HODGERANK: APPLYING COMBINATORIAL HODGE THEORY TO SPORTS RANKING

BY

ROBERT K SIZEMORE

A Thesis Submitted to the Graduate Faculty of

WAKE FOREST UNIVERSITY GRADUATE SCHOOL OF ARTS AND SCIENCES

in Partial Fulfillment of the Requirements

for the Degree of

MASTER OF ARTS

Mathematics

May 2013

Winston-Salem, North Carolina

Approved By:

R. Jason Parsley, Ph.D., Advisor

Sarah Raynor, Ph.D., Chair

Matt Mastin, Ph.D.

W. Frank Moore, Ph.D.

Acknowledgments

There are many people who helped to make this thesis possible. First, I would like to thank my advisor Dr. Jason Parsley, it was at his suggestion that I began studying ranking, Hodge Theory and *HodgeRank* in particular. I would also like to thank my thesis committee: Dr. Raynor, Dr. Moore and Dr. Mastin for all of their feedback. Their comments and suggestions made the final revisions of this document infinitely less painful than it would have been otherwise. I would like to thank Dr. Moore again for all of his help and suggestions regarding actual programming and implementation and for offering to let us use his server. Finally, I want to thank Furman University for inviting Dr. Parsley and me to speak about *HodgeRank* at their Carolina Sport Analytics Meeting this spring.

Table of Contents

List of Figures

Abstract

In this thesis, we examine a ranking method called *HodgeRank*. *HodgeRank* was introduced in 2008 by Jiang, Lim, Yao and Ye as "a promising tool for the statistical analysis of ranking, especially for datasets with cardinal, incomplete, and imbalanced information." To apply these methods, we require data in the form of pairwise comparisons, meaning each voter would have rated items in pairs (A is preferred to B). An obvious candidate for ranking data in the form of pairwise comparisons comes from sports, where such comparisons are very natural (i.e. games between two teams). We describe a simple way in which *HodgeRank* can be used for sports ratings and show how *HodgeRank* generalizes the well-established sports rating method known as *Massey's method*.

The Combinatorial Hodge Theorem, for which *HodgeRank* is named, tells us that the space of possible game results in a given season decomposes into three subspaces: a gradient subspace, a harmonic subspace and a curly subspace. The gradient subspace contains no intransitive game results, that is, no ordinal or cardinal relations of the form $A < B < C < A$, where '$A < B$' indicates team B beating team A. If there are no intransitive relations, it is straightforward to obtain a global ranking of the teams. To this end, *HodgeRank* projects our data onto this subspace of consistent game results. From this projection, we can determine numerical ratings of each team. The residual, which lies in the harmonic and curly subspace, captures these inconsistencies and intransitive relations in the data and so a large residual may indicate a less reliable rating. In a sports context, this may mean that upsets would be more likely or that there is more parity within the league.

Chapter 1: Introduction

It is well known from voting theory that voter preferences may be plagued with inconsistencies and intransitive preference relations. Suppose that we have polled some number of voters and asked them to rate three candidates, by comparing two candidates at a time. That is, we may ask "do you prefer candidate A or candidate B?", "candidate B or candidate C?", etc. It may be the case that voters prefer candidate A to candidate B, candidate B to candidate C, but still prefer candidate C to candidate A, giving us the intransitive preference relation $A < B < C < A$. These same intransitive relations arise in other ranking contexts as well. In sports, a typical example of such an inconsistency could be a cyclic relation of the form "Team A beats Team B beats Team C beats Team A".

When studying a ranking problem, it is often the case that a graph structure can be assigned to the dataset. Suppose we want to rank sports teams, and our data is a list of game results between the teams we wish to rank. We may assign a graph structure to this dataset by assigning vertices to represent each team, and letting edges represent games played, so that two teams have an edge between them if they have played each other at least once. We can give the graph further structure by assigning to each edge a direction and weight. Technically, we think of this as being a skew-symmetric function on pairs of teams, but we usually represent this as a weighted, directed graph. For example, each edge may be given a weight specifying the score difference (winning team score - losing team score) for the game it represents, and a direction specifying (pointing towards) the winning team. When performing actual computations, we generally store these graphs as matrices, the graphs are used mainly

to illustrate the general principles.

$$X = \begin{pmatrix} 0 & -2 & 1 & 0 \\ 2 & 0 & 3 & 1 \\ -1 & -3 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix} \quad \longleftrightarrow$$

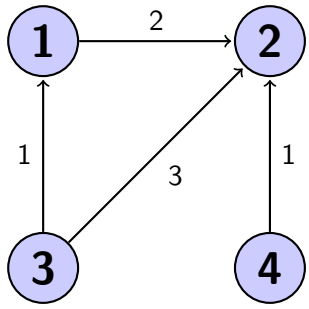

Figure 1.1: A skew-symmetric matrix and the "flow" it induces on a graph

In the case of our weighted directed graph, an inconsistency is a closed path (a path starting and ending at the same vertex) along which the scores raised and lowered along the path are nonzero. That is, if we add or subtract the weights along each edge in the closed path, the "net weight" along the path is nonzero. You may convince yourself that the above graph has no inconsistencies.

In this thesis, we examine a ranking method known as *HodgeRank* which exploits the topology of ranking data to both obtain a global ranking and measure the inherent inconsistency in our dataset. By treating our data as a special type of function on graph complexes, combinatorial Hodge Theory allows us to decompose our data into a cyclic, locally noncyclic, and a globally noncyclic component. We can then formulate a least squares problem to project our data onto the subspace of (globally) noncyclic data. From the least squares projection, the "consistent" component of our data, it is straightforward to obtain a global ranking, given that the underlying graph is connected. However, the novelty of *HodgeRank* is its use of the least squares residual. The residual captures the inconsistencies in the underlying data, and so can be used to analyze the "rankability" of the data. If the residual is large, it suggests

that our least squares ranking may not be very meaningful, that is, the original data had too many inconsistencies.

Suppose we have already assigned a simple graph structure to our data, along with a function specifying the weights and directions along each edge. A simple graph is an instance of a more general, but still relatively simple topological space, known as a simplicial complex. What makes simplicial complexes desirable is that we do not need to consider them as geometric objects to ascertain information about their topologies. That is, simplicial complexes can be a viewed as purely combinatorial objects, i.e., vertices, pairs of vertices (edges), triples of vertices (triangles), etc. By modeling our data as (co)chains on an abstract simplicial complex, we can use some powerful tools from algebraic topology to attack our ranking problem. The combinatorial Hodge Decomposition Theorem tells us that the space in which our data lives decomposes into a acyclic component with no relations of the form $A < B < C < A$ and a space of cyclic rankings where we do get such intransitivities. In general, our data will not lie completely in the subspace of acyclic (consistent) rankings, but we formulate a least square problem to find the nearest match that does.

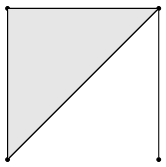

Figure 1.2: The underlying simplicial complex for the graph in Figure 1

The thesis has three parts. In the first chapter we discuss the Hodge Decomposition Theorem for vector fields on domains in $\mathbb{R}^3$. We prove some preliminary

results about the Laplace equation, the Dirichlet and Neumann problems, and the Biot-Savart law. We then state the Hodge Decomposition theorem for vector fields and prove some of various decompositions.

The second chapter is a largely self-contained exposition on the mathematical preliminaries behind *HodgeRank*. We cover the necessary material from graph theory, introduce edge flows, the combinatorial gradient, divergence and curl. We discuss some of the basic geometric and topological properties of simplicial complexes, we include some material that will not be used directly in our applications, but will help motivate the more abstract concepts. Next, we introduce abstract simplicial complexes, (co)chains and simplicial (co)homology groups which will be the objects of interest in *HodgeRank*. Finally, we will prove the combinatorial Hodge theorem, from which *HodgeRank* is named.

In the third chapter of our exposition, we show how the mathematics in the first section can be used to solve ranking problems. We describe and justify *HodgeRank* analytically, and demonstrate how *HodgeRank* can be used in practice. Two particular applications that we are interested in are NCAA Division I Basketball and NCAA Division I Football. The underlying graph structures of each dataset are interesting in their own ways. The football data is more sparse, since each team plays only 12-13 games, only 7-8 of which are conference games, so that even the conferences do not form cliques. In Division I basketball each team plays every team in its conference at least once, so the conferences form cliques in our graph. Since sparsity in the graph often manifests itself as non-trivial homology in the corresponding graph complex, *HodgeRank* may provide unique insights to how the geometry affects the reliability of our rankings, and allow us to isolate certain inconsistencies in the data.

Chapter 2: Hodge Theory for Vector Fields

2.1 Prerequisites

2.1.1 The Dirichlet and Neumann Problems

Let Ω denote a compact subset of $\mathbb{R}^3$ with k connected components $\Omega_1, \ldots, \Omega_k$ with smooth boundaries $\partial\Omega_i$. Given a sufficiently smooth (at least C^2) function ϕ defined on Ω, the **Laplacian** is defined as the operator that acts on ϕ such that:

$$\Delta\phi = \frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial y^2} + \frac{\partial^2\phi}{\partial z^2}.$$

For any given function f defined on Ω, **Poisson's equation** refers the second-order partial differential equation $\Delta\phi = f$. In the special case where $f = 0$, the Poisson equation is referred to as **Laplace's equation**. The solutions of Laplace's equation are called **harmonic functions**. We will want to "construct" solutions of the Poisson and Laplace equations that satisfy certain boundary conditions. To this end, we will need several results regarding the following two problems.

1. **The Dirichlet Problem.** Given a function f defined on Ω, and a function g defined on $\partial\Omega$, find a function ϕ on Ω satisfying

   $$\Delta\phi = f \text{ on } \Omega \text{ and } \phi = g \text{ on } \partial\Omega.$$

   The requirement that $\phi = g$ on $\partial\Omega$ is called a *Dirichlet boundary condition*.

2. **The Neumann Problem.** Given a function f defined on Ω, and a function g

defined on $\partial\Omega$, find a function ϕ on Ω satisfying

$$\Delta\phi = f \text{ on } \Omega \text{ and } \frac{\partial\phi}{\partial n} = g \text{ on } \partial\Omega.$$

The requirement that $\frac{\partial\phi}{\partial n} = g$ on $\partial\Omega$ is called a *Neumann boundary condition.*

It turns out that in order to show the existence of solutions to the Dirichlet/Neumann problem for Poisson's equation ($\Delta\phi = f$), we need only show existence of solutions for the Laplace equation ($\Delta\phi = 0$) [8]. However, before we can do this we will need to develop a few preliminary results regarding harmonic functions, i.e., solutions of Laplace's equation. The first two theorems may be familiar from vector calculus.

Theorem 2.1. ***(The Divergence Theorem/Gauss's Theorem)*** *Let $\vec{v}$ be a vector field that is C^∞ on Ω. Then,*

$$\int_\Omega \nabla \cdot \vec{v} \text{ d(vol)} = \int_{\partial\Omega} \vec{v} \cdot n \text{ d(area)}.$$

Corollary 2.1.1. ***(Gauss's Theorem for Gradients)*** *If $\vec{v} = \nabla\phi$ then:*

$$\int_\Omega \Delta\phi \text{ d(vol)} = \int_{\partial\Omega} \nabla\phi \cdot n \text{ d(area)} = \int_{\partial\Omega} \frac{\partial\phi}{\partial n} \text{ d(area)}$$

We state Gauss's Theorem without proof, as it is a familiar theorem from vector calculus. For a proof of Stokes' theorem, the generalized version of Gauss's Theorem, see [20] or [23]. Next, we derive two identities that will be needed to solve the Dirichlet and Neumann problems.

Theorem 2.2. ***(Green's Identities)*** *Let u and v be functions that are C^2 on Ω.*

1. ***Green's First Identity***

$$\int_{\partial\Omega} v\frac{\partial u}{\partial n} \text{ d(area)} = \int_\Omega (v\Delta u + \nabla u \cdot \nabla v) \text{ d(vol)}$$

2. ***Green's Second Identity***

$$\int_{\partial\Omega}\left(v\frac{\partial u}{\partial n}-u\frac{\partial v}{\partial n}\right)\ \mathrm{d}(\text{area})=\int_{\Omega}(v\Delta u-u\Delta v)\ \mathrm{d}(\text{vol})$$

Proof. 1. If we let ϕ be the vector field given by $\phi = v\nabla u$, then we have:

$$\int_{\partial\Omega}\phi\cdot n\ \mathrm{d}(\text{area})=\int_{\partial\Omega}v\nabla u\cdot n\ \mathrm{d}(\text{area})=\int_{\partial\Omega}v\frac{\partial u}{\partial n}\ \mathrm{d}(\text{area}),$$

and also,

$$\int_{\Omega}\nabla\cdot\phi\ \mathrm{d}(\text{vol})=\int_{\Omega}\nabla\cdot v\nabla u\ \mathrm{d}(\text{vol})=\int_{\Omega}(v\Delta u+\nabla u\cdot\nabla v)\ \mathrm{d}(\text{vol}).$$

Green's first identity then follows from the Divergence Theorem:

$$\int_{\partial\Omega}v\frac{\partial u}{\partial n}\ \mathrm{d}(\text{area})=\int_{\Omega}(v\Delta u+\nabla u\cdot\nabla v)\ \mathrm{d}(\text{vol}).$$

2. Switching the roles of v and u in Green's first identity gives us:

$$\int_{\partial\Omega}u\frac{\partial v}{\partial n}\ \mathrm{d}(\text{area})=\int_{\Omega}(u\Delta v+\nabla v\cdot\nabla u)\ \mathrm{d}(\text{vol}).$$

Subtracting this from the earlier equation:

$$\int_{\partial\Omega}v\frac{\partial u}{\partial n}\ \mathrm{d}(\text{area})=\int_{\Omega}(v\Delta u+\nabla u\cdot\nabla v)\ \mathrm{d}(\text{vol}),$$

gives us Green's second identity:

$$\int_{\partial\Omega}\left(v\frac{\partial u}{\partial n}-u\frac{\partial v}{\partial n}\right)\ \mathrm{d}(\text{area})=\int_{\Omega}(v\Delta u-u\Delta v)\ \mathrm{d}(\text{vol}).\quad\square$$

The next theorem, which is a consequence of the Divergence Theorem, gives a necessary condition for the Neumann problem to be well-posed. If the functions f, g (as in, $\Delta u = f$ on Ω and $u = g$ on $\partial\Omega$) do not satisfy this "compatibility condition", then we know right away there are no solutions to the corresponding Neumann problem.

Theorem 2.3. ***(Compatibility Condition for the Neumann Problem)*** *Given functions f defined on Ω and g defined on $\partial\Omega$. If there exists a function ϕ on Ω satisfying*

$$\Delta\phi = f \text{ on } \Omega \text{ and } \frac{\partial\phi}{\partial n} = g \text{ on } \partial\Omega,$$

then we have

$$\int_{\Omega_i} f \text{ d(vol)} = \int_{\partial\Omega_i} g \text{ d(area)}$$

for each connected component Ω_i of Ω.

Proof. Suppose that ϕ is a solution of the Neumann problem with

$$\Delta\phi = f \text{ on } \Omega \text{ and } \frac{\partial\phi}{\partial n} = g \text{ on } \partial\Omega,$$

then consider the integral of g along the boundary of any component Ω_i of Ω:

$$\int_{\partial\Omega_i} g \text{ d(area)} = \int_{\partial\Omega_i} \frac{\partial\phi}{\partial n} \text{ d(area)} = \int_{\partial\Omega_i} \nabla\phi \cdot n \text{ d(area)}.$$

Applying the Divergence theorem to Ω_i we have:

$$\int_{\partial\Omega_i} \nabla\phi \cdot n \text{ d(area)} = \int_{\Omega_i} \Delta\phi \text{ d(vol)},$$

by hypothesis ϕ is a solution of the Neumann problem and therefore

$$\int_{\partial\Omega_i} g \text{ d(area)} = \int_{\Omega_i} \Delta\phi \text{ d(vol)} = \int_{\Omega_i} f \text{ d(vol)}. \quad \square$$

Corollary 2.1.2. ***(Laplace's Equation)*** *Suppose that u is a solution of the Neumann problem for Laplace's equation with boundary conditions $\partial\phi/\partial n = g$ on $\partial\Omega$, then we must have*

$$\int_{\partial\Omega_i} \frac{\partial u}{\partial n} \text{ d(area)} = \int_{\partial\Omega_i} g \text{ d(area)} = 0$$

for each component Ω_i of Ω.

Proof. Setting $\Delta u = f = 0$ in the previous theorem shows the result. □

2.1.2 Harmonic Functions

The next two lemmas, will be referred to as the "mean-value property" and "maximum principle", shows two important properties of harmonic functions. These two results will be useful in proving later theorems.

