$. By the definition of C we have $\langle\Gamma, \delta(p)\rangle \geq 0$ in $\mathcal{K} \setminus C$. Construct a new chain complex by adding an additional basis element Σ to $C_d(\mathcal{K})$ such that $\partial\Sigma = -\gamma$. By construction $\langle\Sigma, \delta(p)\rangle = -p(\gamma) = 1$, hence we have $\langle\Gamma + \Sigma, \delta(p)\rangle > 0$. However, $\Gamma + \Sigma$ is a d-cycle and $\delta(p)$ is a d-coboundary, so the Hodge decomposition ensures that they are orthogonal. Hence, $\langle\Gamma + \Sigma, \delta(p)\rangle = 0$, a contradiction. □

To conclude the section we will show that computing a minimum combinatorial γ-cut is NP-hard. As we will see in Section 5.2 minimum topological γ-cuts can be computed with linear programming. Our hardness result holds for both the directed and undirected cases. Our hardness result is a reduction from the well-known NP-hard hitting set problem which we will now define. Given a set S and a collection of subsets $\Sigma = (S_1, \ldots, S_n)$ where $S_i \subseteq S$ the hitting set problem asks to find the smallest subset $S' \subseteq S$ such that $S' \cap S_i \neq \emptyset$ for all S_i. We call such a subset S' a *hitting set* for (S, Σ).

Theorem 14. *Let $\mathcal{K}$ be a d-dimensional simplicial complex and γ be a null-homologous $(d-1)$-cycle. Computing a minimum combinatorial γ-cut is NP-hard for $d \geq 2$.*

Proof. Our proof is a reduction from the hitting set problem. First we consider the case when $d = 2$ then we generalize to any $d \geq 2$. Let S be a set and $\Sigma = (S_1, \ldots, S_n)$ where each $S_i \subseteq S$. We construct a 2-dimensional simplicial complex $\mathcal{K}$ from S and Σ in the following way. For each $S_i \in \Sigma$ construct a triangulated disk $\mathcal{D}_i$ such that $\partial\mathcal{D}_i = \gamma$. That is, each $\mathcal{D}_i$ shares the common boundary γ. To accomplish this we construct each $\mathcal{D}_i$ by beginning with a single triangle t with $\partial t = \gamma$ and repeatedly adding a new vertex in the center of some triangle with edges connecting it to every vertex in that triangle. By this process we can construct a disk containing any odd number of triangles as each step increments the number of triangles in the disk by two. Moreover, at each step the boundary of the disk is always γ. We construct each disk $\mathcal{D}_i$ such that $\mathcal{D}_i$ consists of one triangle $t_{i,s}$ for each element $s \in S_i$ and potentially one extra triangle t'_i in the case that $|S_i|$ is even. Next, for each $s \in S$ and S_i with $s \in S_i$, we construct the quotient space by identifying each $t_{i,s}$ into a single triangle. A minimum combinatorial γ-cut C must contain exactly one triangle from each $\mathcal{D}_i$ and without loss of generality we can assume C does not contain any t'_i. If $t'_i \in C$ then by minimality it is the only triangle in $C \cap \operatorname{supp}(\mathcal{D}_i)$ and we can swap it with any other triangle in $\mathcal{D}_i$ without increasing the size of the cut. By construction C is a hitting set for (S, Σ) since each $C \cap \operatorname{supp}(\mathcal{D}_i) \neq \emptyset$ for all $\mathcal{D}_i$.

To perform the above construction in higher dimensions we simply start with a single d-simplex $\sigma = [v_1, \ldots, v_{d+1}]$ with boundary $\partial\sigma = \gamma$. We subdivide σ in the following way: add an additional vertex v_{d+2} and replace σ with the d-simplices $\sigma_i := [v_1, \ldots, v_{i-1}, v_{d+2}, v_{i+1}, \ldots, v_{d+1}]$ for each $1 \leq i \leq d+1$. At each step we increase the number of d-simplices by d; moreover, at each step the complex remains homeomorphic to a d-dimensional disk. Our final complex will have at most d extra d-simplices so for any fixed dimension d the size of the complex is within a constant factor of the size of the given hitting set instance. It remains to show that the subdivision process does not change the boundary of the disk. To accomplish this we will show that $\partial\sigma = \sum_{i=1}^{d+1} \partial\sigma_i$.

For each σ_i the boundary $\partial\sigma_i$ has d terms in its sum; each term is a $(d-1)$-simplex which consists of d vertices. Consider the matrix A such that the entry $A_{i,j}$ contains the jth term of $\partial\sigma_i$. Note that no term in the jth column will contain the vertex v_j. Also note that the sum of the diagonal of A is equal to $\partial\sigma$. Below is an example for $d = 3$.

$$\begin{bmatrix} v_2v_3v_4 & -v_5v_3v_4 & v_5v_2v_4 & -v_5v_2v_3 \\ v_5v_3v_4 & -v_1v_3v_4 & v_1v_5v_4 & -v_1v_5v_3 \\ v_2v_5v_4 & -v_1v_5v_4 & v_1v_2v_4 & -v_1v_2v_5 \\ v_2v_3v_5 & -v_1v_3v_5 & v_1v_2v_5 & -v_1v_2v_3 \end{bmatrix}$$

The matrix A is almost symmetric. The entries $A_{i,j}$ and $A_{j,i}$ contain the same vertices but possibly differ up to a sign or a permutation of the vertices. We want to show that $A_{i,j} = -A_{j,i}$ so that the sum $\sum_{i=1}^{d+1} \partial\sigma_i$ only contains diagonal which is equal to $\partial\sigma$. The ordering of vertices in $A_{i,j}$ differs from $A_{j,i}$ by the placement of v_{d+2}, which we now characterize. Let $v_{i,j}$ denote the vertex in σ_i that is not included in the term $A_{i,j}$. Without loss of generality assume $A_{i,j}$ is in the upper triangle. Then v_{d+2} is in the ith position of $A_{i,j}$ because $v_{i,j}$ appears after it in the ordering of σ_i. It follows that v_{d+2} is in the $(j-1)$th position in the term $A_{j,i}$ since $v_{j,i}$ must appear before v_{d+2} in the ordering of σ_j. So, the position of v_{d+2} in $A_{i,j}$ and $A_{j,i}$ differs by $|i-j+1|$ and this is the number of transpositions needed to permute $A_{i,j}$ into $A_{j,i}$. When $i \equiv j \mod 2$ the terms $A_{i,j}$ and $A_{j,i}$ have the same sign, but differ by an odd permutation so $A_{i,j} = -A_{j,i}$. Similarly, when $i \not\equiv j \mod 2$ the terms $A_{i,j}$ and $A_{j,i}$ have opposite signs, but differ by an even permutation so $A_{i,j} = -A_{j,i}$ which concludes the proof. □

5.2 Linear programming

In this section we model max-flow and min-cut for simplicial complexes as linear programming problems. In Section 5.2.1 we prove a duality theorem reminiscent of the well-known theorem for graphs. However, our theorem comes with a caveat. A max-flow is dual to a directed combinatorial cut arising from a topological cut. That is, a max-flow is dual to a coboundary $\delta(p) = \sum_{i=1}^{n_d} \alpha_i\sigma_i$ minimizing the quantity $\|\delta^-(p)\| = \sum_{\alpha_i<0} |\alpha_i|$. This is the coboundary, or topological cut, which minimizes the absolute value its negative coefficients. In the case of graphs this is equivalent to a min-cut. This is due to the fact that the boundary matrix of a graph is always totally unimodular. If we restrict our domain to topological cuts the linear program finds a minimum directed cut. However, this is not the case for combinatorial cuts.

In Section 5.2.2 we show that there exist simplicial complexes with integral capacities whose maximum flow value is fractional. This is unsurprising given that Theorem 2 tells us that the boundary matrix any simplicial complex with relative torsion cannot be totally unimodular. However, the explicit example we construct is useful in the proof of Theorem 16 which states that computing an integral max-flow is NP-hard.

In Section 5.2.3 we show that there may exist combinatorial cuts of smaller magnitude than what is returned by the linear program. Note that in this case we are viewing the output of the linear program as a combinatorial cut. This statement should also be unsurprising given the lack of total unimodularity in the boundary matrix. In the unweighted case the size of a combinatorial cut of a cycle is upper bound by the number of chains bounding the cycle, which is integral. However, in the presence of relative torsion the linear program may find a coboundary with fractional coefficients which allows for the possibility of a coboundary whose support exceeds the number of bounding chains of the cycle. We show an explicit example where this happens. Finally, in Theorem 17 we show that when the complex is relative torsion-free the linear program returns a minimum combinatorial cut, which gives us the same duality as the case for graphs.

5.2.1 Max-flow min-cut

A ***simplicial flow network*** is a tuple $(\mathcal{K}, c, \gamma)$ where $\mathcal{K}$ is an oriented d-dimensional simplicial complex, c is the ***capacity function*** which is a non-negative function $c\colon \mathcal{K}_d \to \mathbb{R}$, and γ is a null-homologous $(d-1)$-cycle. In a simplicial flow network we work with real coefficients; that is, we consider the chain groups $C_k(\mathcal{K}, \mathbb{R})$. In order to utilize the Hodge decomposition (Theorem 1) we modify $C_d(\mathcal{K})$ by adding an additional basis element Σ such that $\partial\Sigma = -\gamma$. Moreover, we extend the capacity function such that $c(\Sigma) = \infty$. This allows us to work with circulations instead of flows while leaving the solution unchanged. The notation n_d will refer to the number of basis elements in $C_d(\mathcal{K}, \mathbb{R})$ which is now one more than the number of d-simplices in the underlying simplicial complex.

The goal of the maximum flow problem is to find a d-chain f obeying the capacity constraints such that $\partial f = k\gamma$ where $k \in \mathbb{R}$ is maximized. Equivalently, we find a d-cycle f which maximizes $f(\Sigma)$. The linear program for the max-flow problem in a simplicial flow network is identical to the familiar linear program for graphs, but expressed in terms of the coboundary operator. In a graph, conservation of flow at a vertex v is the constraint $\delta_1(v) \cdot f = 0$; to formulate the linear program in higher dimensions we simply replace vertices with $(d-1)$-simplices. The Hodge decomposition states that cycles are orthogonal to coboundaries, so conservation of flow ensures that f is indeed a cycle. We now state the linear program for max-flow in a simplicial flow network.

$$\begin{array}{lll} \text{maximize} & f(\Sigma) & \\ \text{subject to} & \delta(\tau)\cdot f = 0 & \text{for each } \tau \in \mathcal{K}_{d-1} \\ & 0 \leq f(\sigma) \leq c(\sigma) & \text{for each } \sigma \in \mathcal{K}_d \end{array} \tag{LP1}$$

We dualize LP1 to obtain a generalization of the minimum cut problem in directed graphs. To make the dualization more explicit we will write out LP1 in matrix form: maximize $s \cdot f$ subject to $Af \leq b$ and $f \geq 0$, where we have

$$A = \begin{bmatrix} \partial \\ -\partial \\ I_{n_d} \end{bmatrix}, \; b = \begin{bmatrix} 0_{n_{d-1}} \\ 0_{n_{d-1}} \\ c \end{bmatrix}, \; s = \begin{bmatrix} 0_{n_d - 1} \\ 1 \end{bmatrix}.$$

The matrix A has dimension $(2n_{d-1} + n_d) \times n_d$. In our notation I_k is the $k \times k$ identity matrix and 0_k is the $k \times 1$ column vector consisting of all zeros. Since the value of the flow is equal to $f(\Sigma)$ the vector s is all zeros except for the final entry which is indexed by Σ and receives an entry equal to one. Further, c is the $n_d \times 1$ capacity vector indexed by the d-simplices such that the entry indexed by σ has value equal to $c(\sigma)$.

We can now state the dual program in matrix form: minimize $y \cdot b$ subject to $y^T A \geq s$ and $y \geq 0$. The vector y is a $(2n_{d-1} + n_d) \times 1$ column vector indexed by both the $(d-1)$-simplices and the d-simplices. However, only the entries indexed by d-simplices contribute to the objective function since b is zero everywhere outside of the capacity constraints. We will denote the truncated vector consisting of entries indexed by d-simplices by y_d and the entry corresponding to the d-simplex $\sigma \in \mathcal{K}_d$ will be denoted by $y_d(\sigma)$. Similarly we have two truncated vectors y^1_{d-1} and y^2_{d-1} corresponding to the entries indexed by the $(d-1)$-simplices. Moreover, the rows of $y^T A \geq s$ are in the form

$$(y^1_{d-1} - y^2_{d-1})^T \partial + y_d \geq s.$$

For simplicity we define $y_{d-1} = y^1_{d-1} - y^2_{d-1}$ and write $y_{d-1}(\tau)$ for the entry indexed by the $(d-1)$-simplex τ. Putting this together, we state the dual linear program as follows.

$$\begin{aligned}
\text{minimize} \quad & \sum_{\sigma \in \mathcal{K}^d} y_d(\sigma) c(\sigma) \\
\text{subject to} \quad & y_{d-1} \cdot \partial\sigma + y_d(\sigma) \geq 0 \quad \text{for each } \sigma \in \mathcal{K}^d \\
& y_{d-1} \cdot \partial\Sigma + y_d(\Sigma) = 1 \\
& y_d \geq 0
\end{aligned} \tag{LP2}$$

Note the strict equality in the second constraint does not follow from the duality. However, we can assume a strict equality since if $y_{d-1} \cdot \partial\Sigma + y_d(\Sigma) > 1$ we can multiply $[y_{d-1}, y_d]^T$ by some scalar $\epsilon < 1$ to make the inequality tight. This multiplication only decreases the value of $\sum y_d(\sigma)c(\sigma)$ so it does not change the optimal solution.

In the case of graphs LP2 has dual variables for vertices and edges. Moreover, there exists an integral solution such that each vertex is either assigned a 0 or a 1 since a graph cut is a partition of the vertices. The second inequality requires $y_0(s) = 1$ and $y_0(t) = 0$. To see this, when solving an st-cut on a graph, the basis element Σ is an edge with $\partial\Sigma = s - t$, and $y_1(\Sigma) = 0$ otherwise the solution is infinite. This naturally defines a partition of the vertices: S containing vertices assigned a 1, and T containing vertices assigned a 0. The constraints force an edge to be assigned a 1 if it leaves S and enters T, otherwise it is assigned a 0. This solution can be interpreted as a 0-cochain p with $p(st) = 1$, or in the notation of our definition of a simplicial cut: $p(t - s) = -1$. Further, $y_1(e) = 1$ for every edge e that is negative on $\delta(p)$ and a 0 otherwise, hence y_1 fits our definition of a directed st-cut in a 1-complex. We will show the same result holds in higher dimensions; that is, y_d is a directed γ-cut arising from the $(d-1)$-cochain y_{d-1}.

Lemma 14. *Let $y = [y_{d-1}, y_d]^T$ be an optimal solution to LP2. The set* $\mathrm{supp}(y_d)$ *is a directed combinatorial γ-cut.*

Proof. Note that y_{d-1} can be interpreted as a $(d-1)$-cochain. Since $c(\Sigma) = \infty$ we have $y_d(\Sigma) = 0$, otherwise the solution is infinite. It follows from the second constraint that $-y_{d-1}(\gamma) = y_{d-1}(-\gamma) = y_{d-1} \cdot \partial\Sigma = 1$, hence $y_{d-1}(\gamma) = -1$ making y_{d-1} a γ-cut. Expanding $\delta(y_{d-1})$ into a linear combination of d-simplices $\sum \alpha_i \sigma_i$ and applying the first inequality constraint gives us the set equality $\mathrm{supp}(y_d) = \{\sigma_i \mid \alpha_i < 0\}$ since $y_d(\sigma_i) = 0$ precisely when $\alpha_i > 0$. Thus, the result follows from Lemma 1. ∎

Lemma 15. *Let p be a γ-cut with coboundary $\delta(p) = \sum \alpha_i \sigma_i$ and let $\delta(p)^- = \sum_{\alpha_i < 0} \alpha_i \sigma_i$. The vector $[p, -\delta(p)^-]^T$ is a finite feasible solution to LP2.*

Proof. We view the cochains p and $-\delta(p)^-$ as vectors and let them take the roles of y_{d-1} and y_d, respectively. The first constraint is satisfied since for all $\sigma \in \mathcal{K}_d$ we have $-\delta(p)^-(\sigma) = -p \cdot \partial(\sigma)$ when $p \cdot \partial(\sigma)$ is negative, and $-\delta(p)^-(\sigma) = 0$ otherwise. The second constraint and the finiteness of the solution is satisfied by the fact that $p(\gamma) = -1$. □

Lemma 14 tells us that a solution to LP2 yields a directed combinatorial γ-cut. Recall, by Proposition 2 every γ-cut p yields a directed combinatorial γ-cut by taking the coboundary $\delta(p) = \sum \alpha_i \sigma_i$ and considering the negative components $\delta(p)^- = \{\sigma_i \mid \alpha_i < 0\}$. By Lemma 15 $\delta(p)^-$ is a feasible solution to LP2; the cost of this solution is $c \cdot \delta(p)^-$. The coefficients α_i need not always equal one; hence in general we have $\|C\| \neq c \cdot \delta(p)^-$. It follows that LP2 need not return a minimum directed combinatorial γ-cut. In Theorem 17 we will give conditions describing when LP2 returns a directed combinatorial γ-cut. To conclude the section we state our main theorem about LP2 whose proof is immediate from Lemmas 14 and 15.