Lemma 2.1.1. ***(Mean-Value Property of Harmonic Functions)*** *Let u be harmonic on a domain D. Then, for any $x \in D$ and any ball $\overline{B_r(x)} \subset D$, we have:*

$$u(x) = \frac{1}{4\pi r^2} \int_{\partial B_r(x)} u \text{ d(area)} = \frac{1}{\frac{4}{3}\pi r^3} \int_{B_r(x)} u \text{ d(vol)}.$$

Proof. For simplicity suppose that $x = 0$, since if the ball is not centered at the origin, we may make a change variables $x \mapsto x_0 + y$. Let $r > 0$ such that $B_r(0) \subset \Omega$. We let B_ϵ, B_r denote $B_\epsilon(0)$ and $B_r(0)$ respectively. Since u is harmonic on the closure of B, then by the previous corollary we have:

$$\int_{\partial B_r} \frac{\partial u}{\partial n} = 0.$$

Let $\epsilon > 0$ such that $\epsilon < r$. Consider the open set $\Omega = B_r \backslash \overline{B_\epsilon}$. Let $v = 1/r$, then both u and v are harmonic on Ω. By Green's second identity:

$$\int_{\partial\Omega} \left(v\frac{\partial u}{\partial n} - u\frac{\partial v}{\partial n} \right) = \int_\Omega \left(v\cancelto{0}{\Delta u} - u\cancelto{0}{\Delta v} \right) = 0$$

By our construction $\partial\Omega = \partial B_r \cup \partial B_\epsilon$. The outward normal $\mathbf{n}$ for $\partial\Omega$, is $\vec{r}/r$ on ∂B_r and $\vec{\epsilon}/\epsilon$ on ∂B_ϵ. Thus,

$$\frac{\partial v}{\partial n} = \nabla v \cdot n = \nabla \left(\frac{1}{r} \right) \cdot \frac{\vec{r}}{r} = \left(\frac{x}{r^3}, \frac{y}{r^3} \right) \cdot \left(\frac{x}{r}, \frac{y}{r} \right) = \frac{x^2 + y^2}{r^4} = \frac{r^2}{r^4} = \frac{1}{r^2}.$$

Similarly, $\partial v/\partial n = 1/\epsilon$ on ∂B_ϵ. Thus,

$$\begin{aligned} 0 &= \int_{\partial\Omega} \left(v\frac{\partial u}{\partial n} - u\frac{\partial v}{\partial n} \right) \\ &= \int_{\partial B_r} \left(v\frac{\partial u}{\partial n} - u\frac{\partial v}{\partial n} \right) + \int_{\partial B_\epsilon} \left(v\frac{\partial u}{\partial n} - u\frac{\partial v}{\partial n} \right) \\ &= \int_{\partial B_r} \left(\frac{1}{r}\frac{\partial u}{\partial n} - u\frac{1}{r^2} \right) + \int_{\partial B_\epsilon} \left(\frac{1}{2\epsilon}\frac{\partial u}{\partial n} + u\frac{1}{\epsilon^2} \right) \end{aligned}$$

Since, $1/r^2$ and $1/r$ are constant on these boundaries, we can pull them out of the integral. Using the corollary to Green's theorem for harmonic functions, give us:

$$0 = \tfrac{1}{r}\int_{\partial B_r} \overset{0}{\overbrace{\tfrac{\partial u}{\partial n}}} - \tfrac{1}{r^2}\int_{\partial B_r} u + \tfrac{1}{\epsilon}\int_{\partial B_\epsilon} \overset{0}{\overbrace{\tfrac{\partial u}{\partial n}}} + \tfrac{1}{\epsilon^2}\int_{\partial B_\epsilon} u = \tfrac{1}{\epsilon^2}\int_{\partial B_\epsilon} u - \tfrac{1}{r^2}\int_{\partial B_r} u.$$

Rearranging the expression on the right side and multiplying both sides by $1/4\pi$ gives us:

$$\frac{1}{4\pi\epsilon^2}\int_{\partial B_\epsilon} u = \frac{1}{4\pi r^2}\int_{\partial B_r} u.$$

Thus, the average of u over the balls B_ϵ and B_r are the same. This holds for any ϵ such that $0 < \epsilon < r$, so if we take $\epsilon \to 0$, then this average approaches the value $u(0)$ as a consequence of continuity of the integral $\int_{\partial B_\epsilon} u$:

$$u(0) = \lim_{\epsilon\to 0} \tfrac{1}{4\pi\epsilon^2}\int_{\partial B_\epsilon} u = \lim_{\epsilon\to 0} \tfrac{1}{4\pi r^2}\int_{\partial B_r} u = \tfrac{1}{4\pi r^2}\int_{\partial B_r} u.$$

To prove the mean-value theorem for the solid ball B_r, we multiply both sides of the above expression by $4\pi r^2 dr$ and integrate from 0 to r:

$$4\pi\rho^2 u(0) = \int_{\partial B_\rho} u \;\longrightarrow\; \int_0^r 4\pi\rho^2 u(0)d\rho = \int_0^r \int_{\partial B_\rho} u d\rho \;\longrightarrow\; \frac{4}{3}\pi r^3 u(0) = \int_{B_r} u.$$

Thus,

$$u(0) = \frac{1}{\frac{4}{3}\pi r^3}\int_{B_r} u. \quad \square$$

Theorem 2.4. ***(The Maximum Principle)*** *Suppose Ω is a connected, open set in $\mathbb{R}^3$. If the real-valued function u is harmonic on Ω and $\sup_{x\in\Omega} u(x) = m < \infty$, then either $u(x) < m$ for all $x \in \Omega$ or $u(x) = m$ for all $x \in \Omega$.*

Proof. Since u is harmonic on Ω, it is continuous on Ω. Since u is continuous and $\{m\}$ is a closed set, the set $M = u^{-1}(\{m\})$ is closed in Ω. Let $x \in M$, then $x \in \Omega$ so there exists some $r > 0$ such that $\overline{B_r(x)} \subset \Omega$. By the mean-value theorem we have:

$$m = u(x) = \frac{1}{\frac{4}{3}\pi r^3} \int_{B_r(x)} u \, \mathrm{d(vol)},$$

so the average of u over the ball $B_r(x)$ is m, but $u \leq m$ on $B_r(x)$ so we must have $u(x) = m$ for all $x \in B_r(x)$. Thus, $B_r(x) \subset M = u^{-1}(\{m\})$, so the set M must be open. Since Ω is a connected, so the only subsets of Ω which are both open and closed are Ω and $\varnothing$. $M \subset \Omega$ is both closed and open, so $M = u^{-1}(\{m\}) = \Omega$ or $\varnothing$. This is what we wanted to show. □

Corollary 2.1.3. ***Compact Maximum Principle*** *Suppose Ω is a domain in $\mathbb{R}^3$ such that $\overline{\Omega}$ is compact. Let u be harmonic on Ω and continuous on $\partial\Omega$. Then the maximum value of u on $\overline{\Omega}$ is acheived on $\partial\Omega$.*

Proof. Since the function u is harmonic on the compact set $\overline{\Omega}$, it achieves its maximum value on $\overline{\Omega}$ by the extreme value theorem. So, the maximum is acheived either on Int Ω or $\partial\Omega$. If the maximum is achieved at an interior point then u must be constant throughout Int Ω by the previous theorem. If this is the case, then u is also constant on $\overline{\Omega}$ by our additional continuity assumption. Thus, the maximum is also attained on $\partial\Omega$. □

We note that there is a corresponding minimum principle for harmonic functions (consider the maximum principle on the harmonic function $-u$). An important consequence of these minimum/maximum principles is that any harmonic function that

vanishes on $\partial\Omega$ must be zero throughout Ω. The next theorem, which is a direct consequence of the maximum principle, shows that if a solution to the Dirichlet problem exists, then then it is unique.

Theorem 2.5. ***(Uniqueness of Solutions to the Dirichlet Problem)*** *Given a function f defined on Ω, and a function g defined on $\partial\Omega$, if a solution ϕ to the Dirichlet problem:*

$$\Delta\phi = f \text{ on } \Omega \text{ and } \phi = g \text{ on } \partial\Omega,$$

exists, then it is unique.

Proof. Suppose that ϕ, φ are two solutions of the Dirichlet problem stated above. Then define a function ω on Ω by $\omega = \phi - \varphi$. We can see that ω is harmonic on Ω since

$$\Delta\omega = \Delta(\phi - \varphi) = \Delta\phi - \Delta\varphi = f - f = 0.$$

Also, $\omega = \phi - \varphi = g - g = 0$ on $\partial\Omega$. So, ω satisfies

$$\Delta\omega = 0 \text{ on } \Omega \text{ and } \omega = 0 \text{ on } \partial\Omega,$$

therefore by the maximum/minimum principle for harmonic functions we have:

$$0 = \inf_{x\in\partial\Omega} \omega(x) \leq \omega(x) \leq \sup_{x\in\partial\Omega} \omega(x) = 0 \text{ for all } x \in \Omega.$$

Thus, $\omega(x) = 0$ for all $x \in \Omega$ and so $\phi(x) - \varphi(x) = 0$ for all $x \in \Omega$. Therefore, $\phi = \varphi$. □

Next, we introduce a function Φ, called the fundamental solution of Laplace's equation. The defining property of this solution is that Φ is harmonic everywhere except for one point $\xi \in \text{Int } \Omega$, and that $\Delta\Phi = \delta(\xi)$. Here, δ is the Dirac-delta function, which we will describe in the next section.

2.1.3 The Fundamental Solution of Laplace's Equation

Before our next definition, we introduce an important generalized function called the *Dirac delta function*, or simply *delta function*, which will be denoted by δ. A full discussion of the Dirac delta function is beyond the scope of this thesis (see [8] or [9]), but we will note a few of its properties and try to give an informal overview of the concept. In physics, a delta function is often used to represent the mass-density of a point particle with unit mass. A point-particle is a particle with no volume, whose mass is all concentrated at one point. For a point particle located at $\xi \in \mathbb{R}^3$, this would imply that $\delta(x, \xi)$ is infinite at $x = \xi$:

$$\delta(x,\xi) = \begin{cases} 0 & \text{if } x \neq \xi \\ \infty & \text{if } x = \xi \end{cases}.$$

The first thing to note about the delta function is that it not actually a function in the usual sense, since it would be undefined at the point ξ. Secondly, we are treating the delta function as having two arguments, this is simply for convenience, we let ξ denote the location of the point-particle, or the location of singularity (if we drop the physical interpretation), and x be the argument point.

The other crucial property of the delta function is how it acts on a test function $f(x)$. Suppose that our singularity, or point particle, is located at ξ within Ω. Recall that the mass contained within Ω is given by the integral:

$$\text{Mass} = \int_\Omega (\text{Density}) \ \ \text{d(vol)},$$

If our point-mass is located within Ω we should expect that δ satisfies:

$$\int_\Omega \delta(x,\xi) \ \text{d(vol)} = \begin{cases} 0 & \text{if } \xi \notin \Omega \\ 1 & \text{if } \xi \in \Omega. \end{cases}$$

If $\xi \in \Omega$, then convolving a function f by a delta function centered at ξ produces the function's value at ξ.

$$\int_\Omega f(x)\delta(x,\xi)\ \mathrm{d(vol)} = f(\xi).$$

Next, we introduce a function known as the fundamental solution Φ of Laplace's equation, which satisfies the relation $\Delta\Phi = \delta(x,\xi)$ for some ξ.

Definition 1. (Fundamental Solution of Laplace's Equation) The **fundamental solution** to Laplace's equation, is the function Φ defined on $\mathbb{R}^3 - \{0\}$ given by:

$$\Phi(x) = \frac{1}{4\pi}\frac{1}{|x|},$$

or in spherical coordinates:

$$\Phi(r) = \frac{1}{4\pi r}.$$

Physically, we may think of the fundamental solution to Laplace's equation as representing the electric potential of a unit negative charge placed at the origin.

It will be convenient to move the singularity ξ of Φ. We will now treat Φ as a function of two variables, $\Phi(x,\xi)$, where x is the argument point and ξ is the parameter point (the point of singularity):

$$\Phi(x,\xi) = \Phi(x-\xi) = \frac{1}{4\pi}\frac{1}{|x-\xi|}.$$

Our next lemma is simply a point of convenience, following the approach in [8] and [9] in particular. We show that we need only consider the Dirichlet/Neumann problems for Laplace's equation, since if solutions exist for Laplace's equation, then solutions exist for Poisson's equation. For this proof, we will consider convolutions of the fundamental solution.

Lemma 2.1.2. ***(Reduction of Poisson equation to Laplace's Equation)*** *To show there exists solutions of the Dirichlet/Neumann problems for the Poisson equation $\Delta u = f$, with boundary conditions $u = g$, or $\partial u/\partial n = g$, it is enough to show there are solutions to the Laplace equation $\Delta u = 0$ satisfying such boundary conditions.*

Proof. Consider the following three Dirichlet problems:

1. $\Delta u = f$ on Ω, and $u = g$ on $\partial\Omega$,
2. $\Delta v = f$ on Ω, and $v = 0$ on $\partial\Omega$,
3. $\Delta w = 0$ on Ω, and $w = g$ on $\partial\Omega$.

Clearly, if (2) and (3) have solutions v and w, then $u = v + w$ is a solution of (1). We want show that if (3) has a solution, then (1) has a solution u, so by our previous remark it is enough to show that (3) having a solution implies (2) has a solution. Suppose that we can solve (3), and we want to solve (2). First, extend f to the function f' by requiring f' to be zero outside of Ω, that is:

$$f'(x) = \begin{cases} f(x) & \text{if } x \in \Omega \\ 0 & \text{if } x \notin \Omega \end{cases}.$$

Then we define the function v' as the convolution $v' = f' * \Phi$, i.e.,

$$v'(\xi) = \int_\Omega f(x)\Phi(x - \xi)\ \mathrm{d}(\mathrm{vol_x}) = \int_\Omega f(x)\Phi(x, \xi)\ \mathrm{d}(\mathrm{vol_x}).$$

We note that $\Delta v' = f$:

$$\Delta v'(\xi) = \int_\Omega f(x)\Delta\Phi(x, \xi)\ \mathrm{d}(\mathrm{vol_x}) = \int_\Omega f(x)\delta(x, \xi)\ \mathrm{d}(\mathrm{vol_x}) = f(\xi),$$

where $\delta(x, \xi)$ is the Dirac delta function at ξ. Let w be the solution of (3) with $g = v'$, and define $v = v' - w$, then v solves (2).

Consider the Neumann problems:

1. $\Delta u = f$ on Ω, and $\frac{\partial u}{\partial n} = g$ on $\partial\Omega$,
2. $\Delta v = f$ on Ω, and $\frac{\partial v}{\partial n} = 0$ on $\partial\Omega$,
3. $\Delta w = 0$ on Ω, and $\frac{\partial w}{\partial n} = g$ on $\partial\Omega$.

As in the proof for the Dirichlet problem, it is enough to show that if (3) has a solution, then (2) has a solution. Suppose that the Neumann problem (3) can be solved. Let v' be defined as above, and let w be the solution of (3) given by:

$$\Delta w = 0 \text{ on } \Omega, \text{ and } \frac{\partial w}{\partial n} = \frac{\partial v'}{\partial n} \text{ on } \partial\Omega,$$

then as before, $v = v' - w$ solves (2). □

Theorem 2.6. ***(Representation Formula)*** *Any harmonic function $u(x)$ can be represented as an integral over the boundary of Ω. If $\Delta u = 0$ in Ω, then*

$$u(\xi) = \iint_{\partial\Omega} \left[-u(x) \frac{\partial \Phi(x, \xi)}{\partial n} + \Phi(x, \xi) \frac{\partial u}{\partial n} \right] \text{ d(area)}$$

Proof. For simplicity, we assume that Ω contains the origin and set $\xi = 0$. Let $\epsilon > 0$ and let u be harmonic on Ω. We recall that that function $\Phi(x, 0)$ has a discontinuity at the origin, but is C^∞ on the set Ω_ϵ, which we will define as Ω with an ϵ-ball removed around the origin:

$$\Omega_\epsilon \equiv \Omega - B_\epsilon(0) = \{x \in \Omega : |x| \geq 0\}.$$

Applying Green's second identity to the functions $\Phi(x, 0)$ and u on the domain Ω_ϵ gives us:

$$\int_{\partial\Omega_\epsilon} \left(\Phi(x, 0) \frac{\partial u(x)}{\partial n} - u(x) \frac{\partial \Phi(x, 0)}{\partial n} \right) \text{ d(area)} = 0,$$

where the right-hand side vanishes since u and Φ are harmonic on Ω_ϵ. We note that $\partial\Omega_\epsilon = \partial\Omega \cup \partial B_\epsilon(0)$, with the outward normal of Ω_ϵ being the inward normal of $B_\epsilon(0)$ on $\partial\Omega_\epsilon \cap \partial B_\epsilon(0)$. Thus,

$$\begin{aligned} 0 &= \int_{\partial\Omega_\epsilon} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} \\ &= \int_{\partial\Omega} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} - \int_{\partial B_\epsilon(0)} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} \end{aligned}$$

Rearranging this equality gives us:

$$\int_{\partial\Omega} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} = \int_{\partial B_\epsilon(0)} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)},$$

so it is sufficient to show that the right hand side goes to $u(0)$ as $\epsilon \to 0$. Since $\xi = 0$, it will be convenient to switch to spherical coordinates. Then, on $\partial B_\epsilon(0)$, $\partial/\partial n = -\partial/\partial r$ and clearly, $\Phi = 1/4\pi r = 1/4\pi\epsilon$. So, the right-hand side of the above equation is equal to:

$$\begin{aligned} \int_{\partial B_\epsilon(0)} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} &= \int_{\partial B_\epsilon(0)} \left(u\frac{\partial}{\partial r}\left(\frac{1}{4\pi r}\right) - \left(\frac{1}{4\pi r}\right)\frac{\partial u}{\partial r}\right) \text{ d(area)} \\ &= \int_{\partial B_\epsilon(0)} \left(u\left(\frac{1}{4\pi r^2}\right) - \left(\frac{1}{4\pi r}\right)\frac{\partial u}{\partial r}\right) \text{ d(area)} \\ &= \frac{1}{4\pi\epsilon^2}\int_{\partial B_\epsilon(0)} u \text{ d(area)} - \frac{1}{4\pi\epsilon}\int_{\partial B_\epsilon(0)} \frac{\partial u}{\partial r} \text{ d(area)} \\ &= u(0) + \epsilon\overline{\frac{\partial u}{\partial r}}, \end{aligned}$$

In the last step, we have used the mean-value property on the harmonic function u. We let $\overline{\partial u/\partial r}$ denote the average of $\partial u/\partial r$ on $\partial B_\epsilon(0)$. The function $\partial u/\partial r$ is continuous on the compact domain $B_\epsilon(0)$ and therefore bounded on $B_\epsilon(0)$, so as $\epsilon \to 0$:

$$\int_{\partial B_\epsilon(0)} \left(\Phi\frac{\partial u}{\partial n} - u\frac{\partial\Phi}{\partial n}\right) \text{ d(area)} = u(0),$$

this shows the result. □

2.1.4 Green's function and the Neumann function

For now, let us assume the existence of solutions to the Dirichlet/Neumann problems. We can define two special functions called the Green's function and the Neumann function that can be used to characterize the solutions of the Dirichlet problem and the Neumann problem respectively. We will introduce a third special function called the kernel function, a hybrid of the Green's and Neumann functions, that can be used in conjunction with the representation formula to represent the solution of both the Dirichlet and Neumann problems.

Definition 2. Let ξ, ξ_1, ξ_2 denote distinct points in the interior of Ω.

1. The **Green's function** $G = G(x, \xi)$ of Laplace's equation for the domain Ω is the function:

$$G(x, \xi) = \Phi(x, \xi) - u_\Phi(x),$$

   where u_Φ is the solution of the Dirichlet problem with boundary condition $u_\Phi(x, \xi) = \Phi(x)$ on $\partial\Omega$. This implies that the Green's function satisfies $G(x, \xi) = 0$ on $\partial\Omega$.

2. The **Neumann function** $N = N(x, \xi_1, \xi_2)$ of Laplace's equation for the domain Ω is the function:

$$N(x, \xi_1, \xi_2)) = \Phi(x, \xi_1) - \Phi(x, \xi_2) - v_\Phi(x),$$

   where v_Φ is *a* solution of the Neumann problem with the boundary condition:

$$\frac{\partial v_\Phi(x)}{\partial n} = \frac{\partial \Phi(x, \xi_1)}{\partial n} - \frac{\partial \Phi(x, \xi_2)}{\partial n},$$

for x $\in \partial\Omega$. This implies that the Neumann function satisfies $\partial N/\partial n = 0$. Note that v_Φ is not unique; adding a constant term yields another solution satisfying the desired boundary conditions.

The second singularity ξ_2 in the Neumann function is needed to guarantee that v_Φ satisfies the compatibility condition in Corollary 2.1.2. This can be shown as follows:

$$\begin{aligned}
\int_{\partial\Omega} \frac{\partial v_\Phi}{\partial n} \text{ d(area)} &= \int_{\partial\Omega} \left(\frac{\partial \Phi(x,\xi_1)}{\partial n} - \frac{\partial \Phi(x,\xi_2)}{\partial n} \right) \text{ d(area)} \\
&= \int_{\partial\Omega} \Delta \left(\Phi(x,\xi_1) - \Phi(x,\xi_2) \right) \text{ d(vol)} \\
&= \int_{\partial\Omega} \delta(x,\xi_1) \text{ d(vol)} - \int_{\partial\Omega} \delta(x,\xi_2) \text{ d(vol)} \\
&= 1 - 1 = 0
\end{aligned}$$

Next, we will define the kernel function. First, we need to lay out some simplifying conditions that make this definition possible. Recall that for any choice of interior points ξ_1, ξ_2, the Neumann function is given by:

$$N(x,\xi_1,\xi_2)) = \Phi(x,\xi_1) - \Phi(x,\xi_2) - v_\Phi(x),$$

where the harmonic function v_Φ is unique up to an additive constant. For the sake of simplicity, we will make the assumption that the origin lies within Ω, and we will take $\xi_2 = 0$. We will restrict our consideration to particular solutions v_Φ that are zero at the origin, i.e. $v_\Phi(0) = 0$, by introducing an implicit additive constant term.

Definition 3. With the above conditions imposed, the harmonic **kernel function** is given by:

$$K(x,\xi) = N(x;\xi,0) - G(x,\xi) + G(x,0) - \kappa(\xi),$$

where $\kappa(\xi)$ is an additive constant adjusted so that $K(0,\xi) = 0$.

Lemma 2.1.3. *The harmonic kernel function satisfies*

1. $K(x,\xi) = N(x,\xi,0) - \kappa(\xi)$
2. $\frac{\partial K(x,\xi)}{\partial n} = -\frac{\partial G(x,\xi)}{\partial n}$

for $x \in \partial\Omega$.

Proof. The equality in (1) follows from the definition of the Green's function, since they satisfy

$$G(x,\xi) = 0 = G(x,0),$$

for $x \in \partial\Omega$. We recall from the definition that the Neumann function satisfies

$$\partial N/\partial n = 0.$$

When taking the normal derivative of $K(x,\xi)$, the additive constant $\kappa(\xi)$ vanishes and gives us the equality in (2). □

Theorem 2.7. *Given the existence of the harmonic Green's function in* Ω*, the solution to the Dirichlet problems with the boundary condition* $u = f$ *on* $\partial\Omega$*, is as follows:*

$$u(\xi) = -\int_{\partial\Omega} u(x)\frac{\partial G(x,\xi)}{\partial n}\ \mathrm{d(area)} = -\int_{\partial\Omega} f(x)\frac{\partial G(x,\xi)}{\partial n}\ \mathrm{d(area)}$$

Proof. Suppose we want to solve $\Delta u = 0$ subject to the condition that $u(x) = f(x)$ on the boundary of Ω. We work backwards from the representation formula (Theorem 2.1.3). Any solution would necessarily satisfy the following equation:

$$u(\xi) = \int_{\partial\Omega}\left[-u(x)\frac{\partial\Phi(x,\xi)}{\partial n} + \Phi(x,\xi)\frac{\partial u}{\partial n}\right]\ \mathrm{d(area)}$$

Recall that the Green's function, $G(x,\xi)$, is defined by $G(x,\xi) = \Phi(x,\xi) - u_\Phi(x)$, where u_Φ is harmonic and is equal to $\Phi(x,\xi)$ on the boundary of Ω. Thus, using Green's second identity on the harmonic functions u, u_Φ yields:

$$0 = \int_{\partial\Omega} \left(u\frac{\partial u_\Phi}{\partial n} - u_\Phi \frac{\partial u}{\partial n} \right) \text{ d(area)}.$$

Adding this equation to the previous equation gives us:

$$\begin{aligned} u(\xi) &= \int_{\partial\Omega} \left[-u(x) \left(\frac{\partial \Phi(x,\xi)}{\partial n} + \frac{\partial u_\Phi(x)}{\partial n} \right) - \left(\Phi(x,\xi) - u_\Phi(x) \right) \frac{\partial u}{\partial n} \right] \text{ d(area)} \\ &= \int_{\partial\Omega} \left[-u(x)\frac{\partial G(x,\xi)}{\partial n} + G(x,\xi)\frac{\partial u}{\partial n} \right] \text{ d(area)} \\ &= -\int_{\partial\Omega} u(x)\frac{\partial G(x,\xi)}{\partial n} \text{ d(area)} \end{aligned}$$

In the last step, we used the fact that $G(x,\xi)$ vanishes on $\partial\Omega$.

□

Theorem 2.8. *Given the existence of the harmonic Neumann function in Ω, the solution to the Neumann problem with the boundary condition $\partial u/\partial n = g$, is as follows:*

$$v(\xi) = \int_{\partial\Omega} N(x,\xi)\frac{\partial u(x)}{\partial n} \text{ d(area)} = \int_{\partial\Omega} N(x,\xi) g(x) \text{ d(area)},$$

Proof. We use a similar, but slightly more complicated argument than for the Dirichlet problem. Recall that the harmonic Neumann function is defined using two singular points ξ_1, ξ_2. We have assumed that our domain contains the origin, and that $\xi_2 = 0$. Furthermore, since solutions are only unique up to additive constant, we may only consider solutions satisfying: $v(0) = 0 = v_\Phi(0)$. Using the representation formula at

points ξ and 0 yields the equations:

$$v(\xi) = \int_{\partial\Omega} \left[-v(x) \frac{\partial \Phi(x,\xi)}{\partial n} + \Phi(x,\xi) \frac{\partial v}{\partial n} \right] \text{ d(area)}$$

$$0 = v(0) = \int_{\partial\Omega} \left[-v(x) \frac{\partial \Phi(x,0)}{\partial n} + \Phi(x,0) \frac{\partial v}{\partial n} \right] \text{ d(area)},$$

subtracting these two equations gives us:

$$v(\xi) = \int_{\partial\Omega} \left[-v(x) \left(\frac{\partial \Phi(x,\xi)}{\partial n} - \frac{\partial \Phi(x,0)}{\partial n} \right) + (\Phi(x,\xi) - \Phi(x,0)) \frac{\partial v}{\partial n} \right] \text{ d(area)}$$

Again, we use Green's second identity on the harmonic functions v, v_Φ:

$$0 = \int_{\partial\Omega} \left(v \frac{\partial v_\Phi}{\partial n} - v_\Phi \frac{\partial v}{\partial n} \right) \text{ d(area)}.$$

Subtracting the earlier equation for $v(\xi)$ from this equation yields:

$$\begin{aligned}
v(\xi) &= \int_{\partial\Omega} \left[-v(x) \left(\frac{\partial \Phi(x,\xi)}{\partial n} - \frac{\partial \Phi(x,0)}{\partial n} \right) + (\Phi(x,\xi) - \Phi(x,0)) \frac{\partial v}{\partial n} \right] \text{ d(area)} \\
&= \int_{\partial\Omega} \Bigg[- v(x) \left(\frac{\partial \Phi(x,\xi)}{\partial n} - \frac{\partial \Phi(x,0)}{\partial n} - \frac{\partial v_\Phi(x)}{\partial n} \right) \\
&\qquad\qquad + (\Phi(x,\xi) - \Phi(x,0) - v_\Phi(x)) \frac{\partial v}{\partial n} \Bigg] \text{ d(area)} \\
&= \int_{\partial\Omega} \left[-v(x) \frac{\partial N(x,\xi,0)}{\partial n} + N(x,\xi,0) \frac{\partial v}{\partial n} \right] \text{ d(area)} \\
&= \int_{\partial\Omega} N(x,\xi,0) \frac{\partial v}{\partial n} \text{ d(area)}.
\end{aligned}$$

Note, that we have used the fact that $\partial N(x,\xi,0)/\partial n = 0$ on $\partial\Omega$. □

Corollary 2.1.4. *Given the existence of the harmonic kernel functions in Ω, the solutions to the Dirichlet and Neumann problems are respectively:*

1.

$$u(\xi) = \int_{\partial\Omega} u(x) \frac{\partial K(x,\xi)}{\partial n} \text{ d(area)} = \int_{\partial\Omega} f(x) \frac{\partial K(x,\xi)}{\partial n} \text{ d(area)},$$

2.

$$v(\xi) = \int_{\partial\Omega} K(x,\xi)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} = \int_{\partial\Omega} K(x,\xi)g(x)\ \mathrm{d(area)}$$

where the boundary condition for the Dirichlet problem is $u = f$ on $\partial\Omega$, and the boundary condition for the Neumann problem is that $\partial u/\partial n = g$.

Proof. Using Lemma 2.1.3, and substituting the kernel function for the Green's function in the statement of Theorem 2.7, yields equation (1). For equation (2), we start with the result of Theorem 2.8:

$$\begin{aligned}
v(\xi) &= \int_{\partial\Omega} N(x,\xi,0)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} \\
&= \int_{\partial\Omega} N(x,\xi,0)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} + 0 \\
&= \int_{\partial\Omega} N(x,\xi,0)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} + \kappa(\xi)\int_{\partial\Omega}\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} \\
&= \int_{\partial\Omega} \left(N(x,\xi,0) - \kappa(\xi)\right)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)} \\
&= \int_{\partial\Omega} K(x,\xi)\frac{\partial u(x)}{\partial n}\ \mathrm{d(area)}.
\end{aligned}$$

Note that in the third step we used the compatibility condition in Corollary 2.1.2. □

2.2 The Biot-Savart Law

In our discussion of Hodge decomposition in $\mathbb{R}^3$, we will make use of an important result from electrodynamics, namely the *Biot-Savart law* for magnetic fields. We treat our compact, connected set Ω as if it were a volume of conductive material with a steady current distribution represented by the vector field $\mathbf{J}(x)$. In electrodynamics, the direction of $\mathbf{J}$ at any point represents the direction of the current (moving electrical charges) and the magnitude is the charge per unit time passing that point in

the prescribed direction. For our purposes, this interpretation is not crucial since we will consider vector fields $\mathbf{J}$ on Ω that do not represent realistic current distributions.

Suppose that Ω is some volume of conductive material containing a current distribution $\mathbf{J}(x)$, then the magnetic field generated by this current distribution is given by the **Biot-Savart Law**:

$$\mathbf{B}(y) = \frac{\mu_0}{4\pi} \int_\Omega \mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \, \mathrm{d}(\mathrm{vol_x})$$

where x denotes the point within Ω whose infinitesimal contribution to $\mathbf{B}$ we wish to consider, and y denotes the point at which we want to calculate $\mathbf{B}$. Integrating over all such x in Ω we get the full contribution to the magnetic field given by the charge distribution $\mathbf{J}$ in Ω.

Now, let us suppose that we are given any arbitrary smooth vector field $\mathbf{J}$ on Ω (that may or may not represent a physically realistic current distribution). We want to compute the curl and divergence of the resulting field $\mathbf{B}$.

Claim 1. For any vector field $\mathbf{J}$ on Ω, the resulting Biot-Savart field $\mathbf{B}$ has zero divergence.

$$\nabla \cdot \mathbf{B} = 0.$$

Proof. Let $\mathbf{J}$ be a vector field defined on the domain Ω. Then the Biot-Savart field at any point y in $\mathbb{R}^3$ is given by:

$$\mathbf{B}(y) = \frac{\mu_0}{4\pi} \int_\Omega \mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \, \mathrm{d}(\mathrm{vol_x}).$$

We take the divergence of this equation, the subscripts are included so there is no

confusion about the variables x and y:

$$\nabla \cdot \mathbf{B}(y) = \frac{\mu_0}{4\pi} \int_\Omega \nabla_y \cdot \left(\mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \right) \, \mathrm{d}(\mathrm{vol_x}).$$

Recall from vector calculus that the divergence of a cross product satisfies the following identity:

$$\nabla \cdot (\mathbf{A} \times \mathbf{B}) = \mathbf{B} \cdot (\nabla \times \mathbf{A}) - \mathbf{A} \cdot (\nabla \times \mathbf{B})$$

Therefore, the expression in the integrand of the above equation can be expanded using this product rule:

$$\nabla_y \cdot \left(\mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \right) = \frac{y-x}{|y-x|^3} \cdot (\nabla_y \times \mathbf{J}(x)) - \mathbf{J}(x) \cdot \left(\nabla_y \times \frac{y-x}{|y-x|^3} \right)$$

Since $\mathbf{J}(x)$ does not depend on y we have $\nabla_y \times \mathbf{J}(x) = 0$. The curl of an inverse square field is zero so we can conclude that

$$\nabla_y \times \frac{y-x}{|y-x|^3} = 0,$$

and therefore

$$\nabla \cdot \mathbf{B}(y) = \frac{\mu_0}{4\pi} \int_\Omega \left[\frac{y-x}{|y-x|^3} \cdot (\nabla_y \times \mathbf{J}(x)) - \mathbf{J}(x) \cdot \left(\nabla_y \times \frac{y-x}{|y-x|^3} \right) \right] \, \mathrm{d}(\mathrm{vol_x}) = 0$$

□

Claim 2. Given a vector field $\mathbf{J}$ on Ω, the curl of the Biot-Savart field is given by the following equation:

$$\nabla \times \mathbf{B}(y) = \begin{cases} \mu_0 \mathbf{J}(y) & \text{if } y \in \Omega \\ 0 & \text{if } y \notin \Omega \end{cases} + \frac{\mu_0}{4\pi} \nabla_y \int_\Omega \frac{\nabla_x \cdot \mathbf{J}(x)}{|y-x|} \, \mathrm{d}(\mathrm{vol_x}) - \frac{\mu_0}{4\pi} \nabla_y \int_{\partial\Omega} \frac{\mathbf{J}(x) \cdot n}{|y-x|} \, \mathrm{d}(\mathrm{area_x})$$

Proof. Let $\mathbf{J}$ be a vector field defined on the domain Ω. Then the curl of the Biot-Savart field at any point y in $\mathbb{R}^3$ is given by:

$$\nabla \times \mathbf{B}(y) = \frac{\mu_0}{4\pi} \int_\Omega \nabla_y \times \left(\mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \right) \, \mathrm{d}(\mathrm{vol_x}).$$

Again, we recall a product rule from vector calculus:

$$\nabla \times (\mathbf{A} \times \mathbf{B}) = (\mathbf{B} \cdot \nabla)\mathbf{A} - (\mathbf{A} \cdot \nabla)\mathbf{B} + \mathbf{A}(\nabla \cdot \mathbf{B}) - \mathbf{B}(\nabla \cdot \mathbf{A})$$