Theorem 15. *Let $y = [y_{d-1}, y_d]^T$ be an optimal solution to LP2. The set* $\operatorname{supp}(y_d)$ *is a directed combinatorial γ-cut such that $y_d = \delta(y_{d-1})^-$. Moreover, y_d minimizes $c \cdot \delta(p)^-$ where p ranges over all γ-cuts.*

5.2.2 Integral solutions

In this section we provide an example of a simplicial flow network with integral capacity constraints and fractional maximum flow. By Theorem 2 such a network must contain some relative torsion. This is achieved by the inclusion of a Möbius strip in our simplicial flow network. Our example will be used later in Theorem 16 showing that computing a maximum integral flow in a simplicial flow network is NP-hard.

Fractional maximum flow We will now explicitly describe a simplicial flow network with integral capacities whose maximum flow value is fractional. Let $\mathcal{M}$ be a triangulated Möbius strip with boundary $\partial\mathcal{M} = 2\alpha + \gamma$ such that two vertices in α have been identified making α a simple cycle. This identification makes γ a figure-eight. Now let $\mathcal{D}$ be a triangulated disk oriented such that $\partial\mathcal{D} = -\alpha$. See Figure 5.1 for an illustration. Call the resulting complex $\mathcal{MD}$. The capacity function c has $c(t) = 1$ for each triangle $t \in \mathcal{MD}$. Now we solve the max-flow problem on $(\mathcal{MD}, c, \gamma)$. Note that for any flow f we have $f(t_1) = f(t_2)$ for all triangles $t_1, t_2 \in \mathcal{M}$; moreover, for all $t_1, t_2 \in \mathcal{D}$ we also have

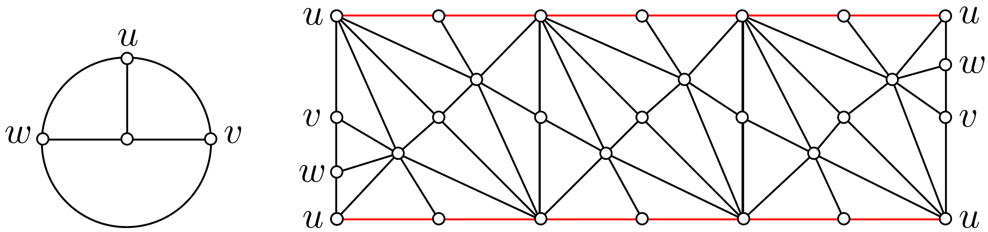

Figure 5.1: A triangulated disk $\mathcal{D}$ (left) and Möbius strip $\mathcal{M}$ (right). The Möbius strip has two points on its boundary identified forming the vertex u. In red we have the input cycle (a figure-eight) γ and we set $\alpha = uwvu$. We orient the complex such that $\partial\mathcal{M} = \gamma + 2\alpha$ and $\partial\mathcal{D} = -\alpha$. The capacity on each simplex in both the disk and Möbius strip is one.

$f(t_1) = f(t_2)$. The value of any flow f on $(\mathcal{MD}, c, \gamma)$ is equal to its value on $\mathcal{M}$, and in order to maintain conservation of flow we must have $f(\mathcal{D}) = 2f(\mathcal{M})$. Now, the capacity constraints imply that the maximum flow f has $f(\mathcal{M}) = 1/2$ and $f(\mathcal{D}) = 1$. We have $\partial f = \frac{1}{2}\partial\mathcal{M} + \partial\mathcal{D} = \frac{1}{2}\gamma + \alpha - \alpha$. Hence, the value of f is equal to 1/2.

Maximum integral flow Given a simplicial flow network $(\mathcal{K}, c, \gamma)$ with integral capacities we consider the problem of finding the maximum integral flow. That is, a d-chain $f \in C_d(\mathcal{K}, \mathbb{Z})$ obeying the capacity constraints such that $\partial f = k\gamma$ where $k \in \mathbb{Z}$ is maximized. We show the problem is NP-hard by a reduction from graph 3-coloring. Our reduction is inspired by a MathOverflow post from Sergei Ivanov showing that finding a subcomplex homeomorphic to the 2-sphere is NP-hard [55]. In the appendix we adapt the proof to show that the high dimensional generalization of computing a directed path in a graph is also NP-hard. Given a graph G we construct a 2-dimensional simplicial flow network whose maximum flow is integral if and only if G is 3-colorable.

Theorem 16. *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network where $\mathcal{K}$ is a 2-complex and c is integral. Computing a maximum integral flow of $(\mathcal{K}, c, \gamma)$ is NP-hard.*

Proof. Let $G = (V, E)$ be a graph. We will construct a simplicial flow network $(\mathcal{K}, c, \gamma)$ such that its maximum flow is integral if and only if G is 3-colorable.

We start our construction with a punctured sphere $\mathcal{S}$ containing $|V| + 1$ boundary components called γ and β_v for each $v \in V$. For each boundary component β_v we construct three disks R_v, B_v, G_v each with boundary $-\beta_v$. These disks represent the three colors in our coloring: red, blue, and green. We refer to these disks as *color disks* and use $\mathcal{C}_v$ to denote an arbitrary color disk associated with v and use $k \in \{r, b, g\}$ to denote an arbitrary color. On each color disk $\mathcal{C}_v$ we add a boundary component for each edge $e = (u, v)$ incident to v. By $\beta_{v,e,k}$ we denote the boundary component corresponding to the vertex v, edge e, and color k. For each edge $e = (u, v)$ and each pair of boundary components β_{u,e,k_u} and β_{v,e,k_v} with $k_u \neq k_v$ we construct a tube with boundary components $-\beta_{u,e,k_u}$ and $-\beta_{v,e,k_v}$ denoted $\mathcal{T}_{e,k_u,k_v}$. When $k_u = k_v = k$ we construct a tube $\mathcal{T}_{e,k,k}$ and puncture it with a third boundary component α and construct a negatively oriented real projective plane $\mathcal{RP}_{e,k}$ with boundary $\partial\mathcal{RP}_{e,k} = -2\alpha$. We call the resulting complex $\mathcal{K}$ and assign a capacity $c(\sigma) = 1$ for every triangle σ in $\mathcal{K}$.

We will show that a maximum integral flow f of $\mathcal{K}$ has value equal to 1 if and only if G is 3-colorable. The following four properties of a maximum integral flow f imply that G is 3-colorable.

- f must assign exactly one unit of flow to each triangle in $\mathcal{S}$ since the value of f is equal to $f(\mathcal{S})$.
- For each vertex $v \in V$ exactly one color disk $\mathcal{C}_v$ is assigned one unit of flow while the other two color disks associated with v are assigned zero units of flow. Otherwise, either conservation of flow is violated or some color disk is assigned a fractional flow value.
- For each edge $e = (u, v) \in E$ exactly one tube $\mathcal{T}_{e,k_u,k_v}$ with $k_u \neq k_v$ must be assigned one unit of flow with all other tubes associated with e assigned zero units of flow. The tube $\mathcal{T}_{e,k_u,k_v}$ assigned one unit of flow is the tube connecting the color disks $\mathcal{C}_v$ and $\mathcal{C}_u$ that are assigned one unit of flow by the previous property.
- f assigns zero flow to every $\mathcal{T}_{e,k,k}$ and $\mathcal{RP}_{e,k}$ since otherwise the triangles in $\mathcal{RP}_{e,k}$ would need to have $1/2$ units of flow assigned to them to maintain conservation of flow.

These four properties imply that the set of color disks $\{\mathcal{C}_v \mid f(\sigma) = 1, \forall \sigma \in \mathcal{C}_v\}$ corresponds to a 3-coloring of G. Conversely, given a 3-coloring of G we assign a flow value of one to each color disk corresponding to the 3-coloring. We extend this assignment to a γ-flow of value one by assigning a flow value of one to $\mathcal{S}$ and the tubes corresponding to the 3-coloring. $\square$

5.2.3 Integral cuts

The goal of this section is to show that for simplicial complexes relative torsion-free in dimension $d-1$ there exists optimal solutions to LP2 whose support is a minimum combinatorial γ-cut. Note that by Theorem 2 a simplicial complex that is relative torsion-free in dimension $d-1$ has a totally unimodular d-dimensional boundary matrix. The total unimodularity is key to our proof. However, we first provide an example of a complex (with relative torsion) whose optimal solution's support does not form a minimum combinatorial γ-cut. Our construction is a slight modification of $\mathcal{MD}$ defined in Section 5.2.2.

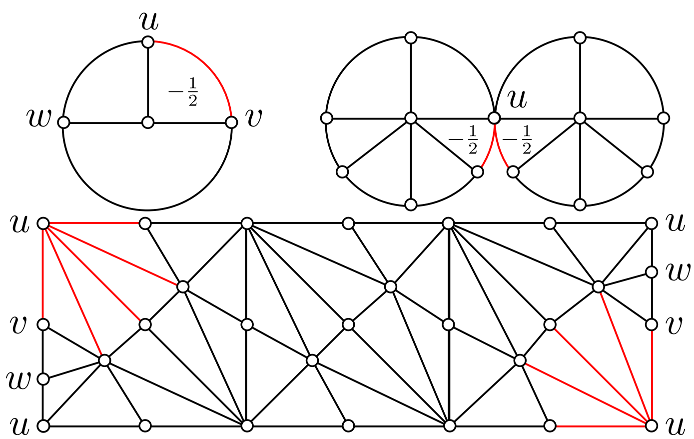

Figure 5.2: The simplicial complex $\mathcal{MD}$ with a wedge sum of two disks $\mathcal{W}$ identified to the figure-eight γ. In red we have a 1-cochain which assigns a value of $-1/2$ to each red edge. The coboundary of the red cochain assigns a value of $-1/2$ to one triangle in $\mathcal{D}$ and a value of $-1/2$ to two triangles in $\mathcal{W}$. The value of the red cochain coincides with the value of the maximum γ-flow. However, its support is not a minimum combinatorial γ-cut. A minimum combinatorial γ-cut picks one triangle from $\mathcal{D}$ and one triangle from $\mathcal{W}$.

Consider the simplicial complex constructed by taking $\mathcal{MD}$ and glueing a wedge sum of two disks $\mathcal{W}$ along the figure-eight γ. That is, $\partial\mathcal{W} = \gamma$. We give every triangle in the resulting complex a capacity equal to one. A maximum γ-flow has value 3/2, so the dual program finds a γ-cut of the same value. One potential optimal solution is a $(d-1)$-cochain whose coboundary assigns a value of $-1/2$ to two triangles in $\mathcal{W}$ and a value of $-1/2$ to one triangle in $\mathcal{D}$. The support of this coboundary has weight equal to three, however a minimal combinatorial γ-cut has weight two by taking only one triangle from $\mathcal{W}$ and one from $\mathcal{D}$. See Figure 5.2 for an illustration.

Now, we show that when $\mathcal{K}$ is relative torsion-free in dimension $d-1$ LP2 has an optimal solution whose support is a minimum directed combinatorial γ-cut. Specifically, we show that a solution existing on a vertex of the polytope defined by the constraints of LP2 is a cochain y_{d-1} with negative coboundary y_d such that $y_d(\sigma) \in \{0,1\}$ for all $\sigma \in \mathcal{K}^d$ hence $\sum y_d(\sigma)c(\sigma) = \|\operatorname{supp}(y_d)\|$. That is, the value of a vertex solution to LP2 is equal to the cost of $\operatorname{supp}(y_d)$ as a directed combinatorial γ-cut.

Theorem 17. *Let $\mathcal{K}$ be d-dimensional simplicial complex that is relative torsion-free in dimension $d-1$ and $y = [y_{d-1}, y_d]^T$ be an optimal vertex solution to the dual program. The set* $\operatorname{supp}(y_d)$ *is a minimum directed combinatorial γ-cut.*

Proof. We can write the constraint matrix of LP2 as the $2n_d \times (n_d + n_{d-1})$ block matrix

$$A = \begin{bmatrix} \delta & I_{n_d} \\ 0_{n_d} & I_{n_d} \end{bmatrix}.$$

Since $\mathcal{K}$ is relative torsion-free in dimension $d-1$ Theorem 2 tells us that ∂_d is totally unimodular; further, we have that $\partial^T = \delta$ is also totally unimodular. Total unimodularity is preserved under the operation of adding a row or column consisting of exactly one component equal to 1 and the remaining components equal to 0, so A is totally unimodular [89, Section 19.4]. We write LP2 as the linear system $Ax \geq b$ where b is a $n_d + n_{d-1}$ dimensional vector with exactly one component equal to 1 and the remaining components equal to 0. Let $y = [y_{d-1}, y_d]^T$ be an optimal vertex solution to LP2. For every $(d-1)$-simplex $\tau \in \mathcal{K}_{d-1}$ we either have $y_{d-1}(\tau) \geq 0$ or $y_{d-1}(\tau) \leq 0$. Let $I'_{n_{d-1}}$ be the matrix whose rows correspond to these inequalities. Note that $I'_{n_{d-1}}$ is a diagonal matrix with entries in $\{-1,1\}$. Now we consider the $(2n_d + n_{d-1}) \times (n_d + n_{d-1})$ dimensional linear system $A'x \geq b'$ where

$$A' = \begin{bmatrix} \delta & I_{n_d} \\ 0 & I_{n_d} \\ I'_{n_{d-1}} & 0 \end{bmatrix}$$

and b' is constructed by appending extra zeros to b. We construct y' from y similarly. Note that A' is totally unimodular and y' is a vertex solution of the system. There exists a vertex v of the polyhedron $P \subseteq \mathbb{R}^{n_d+n_{d-1}}$ corresponding to the linear system such that $A'y' = v \geq b'$ such that $n_{d-1} + n_d$ constraints are linearly independent and tight. Hence, there is an $(n_{d-1} + n_d) \times (n_{d-1} + n_d)$ square submatrix A'' with $A''y' = b''$ where b'' is b' restricted to the tight constraints. We will use Cramer's rule to show that the vertex solution y has components coming from the set $\{-1,0,1\}$. Let $A''_{i,b''}$ be the matrix obtained by replacing the i^{th} column of A'' with b''. By Cramer's rule we compute the i^{th} component of y as $y_i = \frac{\det(A''_{i,b''})}{\det(A'')}$. Since both $A''_{i,b''}$ and A'' are totally unimodular we have $v_i \in \{-1,0,1\}$. Further, we know that A'' is non-singular because it corresponds to linearly independent constraints.

By the above argument we know that an optimal solution y to LP2 has all of its components contained in the set $\{-1,0,1\}$. The constraint $y_d \geq 0$ means that for all d-simplices σ we have $y_d(\sigma) \in \{0,1\}$ and $\sum_{\sigma\in\mathcal{K}^d} y_d(\sigma)c(\sigma) = \|\operatorname{supp}(y_d)\|$. Hence, $\operatorname{supp}(y_d)$ is a minimum directed combinatorial γ-cut. $\square$

5.3 Embedded simplicial complexes

In this section we consider a simplicial flow network $(\mathcal{K}, c, \gamma)$ where $\mathcal{K}$ is a d-dimensional simplicial complex with an embedding into $\mathbb{R}^{d+1}$. Alexander duality implies that $\mathbb{R}^{d+1} \setminus \mathcal{K}$ consists of $\beta_d + 1$ connected components. We call these connected components *voids*; exactly one void is unbounded and we denote the voids by V_i for $1 \leq i \leq \beta_{d+1}$. Given an embedding into $\mathbb{R}^{d+1}$, computing the voids of $\mathcal{K}$ can be done in polynomial time [31]. Further, we assume that the d-simplices are consistently oriented with respect to the voids. The embedding guarantees that every d-simplex σ appears on the boundary of at most two voids; by our assumption if σ appears on the boundary of two voids then it most be oriented positively on one and negatively on the other. We denote the boundary of the void V_i by $\operatorname{Bd}(V_i)$. Every d-simplex contained in the support of some d-cycle is on the boundaries of exactly two voids; it follows that the boundaries of any set of β_d voids is a basis of $H_d(\mathcal{K})$.