Since $\mathbf{J}(x)$ has no y dependence, we have that:

$$\left(\frac{y-x}{|y-x|^3} \cdot \nabla_y \right) \mathbf{J}(x) = 0 = \frac{y-x}{|y-x|^3} (\nabla_y \cdot \mathbf{J}(x)),$$

therefore our equations reduce to the following:

$$\begin{aligned}
\nabla \times \mathbf{B}(y) &= \frac{\mu_0}{4\pi} \int_\Omega \nabla_y \times \left(\mathbf{J}(x) \times \frac{y-x}{|y-x|^3} \right) \, \mathrm{d}(\mathrm{vol_x}) \\
&= \frac{\mu_0}{4\pi} \int_\Omega \left[\mathbf{J}(x) \left(\nabla_y \cdot \frac{y-x}{|y-x|^3} \right) - (\mathbf{J}(x) \cdot \nabla_y) \frac{y-x}{|y-x|^3} \right] \, \mathrm{d}(\mathrm{vol_x}) \\
&= \frac{\mu_0}{4\pi} \int_\Omega \mathbf{J}(x) \left(\nabla_y \cdot \frac{y-x}{|y-x|^3} \right) \, \mathrm{d}(\mathrm{vol_x}) - \frac{\mu_0}{4\pi} \int_\Omega (\mathbf{J}(x) \cdot \nabla_y) \frac{y-x}{|y-x|^3} \, \mathrm{d}(\mathrm{vol_x})
\end{aligned}$$

The first integral simplifies as follows:

$$\frac{\mu_0}{4\pi} \int_\Omega \mathbf{J}(x) \left(\nabla_y \cdot \frac{y-x}{|y-x|^3} \right) \, \mathrm{d}(\mathrm{vol_x}) = \frac{\mu_0}{4\pi} \int_\Omega \mathbf{J}(x) 4\pi \delta^3(x-y) \, \mathrm{d}(\mathrm{vol_x}) = \mu_0 \mathbf{J}(y). \quad (2.1)$$

First, we note that

$$\nabla_y \left(\frac{y-x}{|y-x|^3} \right) = -\nabla_x \left(\frac{y-x}{|y-x|^3} \right),$$

so that if we change ∇_y to ∇_x in (2.1) we get:

$$-\frac{\mu_0}{4\pi} \int_\Omega (\mathbf{J}(x) \cdot \nabla_y) \frac{y-x}{|y-x|^3} \, \mathrm{d}(\mathrm{vol_x}) = \frac{\mu_0}{4\pi} \int_\Omega (\mathbf{J}(x) \cdot \nabla_x) \frac{y-x}{|y-x|^3} \, \mathrm{d}(\mathrm{vol_x}).$$

Secondly, to simplify the integral, we need another product rule:

$$\nabla \cdot (f\mathbf{A}) = f(\nabla \cdot \mathbf{A}) + (\mathbf{A} \cdot \nabla)f \Rightarrow (\mathbf{A} \cdot \nabla)f = \nabla \cdot (f\mathbf{A}) - f(\nabla \cdot \mathbf{A})$$

Applying this product rule and using the divergence theorem gives us:

$$\begin{aligned}
\frac{\mu_0}{4\pi}\int_\Omega (\mathbf{J}(x)\cdot\nabla_x)\frac{y-x}{|y-x|^3}\ \mathrm{d(vol_x)} &= \frac{\mu_0}{4\pi}\int_\Omega \left[\nabla_x \cdot \left(\frac{y-x}{|y-x|^3}\mathbf{J}(x)\right) - \frac{y-x}{|y-x|^3}\left(\nabla_x \cdot \mathbf{J}(x)\right)\right]\ \mathrm{d(vol_x)} \\
&= \frac{\mu_0}{4\pi}\int_\Omega \nabla_x \cdot \left(\frac{y-x}{|y-x|^3}\mathbf{J}(x)\right)\ \mathrm{d(vol_x)} - \frac{\mu_0}{4\pi}\int_\Omega \frac{y-x}{|y-x|^3}\left(\nabla_x \cdot \mathbf{J}(x)\right)\ \mathrm{d(vol_x)} \\
&= \frac{\mu_0}{4\pi}\int_{\partial\Omega} \left(\frac{y-x}{|y-x|^3}\mathbf{J}(x)\cdot n\right)\ \mathrm{d(area_x)} - \frac{\mu_0}{4\pi}\int_\Omega \frac{y-x}{|y-x|^3}\left(\nabla_x \cdot \mathbf{J}(x)\right)\ \mathrm{d(vol_x)}
\end{aligned}$$

Finally, we use the following identity:

$$\nabla_y\left(\frac{1}{|y-x|}\right) = \frac{y-x}{|y-x|^3}$$

Therefore,

$$\int_\Omega (\mathbf{J}(x)\cdot\nabla_x)\frac{y-x}{|y-x|^3}\ \mathrm{d(vol_x)} = \nabla_y \int_\Omega \frac{\nabla_x \cdot \mathbf{J}(x)}{|y-x|}\ \mathrm{d(vol_x)} - \nabla_y \int_{\partial\Omega} \frac{\mathbf{J}(x)\cdot n}{|y-x|}\ \mathrm{d(area_x)}.$$

Substituting this expression back into the earlier equation shows the result. □

2.3 Hodge Decomposition in $\mathbb{R}^3$

In this section, we will devote our efforts to proving the Hodge decomposition for vector fields in $\mathbb{R}^3$. We will borrow heavily from the preceding section on background material. The proof will be broken into several theorems which all together make up the Hodge decomposition theorem. Before stating the Hodge theorem for vector fields, we clarify some of the notation used in this section.

We let Ω denote a compact domain in $\mathbb{R}^3$ with smooth boundary $\partial\Omega$. We let $\mathrm{VF}(\Omega)$ denote the inner-product space of smooth vector fields defined on Ω, with the L^2 inner product given by:

$$\langle V, W\rangle = \int_\Omega V \cdot W \ \mathrm{d(vol)}.$$

Theorem 2.9. ***(Hodge Decomposition Theorem)*** *The space* $\mathrm{VF}(\Omega)$ *is the direct sum of five mutually orthogonal subspaces:*

$$\mathrm{VF}(\Omega) = \ \mathrm{FK} \oplus \ \mathrm{HK} \oplus \ \mathrm{CG} \oplus \ \mathrm{HG} \oplus \ \mathrm{GG}$$

where

$$\begin{aligned}
\mathrm{FK} = \text{fluxless knots} &= \{\nabla \cdot V = 0, V \cdot n = 0, \text{all interior fluxes are } 0\} \\
\mathrm{HK} = \text{harmonic knots} &= \{\nabla \cdot V = 0, \nabla \times V = 0, V \cdot n = 0, \} \\
\mathrm{CG} = \text{curly gradients} &= \{V = \nabla\phi, \nabla \cdot V = 0, \text{all boundary fluxes are } 0\} \\
\mathrm{HG} = \text{harmonic gradients} &= \{V = \nabla\phi, \nabla \cdot V = 0, \phi \text{ is locally constant on } \partial\Omega\} \\
\mathrm{GG} = \text{grounded gradients} &= \{V = \nabla\phi, \phi \text{ vanishes on } \partial\Omega\}.
\end{aligned}$$

with

$$\begin{aligned}
\ker\ \mathit{curl} &= \quad \mathrm{HK} \oplus \ \mathrm{CG} \oplus \ \mathrm{HG} \oplus \ \mathrm{GG} \\
\mathrm{im}\ \mathit{grad} &= \quad \mathrm{CG} \oplus \ \mathrm{HG} \oplus \ \mathrm{GG} \\
\mathrm{im}\ \mathit{curl} &= \ \mathrm{FK} \oplus \ \mathrm{HK} \oplus \ \mathrm{CG} \\
\ker\ \mathit{div} &= \ \mathrm{FK} \oplus \ \mathrm{HK} \oplus \ \mathrm{CG} \oplus \ \mathrm{HG},
\end{aligned}$$

Furthermore,

$$\mathrm{HK} \simeq H_1(\Omega; \mathbb{R}) \simeq H_2(\Omega, \partial\Omega; \mathbb{R}) \simeq \mathbb{R}^{\text{genus of } \partial\Omega}$$

$$\mathrm{HG} \simeq H_2(\Omega; \mathbb{R}) \simeq H_1(\Omega, \partial\Omega; \mathbb{R}) \simeq \mathbb{R}^{(\#\text{ components of } \partial\Omega) - (\#\text{ components of } \Omega)}$$

The Hodge theorem says that the space $\mathrm{VF}(\Omega)$ can be decomposed into 5 subspaces, we prove this in multiple steps. First, we show that $\mathrm{VF}(\Omega)$ can be decomposed into two larger, but simpler subspaces, the knots and the gradients. These two spaces are defined as follows:

$$\begin{aligned} \mathrm{K} = \ & knots & &= \{V \in \ \mathrm{VF}(\Omega) : \nabla \cdot V = 0, V \cdot n = 0\} \\ \mathrm{G} = \ & gradients & &= \{V \in \ \mathrm{VF}(\Omega) : V = \nabla \phi\} \end{aligned}$$

Proposition 2.3.1. *The space* $\mathrm{VF}(\Omega)$ *is the direct sums of the space of knots and the space of gradients:*

$$\mathrm{VF}(\Omega) = \ \mathrm{K} \oplus \ \mathrm{G}$$

Proof. Let V be a smooth vector field on Ω, then the divergence of V defines a smooth function $f = \nabla \cdot V$. Likewise, since Ω has a smooth boundary, we may define a smooth function $g = V \cdot n$ on $\partial\Omega$. Next, we apply the divergence theorem on each component Ω_i of Ω:

$$\int_{\Omega_i} f \ \mathrm{d(vol)} = \int_{\Omega_i} \nabla \cdot V \ \mathrm{d(vol)} = \int_{\partial\Omega_i} V \cdot n \ \mathrm{d(area)} = \int_{\partial\Omega_i} g \ \mathrm{d(area)}$$

We let ϕ be a solution of the Neumann problem $\Delta\phi = f$ on Ω with boundary condition $\partial\phi/\partial n = g$ on $\partial\Omega$. So, given this ϕ we define two vector fields V_1, V_2 with $V_2 = \nabla\phi$ and $V_1 = V - V_2$. It is clear that $V_2 \in G$, so we must show that $V_1 \in K$. Note that on $\partial\Omega$, we have:

$$V_2 \cdot n = \nabla\phi \cdot n = \frac{\partial\phi}{\partial n} = g = V \cdot n$$

Thus, $V_1 \cdot n = (V - V_2) \cdot n = 0$ on $\partial\Omega$. Likewise, on Ω we have

$$\nabla \cdot V_2 = \nabla \cdot \nabla\phi = \Delta\phi = f = \nabla \cdot V$$

and therefore $\nabla \cdot V_1 = 0$, which show that $V_1 \in K$. We have shown that $VF(\Omega)$ is the sum of the subspaces K and G, that is, $VF(\Omega) = K + G$, so next we must show that this is an orthogonal direct sum.

Let V_1 and V_2 be defined as above, then the L^2 inner product of V_1 and V_2 is defined as follows:

$$\langle V_1, V_2 \rangle = \int_\Omega V_1 \cdot V_2 \ \mathrm{d(vol)} = \int_\Omega V_1 \cdot \nabla\phi \ \mathrm{d(vol)}$$

An application of the familiar product rule, $\nabla \cdot (\phi V_1) = (\nabla\phi) \cdot V_1 + \phi(\nabla \cdot V_1)$ gives us:

$$\begin{aligned}
\langle V_1, V_2 \rangle &= \int_\Omega V_1 \cdot \nabla\phi \ \mathrm{d(vol)} \\
&= \int_\Omega (\nabla \cdot (\phi V_1) - \phi(\nabla \cdot V_1) \ \mathrm{d(vol)} \\
&= \int_\Omega \nabla \cdot (\phi V_1) \ \mathrm{d(vol)} \\
&= \int_{\partial\Omega} \phi V_1 \cdot n \ \mathrm{d(area)} = 0.
\end{aligned}$$

Note that we have used the fact that $V_1 \cdot n = 0$ and $\nabla \cdot V_1 = 0$ in our argument. This shows that K and G are orthogonal subspaces and therefore $VF(\Omega) = K \oplus G$. □

Theorem 2.10. *The subspace G of $VF(\Omega)$ is the direct sum of two orthogonal subspaces:*

$$G = DFG \oplus GG.$$

That is,

$$\textit{Gradients} = \textit{Divergence-Free Gradients} \oplus \textit{Grounded Gradients.}$$

Proof. Let V be a vector field such that $V = \nabla\phi$ for some smooth function ϕ defined on Ω. Let ϕ_1 be the solution of the Laplace equation $\nabla\phi_1 = 0$ on Ω with Dirichlet

boundary condition $\phi_1 = \phi$ on $\partial\Omega$, and let be given by $\phi_2 = \phi - \phi_1$. Define $V_1 = \nabla\phi_1$ and $V_2 = \nabla\phi_2$. We see that

$$V = \nabla\phi = \nabla(\phi_1 + \phi_2) = \nabla\phi_1 + \nabla\phi_2 = V_1 + V_2$$

By construction we have that $\nabla \cdot V_1 = \nabla \cdot \nabla\phi = 0$, so $V_1 \in DFG$. Likewise, $\phi_2|_{\partial\Omega} = (\phi - \phi_1)|_{\partial\Omega} = 0$, so $V_2 \in GG$. This shows that $G = DFG + GG$, now we must show that this is a direct sum. Recall that $V_1 \in DFG$ and so $\nabla \cdot V_1 = 0$, and since $V_2 \in GG$ we have $V_2 = \nabla\phi_2$ with $\phi_2 = 0$ on $\partial\Omega$.

$$\begin{aligned}
\langle V_1, V_2\rangle &= \int_\Omega V_1 \cdot V_2 \ \mathrm{d(vol)} = \int_\Omega \nabla\phi_1 \cdot \nabla\phi_2 \ \mathrm{d(vol)} \\
&= \int_\Omega (\nabla \cdot (\phi_2 \nabla\phi_1) - \phi_2 \Delta\phi_1) \ \mathrm{d(vol)} = \int_\Omega \nabla \cdot (\phi_2 \nabla\phi_1) \ \mathrm{d(vol)} \\
&= \int_{\partial\Omega} (\phi_2 \nabla\phi_1) \cdot n \ \mathrm{d(area)} = 0.
\end{aligned}$$

Note that we have used the facts that $\Delta\phi_1 = 0$ and $\phi_2 = 0$ on $\partial\Omega$. This shows that DFG and GG are orthogonal subspaces and therefore $G = DFG \oplus GG$. $\square$

Chapter 3: Hodge Theory on Graphs and Simplicial Complexes

In this section, we outline the mathematics behind *HodgeRank*. We begin by describing simple graphs and special functions called *flows* by which we will represent our ranking data. We show how simple graphs and flows generalize to simplicial complexes and chains. Finally, we prove the Hodge Decomposition theorem for chains.

3.1 Graphs

In this subsection we review some basic definitions and terminology from graph theory. We will occasionally adopt nonstandard notation, in order to more closely follow the original paper [14], and so that the related concepts in simplicial topology and graph theory share an appropriately related notation. Our notation does not change the meaning or validity of the concepts and results.

Throughout this section, V will denote a nonempty finite set, typically representing the vertex set of a graph or simplicial complex. Following [14], we will use the notation $\binom{V}{k}$ to denote the set of k-element of subsets of V, and the notation V^k to denote the set of k-tuples of elements of V.

3.1.1 Basic Definitions

A **(simple) graph** G is an ordered pair (V, E) where V is a nonempty finite set and $E \subset \binom{V}{2}$, that is, E is a collection of two-element subsets of V. The set V is called the **vertex set** of the graph G, the individual elements of V are called **vertices**. If the vertex set of G has n elements we call G a **graph on n vertices**. The set E

is called the **edge set** of G, and the individual elements of E are called **edges**. If $\{i, j\} \in E$, we say that the vertices i and j are **adjacent**, and that $\{i, j\}$ is an edge between i and j.

Since edges are defined to be sets and not ordered pairs, we do not make the distinction between the edges $\{i, j\}$ and $\{j, i\}$. The set definition also precludes the possibility of *loops*, that is, no vertex can be adjacent to itself. Furthermore, since $E \subseteq \binom{V}{2}$, so we do not allow multiples edges, that is $\{i, j\}$ cannot appear twice in E. Graphs that do not contain multiple edges or loops are called *simple*. In Figure 3.1 we see a typical depiction of a graph, here the vertices are the labeled nodes and the edges are lines connecting two nodes. The actual graph represented is $G = (V, E)$ with the vertex set $V = \{1, 2, 3, 4\}$ and the edge set $E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{3, 4\}\}$.

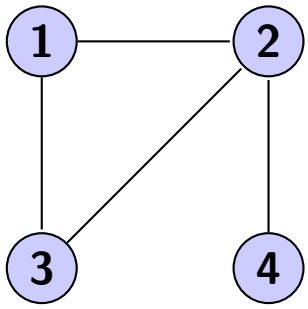

Figure 3.1: A typical depiction of a simple graph.

We say that a graph G is a **complete graph** if for each distinct pair of vertices $i, j \in V$, the edge $\{i, j\}$ is in E. We say that the graph $\bar{G} = (\bar{V}, \bar{E})$ is a **subgraph** of the graph G if $\bar{V} \subseteq V$ and $\bar{E} \subseteq E$. A **clique** of G is a nonempty, complete subgraph of G. If K is a clique of G whose vertex set contains k vertices, we call K a **k-clique** of G. Technically, the cliques of G are themselves graphs, and so should be formally defined as an ordered pair (V, E), but for simplicity we will often identify a clique of

G with its vertex set V, since it will be implicitly understood that the edge set will be $\binom{V}{2}$.