In order to state our theorems we need one additional assumption on $\mathcal{K}$. We assume there exists some void V_i containing two unit γ-flows Γ_1, Γ_2 whose supports partition $\operatorname{Bd}(V_i)$: $\operatorname{supp}(\Gamma_1) \cap \operatorname{supp}(\Gamma_2) = \emptyset$ and $\operatorname{supp}(\Gamma_1) \cup \operatorname{supp}(\Gamma_2) = \operatorname{Bd}(V_i)$. This assumption makes our problem analogous to an st-flow network in a planar graph such that s and t appear on the same face. The existence of two unit γ-flows partitioning the boundary is analogous to the two disjoint st-paths on the boundary of the face. It will be convenient to take the negation of Γ_1 and treat it as a unit $(-\gamma)$-flow; otherwise the assumption conflicts with the assumed consistent orientation. This is equivalent as it does not change the support of the flow, so for the rest of the section we will take Γ_1 to be a unit $(-\gamma)$-flow.

From $\mathcal{K}$ we construct its directed dual graph $\mathcal{K}^*$ as follows. Each void becomes a vertex of $\mathcal{K}^*$. Each d-simplex on the boundary of two voids becomes an edge; since we assumed the d-simplices are consistently oriented we direct the dual edge from the negatively oriented void to the positively oriented void. The remaining d-simplices only appear on one void and become loops in $\mathcal{K}^*$. For a d-simplex σ on the boundary of voids u and v we denote its corresponding dual edge $\sigma^* = (u^*, v^*)$ and we weight the edges by the capacity function: $c^*(\sigma^*) = c(\sigma)$. Let v_i^* be the vertex dual to the void whose boundary is partitioned by $\operatorname{supp}(\Gamma_1)$ and $\operatorname{supp}(\Gamma_2)$. We split v_i^* into two new vertices denoted s^* and t^*. The edges incident to v_i^* whose dual d-simplices were contained in $\operatorname{supp}(\Gamma_1)$ become incident to s^*, and the edges whose dual d-simplices were contained in $\operatorname{supp}(\Gamma_2)$ become incident to t^*. We add the directed edge (t^*, s^*) and set its capacity to infinity; $c^*((t^*, s^*)) = \infty$. Returning to the analogy of a planar graph with s and t on the same face, splitting v_i^* is analogous to adding an additional edge from t to s which splits their common face into two. However, for our purposes we are only concerned with the algebraic properties of the construction and do not actually need to modify the simplicial complex.

We need to update the chain complex associated with $\mathcal{K}$ to account for the voids and the splitting of v_i^*. We add an additional basis element Σ to $C_d(\mathcal{K})$ such that $\partial\Sigma = \gamma$ and give it infinite capacity; $c(\Sigma) = \infty$. In our construction Σ is dual to the edge (t^*, s^*). In our planar graph analogy Σ plays the role of an edge from t to s drawn entirely in the outer face; to make this precise we will need to add an additional chain group $C_{d+1}(\mathcal{K})$. We add each void V_j with $j \neq i$ as a basis element of $C_{d+1}(\mathcal{K})$ and define the boundary operator as $\partial_{d+1} V_j = \sum_{\sigma \in \mathrm{Bd}(v)} (-1)^{k_\sigma} \sigma$ where $k_\sigma = 0$ if σ is oriented positively on V_j and $k_\sigma = 1$ if σ is oriented negatively on V_j. Next we add additional basis elements S and T whose boundaries are defined by $\partial_{d+1} S = \Gamma_1 + \Sigma$ and $\partial_{d+1} T = \Gamma_2 - \Sigma$. The inclusion of $C_{d+1}(\mathcal{K})$ results in a valid chain complex since by definition the image of ∂_{d+1} under each basis element is a d-cycle. Moreover, in the new complex we have $H_d(\mathcal{K}) \cong 0$ since the boundaries of the voids generate $H_d(\mathcal{K})$.

Given our new chain complex we can extend the dual graph $\mathcal{K}^*$ to a dual complex; this construction is reminiscent of the dual of a polyhedron. We define the dual complex by the isomorphism of chain groups $C_k(\mathcal{K}^*) \cong C_{d-k+1}(\mathcal{K})$. The dual boundary operator $\partial_k^*\colon C_k(\mathcal{K}^*) \to C_{k-1}(\mathcal{K}^*)$ is the coboundary operator δ_{d-k+2}, and the dual coboundary operator $\delta_k^*\colon C_{k-1}(\mathcal{K}^*) \to C_k(\mathcal{K}^*)$ is the boundary operator ∂_{d-k+2}. The primal boundary operator commutes with the dual coboundary operator, and the primal coboundary operator commutes with the dual boundary operator. We summarize the duality in the following diagram.

$$\begin{array}{ccccccc}
C_{d+1}(\mathcal{K}) & \xleftrightarrow[\delta_{d+1}]{\partial_{d+1}} & C_d(\mathcal{K}) & \xleftrightarrow[\delta_d]{\partial_d} & \dots & \xleftrightarrow[\delta_1]{\partial_1} & C_0(\mathcal{K}) \\
\updownarrow\cong & & \updownarrow\cong & & & & \updownarrow\cong \\
C_0(\mathcal{K}^*) & \xleftrightarrow[\partial_1^*]{\delta_1^*} & C_1(\mathcal{K}^*) & \xleftrightarrow[\partial_2^*]{\delta_2^*} & \dots & \xleftrightarrow[\partial_{d+1}^*]{\delta_{d+1}^*} & C_{d+1}(\mathcal{K}^*)
\end{array}$$

We now have enough structure to state our duality theorems. In Section 5.3.1 we show that computing a max-flow for $(\mathcal{K}, c, \gamma)$ is equivalent to computing a shortest path from s^* to t^* in $\mathcal{K}^*$. In Section 5.3.2 we show that computing a minimum cost γ-cut p is equivalent to computing a minimum cost unit s^*t^*-flow in $\mathcal{K}^*$.

5.3.1 Max-flow / shortest path duality

We compute a shortest path from s^* to t^* in $\mathcal{K}^*$ using a well-known shortest paths linear program. Details on the linear program can be found in [38].

$$\begin{aligned}
\text{maximize} \quad & \operatorname{dist}(t^*) \\
\text{subject to} \quad & \operatorname{dist}(s^*) = 0 \\
& \operatorname{dist}(v^*) - \operatorname{dist}(u^*) \leq c^*((u^*, v^*)) \quad \forall\, (u^*, v^*) \in E
\end{aligned} \tag{LP3}$$

The solution to LP3 is a function $\operatorname{dist}\colon V(\mathcal{K}^*) \to \mathbb{R}$ which maps a vertex to its distance from s^* under the weight function c. By duality, dist is a $(d+1)$-cochain mapping the voids to $\mathbb{R}$. In the following theorem we will show that dist is equivalent to a γ-flow with value equal to $\operatorname{dist}(t^*)$.

Theorem 18. *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network where $\mathcal{K}$ is a d-dimensional simplicial complex embedded into $\mathbb{R}^{d+1}$ with two unit γ-flows whose supports partition the boundary of some void $\operatorname{Bd}(V_i)$. There is a bijection between γ-flows of $(\mathcal{K}, c, \gamma)$ and s^*t^*-paths in $\mathcal{K}^*$ such that the value of a γ-flow equals the length of its corresponding s^*t^*-path.*

Proof. Recall, by our discussion in Section 5.1 there is a value-preserving bijection between γ-flows and γ-circulations in $\mathcal{K}$ and the value of a γ-circulation f is given by $f(\Sigma)$. We can write f as a linear combination of d-cycles and by our construction there is a basis for the d-cycle space given by β_d elements of $\operatorname{im} \partial_{d+1}$. To form our basis we pick every element of $\operatorname{im} \partial_{d+1}$ except $\partial_{d+1} S$. Hence, f can be written as a linear combination $f = \sum_{i=0}^{\beta_d - 1} \alpha_i B_i + \alpha_T B_T$ where $\partial_{d+1} V_i = B_i$ and $\partial_{d+1} B_T = T$. We construct a $(d+1)$-cochain P by the mapping $P(V_i) = \alpha_i$, $P(S) = 0$, $P(T) = \alpha_T$. Dual to P is a 0-cochain on the vertices of $\mathcal{K}^*$ which we call dist. By construction we have $\text{dist}(s^*) = 0$. Since $\mathcal{K}$ is consistently oriented on the voids and f obeys the capacity constraints we have $\text{dist}(v^*) - \text{dist}(u^*) \leq c^*((u^*, v^*))$ for every edge (u^*, v^*). Finally, we have $\text{dist}(t^*) = f(\Sigma)$ since $f(\Sigma) = P(T) - P(S) = \alpha_T$.

Conversely, let dist be some vector satisfying the constraints of LP3. As dist is a cochain on the vertices of $\mathcal{K}^*$ by duality we may view dist as an element of $C_{d+1}(\mathcal{K})$ hence $\partial(\text{dist})$ is a circulation in $\mathcal{K}$ obeying the capacity constraints. Further, we have that the component of $\partial(\text{dist})$ indexed by Σ is equal to $\text{dist}(t^*) - \text{dist}(s^*) = \text{dist}(t^*)$. Hence, $\|\partial(\text{dist})\| = \text{dist}(t^*)$ which completes the proof. $\square$

5.3.2 Min-cut / min-cost flow duality

We begin this section by stating the minimum cost flow problem in graphs. The minimum cost flow problem asks to find the cheapest way to send k units of flow from s to t. An instance of the minimum cost flow problem is a tuple (G, w, c, k) where $G = (V, E)$ is a directed graph, $w, c \in C_1(G)$, and $k \in \mathbb{R}$. The 1-chains represent the weight and capacity of each edge, and k is the demand of the network. The goal of the minimum cost flow problem is to find an st-flow obeying the following constraints.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{e \in E} w(e) f(e) & \\
\text{subject to} \quad & \delta(v) \cdot f = 0 & \forall v \in V \setminus \{s, t\} \\
& 0 \leq f(e) \leq c(e) & \\
& \delta(s) \cdot f = -k & \\
& \delta(t) \cdot f = k &
\end{aligned}
$$

The first two constraints are conservation of flow and capacity constraints. The third and fourth are the demand constraints which say f must send exactly k units of flow from s to t; note that only one of these constraints is necessary. We will compute a minimum directed γ-cut in $\mathcal{K}$ by solving the minimum cost flow problem with $k = 1$ in $\mathcal{K}^*$. We assume there is a weight function $w \colon \mathcal{K}_d \to \mathbb{R}^+$ on the d-skeleton of $\mathcal{K}$, which after dualizing becomes a weight function w^* on the edges of $\mathcal{K}^*$. In the following theorem the capacity function is not needed, so we will assume each edge in $\mathcal{K}^*$ has infinite capacity.

Theorem 19. *Let $\mathcal{K}$ be a d-dimensional simplicial complex embedded into $\mathbb{R}^{d+1}$ with two unit γ-flows whose supports partitions the boundary of some void $\mathrm{Bd}(V_i)$. There is a bijection between γ-cuts p in $\mathcal{K}$ and unit s^*t^*-flows f in $\mathcal{K}^*$ such that $\|p\| = \sum w^*(e)f(e)$.*

Proof. Let p be a γ-cut in $\mathcal{K}$ and let $\bar{p}$ be its negation; that is, $\bar{p}$ is a $(d-1)$-cochain with $\bar{p}(\gamma) = 1$. By construction we have $\delta(\bar{p})(\Sigma) = \bar{p}((\partial\Sigma) = 1$. We define f to be the image of $\delta(\bar{p})$ under the duality isomorphism. Since $\bar{p}^*$ is a 2-chain in $\mathcal{K}^*$ and $f = \partial^*(\bar{p}^*)$ we see that f is a 1-circulation in $\mathcal{K}^*$. By removing the edge $\Sigma^* = (t^*, s^*)$ from f we see that f is an s^*t^*-flow with $\|f\| = 1$ since $f(\Sigma^*) = \delta(\bar{p})(\Sigma) = 1$. Finally, we have $\|p\| = \|p'\| = \sum w(\sigma)\delta(p)(\sigma) = \sum w(\sigma^*)f(\sigma^*)$.

Conversely, let f^* be a 1-circulation in $\mathcal{K}^*$ with $f(\Sigma^*) = 1$. By assumption we have $H_d(\mathcal{K}) \cong 0$ and by duality $H_1(\mathcal{K}^*) \cong 0$. It follows that f^* can be written as a linear combination of boundaries $f^* = \sum \alpha_i B_i^*$ where $B_i^* \in \operatorname{im}\partial_2^*$. Let p^* be a 2-chain with $\partial_2^* p^* = f^*$. Dual to f^* is a d-coboundary $f = \sum \alpha_i B_i$ where $B_i \in \operatorname{im}(\delta_d)$ and dual to p^* is a $(d-1)$-cochain p with $\delta(p) = f$. We have that $\bar{p}$ is a γ-cut since $\bar{p}(\gamma) = \bar{p}(\partial\Sigma) = \delta(\bar{p})(\Sigma) = -f(\Sigma) = -f^*(\Sigma^*) = -1$. Finally, we have $\sum w^*(e)f^*(e) = \sum |w(\sigma)\delta(p)(\sigma)| = \|p\| = \|\bar{p}\|$. □

Corollary 1. *Let $\mathcal{K}$ be a d-dimensional simplicial complex embedded in $\mathbb{R}^{d+1}$ with two unit γ-flows partitioning some $\mathrm{Bd}(V_i)$. There is a polynomial time algorithm computing a minimum directed combinatorial γ-cut.*

Proof. We solve the minimum cost circulation problem in $\mathcal{K}^*$ setting the demand and every capacity constraint equal to one. The resulting flow is dual to a γ-cut p in $\mathcal{K}$. Since the minimum cost circulation is integral we have $\|\operatorname{supp}(\delta(p))\| = \|p\|$. That is, the cost of p as a γ-cut equals the cost of $\operatorname{supp}(\delta(p))$ as a combinatorial γ-cut. □

5.4 Ford-Fulkerson algorithm

In this section we show how the Ford-Fulkerson algorithm can be used to compute a maximum flow of simplicial flow network $(\mathcal{K}, c, \gamma)$. In a simplicial flow network the Ford-Fulkerson algorithm picks out a ***augmenting chain*** at every iteration which is a high dimensional generalization of an augmenting path. As shown in Section 5.2.2 a maximum flow of a simplicial flow network with integral capacities may not be integral, so it is not immediate that Ford-Fulkerson is guaranteed to halt. To remedy this, our implementation of Ford-Fulkerson contains a heuristic reminiscent of the network simplex algorithm. Our heuristic guarantees that at every iteration of Ford-Fulkerson the flow is a solution on a vertex of the polytope defined by the linear program. Hence, our heuristic makes our implementation of Ford-Fulkerson into a special case of the simplex algorithm. It follows that Ford-Fulkerson does halt on a simplicial flow network, but the running time may be exponential. Our heuristic for picking augmenting chains takes $O(n^{\omega+1})$ time since it requires solving $O(n)$ linear systems, each taking $O(n^{\omega})$ time using standard methods [56]. In Section 5.4.3 we show that solving a particular type of linear system reduces to finding an augmenting chain, giving us an imprecise lower bound on the complexity of our heuristic.

We begin this section by defining the concepts of the ***residual complex*** and of ***augmenting chains*** which serve as high dimensional generalizations of the residual graph and augmenting paths used in the Ford-Fulkerson algorithm for graphs. We show that a flow is maximum if and only if its residual complex contains no augmenting chains, generalizing the well-known graph theoretic result. This work is an extension of previous work done by Latorre who showed one direction of the theorem and leaving the other open [70].

5.4.1 The residual complex

We now present our definitions of the residual complex and an augmenting chain.