Suppose we have a graph $G = (V, E)$. You can verify that the collection of 1-cliques of G is just the vertex set V, and the collection of 2-cliques is the edge set E. We will denote the collection of 3-cliques by the set T:

$$T = \{\{i, j, k\} : \{i, j\}, \{j, k\}, \{k, i\} \in E\}. \tag{3.1}$$

The elements of T form triangles in the visual representation of G, hence the name T. In our example graph depicted above, the triple $\{1, 2, 3\}$ is the only three clique.

3.1.2 Paths and Cycles

Let $G = (V, E)$ be a graph on n vertices, ($|V| = n$). A **trail** in G, or more specifically, a $\mathbf{v_1 - v_k}$ **trail** in G is defined to be a finite tuple of vertices $(v_1, \cdots, v_k)$ such that for each $i < k$, $\{v_i, v_{i+1}\} \in E$. Instead of using tuple notation, we will write a trail as $v_1 - v_2 - \cdots - v_k$. We may think of traveling from vertex v_1 to vertex v_k along the edges $\{v_i, v_{i+1}\}$ for $i = 1, \cdots, k-1$, thus forming a "trail" from v_1 to v_k. If $v_1 - v_2 - \cdots - v_{k-1} - v_k$ is a trail in G such that no vertices are repeated, we call $v_1 - v_2 - \cdots - v_{k-1} - v_k$ a **path** or a more specifically a $\mathbf{v_1 - v_k}$ **path**. A trail with $v_1 = v_k$ is called a **circuit**, and a circuit with no repeated vertices (except for $v_1 = v_k$) is called a **cycle**.

The **length** of a trail, path, circuit, or cycle given by $v_1 - v_2 - \cdots - v_{k-1} - v_k$ is $k-1$, i.e. the number of edges traversed from v_1 to v_k. We will be most interested in cycles, although we will prove some simple results regarding circuits, paths and trails in the next section. We note that the smallest possible length of a non-trivial cycle

in G is 3. A 1-cycle consists of a single vertex, a 2-cycle consists of only two vertices and hence the same edge must be traversed twice in order to return to the starting vertex. The set of triples $\{i, j, k\} \subset V$ around which we can form 3-cycles of the form $i - j - k - i$ is just the set T of 3-cliques in G.

3.1.3 Flows on Graphs

Definition 4 (**Edge Flow**). An **edge flow** on G is a function $X : V \times V \to \mathbb{R}$ that satisfies

$$\begin{cases} X(i,j) = -X(j,i) & \text{if } \{i,j\} \in E \\ X(i,j) = 0 & \text{otherwise} \end{cases}.$$

We note X is zero for all pairs that are not adjacent, and in particular $X(i,i) = 0$ since the way we have defined edges does not allow vertices to be self-adjacent. Let $G = (V, E)$ be a graph, and X be an edge flow on G. We can represent X by a skew-symmetric matrix $[X_{ij}]$ with entries given by $X_{ij} = X(i,j)$ (note that $X_{ij} = 0$ if $\{i,j\} \notin E$). If G is a graph on n vertices, and X is a $n \times n$ skew-symmetric matrix, we can determine an edge flow of G by letting $X(i,j) = X_{ij}$. So, the set of edge flows on G is in one-to-one correspondence with the set of $n \times n$ skew-symmetric matrices satisfying:

$$\{X \in \mathbb{R}^{n \times n} : X^T = -X \text{ and } X_{ij} = 0 \text{ if } \{i,j\} \notin E\}$$

Below, we show how a 4×4 skew-symmetric matrix induces an edge flow on our example graph and vice versa:

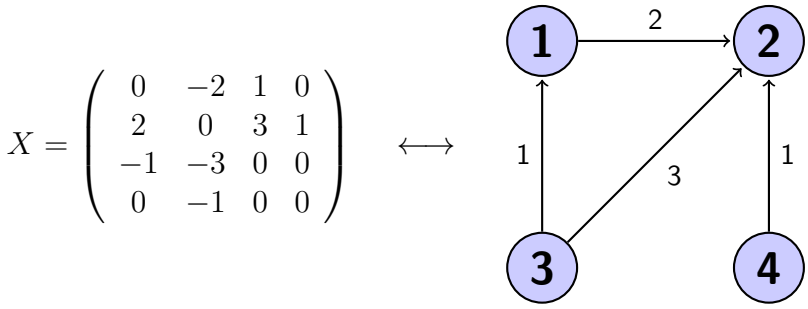

Figure 3.2: A matrix along with its corresponding edge flow.

You may recall from vector calculus that a vector field is called a *conservative field* or a *gradient field* if it is the gradient of a potential function, that is, if the vector field $\vec{F}$ is given by $\vec{F} = \nabla\phi$ for some real-valued function $\phi : \mathbb{R}^n \to \mathbb{R}$. In our graph theoretic setting, a **potential function** s is just a real-valued function on the vertex set V, i.e. $s : V \to \mathbb{R}$. Just as we can identify edge flows with specially structured skew-symmetric matrices, we can identify potential functions on V with vectors in $\mathbb{R}^n$ and write $s(i) = s_i$. Just as in calculus, where a potential function determines a vector field via the gradient operator, a potential function on a vertex set determines an edge flow on a graph via the combinatorial gradient operator, which we define next.

Definition 5. (Combinatorial Gradient) Let G be a graph, and $s \in \mathbb{R}^n$ be a potential function. The **gradient** of s is the the edge flow given by:

$$(\text{grad } s)(i,j) = s_j - s_i \text{ for } \{i,j\} \in E.$$

An edge flow X of the form $X_{ij} = s_j - s_i$ for some potential $s \in \mathbb{R}^n$ is called a **gradient flow**.

The next proposition is a graph theoretic analog of the higher dimensional fundamental theorem of calculus, sometimes called the gradient theorem. Informally,

this theorem states that the integral of a gradient vector field along any path is determined by the value of the potential at the endpoints. Suppose that we have an edge flow X on G and some $v_1 - v_k$ trail in G. Then the flow of X along the trail $v_1 - v_2 - \cdots - v_{k-1} - v_k$, which we informally denote by $X(x_1 - \cdots - x_k)$ is given by

$$X(x_1 - \cdots - x_k) = X(v_1, v_2) + X(v_2, v_3) + \cdots + X(v_{k-2}, v_{k-1}) + X(v_{k-1}, v_k).$$

Proposition 3.1.1. ***(Trail/Path Independence)*** *If X is an gradient flow, i.e. for some $s \in \mathbb{R}^n$, $X_{ij} = s_j - s_i$, then the flow along* ***any*** *v_1-v_k path in G is*

$$s_k - s_i.$$

Proof. Let G be a graph on n vertices, $v_1 - v_2 - \cdots - v_k$ be a trail in G, and X be a gradient flow. Then the flow of X along trail $v_1 - v_2 - \cdots - v_k$ is given by

$$X(v_1, v_2) + X(v_2, v_3) + \cdots + X(v_{k-2}, v_{k-1}) + X(v_{k-1}, v_k).$$

However, since X is a gradient flow, $X(v_i, v_j) = s_{v_j} - s_{v_i}$, thus we get a telescoping sum given by:

$$\begin{aligned} X(v_1, v_2) + X(v_2, v_3) + \cdots + X(v_{k-1}, v_k) = & \\ &= (s_{v_2} - s_{v_1}) + (s_{v_3} - s_{v_2}) + \cdots + (s_{v_k} - s_{v_{k-1}}) \\ &= -s_{v_1} + (s_{v_2} - s_{v_2}) + \cdots + (s_{v_{k-1}} - s_{v_{k-1}}) + s_{v_k} \\ &= s_{v_k} - s_{v_1} \end{aligned}$$

□

It follows immediately from the gradient theorem that the path integral of a gradient field around a closed loop vanishes. An analogous result for flows around circuits follows immediately from path independence:

Corollary 3.1.1. ***(Closed Loops)*** *If X is a gradient flow on G, then the flow around any circuit in G vanishes.*

As you might have suspected, there are many parallels between potential functions and gradients in vector calculus and graph theory beyond the name. You may recall from vector calculus or elementary physics that a potential function (such as gravitational/electrical potential) is only determined uniquely up to a constant. Similarly, in the case of connected graphs, a combinatorial potential is determined uniquely up to an additive constant as we see in the next theorem.

Proposition 3.1.2. ***[Uniqueness of Potentials]*** *Let X be a gradient flow on a connected graph $G = (V, E)$. Then the potential function $s \in \mathbb{R}^n$ such that*

$$X(i,j) = (\text{grad } s)(i,j) = s_j - s_i$$

is uniquely determined up to the addition of a scalar multiple of $[1, \cdots, 1]^T$.

Proof. Let X be a gradient flow on a graph $G = (V, E)$ with $|V| = n$, and let $s = [s_1, \cdots, s_n]^T$ be the potential function such that $X_{ij} = s_j - s_i$ for $\{i, j\} \in E$. Let $c \in \mathbb{R}$ and consider the potential function given by $\bar{s} = [s_1 + c, \cdots, s_n + c]^T$. Then for any i, j we have

$$(\text{grad } \bar{s})(i,j) = \bar{s}_j - \bar{s}_i = (s_j + c) - (s_i + c) = s_j - s_i = (\text{grad } s)(i,j).$$

Next, suppose that s and r are two potential functions whose gradients determine the edge flow X_{ij}. We show that if $r_i = s_i + c$ for some $c \in \mathbb{R}$, then $r_j = s_j + c$ for the same $c \in \mathbb{R}$. Let $i, j \in V$, and let $c = r_i - s_i$, since G is connected, we can find a path between vertices i and j. By the path independence property of gradient flows, since $X = \text{grad } (r)$, the flow along each path is given by $r_i - r_j$. Since $X = \text{grad } (s)$ the flow along this path is also given by $s_i - s_j$. Hence,

$$r_i - r_j = s_i - s_j \implies r_i - s_i = r_j - s_j = c$$

Since G is connected, we may form a similar path between each pair of vertices i and j, hence s and r differ by the addition of $[c, \cdots, c]^T$ □

3.1.4 Curl and Divergence

We can extend the notion of flows to an alternating function defined on the set T of 3-cliques or triangles in G. These *triangular flows* allow us to characterize local inconsistency in graphs. We define a **triangular flow** Φ to be a function $\Phi : V^3 \to \mathbb{R}$ such that for $v_1, v_2, v_3 \in V$ we have

$$\Phi(v_1, v_2, v_3) = \begin{cases} \epsilon_{ijk}\Phi(v_i, v_j, v_k) & \text{if } \{v_1, v_2, v_3\} \in T \\ 0 & \text{if } \{v_1, v_2, v_3\} \notin T \end{cases}.$$

Here ϵ_{ijk} is the Levi-Civita symbol, recall that $\epsilon_{ijk} = 1$ if (ijk) is an even permutation of (123) and $\epsilon_{ijk} = -1$ if (ijk) is an odd permutation of (123). Note that we can store the values of a triangular flow in a skew-symmetric hypermatrix, by identifying Φ_{ijk} with the value $\Phi(i, j, k)$. The **combinatorial curl** or simply **curl** is a linear operator which maps edge flows to triangular flows such that:

Definition 6. (Combinatorial Curl)

$$\text{curl(X)}_{ijk} = \begin{cases} X_{ij} + X_{jk} + X_{ki} & \text{if } \{v_i, v_j, v_k\} \in T \\ 0 & \text{if } \{v_i, v_j, v_k\} \in T \end{cases}$$

In the next lemma, we show that gradient flows have vanishing curl, that is,

$$\text{curl} \ \circ \ \text{grad} \ = 0,$$

just as $\nabla \times \nabla f = 0$ in vector calculus.

Lemma 3.1.1. *Suppose that* $X = \text{grad}\ (s)$ *for some* $s \in \mathbb{R}^n$. *Then,*

$$\text{curl}(X) = 0.$$

Proof. Let $X = \text{grad}\ (s)$ for some $s \in \mathbb{R}^n$, and $\{i, j, k\} \in T$. Then

$$\begin{aligned} \text{curl}(X)_{ijk} &= X_{ij} + X_{jk} + X_{ki} \\ &= (s_i - s_j) + (s_j - s_k) + (s_k - s_i) \\ &= (s_i - s_i) + (s_j - s_j) + (s_k - s_k) = 0. \end{aligned}$$

□

A curl-free edge flow is one for which $\text{curl}(X) = 0$. Intuitively, a curl-free flow has no net flow around a triangular cycle. As we will show later, it is not necessarily the case that $\text{curl}(X) = 0$ implies that $X = \text{grad}\ (s)$ for some $s \in \mathbb{R}^n$.

The final operator we introduce is the combinatorial **divergence** operator, which as might be expected, takes edge flows to functions on vertices. As with curl and gradient there are many similarities between the combinatorial divergence operator and its vector calculus analog.

Definition 7. (Divergence) Let X be an edge flows on a graph $G = (V, E)$, then the divergence is a potential function $\text{div} : V \to \mathbb{R}$ whose values on the vertex i is given by:

$$\text{div}\, X(i) = \sum_{j \text{s.t.} \{i,j\} \in E} X_{ij}.$$

Since $X_{ij} = 0$ if $\{i, j\} \notin E$, the divergence amounts to summing the entries of the i^{th} row of X. Graphically, this amounts to determining the net flow into a the vertex i.

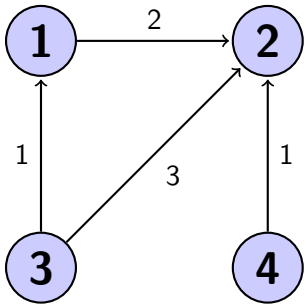

$$\text{div}\, X = \begin{bmatrix} -1 \\ 6 \\ -4 \\ -1 \end{bmatrix}$$

Figure 3.3: An edge flow along with its divergence.

Let X be an edge flow on the graph G. If X has the property that $\operatorname{div} X = 0$ on each vertex, we call the graph G **divergence-free**. Intuitively, this means that the net flow through any vertex is zero. As in vector calculus, a vertex i on which $(\operatorname{div} X)(i) < 0$ is called a **source**, and a vertex j on which $(\operatorname{div} X)(j) > 0$ is called a **sink**. Intuitively, a source has more out-flow than in-flow, and so can be thought of as creating flow out of nothing, hence a "source" of flow. Similarly, a sink has more out-flow than in-flow, reminiscent of a drain on a sink.

Definition 8. (Harmonic Flows) Let X be an edge flow on the graph G. The flow X is a **harmonic flow** if it both curl-free and divergence free.

In the next section, we provide examples of flows that are gradients, curl-free, divergence-free and harmonic, but we should note that more so than the other types, harmonic flows are tied up in the underlying geometry of our graph. Whether or not harmonic flows even exist on our graph depends on homology of the underlying graph complex. We will prove this in the next section, when we introduce simplicial homology and cohomology which generalizes flows on simple graphs.

3.1.5 Examples

Let $G = (V, E)$ be the graph with vertex set $V = \{1, 2, 3, 4, 5, 6\}$ and edge set $E = \{\{1, 2\}, \{1, 3\}, \{1, 5\}, \{1, 6\}\{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}$. Then G has the following planar representation, which we will adopt henceforth:

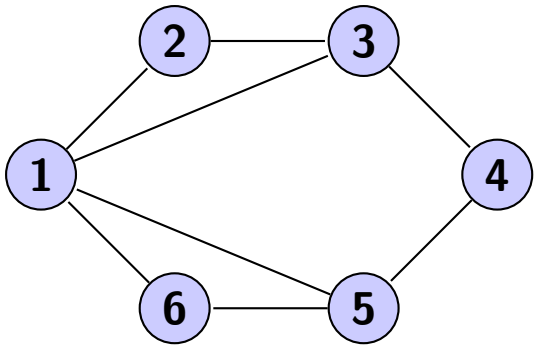

Figure 3.4: An example graph G.

In this section, we will consider several examples of flows on the same underlying graph G. We will try to pick examples that are mutually exclusive, for instance, our example of a curl-free flow will not also be divergence-free, clearly our harmonic flow will be both curl-free and divergence-free by definition.

1. **(Gradient Flow)** Recall that a gradient flow X is an edge flow for which there exists a potential $s \in \mathbb{R}^n$ such that

$$X_{ij} = \begin{cases} s_j - s_i & \text{for } \{i,j\} \in E \\ 0 & \text{otherwise} \end{cases}$$

   We call X the gradient of s and write $X = \text{grad}\ (s)$. Below is an example of a gradient flow on G along with a corresponding potential function s.

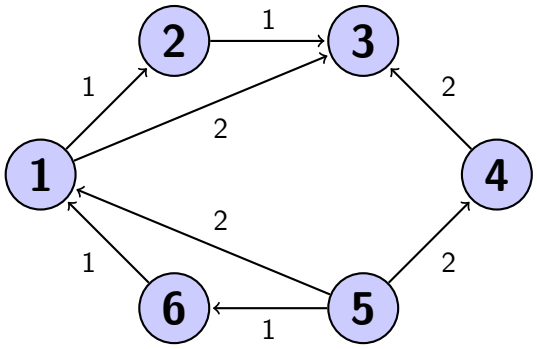

Figure 3.5: A gradient flow on G along with the potential s

Our flow is clearly not divergence-free, but it is necessarily curl-free on each of the triangles $\{1, 2, 3\}$ and $\{1, 5, 6\}$. We can also see that this graph satisfies the path-independence and closed loop properties we expect. That is, the net flow around a closed path or closed trail vanishes, and the net flow along any path between the same two vertices is equal.