Definition 7 (Residual complex)**.** *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network and f be a feasible flow on the network. We define a new simplicial flow network called the **residual complex** to be the tuple $(\mathcal{K}_f, c_f, \gamma)$ constructed in the following way. The d-skeleton of $\mathcal{K}_f$ is the union $\mathcal{K}_d \cup -\mathcal{K}_d$, that is, for each d-simplex σ in $\mathcal{K}$ we add an additional d-simplex $-\sigma$ whose orientation is opposite of σ. $(\mathcal{K}_f)_{d'} = (\mathcal{K})_{d'}$ for dimensions $d' < d$. The **residual capacity function** $c_f \colon (\mathcal{K}_f)_d \to \mathbb{R}$ is given by*

$$c_f(\sigma) = \begin{cases} c(\sigma) - f(\sigma) & \sigma \in \mathcal{K}_d, \\ f(\sigma) & -\sigma \in \mathcal{K}_d. \end{cases}$$

Definition 8 (Augmenting chain)**.** *Let $\mathcal{K}_f$ be a residual complex for the simplicial flow network $(\mathcal{K}, c, \gamma)$. An **augmenting chain** is a d-chain $\Gamma \in C_d(\mathcal{K}_f)$ such that $\Gamma = \sum \alpha_i \sigma_i$ and $\partial\Gamma = \gamma$ with $\alpha_i \geq 0$.*

Note that an augmenting chain need not obey the residual capacity constraint c_f. This is because after finding an augmenting chain the amount of flow sent through the chain will be normalized by the coefficients α_i producing a new chain respecting the capacity constraints. The following two lemmas prove the main result of the section. The first of which was observed by Latorre [70].

Lemma 16 (Latorre [70])**.** *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network. If f is a maximum flow then $\mathcal{K}_f$ contains no augmenting chains.*

Proof. Let $\Gamma = \sum \alpha_i \sigma_i$ be an augmenting chain in $\mathcal{K}_f$ and let $\alpha = \min\{\frac{1}{\alpha_i} c_f(\sigma_i)\}$. Define the new flow as follows

$$f'(\sigma) = \begin{cases} f(\sigma) + \alpha \cdot \alpha_i & \sigma = \sigma_i, \\ f(\sigma) - \alpha \cdot \alpha_i & -\sigma = \sigma_i. \end{cases}$$

We have $f(\sigma_i) + \alpha \cdot \alpha_i \leq f(\sigma_i) + c_f(\sigma_i) = c(\sigma_i)$ and $f(\sigma_i) - \alpha \cdot \alpha_i \geq f(\sigma_i) - c_f(\sigma_i) = 0$ so f' obeys the capacity constraints. To show that f' obeys conservation of flow we compute the following equality

$$\begin{aligned} \sum_{\sigma \in \mathcal{K}_d} f'(\sigma)\partial(\sigma) &= \sum_{\sigma \in \mathcal{K}_d \setminus \mathrm{supp}(\Gamma)} f(\sigma)\partial(\sigma) + \sum_{\sigma_i \in \mathrm{supp}(\Gamma)} (f(\sigma_i) \pm \alpha \cdot \alpha_i)\partial(\sigma_i) \\ &= \alpha \cdot \partial \left(\sum_{\sigma_i \in \mathrm{supp}(\Gamma)} \alpha_i \sigma_i \right) \\ &= 0. \end{aligned}$$

□

Lemma 17. *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network. If f is a flow such that $\mathcal{K}_f$ contains no augmenting chains then f is a maximum flow.*

Proof. By way of contradiction assume that f is not a maximum flow and g is some other flow with higher value than f. We show that the d-chain $g - f$ is an augmenting chain in $\mathcal{K}_f$. First, we note that the boundary of $g - f$ is equal to $(\|g\| - \|f\|)\gamma$ implying $g - f$ obeys conservation of flow, so all that remains is to check that $g - f$ obeys the capacity constraints in $\mathcal{K}_f$. Let σ be a d-simplex in $\mathcal{K}$, there are two cases to consider. First, if $f(\sigma) \leq g(\sigma)$ we have $g(\sigma) - f(\sigma) \leq c(\sigma) - f(\sigma) = c_f(\sigma)$ and the capacity constraint is obeyed. Second, if $g(\sigma) < f(\sigma)$ we interpret this as applying $|g(\sigma) - f(\sigma)|$ flow to $-\sigma$. Hence, $|g(\sigma) - f(\sigma)| = f(\sigma) - g(\sigma) < f(\sigma) = c_f(\sigma)$ which concludes the proof. □

Lemmas 16 and 17 give us the main theorem of the section.

Theorem 20. *Let $(\mathcal{K}, c, \gamma)$ be a simplicial flow network. A flow f is a maximum flow if and only if $\mathcal{K}_f$ contains no augmenting chains.*

5.4.2 Augmenting chain heuristic

In this section we provide a heuristic for the Ford-Fulkerson algorithm that is guaranteed to halt on a simplicial flow network. Our example in Section 5.2.2 shows that a maximum flow may have fractional value, so it's not immediately clear that Ford-Fulkerson halts on all simplicial flow networks. To remedy this our heuristic ensures that at each step the flow corresponds to a vertex of the flow polytope (defined in the next paragraph). As there are a finite number of vertices, and the value of the flow increases at every step, it follows that under this heuristic Ford-Fulkerson must halt. Under our heuristic Ford-Fulkerson becomes a special case of the simplex algorithm. Our heuristic is reminiscent of the network simplex algorithm which maintains a tree at every iteration. See the book by Ahuja, Magnanti, and Orlin for an overview of the network simplex algorithm [3].

We define the ***flow polytope*** of $(\mathcal{K}, c, \gamma)$ to be the polytope $P \subset \mathbb{R}^{n_d}$ defined by the constraints of the maximum flow linear program LP1. A ***vertex*** of the polytope P is any feasible solution to LP1 with at least n_d tight linearly independent constraints. We will ensure that at every step of Ford-Fulkerson our flow f is a vertex of P. To do this we will make sure that the d-simplices corresponding to non-tight constraints of LP1 form an acyclic complex. Some straightforward algebra implies that this condition is enough to make at least n_d constraints tight. Let $\mathcal{H}_f$ be the subcomplex of d-simplices "half-saturated" by f; that is, $\sigma \in \mathcal{H}_f$ if and only if its capacity constraint is a strict inequality: $0 < f(\sigma) < c(\sigma)$. The half-saturated simplices do not make either of their two corresponding constraints tight, while d-simplices not in $\mathcal{H}_f$ make exactly one of their corresponding constraints tight. We require that $\mathcal{H}_f$ be an acyclic complex at each step of Ford-Fulkerson. In the case of graphs, this just means that $\mathcal{H}_f$ is a forest. For a d-dimensional complex it means that $H_d(\mathcal{H}_f) = 0$. Acyclic complexes have been studied by Duval, Klivans, and Martin who show that they share many properties with forests and trees in graphs [36]. The following lemma shows that if $\mathcal{H}_f$ is acyclic then f is a vertex of the flow polytope.

Lemma 18. *Let f be a feasible flow for the d-dimensional simplicial flow network $(\mathcal{K}, c, \gamma)$.*

If the subcomplex of half-saturated d-simplices $\mathcal{H}_f$ is acyclic then f is a vertex of the flow polytope P.

Proof. In order for f to be a vertex of P we need to show that at least n_d of the constraints of LP1 are tight, and that these n_d tight constraints are linearly independent. As f is a flow the $2n_{d-1}$ constraints ensuring conservation of flow are always tight. It follows that we have $n_d - \beta_d$ tight linearly independent conservation of flow constraints corresponding to a basis for $\operatorname{im}\delta_d$. As $\mathcal{H}_f$ is acyclic we have that $|\operatorname{supp}(\mathcal{H}_f)| \leq n_d - \beta_d$ since at least one d-simplex from each basis element of $H_d(\mathcal{K})$ must be missing from $\mathcal{H}_f$. This implies that at most $2(n_d - \beta_d)$ of the $2n_d$ capacity constraints are not tight; equivalently, at least $2n_d - 2(n_d - \beta_d) = 2\beta_d$ of the capacity constraints are tight. Since $\mathcal{H}_f$ is acyclic we can pick a set of β_d d-simplices Σ such that $\dim H_d(\mathcal{K}\setminus\Sigma) = 0$ and $\dim H_{d-1}(\mathcal{K}\setminus\Sigma) = \dim H_{d-1}(\mathcal{K})$. Each d-simplex $\sigma \in \Sigma$ corresponds to some tight capacity constraint, and since removing σ does not change the dimension of H_{d-1} it is not contained in $\operatorname{im}\delta_d$. It follows that the tight capacity constraint corresponding to σ is linearly independent from the $n_d - \beta_d$ conservation of flow constraints. Finally, since each tight capacity constraint corresponds to a unique σ they are all linearly independent from each other. $\square$

At each iteration of Ford-Fulkerson we want to pick an augmenting chain such that the resulting flow leaves $\mathcal{H}_f$ acyclic. It's not clear how to pick such an augmenting chain. However, no matter what augmenting chain we pick we can always repair the flow in a way that the resulting flow leaves $\mathcal{H}_f$ acyclic. We describe our method for repairing the flow in the following lemma.

Lemma 19. *Let f be a feasible flow for the d-dimensional simplicial flow network $(\mathcal{K}, c, \gamma)$. If the subcomplex of half-saturated d-simplices $\mathcal{H}_f$ is not acyclic then in $O(n^{\omega+1})$ time we can construct a new flow f' such that $\mathcal{H}_{f'}$ is acyclic and $\|f\| = \|f'\|$.*

Proof. Let f be a feasible flow such that $\mathcal{H}_f$ is not acyclic. In polynomial time we compute a basis for $H_d(\mathcal{H}_f)$. Let $\Sigma = \sum \alpha_i\sigma_i$ be a basis element of $H_d(\mathcal{H}_f)$ and let $\alpha = \min\{\frac{1}{\alpha_i} c_f(\sigma_i)\}$. As in Lemma 16 we construct the new flow f' by

$$f' = \begin{cases} f(\sigma) + \alpha \cdot \alpha_i & \sigma = \sigma_i \\ f(\sigma) - \alpha \cdot \alpha_i & -\sigma = \sigma_i \end{cases}.$$

The new flow f' saturates some d-simplex in Σ and does not introduce any new d-cycles to $\mathcal{H}_f$ as it only affects the half-saturated edges. We call the new subcomplex of half-saturated simplices $\mathcal{H}_{f'}$ and observe that $\dim H_d(\mathcal{H}_{f'}) < \dim H_d(\mathcal{H}_f)$. Since f' is constructed by adding a d-cycle to f we have that $\|f\| = \|f'\|$. We repeat the process of computing a homology basis for $\mathcal{H}_{f'}$ and saturating some basis element until $\mathcal{H}_{f'}$ is acyclic.

It remains to compute the running time of the above procedure. Computing a homology basis takes $O(n^{\omega})$ time . To repair the flow we need to make at most $O(n)$ homology basis computations, hence the total running time is $O(n^{\omega+1})$. □

To wrap up the section, we state our main theorem whose proof is immediate from Lemmas 18 and 19.

Theorem 21. *Given a simplicial flow network $(\mathcal{K}, c, \gamma)$ we can compute a maximum flow f by using the Ford-Fulkerson algorithm with the following heuristic: at every iteration pick an augmenting chain such that the subcomplex of half-saturated d-simplices $\mathcal{H}_f$ is acyclic.*

5.4.3 Lower bounds

We have shown a heuristic for which given a simplicial flow network $(\mathcal{K}, c, \gamma)$ Ford-Fulkerson is guaranteed to halt with a maximum flow. In the worst case our algorithm runs in exponential time. In this section we focus our attention on the time required to find an augmenting chain. The running time of our heuristic is determined by the time it takes to compute the augmenting chain and repair the flow at each iteration of the algorithm. Computing the augmenting chain takes $O(n^{\omega})$ time by solving the linear system. Repairing the flow takes $O(n^{\omega+1})$ time since it requires $O(n)$ homology basis computations which each take time $O(n^{\omega})$. We show that it is unlikely that this running time can be substantially improved. More specifically, we show that finding a non-negative solution to a linear system $Ax = b$ when A has entries in $\{-1, 0, 1\}$ reduces to computing an augmenting chain for $(\mathcal{K}, c, \gamma)$, hence the complexity of solving a linear system in this form serves as a lower bound on computing an augmenting chain. Given a linear system we construct a 2-complex with a 1-cycle γ such that finding a 2-chain Γ with non-negative coefficients and $\partial\Gamma = \gamma$ is equivalent to finding a non-negative feasible solution to the linear system. Further, the complex $\mathcal{K}$ used in the reduction is relatively torsion-free, so the total unimodularity of its boundary matrix cannot be used to speed up the computation. Our reduction is essentially the same as one used by Chen and Freedman to show homology localization over $\mathbb{Z}_2$ is NP-hard [23]. However, we modify it slightly since we consider coefficients over $\mathbb{R}$. We only give a proof sketch as our reduction is almost identical to that of Chen and Freedman's.

Theorem 22. *Let $Ax = b$ be a linear system where A has entries in $\{-1, 0, 1\}$. In polynomial time we can construct a 2-dimensional, relatively torsion-free, simplicial complex $\mathcal{K}$ and a 1-cycle γ such that if a 2-chain Γ is an augmenting chain for γ then it is a non-negative solution to $Ax = b$.*

Proof. We construct a cell complex $\mathcal{K}$ from A as follows. For each of the m rows we construct a 1-cycle C_i. For each column vector v_j we construct a punctured sphere T_j with boundary components $C_{i,j}$, for each $v_{i,j} = 1$ and $-C_{i,j}$ for each $v_{i,j} = -1$. Define the 1-cycle γ to be b with respect to the basis given by the boundary components of $\mathcal{K}$. By our construction a vector $x = (x_1, \ldots, x_n)$ is a feasible solution to $Ax = b$ if and only if the 2-chain $\sum x_i T_i$ has boundary γ. Hence, computing an augmenting chain for γ is equivalent to computing a non-negative solution to the linear system. It remains to show that $\mathcal{K}$ can be triangulated into a simplicial complex. We refer the reader to [23] for a triangulation. To see that $\mathcal{K}$ is relatively torsion-free we refer to reader to [30] which characterizes relative torsion in 2-complexes by forbidding certain Möbius subcomplexes. $\square$

Chapter 6: Resistance and Capacitance

In this section we introduce the concepts of ***effective resistance*** and ***effective capacitance*** of a cycle γ in a simplicial complex. These concepts are topological generalizations of graph theoretic notions that were originally studied in the context of electrical networks [65]. Intuitively, in a graph the effective resistance between s and t is a measure of "how connected" s and t are to one another, and there are several interesting applications illustrating this. Effective resistance forms a metric on the vertices of graph [68]. It is proportional to the expected length of a random walk beginning and ending at s and including t [21]. It is proportional to the probability that an edge $\{s, t\}$ is included in a random spanning tree [65]. The random subgraph of a graph obtained by keeping an edge $\{s, t\}$ with probability proportional to the effective resistance between s and t approximates the spectrum of the Laplacian of the original graph [91]. Effective capacitance has been defined to be used as a parameter for a quantum algorithm deciding st-connectivity [59].

The concept of effective resistance in simplicial complexes has been studied in previous work, but it seems our treatment of effective capacitance in simplicial complexes is new. Previously there have been two competing definitions for effective resistance in simplicial complexes. Kook and Lee [69] as well as Osting, Palande, and Wang [82] define the effective resistance as a quantity associated with the boundary of a d-simplex. This is consistent with the view that the effective resistance is a quantity associated with a pair of vertices s and t. Note that the d-simplex is not required to exist in the complex; only the boundary is required to be present. Hasen and Ghrist define the effective resistance to be a quantity associated with a null-homologous cycle [49] which is a more general than the previous definition since the boundary of a d-simplex can be made null-homologous by including the d-simplex in the complex. Our definition of effective resistance is equivalent to that of Hasen and Ghrist.

A few results about effective resistance in graphs have been generalized to simplicial complexes. Kook and Lee show that the probability a d-simplex appears in a spanning acyclic subcomplex is proportional to the effective resistance of its boundary [69]. A spanning acyclic subcomplex is a high dimensional generalization of a spanning tree, sometimes called a hypertree, and shares many of the same properties as a spanning tree [36]. Osting, Palande, and Wang show that the random subcomplex obtained by keeping a d-simplex with probability proportional to its boundary's effective resistance results in a complex which approximates the spectrum of the $(d-1)$-up Laplacian [82].

Our definitions of effective resistance and capacitance heavily rely upon our study of generalized flows and cuts presented in Chapter 5. The effective resistance and capacitance of a cycle γ are the sizes of the minimum unit γ-flow and the minimum γ-cut. However, we minimize the flow and cut with respect to the ℓ_2-norm rather than the ℓ_1-norm. In this Section 6.2 we will provie upper bounds on effective resistance and capacitance that are polynomial in both the number of d-simplices and the size of the relative torsion subgroups of the complex. In general the size of a relative torsion subgroup can be exponential in the number of d-simplices. In Section 6.2.1 we provide explict examples of complexes with effective resistance and capacitance that is exponential with respect to the number of d-simplices. In Section 6.3 we show that effective resistance and capacitance can be used to parameterize the query complexity of a quantum algorithm that decides whether or not a cycle is null-homologous in a simplicial complex.