2. **(Curl-Free Flow)** Recall that a flow X is curl-free if for each $\{i, j, k\} \in T$,

$$X_{ij} + X_{jk} + X_{ki} = 0.$$

Graphically, this means the flow along any triangular path in G vanishes.

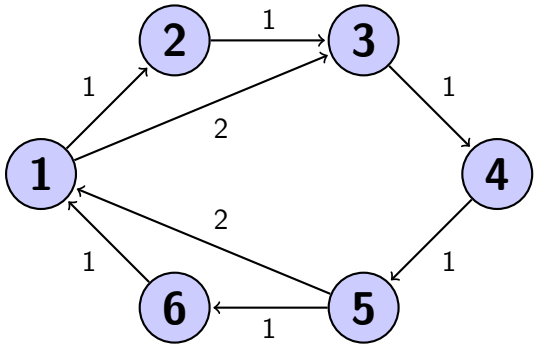

Figure 3.6: A curl-free flow on G.

Although the flow along triangular paths vanishes, the flow along the closed path $1-3-4-5-1$ does not vanish. In general, being curl-free is not enough to ensure that the net flow around any loop is zero.

3. **(Divergence-Free Flow)** An edge flow X is divergence-free if for each vertex i, we have

$$\sum_{j \text{s.t.} \{i,j\} \in E} X_{ij} = 0.$$

That is, the net flow through each vertex is zero, so there are no sinks or sources.

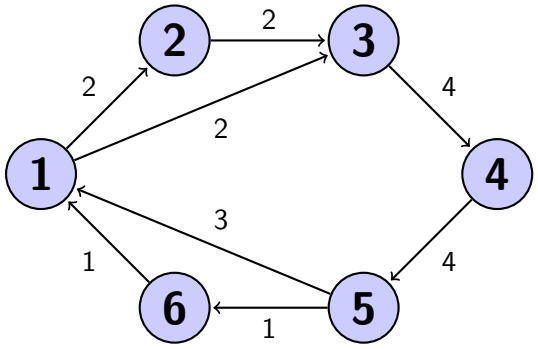

Figure 3.7: A divergence-free flow on G.

A divergence-free flow is unlikely to have vanishing flow around an arbitrary path, although divergence-free flows can also be curl-free.

4. **(Harmonic Flow)** A harmonic flow is a flow that is both curl-free and divergence-free.

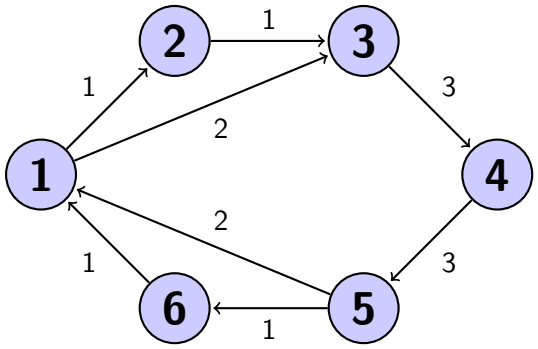

Figure 3.8: A harmonic flow on G.

By definition, the flow around any triangular path vanishes, but there exists longer paths on which the net-flow is non-zero.

3.2 Simplicial Complexes

Next, we introduce an important class of topological objects called simplicial complexes. We can think of simplicial complexes as generalizing the notion of a simple graph to higher dimensions. In our applications to ranking theory, we will work almost exclusively with abstract simplicial complexes rather than geometric simplicial complexes. However, we briefly describe geometrical simplicial complexes, since the geometric object is more intuitive, which may make the connections to graph theory more obvious.

A set of $k+1$ points $\{v_0, v_1, \cdots, v_k\}$ in $\mathbb{R}^n$ is said to be **geometrically independent** if the vectors given by $\{v_1 - v_0, \cdots, v_k - v_0\}$ form a linearly independent set. Thus, two points are geometrically independent if they are distinct, three points are geometrically independent if they are distinct and not colinear, four points are geometrically independent if they are distinct and not coplanar, etc.

Definition 9. (Simplex) If $V = \{v_0, v_1, \cdots, v_n\}$ is a geometrically independent set in $\mathbb{R}^m$, then the **n-simplex** σ **spanned by** $v_1, \cdots, v_n$ is the convex hull of the

points $v_0, \cdots, v_n$. Recall that the convex hull is the intersection of all convex sets containing the points $v_0, \cdots, v_n$, making σ the smallest convex set containing each of these points. It can be shown that the spanning set of a simplex is unique. (See [19].)

To visualize these sets, we recall that a convex set X is a set in which for each pair of points $x, y \in X$, the line segment between x and y is also contained in X. Since a single point is trivially convex, the 0-simplices are just single points. Two points v_0, v_1 are geometrically independent if they are distinct. By definition, the 1-simplex spanned by v_0 and v_1 is the smallest convex set containing v_0 and v_1, so it must contain the line segment between v_0 and v_1. In fact, the 1-simplex spanned by v_0 and v_1 is precisely the line segment between v_0 and v_1. Similarly, the 2-simplex is the smallest convex set containing the three non-colinear points v_0, v_1, v_2. This simplex must contain the line segments between each of v_0, v_1, v_2, which all together form the boundary of a triangle with vertices v_0, v_1, v_2. This simplex is precisely the "filled-in" triangle with vertices v_0, v_1, v_2.

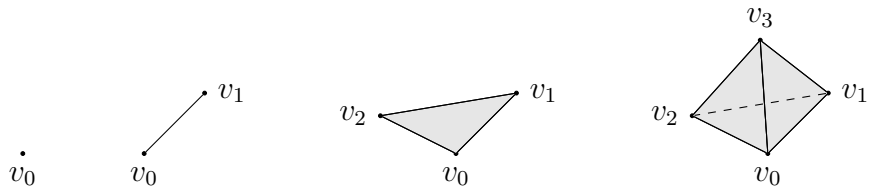

Figure 3.9: Simplices of dimension 0,1,2,3.

The set of points $V = \{v_0, \cdots, v_n\}$ that span σ is called the **vertex set** of σ, the individual points are called **vertices**, and the number n is the **dimension** of σ. The simplex spanned by any subset of V is called a **face** of σ. The faces of σ distinct from σ itself are called **proper faces** of σ; the union of all proper faces of σ is called the **boundary** of σ, denoted Bd σ. The **interior** of σ, denoted Int σ is defined as the set difference Int $\sigma = \sigma -$ Bd σ. The boundaries of the simplices in Figure (3.9)

above are:

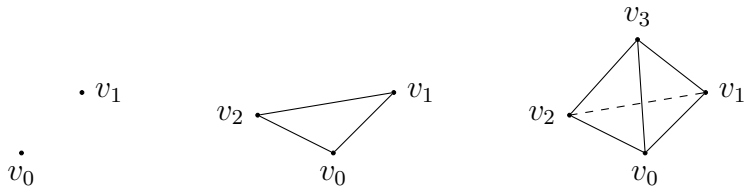

Figure 3.10: Boundaries of various simplices.

Note, that the boundary of 0-simplices are empty. We define the topology of a simplex σ to be the topology it inherits as a subspace of the Euclidean space $\mathbb{R}^m$. We note that the boundary and interior as defined above do not always correspond to the topological boundary and interior of the simplices as subspaces of $\mathbb{R}^m$.

Definition 10. (Simplicial Complex)
A **simplicial complex** K in $\mathbb{R}^m$ is a (finite) collection of simplices in $\mathbb{R}^m$ such that:

(1) Every face of a simplex in K is in K.

(2) The intersection of any two simplices in K is a simplex in K.

We don't necessarily need the qualifier "finite" in this definition, but for our purposes we will only be interested in finite simplicial complexes. In the infinite case, we would need to be more careful in prescribing a sensible topology on our simplicial complex. However, for our purposes this would needlessly complicate our discussion of geometric simplicial complexes.

A subcollection L of a simplicial complex K that contains the faces of all of its simplices (and therefore is also a simplicial complex) is called a **subcomplex** of K. An important example of a subcomplex is the collection of all simplices in K of dimension at most p called the **p-skeleton** of K, denoted by $K^{(p)}$. The 0-skeleton

$K^{(0)}$ of K, is often called the **vertex set** of K. Let $|K|$ denote the subset of $\mathbb{R}^m$ given by the union of all simplices of K. The set $|K|$ is called the **underlying space** of K, we topologize $|K|$ by endowing it with subspace topology inherited from $\mathbb{R}^m$.

3.2.1 Abstract Simplicial Complexes

An **abstract simplicial complex** is a collection K of finite sets such that if $\sigma \in K$, then so is every subset of σ. We say that K is *closed under inclusion*, that is, if $\sigma \in K$ and $\tau \subset \sigma$, then $\tau \in K$. Each element σ of K is called a **simplex**, and the **dimension** of σ is defined to be one less than the number of its elements. The dimension of the simplicial complex K is the highest dimension of any simplex contained in K. If the dimension of σ is p we call σ a **p-simplex**. The **vertex set** V, or sometimes K_0, of K is the union of all 0-simplices (singleton sets) in K. We will not distinguish the 0-simplex $\{v\}$ with the vertex v. If K has dimension n, then for $1 \leq p \leq n$, we let K_p be the collection of p-simplices of K.

Example 1. (Clique Complexes) Let $G = (V, E)$ be a graph on n vertices. For $k = 0, \cdots, n$, we can construct an abstract simplicial complex Σ_G^k called the **k-clique complex of** G as follows. We let the vertex set of Σ_G^k be the vertex set V of G. For any p with $1 \leq p \leq n$, we let Σ_p be the collection of sets of vertices that form a p-clique of G. Note that by our definition, $\Sigma_0 = V$, $\Sigma_1 = E$, and $\Sigma_2 = T$. Then, the k-clique complex of G, Σ_G^k is given by:

$$\Sigma_G^k = \Sigma_k \cup \Sigma_{k-1} \cup \cdots \cup \Sigma_1 \cup \Sigma_0,$$

where the sets Σ_k are the sets of k-simplices in G.

Let σ be a k-simplex with vertex set V. We say two orderings of its vertices are equivalent if they differ by an even permutation. For $k > 0$, there are two equiva-

lence classes of orderings, each of these is called a **orientation** of σ. An **oriented simplex** is a simplex σ together with an orientation of σ. If σ has the vertex set $V = \{v_0, v_1, \cdots, v_n\}$ we use the notation $\sigma = [v_0, \cdots, v_n]$, to denote the oriented simplex σ with the orientation given by even permutations of the tuple $(v_0, v_1, \cdots, v_n)$.

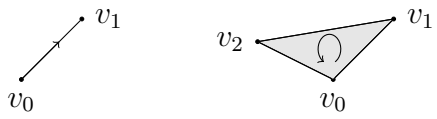

Figure 3.11: The oriented simplices $[v_0, v_1]$ and $[v_0, v_1, v_2]$.

The final subject we discuss from simplicial topology are the chain/cochain groups. Just as simplices generalize simple graphs, chains and cochains will generalize our notion of flows on graphs.

3.2.2 Simplicial Chains

Let K be a simplicial complex. Suppose that we have chosen an orientation for each of the simplices of K. Recall, that we denote the set of (oriented) p-simplices of K by K_p.

Definition 11. (Chains)

A **p-chain** c_p is a formal $\mathbb{R}$-linear sum of oriented p-simplices in K, whose coefficients are skew-symmetric with respect to the orientation of σ. More precisely, if σ' denotes the opposite orientation of the orientated simplex σ, and $c_p(\sigma)$ denotes the coefficient of the oriented simplex σ in the chain c_p, then

$$c_p(\sigma) = -c_p(\sigma') \text{ if } \sigma \text{ is replaced by } \sigma'.$$

The set of p-chains on K is formally denoted by $C_p(K, \mathbb{R})$, however, since we will work exclusively with real coefficients, we will use the simpler notation $C_p(K)$.

If $\sigma \in K_p$ is an oriented p-simplex, we will abuse notation and let σ denote the **elementary p-chain** such that $c(\sigma) = 1$ on σ, $c(\sigma') = -1$ for oppositely oriented σ' and $c(\tau) = 0$ on any other p-simplex τ.

Lemma 3.2.1. ***(Vector Space)*** *The set $C_p(K)$ of p-chains on K is a real vector space. Furthermore, the elementary p-chains $\{\sigma_1, \cdots, \sigma_n\}$ form a basis for $C_p(K)$.*

Proof. Let $c_1, c_2 \in C_p(K)$, and $\alpha \in \mathbb{R}$. We define addition of two chains by adding their coefficients on each simplex, and define scalar multiplication on $C_p(K)$ by scaling each coefficient. That is,

$$(c_1 + c_2)(\sigma) = c_1(\sigma) + c_2(\sigma) \text{ for each } \sigma \in K_p$$

$$(\alpha c_1)(\sigma) = \alpha c_1(\sigma) \text{ for each } \sigma \in K_p$$

If σ' has the opposite orientation of σ then

$$(c_1 + c_2)(\sigma') = c_1(\sigma') + c_2(\sigma') = -c_1(\sigma) - c_2(\sigma) = -(c_1 + c_2)(\sigma'),$$

and similarly

$$(\alpha c_1)(\sigma') = \alpha c_1(\sigma') = -\alpha c_1(\sigma) = -(\alpha c_1)(\sigma).$$

This shows that addition and scalar multiplication operations are well-defined on $C_p(K)$. It remains to show that $\{\sigma_1, \cdots, \sigma_n\}$ is a basis for $C_p(K)$. Let c be an arbitrary chain in $C_p(K)$, then for some $\alpha_1, \cdots, \alpha_n \in \mathbb{R}$ we have:

$$c = \alpha_1\sigma_1 + \cdots + \alpha_n\sigma_n, \text{ where the } \sigma_i \text{ are } p\text{-simplices.}$$

If we replace each simplex σ_i with the elementary chain σ_i, we get an identical expression, now with c written as a linear combination of the elementary chains $\{\sigma_1, \cdots, \sigma_n\}$ Since c was an arbitrary chain, it follows that $\{\sigma_1, \cdots, \sigma_n\}$ spans $C_p(K)$. It remains to show that the set $\{\sigma_1, \cdots, \sigma_n\}$ is linearly independent. Suppose we have written the trivial (zero-valued) p-chain as a linear combination of the σ_i:

$$0 = \alpha_1\sigma_1 + \cdots + \alpha_n\sigma_n, \text{ where the } \sigma_i \text{ are elementary } p\text{-chains.}$$

Each of *chains* σ_j has zero as the coefficient of *simplex* σ_i for $i \neq j$, and so by adding these chains we get the identical expression:

$$0 = \alpha_1\sigma_1 + \cdots + \alpha_n\sigma_n, \text{ where the } \sigma_i \text{ are } p\text{-simplices.}$$

Since the trivial chain has 0 as each of its coefficients, we must have $a_i = 0$ for all i. Hence, the elementary chains $\{\sigma_1, \cdots, \sigma_n\}$ form a basis for $C_p(K)$. □

Note 1. By convention, if p is the largest dimension of any simplex in K, we define the chain groups $C_q(K)$ to be the trivial group $\{0\}$ for $q > p$ and $q < 0$. It is also worth remarking that a 0-simplex has only one orientation, so the skew-symmetry of higher dimensional chains does not appear. So, if K has n different 0-simplices $v_1, \cdots, v_n$, a 0-chain is a sum of the form $c = a_1v_1 + \cdots + a_nv_n$, and the chain group $C_0(K)$ is isomorphic to $\mathbb{R}^n$.

3.2.3 The Boundary Operator

In this section, we introduce an important linear operator on the chain spaces $C_p(K)$ which has the effect of taking an elementary p-chain σ to a $(p-1)$-chain defined on the boundary of σ, so that the orientation of Bd σ is consistent with the orientation of σ.

Definition 12. (Boundary Operator) We define the **boundary operator** $\partial_p : C_p(K) \to C_{p-1}(K)$ as the linear operator which acts on each elementary chain $\sigma = [v_0, \cdots, v_p]$ as follows:

$$\partial_p([v_0, \cdots, v_p]) = \sum_{i=0}^{p} (-1)^p [v_0, \cdots, \widehat{v_i}, \cdots, v_p].$$

Thus, ∂_p acts on the chain $c = \sum \alpha_i \sigma_i$ by

$$\partial_p(c) = \sum_i \alpha_i (\partial_p \sigma_i).$$

To see how the boundary operator works, and why it is named the boundary operator, we illustrate the boundary operator's effect on an elementary chain.

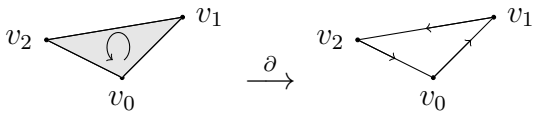

Figure 3.12: The oriented simplex $[v_0, v_1, v_2]$ along with $\partial[v_0, v_1, v_2]$.

Let $\tau \in C_p(K)$. If $\tau = \partial_{p+1}\sigma$ for some $\sigma \in C_{p+1}(K)$ we call τ a **boundary** or more specifically a **p-boundary**. If $\partial\tau = 0$, we say that τ is a **cycle** or **p-cycle**. The next lemma shows that the p-boundaries are a subset of the p-cycles.