6.1 Definitions of Effective Resistance and Capacitance

We now begin with our definitions of effective resistance and capacitance. Our definition of effective resistance requires the notion of a γ-flow defined in Chapter 5. We recall the definition of a ***unit*** γ***-flow***. Throughout this section we work with weighted simplicial complexes and assume there is a weight function $w\colon \mathcal{K}_d \to \mathbb{R}^+$. By W we denote the diagonal matrix whose entries are the weights of the d-simplices. Intuitively, these weights represent the capacitance of an electrical network which has an inverse relationship with the resistance.

Definition 9. *Given a d-dimensional simplicial complex $\mathcal{K}$ with weight function $w\colon \mathcal{K}_d \to \mathbb{R}$ and a $(d-1)$-dimensional null-homologous cycle γ, a* ***unit*** $\boldsymbol{\gamma}$***-flow*** *is a d-chain $f \in C_d(\mathcal{K})$ such that $\partial f = \gamma$.*

Recall that the definition of a γ-flows in Chapter 5 included any d-chain f such that $\partial f = k\gamma$. In this chapter we only consider the case when $k = 1$, hence the name *unit* γ flow. Next, we define the ***flow energy*** of a unit γ-flow which is almost equivalent to the definition of $\|f\|_2$ but accounts for the weight function.

Definition 10. *Given a d-dimensional simplicial complex $\mathcal{K}$ with weight function w :*

*$C_d(\mathcal{K}) \to \mathbb{R}^+$ and a unit γ-flow f, the **flow energy** of f on $\mathcal{K}$ is*

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

Finally, we define the ***effective resistance*** of γ which just minimizes over the flow energy of each unit γ-flow.

Definition 11. *Given a d-dimensional simplicial complex $\mathcal{K}$ and a $(d-1)$-dimensional cycle γ in $\mathcal{K}$ if γ is null-homologous in $\mathcal{K}$ the **effective resistance** of γ in $\mathcal{K}$ is given by the size of the minimum energy unit γ-flow,*

$$\mathcal{R}_\gamma(\mathcal{K}) = \min_{f \in C_d(\mathcal{K}) :\, \partial f = \gamma} \mathsf{J}(f).$$

If γ is not null-homologous in $\mathcal{K}$ we have $\mathcal{R}_\gamma(\mathcal{K}) = \infty$.

Our definition of effective resistance is equivalent but different than the original. Traditionally, the effective resistance of γ is defined as the quadratic form $\gamma^T L^+ \gamma$ where $L = \partial_{d+1} W \delta_d$ is the d-dimensional weighted up-Laplacian; W denotes the $n_d \times n_d$ matrix whose diagonal contains the weights of the d-simplices. By L^+ we denote the Moore-Penrose pseudoinverse of L. In the following theorem we show that the two definitions are equivalent. The equivalence is well-known folklore in graphs, and the proof relies on the fact that $L^+\gamma$ is a minimum energy unit γ-flow.

Theorem 23. *Let $\mathcal{K}$ be a d-dimensional simplicial complex and let γ be a null-homologous $(d-1)$-cycle. The effective resistance of γ in $\mathcal{K}$ has the equality $\mathcal{R}_\gamma(\mathcal{K}) = \gamma^T L^+ \gamma$.*

Proof. We use two well-known properties of the Moore-Penrose pseudoinverse. For any matrix $B = AA^T$ we have $B^+ = (A^T)^+ A^+$, and for any matrix B we have $B^+ v = \arg\min\{u \mid Bu = v\}$. We proceed with the following calculation:

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

from which we conclude that $L^+ = (\partial_{d+1} W^{1/2})^T)^+ (\partial_{d+1} W^{1/2})^+$. Using the fact that for any matrix $(B^T)^+ = (B^+)^T$ we can express the quadratic form as

$$\gamma^T L^+ \gamma = \langle (\partial_{d+1} W^{1/2})^+ \gamma, (\partial_{d+1} W^{1/2})^+ \gamma \rangle,$$

which is the magnitude of the smallest vector f with $\partial_{d+1} W^{1/2} f = \gamma$; we choose f to be the vector minimizing this quantity. By our construction we have that $W^{1/2} f$ is a unit γ-flow. We will prove that $W^{1/2} f$ minimizes the flow energy. From the definition of flow energy we have

$$\mathsf{J}(W^{1/2} f) = (W^{1/2} f)^T W^{-1} (W^{1/2} f) = \|f\|^2.$$

For any unit γ-flow f we have

$$\mathsf{J}(f) = f^T W^{-1} f = (W^{-1/2} f)^T (W^{-1/2} f),$$

so if f is a unit γ-flow of minimum energy then f minimizes $\|W^{-1/2} f\|$ and $f' = W^{-1/2} f$ must be the minimum norm vector such that $\partial_{d+1} W^{1/2} f' = \gamma$. □

Before providing the definition of a ***unit γ-potential*** in a simplicial complex we will begin by reviewing the definition of a unit st-potential a graph, which can be found in [59]. 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. A unit st-potential is function $p\colon V(G) \to \mathbb{R}$ such that $p(s) = 0$, $p(t) = 1$ and for any two vertices u, v in the same connected component we have $p(u) = p(v)$. Viewing p as a cochain we see that its coboundary is zero in H. In a sense the size of its coboundary in G is a measure of "how disconnected" s is from t in H. Intuitively, our definition of a unit γ-potential measures "how far" a cycle γ is from null-homologous in a subcomplex $\mathcal{L}$ of $\mathcal{K}$. We are now ready to present our definitions.

Definition 12. *Let $\mathcal{L} \subset \mathcal{K}$ be simplicial complexes, and let $\gamma \in C_{d-1}(\mathcal{L})$ be a $(d-1)$-cycle such that γ is null-homologous in $\mathcal{K}$ but not $\mathcal{L}$. A **unit γ-potential** in $\mathcal{L}$ is a $(d-1)$-cochain p such that $\delta[\mathcal{L}]p = 0$ and $p(\gamma) = 1$.*

Now we define the notion of potential energy which is used to measure the size of a unit γ-potential.

Definition 13. *Given simplicial complexes $\mathcal{L} \subset \mathcal{K}$ with weight function $w\colon C_d(\mathcal{K}) \to \mathbb{R}$*

*and a γ-potential p in $\mathcal{L}$, the **potential energy** of p on $\mathcal{K}$ is*

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

Just like effective resistance minimizes over flow energy, the effective capacitance of γ is minimized over the potential energy.

Definition 14. *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 null-homologous in $\mathcal{K}$. If γ is not null-homologous in $\mathcal{L}$, the **effective capacitance** of γ in $\mathcal{L}$ is $\mathcal{C}_\gamma(\mathcal{L}) = \min_p \mathcal{J}(p)$ where p is a γ-potential. If γ is null-homologous in $\mathcal{L}$, then $\mathcal{C}_\gamma(\mathcal{L}) = \infty$.*

Unlike unit γ-flows, it is not obvious from the definition that a unit γ-potential even exists. In the following theorem we show that it is indeed the case that they exist.

Theorem 24. *Let $\mathcal{L} \subset \mathcal{K}$ be d-dimensional simplicial complexes whose $(d-1)$-skeletons are equal. Let γ be a $(d-1)$-cycle that is null-homologous in $\mathcal{K}$. There exists a unit γ-potential in $\mathcal{L}$ if and only if γ is not null-homologous in $\mathcal{L}$.*

Proof. First we note that $\ker \delta_d[\mathcal{L}]$ is the orthogonal complement of $\operatorname{im} \delta_d[\mathcal{L}]$ in $C_{d-1}(\mathcal{L}, \mathbb{R})$. Assume that there exists a unit γ-potential p in $\mathcal{L}$, so $p \in \ker \delta_d[\mathcal{L}] = (\operatorname{im} \partial_d[\mathcal{L}])^\perp$. Since $\langle p, \gamma \rangle = 1$ we see that γ has a non-zero component in $(\operatorname{im} \partial_d[\mathcal{L}])^\perp$, so $\gamma \notin \operatorname{im} \partial_d[\mathcal{L}]$.

Conversely, assume that γ is not null-homologous in $\mathcal{L}$. Then γ has some non-zero component coming from $(\operatorname{im} \partial_d[\mathcal{L}])^\perp = \ker \delta_d[\mathcal{L}]$. Let $q = \Pi_{\ker \delta_d} \gamma$ be the projection of γ onto $\ker \delta_d[\mathcal{L}]$, so $\langle q, \gamma \rangle \neq 0$ and $\delta_d[\mathcal{L}] = 0$. It follows that $\frac{q}{\langle q, \gamma \rangle}$ is a unit γ-potential in $\mathcal{L}$. $\square$

6.2 Bounds on Effective Resistance and Capacitance

In this section, we provide upper bounds on the resistance and capacitance of a cycle γ in a simplicial complex $\mathcal{K}$. Our upper bounds are polynomial in the number of d-simplices and the cardinality of the largest torsion subgroup of a relative homology group. 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 with respect to the number of d-simplices. There exist simplicial complexes such that the torsion subgroup of $H_{d-1}(\mathcal{K}, \mathbb{Z})$ has cardinality n while $\mathcal{K}$ only has $O(\log^{1/d} n)$ vertices [79] note that such a complex contains at most $O(\log n)$ d-simplices.

In Theorems 28 and 29 we provide explicit examples of simplicial complexes containing a cycle γ whose effective resistance and capacitance 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. There exist simplicial complexes with no torsion in their homology groups but that do have torsion in their relative homology groups. An example of this is the Möbius strip. The Möbius strip has no torsion, but it has torsion relative to its boundary [30].

Our results rely on a change of basis on the boundary matrix called the 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 25 (Munkres, Chapter 1 Section 11 [77])**.** *There are bases for $C_d(\mathcal{K})$ and $C_{d-1}(\mathcal{K})$ such that the matrix for the boundary operator $\partial_d \colon C_d(\mathcal{K}, \mathbb{Z}) \to C_{d-1}(\mathcal{K}, \mathbb{Z})$ can be expressed in these bases yielding the* ***normal form*** *of the matrix,*

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

where D is a diagonal matrix with entries $d_1, \ldots, d_m$ 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 $d_1, \ldots, d_m$ correspond to the torsion coefficients of $H_{d-1}(\mathcal{K}, \mathbb{Z}) \cong \mathbb{Z}^{\beta_{d-1}} \oplus \mathbb{Z}_{d_1} \oplus \cdots \oplus \mathbb{Z}_{d_m}$,*

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 a set of elementary row and column operations. If ∂ is square, these operations multiply $\det \partial$ by ± 1.

Using Theorem 25, 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 $H_k(\mathcal{L}, \mathcal{L}_0, \mathbb{Z}) = \ker \partial_k[\mathcal{L}, \mathcal{L}_0] / \operatorname{im} \partial_{k+1}[\mathcal{L}, \mathcal{L}_0]$. More information on the relative boundary matrix can be found in [30]. We denote the cardinality of the torsion subgroup of the relative homology group $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$ by $t(\mathcal{L}, \mathcal{L}_0)$. By t we denote the maximum $t(\mathcal{L}, \mathcal{L}_0)$ over all relative homology groups.

Lemma 20. *Let $\partial_d[\mathcal{L}, \mathcal{L}_0]$ be a $k \times k$ relative boundary matrix of $\mathcal{K}$. The magnitude of the determinant of $\partial_d[\mathcal{L}, \mathcal{L}_0]$ is bounded above by the cardinality of the torsion subgroup of $H_{d-1}(\mathcal{L}, \mathcal{L}_0, \mathbb{Z})$,*

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

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 $D = \operatorname{diag}(d_1, \ldots, d_k)$. The determinant is equal to $\pm \prod_{i=1}^{k} d_i$ and by Theorem 25 the torsion subgroup of $H_{d-1}(\mathcal{L}, \mathcal{L}_0)$ is $\mathbb{Z}_{d_1} \oplus \cdots \oplus \mathbb{Z}_{d_k}$ which has cardinality $t(\mathcal{L}, \mathcal{L}_0) = \prod_{i=1}^{k} d_i$. $\square$

Before proving our upper bounds on effective resistance and capacitance we need to state some assumptions on the cycle. Let $\mathcal{K}$ be a d-dimensional simplicial complex with weight function $w \colon \mathcal{K}_d \to \mathbb{R}^+$ and let γ be a null-homologous $(d-1)$-cycle in $\mathcal{K}$. The weight function appears in our definitions of flow energy and potential energy, which implies that the effective resistance and capacitance is dependent on the weight function. In general the weights can be arbitrarily large, which means that the effective resistance and capacitance can also be arbitrarily large. To remedy this we must assume that $w(\sigma) = \Theta(1)$ for each $\sigma \in \mathcal{K}_d$. In our analysis we will treat the treat the simplicial complexes as if they were unweighted; under our assumption on the weights this does not affect the asymptotic analysis.

We need one assumption on the input cycle γ. Recall that $\gamma = \sum_{i=1}^{n_{d-1}} \alpha_i \sigma_i$ for $\sigma_i \in \mathcal{K}_{d-1}$ where each $\alpha_i \in \mathbb{R}$. There are two problems with this representation. Our proof for the upper bound on effective capacitance breaks if a pair of coefficients can become arbitrarily close to one another; that is, $|\alpha_i - \alpha_j| < \epsilon$ for some arbitrarily small ϵ. To remedy this we make the assumption that $\gamma \in C_{d-1}(\mathcal{K}, \mathbb{Z})$ such that each $\alpha_i = \Theta(1)$. Under this assumption we can still test the null-homology of any bounded cycle in $C_{d-1}(\mathcal{K}, \mathbb{Q})$ by multiplying the cycle by an integer. Since our algorithm tests null-homology over the reals multiplying by a scalar does not affect the homology of γ.

The second issue is that the magnitude of γ could be arbitrarily large, which implies that the effective resistance and capacitance could also be arbitrarily large. To remedy this we normalize γ as a preprocessing step to obtain a cycle whose magnitude is one. Let $\hat{\gamma}$ be the input cycle, so $\hat{\gamma} = \sum_{i=1}^{n_{d-1}} \alpha_i \sigma_i$ for $\alpha_i \in \mathbb{Z}$ and $\sigma_i \in \mathcal{K}_{d-1}$. After preprocessing we obtain the cycle $\gamma = \sum_{i=1}^{n_{d-1}} \frac{\alpha_i}{\|\hat{\gamma}\|}\sigma_i$ which is a cycle in $C_{d-1}(\mathcal{K}, \mathbb{Q})$. In Theorems 26 and 27 we will assume γ is in this form. We are now ready to upper bound the effective resistance of a cycle γ in a simplicial complex $\mathcal{K}$.

Theorem 26. *Let $\mathcal{K}$ be a d-dimensional simplicial complex and γ a null-homologous $(d-1)$-cycle in $\mathcal{K}$. The effective resistance of γ is bounded above by $\mathcal{R}_\gamma(\mathcal{K}) = O\left(n^2 t^2\right)$.*

Proof. We consider the subcomplex $\mathcal{H}$ obtained by removing β_d d-simplices from $\mathcal{K}$ such that $\ker(\partial_d[\mathcal{H}]) = 0$; in other words $\mathcal{H}$ is a hypertree. Since the effective resistance minimizes over the energy of all unit γ-flows we have $\mathcal{R}_\gamma(\mathcal{K}) \leq \mathcal{R}_\gamma(\mathcal{H})$. Further, since $\mathcal{H}$ is a hypertree there exists a unique unit γ-flow $f \in C_d(\mathcal{H})$ hence $\mathcal{R}_\gamma(\mathcal{H}) = \|f\|^2$.

The matrix $\partial_d[\mathcal{H}]$ has full column rank, so there exists a non-singular $n_d \times n_d$ submatrix of $\partial_d[\mathcal{H}]$. We can express this submatrix as a relative boundary matrix $\partial_d[\mathcal{H}, \mathcal{H}_0]$ where $\mathcal{H}_0$ is the $(d-1)$-subcomplex corresponding to the rows excluded from the submatrix. We have $\partial_d[\mathcal{H}, \mathcal{H}_0]f = \gamma'$ where γ' is the restriction of γ to the rows of $\partial_d[\mathcal{H}, \mathcal{H}_0]$. We have the inequality $\|\gamma'\| \leq \|\gamma\| = 1$.

We will apply Cramer's rule to upper bound the size of f. Let $f(\sigma)$ denote the component of f indexed by the d-simplex σ. By Cramer's rule we have the equality

$$f(\sigma) = \frac{\det(\partial_d[\mathcal{H}, \mathcal{H}_0]_{\sigma,\gamma'})}{\det(\partial_d[\mathcal{H}, \mathcal{H}_0])}$$

where $\partial_d[\mathcal{H}, \mathcal{H}_0]_{\sigma,\gamma'}$ is the matrix obtained by replacing the column of $\partial_d[\mathcal{H}, \mathcal{H}_0]$ indexed by σ with the vector γ'. Since $\det(\partial_d[\mathcal{H}, \mathcal{H}_0])$ is integral $|\det(\partial_d[\mathcal{H}, \mathcal{H}_0])| \geq 1$ so we drop the denominator and focus on the inequality $|f(\sigma)| \leq |\det(\partial_d[\mathcal{H}, \mathcal{H}_0]_{\sigma,\gamma'})|$. We bound

$|\det(\partial_d[\mathcal{H},\mathcal{H}_0]_{\sigma,\gamma'})|$ by its cofactor expansion,

$$\begin{aligned}|\det(\partial_d[\mathcal{H},\mathcal{H}_0]_{\sigma,\gamma_i})| &= \left|\sum_{i=1}^{n_{d-1}}(-1)^i\cdot\gamma'_i\cdot\det(\partial_d[\mathcal{H},\mathcal{H}_0]^{\gamma',i}_{\sigma,\gamma'})\right|\\ &\le \sum_{i=1}^{n_{d-1}}|\gamma'_i|\cdot t\\ &= O\left(\sqrt{n}\cdot t\right)\end{aligned}$$

where $\partial_d[\mathcal{H},\mathcal{H}_0]^{\gamma',i}_{\sigma,\gamma'}$ denotes the submatrix obtained by removing the column c and removing the ith row and c_i denotes the ith component of c. The first inequality comes from Lemma 20, as $\partial_d[\mathcal{H},\mathcal{H}_0]^{\gamma',i}_{\sigma,\gamma'}$ is the relative boundary matrix $\partial[\mathcal{H}\setminus\{\sigma\},\mathcal{H}_0\cup\sigma_i]$, where σ_i is the $(d-1)$-simplex corresponding to the ith row of $\partial_d[\mathcal{H},\mathcal{H}_0]$. The factor of $\sqrt{n}$ comes from the inequality $\|\gamma'\|\le 1$ and the fact that $\|\gamma'\|_1\le\sqrt{n}\|\gamma'\|_2$, which can be shown using the Cauchy-Schwarz inequality. Finally, we compute the flow energy of f as

$$\begin{aligned}\mathsf{J}(f) &= \sum_{\sigma\in\mathcal{K}^d} f(\sigma)^2\\ &\le \sum_{i=1}^{n_d} nt^2\\ &= O\left(n^2t^2\right).\end{aligned}$$

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

The same argument also applies for any subcomplex $\mathcal{L}\subset\mathcal{K}$ where γ is null-homologous in $\mathcal{L}$ which gives us the following corollary.

Corollary 2. *Let $\mathcal{L}\subset\mathcal{K}$ be a d-dimensional simplicial complex and γ a null-homologous $(d-1)$-cycle in $\mathcal{L}$. The effective resistance of γ in $\mathcal{L}$ is bounded above by $\mathcal{R}_\gamma(\mathcal{L})=O\left(n^2t^2\right)$.*

When γ is the boundary of a d-simplex we only sum over at most d non-zero components of γ' when computing the determinant. This reduces the effective resistance to $O\left(nt^2\right)$, and to $O(n)$ when $\mathcal{K}$ is relative torsion-free. This matches the upper bound of $O(n)$ on the effective resistance of a pair of edges in a graph. The upper bound in graphs can easily be realized by an st-path of length $O(n)$.

Before obtaining our upper bound on the effective capacitance of a cycle we need to prove one lemma. In the following lemma, we provide an upper bound on the largest singular value of the coboundary matrix.

Lemma 21. *The largest singular value of the coboundary matrix δ_{d-1} is $\sigma_{max}(\delta_{d-1}) = O(\sqrt{dn})$.*

Proof. The squared singular values of δ_{d-1} are the eigenvalues of $\delta_{d-1}^T\delta_{d-1} = L$. Thus, $\sigma_{\max}(\delta_{d-1})^2 \leq \sum_i \sigma_i(\delta_{d-1})^2 = \text{trace}(L)$, where the $\sigma_i(\delta_{d-1})$ are the singular values of δ_{d-1}. We can obtain an upper bound on $\sigma_{\max}(\delta_{d-1})$ by computing the trace of L. The diagonal elements of L are the degrees of the $(d-1)$-simplices [46, Proposition 3.3.2]. Each d-simplex is the coface of $d+1$ $(d-1)$-simplices, so summing up the diagonal of L, we find $\text{trace}(L) = O(dn)$. Thus, $\sigma_{\max}(\delta_{d-1}) = O(\sqrt{dn})$. □

We now proceed with our upper bound on effective capacitance, which is a similar argument by Cramer's rule.

Theorem 27. *Let $\mathcal{L} \subset \mathcal{K}$ be d-dimensional simplicial complexes, and let γ be a $(d-1)$-dimensional cycle that is null-homologous in $\mathcal{K}$ but not in $\mathcal{L}$. The effective capacitance of γ in $\mathcal{L}$ is bounded above by $\mathcal{C}_\gamma(\mathcal{L}) = O\left(dn^3t^2\right)$.*

Proof. Let p be a γ-potential. By definition, $\delta[\mathcal{L}]p = 0$ and $p(\gamma) = 1$. These constraints can be expressed in the linear system

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

We remove columns from the system until the system has full column rank. Columns in the system correspond to $(d-1)$-simplices in $\mathcal{L}$. Let $\mathcal{L}_0$ be the $(d-1)$-complex obtained by excluding the dropped columns from the $(d-1)$-skeleton of $\mathcal{L}$. We can express the newly obtained system by the relative coboundary operator $\delta[\mathcal{L}, \mathcal{L}_0]$. Removing the columns does not change the image of the linear system, so there is still a solution p'. The newly obtained linear system is

$$\begin{bmatrix} \\ \delta[\mathcal{L}, \mathcal{L}_0] \\ \\ \gamma' \end{bmatrix} p' = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}$$

where γ' is the restriction of γ^T to the remaining columns. The vector p' is not a γ-potential since it is a vector in $C_{d-1}(\mathcal{L} \setminus \mathcal{L}_0)$, not $C_{d-1}(\mathcal{L})$. However, we can extend p' to be a γ-potential by adding zeros in the entries indexed by $\mathcal{L}_0$. Adding zero-valued entries preserves the length of p'.

The next step is to remove rows from the linear system until it has full row rank. This corresponds to removing d-simplices from t he complex $\mathcal{L}$ to create a subcomplex $\mathcal{L}_1$. The row γ' must always be present since removing it would make p' a non-zero solution to the system, and then the system could not have full rank. We obtain the linear system

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

Let $A = \begin{bmatrix} \delta[\mathcal{L}_1, \mathcal{L}_0]^T & (\gamma')^T \end{bmatrix}^T$ and $b = \begin{bmatrix} 0 & 0 & \cdots & 1 \end{bmatrix}^T$. Note that A is an square matrix and we denote its dimensionality by $m \times m$. We now use Cramer's rule to bound the size of $\|p'\|$. By Cramer's rule, p'_i, the ith entry of p', is

$$p'_i = \frac{\det(A_{i,b})}{\det(A)}.$$

where $A_{i,b}$ is the matrix obtained by replacing the ith column with b. We first lower bound $|\det(A)|$. We can express $\det(A)$ by its cofactor expansion on the row of γ' as

$$\det(A) = \sum_{i=1}^{m} (-1)^i \cdot \gamma'_i \cdot \det(\delta[\mathcal{L}_1, \mathcal{L}_0]_i)$$

where $\delta[\mathcal{L}_1, \mathcal{L}_0]_i$ is $\delta[\mathcal{L}_1, \mathcal{L}_0]$ without the ith column. Each term $\delta[\mathcal{L}_1, \mathcal{L}_0]_i$ is integral as $\delta[\mathcal{L}_1, \mathcal{L}_0]$ is an integral matrix. Moreover, each term $\gamma'_i = \gamma_{j_i}/\|\gamma\|$ where $\gamma_{j_i} = \Theta(1)$ is an integer. By our assumptions we have $\|\gamma\| = O(\sqrt{n})$, and using this fact we can then derive a lower bound on $|\det(A)|$,

$$\begin{aligned}
|\det(A)| &= \left|\sum_{i=1}^{m}(-1)^i \cdot \gamma'_i \cdot \det(\delta[\mathcal{L}_1, \mathcal{L}_0]_i)\right| \\
&= \frac{1}{\|\gamma\|}\left|\sum_{i=1}^{m}(-1)^i \cdot \gamma_{j_i} \cdot \det(\delta[\mathcal{L}_1, \mathcal{L}_0]_i)\right| \\
&\geq \frac{1}{\|\gamma\|} \\
&= \Omega\left(\frac{1}{\sqrt{n}}\right).
\end{aligned}$$

We now upper bound $|\det(A_{i,b})|$. We calculate $\det(A_{i,b})$ with the cofactor expansion on the column replaced by b. As b has 1 in its last entry and 0s everywhere else, the cofactor expansion is $\det(A_{i,b}) = \det(A_{i,b}^{i,\gamma'})$ where $A_{i,b}^{i,\gamma'}$ is the matrix where we dropped the ith column and the row γ' from $A_{i,b}$. The matrix $A_{i,b}^{i,\gamma'}$ is a square submatrix of $\delta[\mathcal{K}]$, so we can bound $|\det(A_{i,b})| \leq t$. Thus, $p'_i = \det(A_{i,b})/\det(A) \leq \sqrt{n} \cdot t$ and

$$\begin{aligned}
\|p'\| &= \sqrt{\sum_{i=1}^{m}(p')_i^2} \\
&\leq \sqrt{n^2t^2} \\
&= nt
\end{aligned}$$

The potential energy of p' is $\|\delta[\mathcal{K}]p'\|^2$. We can bound this using Lemma 21 to obtain $\|\delta[\mathcal{K}]p'\|^2 = O\left(dn^3t^2\right)$. □

When γ is the boundary of a d-simplex we sum over at most d non-zero components in the cofactor expansion which allows us to shave off a factor of $O(n)$ from the upper bound. Under the further assumption that $\mathcal{K}$ is relative torsion-free we obtain an upper bound of $O(n^2)$ which matches the bound in graphs [8].

6.2.1 Lower Bounds

Recall that t could be exponential in the number of d-simplices in $\mathcal{K}$. In this section we prove that our bounds on effective resistance and capacitance are tight by providing examples of simplicial complexes containing cycles whose effective resistance and capacitance are exponential in the number of simplices.

Theorem 28. *There exists a 2-dimensional simplicial complex with $\Theta(n)$ triangles containing a cycle γ such that the effective resistance of γ is $\Theta(2^{2n})$.*

Proof. Let $\mathcal{RP}_\gamma$ denote a simplicial complex homeomorphic to the real projective plane with a disk removed. The boundary of the removed disk is the cycle γ. Hence, we have that the boundary of the complex is $\partial_2 \mathcal{RP}_\gamma = 2\alpha + \gamma$ for some 1-cycle α. We require $\mathcal{RP}_\gamma$ to be triangulated in such a way that $|\operatorname{supp}(\alpha)| = |\operatorname{supp}(\gamma)|$, that is, α and γ contain the same number of edges. Let the constant c denote the number of triangles in $\mathcal{RP}_\gamma$.

We consider a collection of complexes $\mathcal{RP}_{\gamma_0}, \mathcal{RP}_{\gamma_1}, \ldots, \mathcal{RP}_{\gamma_{n-1}}, \mathcal{D}_{\gamma_n}$. Each $\mathcal{RP}_{\gamma_i}$ is constructed in the same way was $\mathcal{RP}_\gamma$ but with disjoint simplices, and $\mathcal{D}_{\gamma_n}$ triangulation of a disk using c triangles with boundary γ_n such that $|\operatorname{supp}(\gamma_n)| = |\operatorname{supp}(\gamma)|$. Each $\mathcal{RP}_{\gamma_i}$ has boundary $\partial_2 \mathcal{RP}_{\gamma_i} = 2\alpha_i + \gamma_i$.

We consider the simplicial complex $\mathcal{K}$ constructed by taking the quotient space under the identification $\alpha_i \sim \gamma_{i+1}$. That is, we glue the boundary component α_i along the boundary component γ_{i+1}. The resulting complex contains a unique unit γ_0-flow f. This is because the value of f must be equal to 1 on $\mathcal{RP}_{\gamma_0}$ (otherwise f is not a unit γ_0-flow) and this completely determines the assignments on each other $\mathcal{RP}_{\gamma_i}$ and $\mathcal{D}_{\gamma_n}$ as the boundary components must cancel each other out.

We now compute the effective resistance of γ_0 by explicitly describing the unit γ_0-flow f. The flow will have the property that $f(\sigma_i) = f(\sigma_j)$ for each $\sigma_i, \sigma_j \in \mathcal{RP}_{\gamma_k}$ (and similarly for $\mathcal{D}_{\gamma_n}$). That is, f is uniform on each subcomplex in our original collection. For each σ in the subcomplex indexed by γ_i we set $f(\sigma) = (-1)^i \cdot 2^i$. To see that this is indeed a unit γ_0-flow we compute the boundary:

$$\begin{aligned}
\partial f &= \left(\sum_{i=0}^{n-1} (-1)^i \cdot 2^i \partial \mathcal{RP}_{\gamma_i} \right) + (-1)^n \cdot 2^n \partial \mathcal{D}_{\gamma_n} \\
&= \left(\sum_{i=0}^{n-1} (-1)^i \cdot 2^i (\gamma_i + 2\gamma_{i+1}) \right) + (-1)^n \cdot 2^n \gamma_n \\
&= \gamma_0.
\end{aligned}$$

Finally, we compute the flow energy of f:

$$\mathsf{J}(f) = \sum_{i=0}^{n} c \cdot (2^i)^2 = \frac{c}{3}\left(2^{2n+2} - 1\right) = \Theta(2^{2n}).$$

□

Now we use a very similar construction to provide an example of a simplicial complex with a cycle whose effective capacitance is exponential.

Theorem 29. *There exists a 2-dimensional simplicial complex with* $\Theta(n)$ *triangles containing a cycle* γ *and a subcomplex such that the effective capacitance of* γ *in the subcomplex is* $\Theta(2^{2n})$.

Proof. As in Theorem 28 we let $\mathcal{RP}_\gamma$ be a triangulation of the real projective plane with a disk removed, we consider the same sequence $\mathcal{RP}_{\gamma_0}, \mathcal{RP}_{\gamma_1}, \ldots, \mathcal{RP}_{\gamma_{n-1}}, \mathcal{D}_{\gamma_n}$, but under the identifications $\gamma_i \sim \alpha_{i+1}$ and $\gamma_{n-1} \sim \gamma_n$. We denote the obtained quotient space by $\mathcal{K}$. Note that α_0 is a null-homologous cycle with a unique unit γ-flow given by

$$f = \left(\sum_{i=0}^{n-1} (-1)^i \cdot \frac{1}{2^{i+1}} \mathcal{RP}_{\gamma_i}\right) + (-1)^n \cdot \frac{1}{2^{n+1}} \mathcal{D}_{\gamma_n},$$

which can be verified similarly to Theorem 28.

Let T be a triangle in $\mathcal{D}_{\gamma_n}$ and let p be a unit α_0-potential in $\mathcal{K} \setminus T$. We will express p as a collection of flows in graphs $G_1, \ldots, G_n$ with each G_i being constructed the dual graph of $\mathcal{RP}_{\gamma_i}$ and G_n being constructed from the dual of $\mathcal{D}_{\gamma_n} \setminus T$. Each G_i is the dual graph of $\mathcal{RP}_{\gamma_i}$ but with an additional vertex v_{γ_i} representing the boundary component γ_i; it is adjacent to each vertex dual to a triangle incident to γ_i. The graph G_n is constructed the same way from $\mathcal{D}_{\gamma_n}$.

We claim that when $\delta[\mathcal{K} \setminus T](p)$ is restricted to a single $\mathcal{RP}_{\gamma_i}$ (or $\mathcal{D}_{\gamma_n}$) it is dual to a multi-source single-sink flow in G_i. The set of source vertices S_i is the set of vertices dual to the triangles incident to α_i, and the sink is v_{γ_i}. We denote this flow by p_i, and further we claim that the flow has value $\|p_i\| = 2^{i+1}$. Next, we will prove our claims by induction, beginning with G_0 as the base case.