Lemma 3.2.2. ***(Closedness)*** *The boundary operators on the spaces $C_p(K)$ satisfy the following closedness relation:*

$$\partial_{k-1} \circ \partial_k = 0.$$

Proof. It is enough to show the results for elementary chains. Let $\sigma = [v_0, \cdots, v_k] \in C_p(K)$ be an elementary chain, then $\partial_k(\sigma) \in C_{k-1}(K)$ is given by:

$$\partial_k([v_0, \cdots, v_k]) = \sum_{i=0}^{k} (-1)^i [v_0, \cdots, \widehat{v_i}, \cdots, v_k].$$

So then $\partial_{k-1}(\partial_k \sigma)$ is given in terms of $\partial_k \sigma$ as follows:

$$\partial_{k-1}(\partial_k[v_0, \cdots, v_k]) = \sum_{j=0}^{k-1} (-1)^j \partial_k([v_0, \cdots, \widehat{v_j}, \cdots, v_k]).$$

The term $\partial_k([v_0, \cdots, \widehat{v_j}, \cdots, v_k])$ is also a summation, so we let i index the inner summation and j index the outer summation. We can write

$$\begin{aligned}\partial_{k-1}(\partial_k[v_0, \cdots, v_k]) &= \sum_{j=0}^{k-1}(-1)^j \partial_k([v_0, \cdots, \widehat{v_j}, \cdots, v_k]) \\ &= \sum_{i<j}(-1)^j(-1)^i[\cdots, v_i, \cdots, v_j, \cdots] \\ &\quad + \sum_{j<i}(-1)^i(-1)^{j-1}[\cdots, v_j, \cdots, v_i, \cdots]\end{aligned}$$

All elementary chains in the full summation over i and j are of the form:

$$[v_0, \ldots, \widehat{v_m}, \ldots, \widehat{v_n}, \ldots, v_k] \text{ for } m < n.$$

Each term appears twice, once with $j = m$ and $i = n$ and once with $i = m$ and $j = n$ (and since $j < i$, we have an opposite sign). So, we may rearrange the summation to so that the terms vanish in pairs. □

The ultimate goal is to develop the Combinatorial Hodge Decomposition theorem, which states that the space $C_p(K)$ decomposes orthogonally as the direct sum of three subspaces. However, before we can discuss orthogonality, we need to define an inner product on $C_p(K)$. If $c_1 = \alpha_1\sigma_1 + \cdots + \alpha_n\sigma_n$ and $c_2 = \beta_1\sigma_1 + \cdots + \beta_n\sigma_n$ we define $\langle c_1, c_2 \rangle_p$ by:

$$\langle c_1, c_2 \rangle_p = \sum_{i=1}^{n} \alpha_i\beta_i. \tag{3.2}$$

In our ranking applications, we will frequently use a weighted inner product

$$\langle c_1, c_2 \rangle_{p,w} = \sum_{i=1}^{n} w_i\alpha_i\beta_i, \tag{3.3}$$

where each $w_i > 0$.

3.2.4 Laplacians

Since $C_k(K)$ and $C_{k-1}(K)$ are both inner product spaces and $\partial_k : C_k(K) \to C_{k-1}(K)$, we can define the adjoint operator $\partial_k^* : C_{k-1}(K) \to C_k(K)$. We recall from linear algebra, that for any $\sigma \in C_k$, the adjoint of ∂ satisfies

$$\langle \partial_k \sigma, \tau \rangle_{k-1} = \langle \sigma, \partial_k^* \tau \rangle_k \text{ for all } \tau,$$

where $\langle \cdot, \cdot \rangle$ may be a weighted inner product. We derive ∂_k^* explicitly in the special cases of $k = 1, 2$, since these we will be useful in our applications. Suppose our complex K has the vertex set $V = \{v_1, \cdots, v_n\}$, and that $\sigma \in C_1(K)$ and $\tau \in C_0(K)$ are elementary chains. That is, $\sigma = [v_i, v_j]$ and $\tau = v_k$ for some $1 \leq i, j, k \leq n$, then we have:

$$\begin{aligned} \langle \partial_1 \sigma, \tau \rangle_0 = \langle \sigma, \partial_1^* \tau \rangle_1 &\implies \langle \partial_1 [v_i, v_j], v_k \rangle_0 = \langle [v_i, v_j], \partial_1^* v_k \rangle_1 \\ &\implies \langle v_j - v_i, v_k \rangle_0 = \langle [v_i, v_j], \partial_1^* v_k \rangle_1 . \end{aligned}$$

Hence,

$$\langle [v_i, v_j], \partial_1^* v_k \rangle_1 = \begin{cases} 1 & \text{if } k = j \\ -1 & \text{if } k = i \\ 0 & \text{otherwise} \end{cases} .$$

But $\langle [v_i, v_j], \partial_1^* v_k \rangle_1$ is just the coefficient of the oriented simplex $[v_i, v_j]$ in the 1-chain $\partial_1^* v_k$. Since $-[v_j, v_i] = [v_i, v_j]$ it follows that

$$\partial_1^* v_k = \sum_{i \text{ s.t } [v_i, v_k] \in K_1} [v_i, v_k].$$

Next, suppose that $\sigma \in C_2(K)$ and $\tau \in C_1(K)$ are elementary chains, write $\sigma = [v_i, v_j, v_k]$ and $\tau = [v_p, v_q]$ for some $1 \leq i, j, k, p, q \leq n$. So, as before we have

$$\begin{aligned} \langle \partial_2 \sigma, \tau \rangle_1 = \langle \sigma, \partial_2^* \tau \rangle_2 &\implies \langle \partial_2 [v_i, v_j, v_k], [v_p, v_q] \rangle_1 = \langle [v_i, v_j, v_k], \partial_2^* [v_p, v_q] \rangle_2 \\ &\implies \langle [v_j, v_k] - [v_i, v_k] + [v_i, v_j], [v_p, v_q] \rangle_0 = \langle [v_i, v_j, v_k], \partial_1^* [v_p, v_q] \rangle_1 . \end{aligned}$$

As before $\langle [v_i, v_j, v_k], \partial_1^*[v_p, v_q] \rangle_1$ is the coefficient of $[v_i, v_j, v_k]$ in the chain $\partial_1^*[v_p, v_q]$. If follows that

$$\partial_1^*[v_p, v_q] = \sum_{r \text{ s.t } [v_r, v_p, v_q] \in K_2} [v_r, v_p, v_q].$$

Definition 13. (Laplacians)

Let K be an abstract simplicial complex. The k-dimensional **combinatorial Laplacian** is the operator Δ_k on $C_k(K)$ defined by:

$$\Delta_k = \partial_{k+1} \circ \partial_{k+1}^* + \partial_k^* \circ \partial_k.$$

As we remarked earlier, if K is a n-dimensional simplicial complex, then the vector spaces $C_p(K)$ are trivial for $p < 0$ and $n > p$. So, by convention $\partial_0 : C_0(K) \to C_{-1}(K)$ is the trivial homomorphism $c \mapsto 0$. Hence, Δ_0, which is often referred to as the **graph Laplacian** is given by $\Delta_0 = \partial_1 \circ \partial_1^*$. To derive Δ_0, let the vertex set of K be $V = \{v_1, \cdots, v_n\}$. The j^{th} column of Δ_0 is given by $\Delta_0 v_j = \partial_0(\partial_0^* v_j)$, so $[\Delta_0]_{ij}$ is just the coefficient of v_i in this sum. We isolate this coefficient by taking an inner product:

$$[\Delta_0]_{ij} = \langle \partial_0(\partial_0^* v_j), v_i \rangle_0 .$$

From our derivations, this is given by:

$$\partial_0(\partial_0^* v_j) = \partial_0 \left(\sum_{k \text{ s.t } [v_k, v_j] \in K_1} [v_k, v_j] \right) = \sum_{k \text{ s.t } [v_k, v_j] \in K_1} \partial_0[v_k, v_j] = \sum_{k \text{ s.t } [v_k, v_j] \in K_1} (v_j - v_k).$$

Hence,

$$[\Delta_0]_{ij} = \left\langle \sum_{k \text{ s.t } [v_k, v_j] \in K_1} (v_j - v_k), v_i \right\rangle_0 .$$

If $i \neq j$ all terms vanish except for the one corresponding to $k = i$. If $i = j$ we get a sum that essentially counts the number of 1-simplices j is a face of:

$$[\Delta_0]_{ij} = \langle \partial_0(\partial_0^* v_j), v_i \rangle_0 = \begin{cases} -1 & \text{if } \{v_i, v_j\} \in K_1 \\ \#\{j : \{v_i, v_j\} \in K_1\} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

If we think of the 1-skeleton of K as a simple graph, then $[\Delta_0]$ is more simply given by

$$[\Delta_0]_{ij} = \begin{cases} -1 & \text{if } \{v_i, v_j\} \in E \\ \deg(v_i) & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

A similar calculation with the weighted inner product in (3.3) give us the Laplacian:

$$[\Delta_0]_{ij} = \begin{cases} -w_{ij} & \text{if } \{v_i, v_j\} \in E \\ \sum_k w_{ik} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

3.2.5 Homology with Real Coefficients

Let K denote a k-dimensional simplicial complex, and let ∂_p, ∂_p^* denote the p^{th} boundary operator and its adjoint. Then we get the following sequence of vector spaces and homomorphisms

$$C_0 \xrightarrow{\partial_0} C_1 \xrightarrow{\partial_1} C_2 \xrightarrow{\partial_2} C_3 \cdots C_{k-1} \xrightarrow{\partial_{k-1}} C_k,$$

for each p satisfying $1 \leq p < k$. If we consider the adjoint to the boundary operator, we get a similar sequence with the directions of the mappings reversed:

$$C_0 \xleftarrow{\partial_0^*} C_1 \xleftarrow{\partial_1^*} C_2 \xleftarrow{\partial_2^*} C_3 \cdots C_{k-1} \xleftarrow{\partial_{k-1}^*} C_k.$$

By Lemma 3.2.2, we know that $\partial_p \circ \partial_{p-1} = 0$ and $\partial_{p-1}^* \circ \partial_p^* = 0$. We call such a sequence of vector spaces and homomorphisms a *chain complex*. Recall that the

closedness property implies that $\text{im}\partial_{p-1} \subseteq \ker \partial_p$ and $\text{im}\partial_p^* \subseteq \ker \partial_{p-1}^*$, then the p^{th} homology group is defined as follows.

Definition 14. (Homology Groups) Let K denote a k-dimensional simplicial complex, and let ∂_p, ∂_p^* denote the p^{th} boundary operator and its adjoint. Then the p^{th} **homology group** $H_p(K)$ of K is given by:

$$H_p(K) = \ker \partial_p \Big/ \text{im}\partial_{p-1}.$$

We will informally refer to $H_p(K)$ as the p-dimensional homology of K.

Intuitively, the p-dimensional homology measures the number of p-dimensional holes in K that do not arise from the removal of $(p+1)$-dimensional faces. We will show in the last section that the one-dimensional homology of a graph G, thought of as a clique-complex, determines whether G admits harmonic flows.

3.2.6 The Combinatorial Hodge Decomposition Theorem

Theorem 3.1. ***The Combinatorial Hodge Decomposition Theorem*** *The space $C_k(K, \mathbb{R})$ admits an orthogonal decomposition*

$$C_k(K, \mathbb{R}) = \text{im}(\partial_k^*) \oplus \ker(\Delta_k) \oplus \text{im}(\partial_{k+1}),$$

with

$$\ker(\Delta_k) = \ker(\partial_{k+1}^*) \cap \ker(\partial_k).$$

Proof. Let K be a simplical complex. Recall that ∂_k is a linear operator from $C_k(K)$ to $C_{k-1}(K)$, so that the adjoint of ∂_k is a linear operator $\partial_k^* : C_{k-1}(K) \to C_k(K)$. By Lemma A.1 we know that the space of k-dimensional chains $C_k(K)$ decomposes orthogonally as follows:

$$C_k(K) = \text{im}(\partial_k^*) \oplus \ker(\partial_k). \tag{3.4}$$

Since $\partial_k \circ \partial_{k+1} = 0$ we know that $\text{im}(\partial_{k-1}) \subseteq \ker(\partial_k)$, so applying Lemma A.1, this time thinking of ∂_{k-1} as a linear operator on $\ker(\partial_k)$ gives us

$$\ker \partial_k = \text{im}\partial_{k-1} \oplus [\ker \partial_{k-1} \cap \ker \partial_k]. \tag{3.5}$$

Thus,

$$C_k(K, \mathbb{R}) = \text{im}(\partial_k^*) \oplus \ker(\Delta_k) \oplus \text{im}(\partial_{k+1}). \tag{3.6}$$

It only remains to show that $\ker(\Delta_k) = \ker(\partial_{k+1}^*) \cap \ker(\partial_k)$. From Lemma (A.1), we know that $C_k(K, \mathbb{R})$ decomposes as:

$$\begin{aligned} C_k(K, \mathbb{R}) &= \text{im}(\partial_k^*) \oplus \ker(\partial_k) \\ C_k(K, \mathbb{R}) &= \text{im}(\partial_{k+1}) \oplus \ker(\partial_{k+1}^*) \end{aligned} \tag{3.7}$$

Hence, by (3.6) we know that:

$$\begin{aligned} \ker(\partial_k) &= \ker(\Delta_k) \oplus \text{im}(\partial_{k+1}) \\ \ker(\partial_{k+1}^*) &= \text{im}(\partial_k^*) \oplus \ker(\Delta_k). \end{aligned}$$

Taking intersections with the equations (3.7) in mind, we see that

$$\ker(\Delta_k) = \ker(\partial_{k+1}^*) \cap \ker(\partial_k)$$

□

Corollary 3.2.1. *The p-dimensional homology of K is isomorphic to the kernel of the p-dimensional Laplacian, i.e.*

$$H_p(K) \cong \ker(\Delta_p).$$

In particular, if K has trivial p-dimensional homology, then Δ_p is injective.

Proof. The previous theorem combined with equation (3.5) implies that

$$\ker \partial_k = \text{im}\partial_{k-1} \oplus \ker(\Delta_k).$$

From Lemma (A.1.1), we have

$$\ker(\Delta_k) \cong \ker \partial_p \Big/ \mathrm{im}\partial_{p-1} = H_p(K).$$

□

3.3 Cochains and Cohomology

As was mentioned earlier, the goal of this chapter was to generalize graphs and edge flows. It should be clear how simplicial complexes generalize graphs, in fact, some texts (see [17]) define simple graphs to be one-dimensional simplicial complexes. However, it may not be clear whether we have made progress towards a suitable generalization for edge flows or the gradient, curl and divergence operators. In this section, we introduce cochains, the coboundary operator, the objects that generalize flows and div/grad/curl respectively, and rephrase the combinatorial Hodge Decomposition Theorem in this new framework.

Fortunately, there is a nice duality between homology and cohomology in the case of finite-dimensional vector spaces such as our chain groups $C_p(K)$. So it will not be necessary to reprove every theorem, but we will at least explain the duality and restate the most relevant results.

Definition 15. (Cochains) Let K be a k-dimensional simplicial complex. The vector space $C^p(K)$ of **p-dimensional cochains** of K is defined to be

$$C^p(K) = \mathrm{Hom}(C_p(K), \mathbb{R}).$$

The **p-th coboundary map** $\delta_p : C^p \to C^{p+1}$. is defined to be the dual of the boundary map $\partial_{p+1} : C_{p+1}(K) \to C_p(K)$.

We see that the space of p-dimensional cochains is just the dual space of the space of p-dimensional chains, so we know that $C_p(K) \cong C^p(K)$ by Lemma (A.1.2). If

$\{\sigma_1, \cdots, \sigma_n\}$ is a basis of $C_p(K)$, then a basis for $C^p(K)$ is the set $\{\sigma^1, \cdots, \sigma^n\}$ of cochains given by:

$$\sigma^i(\tau) = \langle \sigma_i, \tau \rangle_p ,$$

where $\langle \cdot, \cdot \rangle_p$is the inner product from 3.2. That is σ^i maps τ to the coefficient of σ_i in the expression for τ. We call the σ^i an **elementary p-cochain**. In general, any p-chain τ_p determines a p-cochain τ^p which is defined by

$$\tau^p(\sigma) = \langle \tau_p, \sigma \rangle_p ,$$

likewise the p-cochain τ^* can be written as $\tau^*(\sigma) = \langle \tau, \sigma \rangle_p$ for some p-chain τ. So, we can define a vector space isomorphism $\varphi : C_p(K) \to C^p(K)$ such that φ maps the chain σ to the linear functional that maps $\tau \mapsto \langle \sigma, \tau \rangle_p$.