Since p is a unit α_0-potential we have $p(\alpha_0) = 1$, so $\delta(p)(S_i^*) = 2$ since there are two triangles incident to each edge in α_0. Let $\overrightarrow{S}_i$ denote the set of edges leaving S_i; we have $p_i(\overrightarrow{S_i^*}) = 2$ which is the amount of flow p_i sends out of S_i. For any vertex $v \notin S_i \cup \{v_{\gamma_i}\}$

the net flow leaving v must be equal to zero, otherwise its dual triangle v^* has $\delta(p)(v^*) \neq 0$ contradicting the assumption that p is a unit α_0-potential in $\mathcal{K} \setminus T$. Thus, we conclude that the amount of flow entering v_{γ_i} is equal to 2 which proves the base case.

Next, we assume that p_{i-1} is a flow in G_{i-1} from S_{i-1} to v_{γ_i} of value 2^i. Consider the set of triangles S_i^* which are dual to the source S_i. The inductive hypothesis ensures that $p(\alpha_i) = 2^i$, but $\delta[\mathcal{K} \setminus T](p) = 0$, so we must have that $\delta[\mathcal{RP}_{\gamma_i}](p)(S_i) = -2^{i+1}$ since each edge in α_i is incident to two triangles in $\mathcal{RP}_{\gamma_i}$. Hence, S_i has a net flow of 2^{i+1} exiting it in p_i. By the same reasoning as the previous paragraph we see that any vertex $v \notin S_i \cup \{v_{\gamma_i}\}$ must have a net flow of zero exiting it. Thus, we have a flow from S_i to γ_i of value 2^{i+1}. Note that the same reasoning proves the analogous result for $\mathcal{D}_{\gamma_n}$. In particular, we have that the net flow entering v_{γ_n} is 2^{n+1}.

The above argument shows that $\delta[\mathcal{K} \setminus T](p)(t) = 0$ and that $\delta(p)(T) = 2^{n+1}$. Hence, the potential energy of p is $\mathcal{J}(p) = \Theta(2^{2n})$. $\square$

6.3 Applications to Quantum Algorithms

In this section we provide an example of an algorithm whose complexity can be parameterized by the effective resistance and capacitance of a cycle. We consider the following problem: 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 $\mathcal{L} \subseteq \mathcal{K}$, decide whether or not γ is null-homologous in $\mathcal{L}$. This problem is a high dimensional generalization of testing st-connectivity in a subgraph.

We present a quantum algorithm based on the span program model solving this problem. Our algorithm is an adaption of the quantum algorithm deciding st-connectivity of Belovs and Reichardt [8] which has a query complexity of $O(n^{3/2})$, which can be obtained by parameterizing the query complexity in terms of the effective resistance and capacitance of $\{s, t\}$.

Our algorithm has a query complexity of $O\left(\sqrt{\mathcal{R}_{\max}(\gamma) \cdot \mathcal{C}_{\max}(\gamma)}\right)$ where $\mathcal{R}_{\max}(\gamma)$ and $\mathcal{C}_{\max}(\gamma)$ are the maximum effective resistance and capacitance of γ over all subcomplexes of $\mathcal{K}$. Our bounds on effective resistance and capacitance obtained in Theorems 26 and 27 hold for any subcomplex, so we obtain a query complexity of $O\left(n^{5/2}t^2\right)$. Under the assumptions that γ is the boundary of a d-simplex and that $\mathcal{K}$ is relative torsion-free we obtain the query complexity $O(n^{3/2})$, matching the query complexity of st-connectivity.

These assumptions always hold in the case of st-connectivity in graphs.

We will begin with a short introduction to span programs. We will present only the necessary definitions and theorems needed to obtain our quantum query algorithm. Then we will show that the problem of deciding whether or not γ is null-homologous fits into the span program model which will imply the existence of our quantum query algorithm.

In the spirit of quantum computing we will adopt ***bra-ket notation*** for vectors. A column vector v will be written as a ***ket*** denoted as $|v\rangle$. The corresponding row vector v^T will be written as a ***bra*** denoted as $\langle v|$. The inner product of two vectors u, v will be denoted as $\langle u|v\rangle$.

6.3.1 A brief introduction to span programs

Span programs were first defined by Karchmer and Wigderson [61] and were first used for quantum algorithms by Reichardt and Špalek [86]. Intuitively, a span program is a model of computation which encodes a boolean function $f\colon \{0,1\}^n \to \{0,1\}$ 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 30.)

Definition 15. *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* $A\colon \mathcal{H} \to \mathcal{U}$*.*

For every string $x = (x_1, \ldots, x_n) \in \{0,1\}^n$ *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* $A(x) = A\Pi_{\mathcal{H}(x)}\colon \mathcal{H} \to \mathcal{U}$ *where* $\Pi_{\mathcal{H}(x)}$ *is the projection of* $\mathcal{H}$ *onto* $\mathcal{H}(x)$*.*

The quantum query complexity of evaluating $\mathcal{P}$ depends on the sizes of the positive and negative witnesses, which we now define. Let $\mathcal{P}$ be a span program and let $x \in \{0,1\}^n$ be a binary string. 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 the size of the smallest positive witness of x given by its ℓ_2 norm,

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

If no positive witness exists for x, then $w_+(x,\mathcal{P}) = \infty$. If there is a positive witness for x, then x is a ***positive instance***. That is, there exists a solution to the linear system contained entirely in the subspace $\mathcal{H}(x)$ if and only if there exists a positive witness for x.

A ***negative witness*** for x is a linear map $\langle w| : \mathcal{U} \to \mathbb{R}$ such that $\langle w|A\Pi_{\mathcal{H}(x)} = 0$ and $\langle w|\tau\rangle = 1$. Similarly, the ***negative witness size*** of x is the minimum size over all negative witnesses given by the value $\|A^T|w\rangle\|$ or equivalently,

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

If no negative witness exists for x, then $w_-(x,\mathcal{P}) = \infty$. If there is a negative witness for x, then x is a ***negative instance***. If there exists a negative instance of x then there is no solution to the linear system $A|y\rangle = |\tau\rangle$ with $|y\rangle \in \mathcal{H}(x)$ since $\langle w|$ is orthogonal to the image of $A\Pi_{\mathcal{H}(x)}$. Conversely, if there exists no solution to $Ay = |\tau\rangle$ in $\mathcal{H}(x)$ then we can write $|\tau\rangle = |\tau_1\rangle + |\tau_2\rangle$ with $|\tau_1\rangle \in \operatorname{im} A\Pi_{\mathcal{H}(x)}$ and $|\tau_2\rangle \in (\operatorname{im} A\Pi_{\mathcal{H}(x)})^\perp$ and, after scaling by an appropriate constant, $|\tau_2\rangle$ is a negative witness.

A string $x \in \{0,1\}^n$ will either be a positive or negative instance of $\mathcal{P}$. A span program $\mathcal{P}$ ***decides*** the function $f\colon \{0,1\}^n \to \{0,1\}$ if $f(x) = 1$ when x is a positive instance and $f(x) = 0$ when x is a negative instance. We define the positive and negative witness sizes for the span program $\mathcal{P}$ with respect to the boolean function f to be $W_+(f,\mathcal{P}) = \max_{x\in f^{-1}(1)} w_+(x,\mathcal{P})$ and $W_-(f,\mathcal{P}) = \max_{x\in f^{-1}(0)} w_-(x,\mathcal{P})$. Since we are maximizing over the positive and negative witnesses the values of $W_+(f,\mathcal{P})$ and $W_-(f,\mathcal{P})$ are finite. Reichardt [87] showed that the query complexity of a span program is a function of the positive and negative witness sizes of the program.

Theorem 30 (Reichardt [87])**.** *Let $f : \{0,1\}^n \to \{0,1\}$. Let $\mathcal{P}$ be a span program that decides f. There is a bounded error quantum algorithm that decides f with query complexity $O\left(\sqrt{W_+(f,\mathcal{P})W_-(f,\mathcal{P})}\right)$.*

Reichardt's algorithm in *bounded error* in the sense that it returns the correct answer with probability $2/3$. We will not go into details on the analysis of the complexity of a span program, however we give a brief overview here following the work of Belovs and Reichardt [8, Theorem 13]. The query model evaluates the complexity of a quantum algorithm by the number of times it queries an oracle. The oracle our quantum algorithm will query is a unitary transformation $U = R_{\ker A} R_{\mathcal{H}(x)}$ where $R_{\ker A} = 2\Pi_{\ker A} - I$ and $R_{\mathcal{H}(x)} = 2\Pi_{\mathcal{H}(x)} - I$ are reflections around the subspaces $\ker A$ and $\mathcal{H}(x)$. By Π_S we denote the orthogonal projection onto the subspace S. To evaluate the span program we perform phase estimation [67] on U with precision $\Theta\left(1/\sqrt{W_+(f,\mathcal{P})W_-(f,\mathcal{P})}\right)$ and accept if and only if the measured phase is zero. The query complexity follows from the fact that performing phase estimation to precision ϵ requires $\Theta(1/\epsilon)$ queries to U. Time efficient implementations of the algorithm can be obtained by constructing a polynomial sized quantum circuits for U.

6.3.2 A span program for deciding null-homology

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 [61] and used to develop quantum algorithms in [8, 15, 59, 60]. Let $\mathcal{K}$ be a d-dimensional simplicial complex. Let $|\gamma\rangle \in C_{d-1}(\mathcal{K})$ be a $(d-1)$-cycle. Let n be the number of d-simplices in $\mathcal{K}$. Order the d-simplices $\{\sigma_1, \dots, \sigma_n\}$. Let $\langle w| \colon \mathcal{K}_d \to \mathbb{R}$ be a weight function on the d-simplices. We define a span program over the strings $\{0,1\}^n$ in the following way.

1. $\mathcal{H} = C_d(\mathcal{K})$, with $\mathcal{H}_{i,1} = \mathrm{span}\{\sigma_i\}$ and $\mathcal{H}_{i,0} = \{0\}$.
2. $\mathcal{U} = C_{d-1}(\mathcal{K})$
3. $|\tau\rangle = |\gamma\rangle$
4. $A = \partial_d W^{1/2} \colon C_d(\mathcal{K}) \to C_{d-1}(\mathcal{K})$

We denote the above span program by $\mathcal{P}_{\mathcal{K}}$. Let $x \in \{0,1\}^N$ be a binary string. We define the subcomplex $\mathcal{K}(x) \coloneqq \mathcal{K}^{d-1} \cup \{\sigma_i : x_i = 1\}$. That is, $\mathcal{K}(x)$ contains the d-simplices σ_i such that $x_i = 1$. The subcomplex $\mathcal{K}(x)$ corresponds to the subspace $\mathcal{H}(x)$.

There exists a solution to the linear system $\partial_d W^{1/2}\Pi_{\mathcal{K}(x)}|f\rangle = |\gamma\rangle$ if and only if the cycle γ is null-homologous in $\mathcal{K}(x)$. If $\partial_d[\mathcal{K}(x)]|f\rangle = |\gamma\rangle$ then $W^{1/2}|f\rangle$ is a solution. Conversely, if $\partial_d W^{1/2}\Pi_{\mathcal{K}(x)}|f\rangle = \gamma$ then $\partial_d W^{-1/2}|f\rangle = |\gamma\rangle$. Hence, $|\gamma\rangle$ is null-homologous 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 $f : \{0,1\}^n \to \{0,1\}$ where $f(x) = 1$ if and only if γ is a null-homologous cycle in the subcomplex $\mathcal{K}(x)$.

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

Lemma 22. *Let $x \in \{0,1\}^N$ be a positive instance. There is a bijection between positive witnesses $|w_+\rangle$ for x and unit γ-flows $|f\rangle$ in $\mathcal{K}(x)$. Moreover, the positive witness size is equal to the effective resistance of $|\gamma\rangle$ 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 $\partial_d W^{1/2}|w_+\rangle = |\gamma\rangle$. We construct a unit γ-flow $|f\rangle$ in $\mathcal{K}(x)$ by $|f\rangle = W^{1/2}|w_+\rangle$; $|f\rangle$ is indeed a unit γ-flow as $\partial_d|f\rangle = \partial_d W^{1/2}|w_+\rangle = |\gamma\rangle$. Moreover, $|w_+\rangle = W^{-1/2}|f\rangle$. The flow energy of $|f\rangle$ is

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

Hence, the flow energy of $|f\rangle$ equals the witness size of x.

Conversely, let $|f\rangle$ 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 $|f\rangle$ equals the positive witness size of x. □

Lemma 23. *Let $x \in \{0,1\}^N$ be a negative instance. There is a bijection between negative witnesses $\langle w_-|$ for x and unit γ-potentials $\langle p|$ in $\mathcal{K}(x)$. Moreover, the negative witness size is equal to the effective capacitance of $|\gamma\rangle$ in $\mathcal{K}(x)$; that is, $w_-(x, \mathcal{P}_\mathcal{K}) = C_\gamma(\mathcal{K}(x))$.*

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 $(d-1)$-cochain; there exists some $\langle p| \in C^{d-1}(\mathcal{K})$ such that $\langle p| = \langle w_-|$. Since $\langle w_-|\gamma\rangle = 1$ we immediately obtain the equality $\langle p|\gamma\rangle = 1$. To show that $\langle p|$ is a unit γ-potential we must show that the coboundary of $\langle p|$ is zero in $\mathcal{K}(x)$. By the definition of a negative witness we have

$$\begin{aligned} 0 &= \langle w_-|\partial_d W^{1/2}\Pi_{\mathcal{K}(x)} \\ &= \langle p|\partial_d W^{1/2}\Pi_{\mathcal{K}(x)} \\ &= \langle \delta_d(p)|W^{1/2}\Pi_{\mathcal{K}(x)}. \end{aligned}$$

Since $W^{1/2}$ 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 $\langle p|$ we have

$$\begin{aligned}
\|\langle w_-|\partial_d W^{1/2}\|^2 &= \langle p\partial_d W^{1/2}|p\partial_d W^{1/2}\rangle \\
&= \langle W^{1/2}\delta_d(p)|W^{1/2}\delta_d(p)\rangle \\
&= \sum_{\sigma\in\mathcal{K}_d} \langle \delta_d(p)|\sigma\rangle^2 w(\sigma) \\
&= \mathcal{J}(p).
\end{aligned}$$

Conversely, let $\langle p|$ be a unit γ-potential for $\mathcal{K}(x)$ we construct a negative witness for x by setting $\langle w_-| := \langle p|$. Since the coboundary of $\langle p|$ 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 $\langle w_-|\partial_d W^{1/2}\Pi_{\mathcal{K}(x)} = 0$ by the reasoning in the previous paragraph. Also by the previous paragraph we have that the potential energy of $\langle p|$ is equal to the negative witness size of $\langle w_-|$ which concludes the proof. □

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

Theorem 31. *Given a d-dimensional simplicial complex $\mathcal{K}$, a $(d-1)$-dimensional cycle γ that is null-homologous in $\mathcal{K}$, a witness string $x \in \{0,1\}^n$, 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 $\mathcal{K}(x)$ whose quantum query complexity is $O\left(\sqrt{\mathcal{R}_{\max}(\gamma)\mathcal{C}_{\max}(\gamma)}\right)$, where $\mathcal{R}_{\max}$ is the maximum effective resistance of γ in any subcomplex $\mathcal{K}(y)$ and $\mathcal{C}_{\max}$ is the maximum effective capacitance γ in any subcomplex $\mathcal{L} \subseteq \mathcal{K}$.*

Proof. By Theorem 30, 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 $W_-(f,\mathcal{P}_{\mathcal{K}}) = \max_{x\in f^{-1}(0)} \mathcal{C}_\gamma(\mathcal{K}(x)) = \mathcal{C}_{\max}(\gamma)$. □

By Theorems 26 and 27 we obtain an upper bound on the query complexity parameterized by the number of simplices and the cardinality of the torsion subgroups of the relative homology groups.

Theorem 32. *Let $\mathcal{K}$ be a d-dimensional simplicial complex, γ a $(d-1)$-dimensional null-homologous cycle in $\mathcal{K}$, and $\mathcal{K}(x)$ a d-dimensional subcomplex $\mathcal{K}(x) \subseteq \mathcal{K}$. There exists a* *quantum algorithm deciding whether or not γ is null-homologous in $\mathcal{K}(x)$ whose quantum query complexity is $O\left(n^{5/2}t^2\right)$.*

Finally, we state the query complexity under some assumptions that arise in the case of st-connectivity in graphs. In graphs we have that $\gamma = t - s$, hence the support of γ is equal to 2. Under the assumption that the support of γ is bounded above by $O(d)$ we can shave off a factor of n from both the flow energy and potential energy of any unit γ-flow and unit γ-potential. Further, graphs do not contain torsion, so we make the additional assumption that $\mathcal{K}$ is relative torsion-free. Under these assumptions our query complexity matches the query complexity arising from the span program deciding st-connectivity.

Corollary 3. *Let $\mathcal{K}$ be a d-dimensional simplicial complex, γ be a $(d-1)$-dimensional cycle that is null-homologous in $\mathcal{K}$, and $\mathcal{K}(x) \subseteq \mathcal{K}$ be a d-dimensional subcomplex for a fixed d. Further assume that $\mathcal{K}$ is relative torsion-free and that $|\operatorname{supp}(\gamma)| = O(d)$. There exists a quantum algorithm deciding whether or not γ is null-homologous in $\mathcal{K}(x)$ whose quantum query complexity is $O\left(n^{3/2}\right)$.*

Bibliography

[1] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $O(\sqrt{\log n})$ approximation algorithms for min uncut, min 2CNF deletion, and directed cut problems. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 573–581, New York, NY, USA, 2005. ACM.

[2] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(09):3821–3851, September 2006.

[3] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications.* Prentice-Hall, Inc., USA, 1993.

[4] J. W. Alexander. A proof and extension of the Jordan-Brouwer separation theorem. *Transactions of the American Mathematical Society*, 23(4):333–349, 1922.

[5] Josh Alman and Virginia Vassilevska Williams. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 522–539. Society for Industrial and Applied Mathematics, January 2021.

[6] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), April 2009.

[7] Per Austrin and Subhash Khot. A simple deterministic reduction for the gap minimum distance of code problem. In *Proceedings of the 38th International Colloquim Conference on Automata, Languages and Programming - Volume Part I*, ICALP'11, pages 474–485, Berlin, Heidelberg, 2011. Springer-Verlag.

[8] 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.

[9] E. Berlekamp, R. McEliece, and H. van Tilborg. On the inherent intractability of certain coding problems (corresp.). *IEEE Transactions on Information Theory*, 24(3):384–386, May 1978.

[10] Arnab Bhattacharyya, Ameet Gadekar, Suprovat Ghoshal, and Rishi Saket. On the hardness of learning sparse parities. In *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, pages 11:1–11:17, 2016.

[11] Glencora Borradaile, Erin Wolf Chambers, Kyle Fox, and Amir Nayyeri. Minimum Cycle and Homology Bases of Surface Embedded Graphs. In Sándor Fekete and Anna Lubiw, editors, *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 23:1–23:15, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[12] Glencora Borradaile and Philip Klein. An $O(n \log n)$ algorithm for maximum *st*-flow in a directed planar graph. *J. ACM*, 56(2), April 2009.

[13] Glencora Borradaile, William Maxwell, and Amir Nayyeri. Minimum Bounded Chains and Minimum Homologous Chains in Embedded Simplicial Complexes. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:15, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[14] Oleksiy Busaryev, Sergio Cabello, Chao Chen, Tamal K. Dey, and Yusu Wang. Annotating simplices with a homology basis and its applications. In Fedor V. Fomin and Petteri Kaski, editors, *Algorithm Theory – SWAT 2012*, pages 189–200, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[15] 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.

[16] J. Carvalho, Mikael Vejdemo-Johansson, Danica Kragic, and Florian Pokorny. An algorithm for calculating top-dimensional bounding chains. *PeerJ Computer Science*, 4:e153, 05 2018.

[17] Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. In *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, SCG '09, pages 377–385, New York, NY, USA, 2009. ACM.

[18] Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Minimum cuts and shortest homologous cycles. In *Proc. 25th Ann. Symp. Comput. Geom.*, pages 377–385, 2009.

[19] Erin W. Chambers, Jeff Erickson, and Amir Nayyeri. Homology flows, cohomology cuts. *SIAM Journal on Computing*, 41(6):1605–1634, 2012.

[20] Erin Wolf Chambers and Mikael Vejdemo-Johansson. Computing minimum area homologies. *Comput. Graph. Forum*, 34(6):13–21, September 2015.

[21] A. K. Chandra, P. Raghavan, W. L. Ruzzo, and R. Smolensky. The electrical resistance of a graph captures its commute and cover times. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, page 574–586, New York, NY, USA, 1989. Association for Computing Machinery.

[22] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Comput. Complex.*, 15(2):94–114, June 2006.

[23] Chao Chen and Daniel Freedman. Hardness results for homology localization. *Discrete & Computational Geometry*, 45(3):425–448, Apr 2011.

[24] A.V. Chernavsky and V.P. Leksine. Unrecognizability of manifolds. *Annals of Pure and Applied Logic*, 141(3):325 – 335, 2006. Papers presented at the Second St. Petersburg Days of Logic and Computability Conference on the occasion of the centennial of Andrey Andreevich Markov, Jr.

[25] R.J. Daverman and G. Venema. *Embeddings in Manifolds*. Graduate studies in mathematics. American Mathematical Society, 2009.

[26] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. Embeddability in $\mathbb{R}^3$ is NP-Hard. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, page 1316–1329, USA, 2018. Society for Industrial and Applied Mathematics.

[27] 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.

[28] Tamal K. Dey. *Computing Height Persistence and Homology Generators in $\mathbb{R}^3$ Efficiently*, pages 2649–2662.

[29] Tamal K. Dey and Sumanta Guha. Algorithms for manifolds and simplicial complexes in Euclidean 3-space (preliminary version). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 398–407, New York, NY, USA, 1996. Association for Computing Machinery.

[30] 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.

[31] 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.

[32] Tamal K. Dey, Tianqi Li, and Yusu Wang. Efficient algorithms for computing a minimal homology basis. In Michael A. Bender, Martín Farach-Colton, and Miguel A. Mosteiro, editors, *LATIN 2018: Theoretical Informatics*, pages 376–398, Cham, 2018. Springer International Publishing.

[33] Rod G. Downey, Michael R. Fellows, Alexander Vardy, and Geoff Whittle. The parametrized complexity of some fundamental problems in coding theory. *SIAM J. Comput.*, 29(2):545–570, October 1999.

[34] 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, pages 135–144, New York, NY, USA, 2011. ACM.

[35] Art M. Duval, Caroline J. Klivans, and Jeremy L. Martin. Cuts and flows of cell complexes. *Journal of Algebraic Combinatorics*, 41:969–999, 2015.

[36] Art M. Duval, Caroline J. Klivans, and Jeremy L. Martin. *Simplicial and cellular trees*, pages 713–752. Springer International Publishing, Cham, 2016.

[37] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. Applied Mathematics. American Mathematical Society, 2010.

[38] Jeff Erickson. *Algorithms*. `http://algorithms.wtf`, 2019.

[39] Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 1166–1176, Philadelphia, PA, USA, 2011. Society for Industrial and Applied Mathematics.

[40] Jeff Erickson and Amir Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. In *Proc. 22nd Ann. ACM-SIAM Symp. Discrete Algorithms*, pages 1166–1176, 2011.

[41] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, page 1038–1046, USA, 2005. Society for Industrial and Applied Mathematics.

[42] Arthur F. Veinott, Jr. and George B. Dantzig. Integral extreme points. *SIAM Review*, 10(3):371–372, July 1968.

[43] Ralph H. Fox. On the imbedding of polyhedra in 3-space. *Annals of Mathematics*, pages 462 – 470, 1948.

[44] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs with applications. *SIAM J. Comput.*, 16(6):1004–1004, 1987.

[45] Robert Ghrist and Sanjeevi Krishnan. A topological max-flow-min-cut theorem. *2013 IEEE Global Conference on Signal and Information Processing*, pages 815–818, 2013.

[46] Timothy Goldberg. Combinatorial laplacians of simplicial complexes. B.S. Thesis, Annandale-on-Hudson, NY, 2002.

[47] Joshua Grochow and Jamie Tucker-Foltz. Computational topology and the unique games conjecture. 06 2018.

[48] gyashfe (https://mathoverflow.net/users/75344/gyashfe). Planar duality generalized to embedded simplicial complexes. MathOverflow. URL:https://mathoverflow.net/q/357833 (version: 2020-04-19).

[49] Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4):315–358, Aug 2019.

[50] Refael Hassin. Maximum flow in (s, t) planar networks. *Inf. Process. Lett.*, 13:107, 1981.

[51] Refael Hassin and Donald B. Johnson. An $O(n \log^2 n)$ algorithm for maximum flow in undirected planar networks. *SIAM J. Comput.*, 14(3):612–624, 1985.

[52] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[53] John Hopcroft and Robert Tarjan. Efficient planarity testing. *J. ACM*, 21(4):549–568, October 1974.

[54] J. D. Horton. A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM Journal on Computing*, 16(2):358–366, April 1987.

[55] Sergei Ivanov (https://mathoverflow.net/users/4354/sergei ivanov). computational complexity. MathOverflow. URL:https://mathoverflow.net/q/118428 (version: 2013-01-09).

[56] Oscar Ibarra, Shlomo Moran, and Roger Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms*, 3:45–56, 03 1982.

[57] Alon Itai and Yossi Shiloach. Maximum flow in planar networks. *SIAM J. Comput.*, 8:135–150, 1979.

[58] Giuseppe F. Italiano, Yahav Nussbaum, Piotr Sankowski, and Christian Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. In *Proc. 43rd Ann. ACM Symp. Theory Comput.*, pages 313–322, 2011.

[59] 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.

[60] Stacey Jeffery and Shelby Kimmel. Quantum algorithms for graph connectivity and formula evaluation. *Quantum*, 1:26, August 2017.

[61] M. Karchmer and A. Wigderson. On span programs. In *[1993] Proceedings of the Eigth Annual Structure in Complexity Theory Conference*. IEEE Comput. Soc. Press.

[62] Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna Paluch. A faster algorithm for minimum cycle basis of graphs. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming*, pages 846–857, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

[63] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM J. Comput.*, 37(1):319–357, April 2007.

[64] Subhash A. Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative-type metrics into ℓ_1. *J. ACM*, 62(1):8:1–8:39, March 2015.

[65] G. Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148:497–508, 1847.

[66] Danil Kirsanov and Steven Gortler. A discrete global minimization algorithm for continuous variational problems. *ArXiv*, 2004.

[67] A. Yu. Kitaev. Quantum measurements and the Abelian Stabilizer Problem. *arXiv e-prints*, pages quant–ph/9511026, November 1995.

[68] D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, Dec 1993.

[69] Woong Kook and Kang-Ju Lee. Simplicial networks and effective resistance. *Advances in Applied Mathematics*, 100:71–86, 2018.

[70] Fabian Latorre. The maxflow problem and a generalization to simplicial complexes. *arXiv*, 1212.1406, 2012.

[71] Lek-Heng Lim. Hodge Laplacians on graphs. *SIAM Review*, 63(3):685–715, 2020.

[72] Saunders Mac Lane. A combinatorial condition for planar graphs. *Fundamenta Mathematicae*, 28(1):22–32, 1937.

[73] Pasin Manurangsi and Luca Trevisan. Mildly exponential time approximation algorithms for vertex cover, balanced separator and uniform sparsest cut. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, pages 20:1–20:17, 2018.

[74] A. Markov. Insolubility of the problem of homeomorphy. *Doklady Akademii Nauk SSSR*, 121, 01 1958.

[75] Jiří Matoušek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, page 855–864, USA, 2009. Society for Industrial and Applied Mathematics.

[76] Sarah Morell, Ina Seidel, and Stefan Weltge. Minimum-cost integer circulations in given homology classes. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2725–2739. Society for Industrial and Applied Mathematics, January 2021.

[77] James R. Munkres. *Elements of Algebraic Topology.* Addison Wesley Publishing Company, 1984.

[78] A. Nabutovsky. Einstein structures: Existence versus uniqueness. *Geometric and Functional Analysis*, 5(1):76–91, January 1995.

[79] Andrew Newman. Small simplicial complexes with prescribed torsion in homology. *Discrete & Computational Geometry*, 62(2):433–460, March 2018.

[80] P. V. O'Neil. A short proof of mac lane's planarity theorem. *Proceedings of the American Mathematical Society*, 37(2):617–617, February 1973.

[81] James B. Orlin. Max flows in $O(nm)$ time, or better. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, pages 765–774, 2013.

[82] Braxton Osting, Sourabh Palande, and Bei Wang. Spectral sparsification of simplicial complexes for clustering and label propagation, 2019.

[83] Christos Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 229–234, New York, NY, USA, 1988. ACM.

[84] Zhaoxia Qu. On Tutte's algorithm for recognizing binary graphic matroids. Master's thesis, The Pennsylvania State University, 2002.

[85] Igor Razgon and Barry O'Sullivan. Almost 2-SAT is fixed-parameter tractable. *J. Comput. Syst. Sci.*, 75(8):435–450, December 2009.

[86] Ben Reichardt and Robert Spalek. Span-program-based quantum algorithm for evaluating formulas. *Theory of Computing*, 8(1):291–319, 2012.

[87] 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.

[88] John Reif. Minimum s-t cut of a planar undirected network in $O(n \log^2 n)$ time. *SIAM J. Comput.*, 12:71–81, 1983.

[89] Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., USA, 1986.

[90] Primoz Skraba, Gugan Thoppe, and D. Yogeshwaran. Randomly weighted d-complexes: Minimal spanning acycles and persistence diagrams. *The Electronic Journal of Combinatorics*, 27(2), April 2020.

[91] Daniel Spielman. *Spectral and Algebraic Graph Theory*. 2019.

[92] John M. Sullivan. *A Crystalline Approximation Theorem for Hypersurfaces*. PhD thesis, Princeton University, 1990.

[93] W. T. Tutte. An algorithm for determining whether a given binary matroid is graphic. *Proceedings of the American Mathematical Society*, 11(6):905–917, 1960.

[94] A. Vardy. The intractability of computing the minimum distance of a code. *IEEE Trans. Inf. Theor.*, 43(6):1757–1766, November 1997.

[95] I Volodin, V Kuznetsov, and A Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Russian Mathematical Surveys*, 29:71, 10 2007.

[96] Hassler Whitney. On the abstract properties of linear dependence. *American Journal of Mathematics*, 57(3):509–533, 1935.

APPENDICES

.1 Directed paths in simplicial complexes

In a directed graph an st-path is equivalent to an 1-chain $P = \sum \alpha_i e_i$ with $\partial P = t - s$ such that $\alpha_i \in \{0, 1\}$. The generalization to simplicial complexes is straightforward: we define a d-dimensional directed γ-path in a d-complex $\mathcal{K}$ to be a d-chain $P = \sum \alpha_i \sigma_i$ such that $\partial P = \gamma$ and $\alpha_i \in \{0, 1\}$ for all i. We will now show that computing a directed γ-path is NP-complete for $d \geq 2$. The reduction from graph 3-coloring is a slight adaptation of the proof of Theorem 16.

Theorem 33. *Computing a directed γ-path in a d-dimensional simplicial complex is NP-complete for $d \geq 2$.*

Proof. First, to show that computing a directed γ-path is in NP we note that we can check the coefficients and the boundary of a 2-chain in polynomial time. To prove NP-hardness we give a reduction from graph 3-coloring.

Given a graph $G = (V, E)$ we construct a 2-complex $\mathcal{K}$ with a boundary component γ such that there exists a directed γ-path if and only if G is 3-colorable. First, construct a punctured sphere with $|V| + 1$ boundary components. One of these boundary components is γ, the remaining $|V|$ are in bijection with the vertices of G and we will denote the component corresponding to the vertex v as γ_v. For each vertex v we construct three additional punctured spheres v_r, v_b, v_g each with $\deg(v) + 1$ boundary components. These punctured spheres correspond to the three potential colors of v: red, blue, and green and we refer to them as the *color surfaces* of v. We glue v_r, v_b, v_g to γ_v each along some boundary component. For each edge $e = (u, v)$ we construct nine tubes with two boundary components. Each tube connects a color surface of u to a color surface of v, for example the tube $\mathcal{T}_{r,b}$ connects the red color surface of u with the blue color surface of v. We orient the complex such that each 2-chain with boundary γ and coefficients in $\{0, 1\}$ is an oriented manifold with boundary γ. For each tube connecting two color surfaces of the same color (for example, $\mathcal{T}_{r,r}$) we invert the orientation of one simplex, such that any bounding chain for γ including the inverted simplex will have to assign it a coefficient of -1.

Any 2-chain with boundary γ and coefficients in $\{0, 1\}$ is a surface as for each vertex it must contain exactly one color surface and for each edge it must contain exactly one tube. Moreover, if any tube connecting two color surfaces of the same color is contained in the solution it must contain some simplex a negative coefficient and is not an directed γ-path. It follows that G is 3-colorable if and only if there is some directed γ-path. $\square$