Definition 16. (Coboundary Operator) The **coboundary operator** $\delta_p : C^p \to C^{p+1}$ is defined to be the dual operator to the boundary operator $\partial_{p+1} : C_{p+1}(K) \to C_p(K)$. Let τ^p be the cochain corresponding to $\tau_p \in C_p(K)$, then for any $\sigma \in C_{p+1}(K)$:

$$\delta\tau^p(\sigma) = \tau^p(\partial_{p+1}\sigma) = \langle \tau_p, \partial_{p+1}\sigma \rangle . \tag{3.8}$$

From our the definition of adjoint we know that if $\partial^*_{p+1}\tau$, then for all $\sigma \in C_{p+1}(K)$:

$$\left\langle \partial^*_{p+1}\tau, \sigma \right\rangle = \langle \tau, \partial_{p+1}\sigma \rangle$$

So if we let $\omega_p = \partial^*_{p+1}\tau$, we expect $\delta\tau^p = \omega^p$ If $\varphi : C_p(K) \to C^p(K)$ is the isomorphism sending $\tau \mapsto \tau^*$, we get the following commutative diagram between chain complexes:

$$\begin{array}{ccccccccc}
\cdots & \xrightarrow{\partial^*_{p-1}} & C_p(K) & \xrightarrow{\partial^*_p} & C_{p+1}(K) & \xrightarrow{\partial^*_{p+1}} & \cdots \\
 & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \varphi} & & \\
\cdots & \xrightarrow[\delta_{p-1}]{} & C^p(K) & \xrightarrow[\delta_p]{} & C^{p+1}(K) & \xrightarrow[\delta_{p+1}]{} & \cdots
\end{array}$$

If we look at the adjoint of δ, and the boundary operator ∂, we get a similar reversed diagram

$$\begin{array}{ccccccc}
\cdots & \xleftarrow{\partial_{p-1}} & C_p(K) & \xleftarrow{\partial_p} & C_{p+1}(K) & \xleftarrow{\partial_{p+1}} & \cdots \\
 & & \downarrow{\scriptstyle \varphi} & & \downarrow{\scriptstyle \varphi} & & \\
\cdots & \xleftarrow[\delta^*_{p-1}]{} & C^p(K) & \xleftarrow[\delta^*_p]{} & C^{p+1}(K) & \xleftarrow[\delta^*_{p+1}]{} & \cdots
\end{array}$$

From this duality, we get a cochain analog of each theorem for chains proved in this chapter. The closedness of the boundary operator ∂ implies the closedness of the adjoint ∂^*, our isomorphism φ also gives us

Theorem 3.2. ***(Closedness)*** *The coboundary operators on the spaces $C^k(K, \mathbb{R})$ satisfy the following relation:*

$$\delta_{k+1} \circ \delta_k = 0.$$

We can define a cochain Laplacian operator in the same way the combintorial Laplacian was defined.

Definition 17. Let K be an abstract simplicial complex. We can define a k-dimensional **Laplacian** Δ^k on the cochain space $C^k(K, \mathbb{R})$ to be the operator:

$$\Delta^k = \delta^*_k \circ \delta_k + \delta_{k-1} \circ \delta^*_{k-1}$$

We can also define p-dimensional cohomology, in the same fashion that p-dimensional homology was defined.

Definition 18. (Cohomology) Let K denote a k-dimensional simplicial complex, and let δ_p denote the p^{th} coboundary operator. Then the p^{th} **cohomology group**

$H^p(K)$ of K is given by:

$$H^p(K) = \ker \delta_p \Big/ \mathrm{im}\delta_{p-1}.$$

As may be expect, $H_p(K) \cong H^p(K)$, in fact $\varphi : C_p(K) \to C^p(K)$ induces an isomorphism $\overline{\varphi} : H_p(K) \to H^p(K)$ by defining $\overline{\varphi}$ to be φ acting on coset representatives of $H_p(K)$. We also get a Hodge Decomposition Theorem for cochains,

Theorem 3.3. ***(The Combinatorial Hodge Decomposition Theorem)*** *The space $C^k(K, \mathbb{R})$ admits an orthogonal decomposition*

$$C^k(K, \mathbb{R}) = \mathrm{im}(\delta_{k-1}) \oplus \ker(\Delta^k) \oplus \mathrm{im}(\delta_k^*),$$

with

$$\ker(\Delta^k) = \ker(\delta_k) \cap \ker(\delta_{k-1}^*).$$

Furthermore,

$$\ker(\Delta^k) \cong H^k(K).$$

To expedite our exposition, we could have started with cohomology and perhaps ignored homology altogether, however we believe the geometric intuition of the boundary operator is more apparent than the coboundary operator.

3.3.1 Simple Graphs Revisited

Throughout this section we will not distinguish the oriented k-simplex $[v_{i_1}, \cdots, v_{i_k}]$ with the elementary k-chain $[v_{i_1}, \cdots, v_{i_k}]$. We will denote the elementary k-cochain which takes a value of 1 on the chain $[v_{i_1}, \cdots, v_{i_k}]$ by $[v^{i_1}, \cdots, v^{i_k}]$.

In §3.1 we defined a **(simple) graph** G to be an ordered pair (V, E), with V a nonempty finite set and $E \subset \binom{V}{2}$. From G we can define a simplicial complex K, called the 3 clique-complex of G given by:

$$K = T \cup E \cup V.$$

An edge flow on a graph G was defined to be a function $X : V \times V \to \mathbb{R}$ that satisfies

$$\begin{cases} X(v_i, v_j) = -X(v_j, v_i) & \text{if } \{v_i, v_j\} \in E \\ X(v_i, v_j) = 0 & \text{otherwise} \end{cases}.$$

Just as an edge flow is completely determined by its value on each edge oriented in some way, a 1-cochain is completely determined by its value on elementary 1-chains. Hence, we can identify the edge flow X with the cochain $\tilde{X} \in C^1(K)$ given by:

$$\tilde{X}([i, j]) = X_{ij} \text{ for each oriented simplex } [i, j] \text{ in } E.$$

Alternatively, we can write this as a linear combination of elementary 1-cochains:

$$\tilde{X} = \sum_{[v_i, v_j] \in E} X_{ij}[v^i, v^j]$$

A 2-cochain Θ can be similarly identified with a triangular flow Φ, so that the value of Θ on the elementary 2-chain $[v_i, v_j, v_k]$ is given by

$$\Theta([v_i, v_j, v_k]) = \Phi(v_i, v_j, v_k) = \Phi_{ijk}.$$

As before, we can write this as a linear combination of elementary 2-cochains:

$$\Theta = \sum_{[v_i, v_j, v_k] \in E} \Phi_{ijk}[v^i, v^j, v^k].$$

Likewise, a 1-cochain τ can be identified with a score function s, so that the value of τ on the elementary 1-chain $[i] = \{i\}$ is

$$\tau([i]) = s(i),$$

written as a sum this becomes:

$$\tau = \sum_{v_i \in V} s_i[v^i].$$

The cochain notation is awkward and will be largely replaced by the nicer graph theoretic notation in later sections. We have shown that we can identify the cochain spaces $C^0(K), C^1(K), C^2(K)$ with score functions, edge flows and triangular flows respectively. We show that the coboundary operators δ_0, δ_1 generalize grad and curl respectively. Suppose we have a score vector $s \in \mathbb{R}^n$, then the value of grad (s) on (i, j) is $s_j - s_i$. Write the score function s as a 0-chain:

$$s = \sum_{i=1}^{n} s_i [v^i].$$

We show that the value of δs on $[v_m, v_n]$ is $s_m - s_n$ which is actually $-$grad (s). From our definition of the coboundary operator we have

$$\delta_0 s = \delta_0 \sum_{i=1}^{n} s_i v^i = \sum_{i=1}^{n} s_i \delta_0 v^i = \sum_{i=1}^{n} s_i \sum_{i \text{ s.t } [v_i, v_k] \in E} [v^i, v^k].$$

Evaluating at $[v_m, v_n]$ we see that all terms vanish except for $s_m[v^m, v^n] + s_n[v^n, v^m]$, which evaluates to $s_m - s_n$. If we are given an edge flow X, then the value of curl X on (i, j, k) is $X_{ij} + X_{jk} + X_{ki}$. If we write X as a 1-chain in the fashion described earlier:

$$X = \sum_{[v_i, v_j] \in E} X_{ij} [v^i, v^j].$$

So that,

$$\delta_1 X = \sum_{[v_i, v_j] \in E} X_{ij} \delta_1 [v^i, v^j] = \sum_{[v_i, v_j] \in E} X_{ij} \sum_{k \text{ s.t } [v_i, v_j, v_k] \in K_2} [v_i, v_j, v_k].$$

Before we evaluate $\delta_1 X$ at $[v_p, v_q, v_r]$, we see that all but three terms vanish, namely

$$V_{pq}[v^p, v^q, v^r] + V_{pr}[v^p, v^r, v^q] + V_{qr}[v^q, v^r, v^p].$$

Evaluating at $[v_p, v_q, v_r]$ gives us:

$$V_{pq} - V_{pr} + V_{qr} = V_{pq} + V_{rp} + V_{qr} = V_{pq} + V_{qr} + V_{rp}.$$

The last operator we consider is the divergence. Recall that if we are given an edge flow X, then $\operatorname{div} X$ is a score function whose value on v_j is given by $\sum_i X_{ij}$. We write X out as a 1-chain,

$$X = \sum_{[v_i,v_j]\in E} X_{ij}[v^i, v^j].$$

The adjoint δ_0^*, which is dual to the boundary operator, is then given by

$$\delta_0^* X = \sum_{[v_i,v_j]\in E} X_{ij}\delta_0^*[v^i, v^j] = \sum_{[v_i,v_j]\in E} X_{ij}(v^j - v^i).$$

Evaluating at v_j then gives us

$$\sum_{i \text{ s.t. } [v_i,v_j]\in E} X_{ij}.$$

One last thing to note is that while $\delta_0^* X$ and $\delta_1 X$ gave us the standard formulas for curl and div, $\delta_0 s$ gave us $-\text{grad }(s)$. The reason for this is that grad and div are actually negative adjoints of one another, that is, $\text{grad}^* = -\operatorname{div}$. With this discussion in mind, we can rephrase the Hodge Decomposition Theorem in terms of edge flows.

Theorem 3.4. ***(The Combinatorial Hodge Decomposition Theorem)*** *The space of* edge flows $C_1(G)$ *on a graph* $G = (V, E)$ *admits an orthogonal decomposition*

$$C_1(G) = \operatorname{im}(\text{grad }) \oplus \ker(\Delta^k) \oplus \operatorname{im}(\text{curl}^*),$$

with

$$\ker(\Delta^k) = \ker(\text{curl}) \cap \ker(\text{div}).$$

What this means is that any edge flow on a graph G decomposes orthogonally into a gradient component, a harmonic component, and a "curly" component. Since both gradient flows and harmonic flows are curl-free, we call the third subspace "curly" flows since it is the only component with nonzero curl. We show in the next chapter how this theorem allows us to analyze ranking data.

Chapter 4: Ranking

4.1 Pairwise Comparisons

We now describe a generic situation in which one may use the methods described in this paper. Suppose we wish to rank a set of n items $V = \{1, \cdots, n\}$, each of which have been rated by some subset of m voters $\Lambda = \{\alpha_1, \cdots, \alpha_m\}$. Here, voters and alternatives are used in a general sense. In various contexts, our alternatives could be movies, restaurants, websites, sports teams, candidates for mayor, and our voters could be traditional voters, hyperlinks, matches (games played between teams), or movie reviews.

In order to apply *HodgeRank* to a given dataset, we require that our data be given in the form of **pairwise comparisons**, that is, each voter would have compared some pairs of alternatives. These could either be *ordinal* pairwise comparisons (A is preferred to B), or *cardinal* pairwise comparisons (A is preferred to B by some value). A simple example of ranking data in the form of ordinal pairwise comparisons is a blind food tasting (e.g., the "Pepsi Challenge"), where voters are asked to choose between two unknown foods. An example of a cardinal pairwise comparison, and one which we will revisit later, is a sports match (e.g., a basketball game) along with margin of victory. Each game result can be thought of as a pairwise comparison between the teams involved, and the margin of victory acts as a cardinal rating.

We will work exclusively with data already in the form of pairwise comparisons, but it is generally not difficult to convert other forms of ranking data into pairwise comparisons. We may have data in the form of *cardinal scores*, in which each voter

assigns a rating to some subset of items. For example, the dataset released for the Netflix prize is given in the form of user ratings, where Netflix movies are given ratings of 1-5 stars. Another standard form of ranking data is *ordinal preferences*, in which each voters lists the alternatives in order of preference. In both cases, cardinal scores and ordinal preferences can be easily converted to pairwise comparisons. For example, if voter α has assigned items i and j cardinal scores s_i, s_j respectively, voter α has indirectly compared i, j pairwise, and we can say the degree of preference for item i over item j is $s_i - s_j$. A similar process can be used to convert ordinal rankings to pairwise comparisons.

Perhaps the most natural and familiar examples of pairwise comparisons comes from sports. In most team sports, each game played can be thought of as a direct comparison of the two teams. We can treat the outcome of the match as either an *ordinal* pairwise comparison (i.e., winning team > losing team) or a *cardinal* pairwise comparison (i.e., winning team > losing team by n points). It should be noted that there are more sophisticated ways to form these comparisons, but these will suffice for the purposes of illustrating *HodgeRank*.

4.1.1 Pairwise Comparison Matrices

Recall that $V = \{1, \cdots, n\}$ and $\Lambda = \{\alpha_1, \cdots, \alpha_m\}$ denote the set of voters and alternatives, respectively. Suppose that we are given ranking data already in the form of pairwise comparisons. For each voter $\alpha \in \Lambda$, we can form two related matrices W^α and Y^α called the *weight matrix* and *pairwise comparison matrix* for voter α. The weight matrix W^α is the symmetric n-by-n matrix defined by:

$$[W_{ij}^\alpha] = \begin{cases} 1 & \text{if } \alpha \text{ made a pairwise comparison for } \{i, j\} \\ 0 & \text{otherwise} \end{cases}.$$

We may think of W^α as keeping track of which comparisons have been made by voter α. Note that the weight matrix W^α is symmetric. The pairwise comparison matrix Y^α is given by:

$$[Y_{ij}^\alpha] = \begin{cases} \text{degree of preference of } i \text{ over } j & \text{if } W_{ij}^\alpha = 1 \\ 0 & \text{if } W_{ij}^\alpha = 0 \end{cases}$$

If $Y_{ij}^\alpha > 0$, then voter α prefers i over j by the amount $|Y_{ij}^\alpha|$. So we should expect that α prefers j over i by the negative of this amount:

$$Y_{ij}^\alpha = -Y_{ji}^\alpha.$$

Hence, the pairwise comparison matrix Y^α is skew-symmetric.

Next, we describe the perhaps the simplest way one can aggregate pairwise comparisons over voters. In essence, we are averaging comparisons over all voters such that our weight matrix W reflects the total number of comparisons made by all voters and that our pairwise comparison matrix Y contains the average degree of preference for each alternative. Define the weight matrix W to be the matrix $[W_{ij}]$ where

$$W_{ij} = \sum_{\alpha \in \Lambda} W_{ij}^\alpha,$$

and the pairwise comparison matrix Y to be the matrix Y_{ij} given by:

$$Y_{ij} = \frac{1}{W_{ij}} \sum_{\alpha \in \Lambda} Y_{ij}^\alpha.$$

The matrix W is symmetric since each of the matrices W_{ij}^α are symmetric:

$$W_{ji} = \sum_{\alpha \in \Lambda} W_{ji}^\alpha = \sum_{\alpha \in \Lambda} W_{ij}^\alpha = W_{ij}.$$

Similarly, the matrix Y is skew-symmetric since each of the Y_{ij}^{α} are skew-symmetric:

$$Y_{ji} = \frac{1}{W_{ji}} \sum_{\alpha \in \Lambda} Y_{ji}^{\alpha} = \frac{1}{W_{ij}} \sum_{\alpha \in \Lambda} (-1) Y_{ij}^{\alpha} = \frac{-1}{W_{ij}} \sum_{\alpha \in \Lambda} Y_{ij}^{\alpha} = -Y_{ij}.$$

Henceforth, W and Y will be assumed to have been aggregated in the fashion described above. See [14] for other examples of aggregation.

4.1.2 Pairwise Comparison Graphs

Suppose we have a weight matrix W and a pairwise comparison matrix Y representing our ranking data, perhaps aggregated over the voters. Then we can define an underlying graph $G = (V, E)$, called the **pairwise comparison graph**. As before, let V denote the set of alternatives to be ranked, then we can define the edge set E to be the set of pairs where comparisons have been made:

$$E = \left\{ \{i, j\} \in \binom{V}{2} : W_{ij} > 0 \right\}.$$

Our pairwise comparison matrix Y then induces an edge flow on our graph G by defining $Y(i, j) = Y_{ij}$. The reader can verify that that Y satisfies the properties of an edge flow. Below is an example of a skew-symmetric matrix Y inducing an edge flow on a graph, we may think of Y as representing a pairwise comparison matrix.

$$Y = \begin{pmatrix} 0 & -1 & 2 & 0 & 2 & -2 \\ 1 & 0 & 3 & 0 & 0 & 0 \\ 2 & -3 & 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 1 & 0 \\ -2 & 0 & 0 & -1 & 0 & -2 \\ 2 & 0 & 0 & 0 & 2 & 0 \end{pmatrix} \quad \longleftrightarrow$$

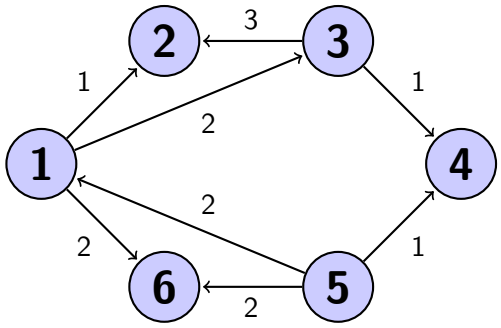

Figure 4.1: Pairwise comparison matrix and corresponding graph

Now that we have represented our data as a pairwise comparison graph (i.e., an edge flow), we are free to use the results from the previous chapter to analyze our data. In the next section, we revisit these results in a ranking context and show how Hodge theory can be used to rate our set of alternatives V.

4.2 Hodge Rank

The fundamental motivation behind *HodgeRank* is to examine the notion of consistency in our ranking data. The example we used to introduce this notion was the *Condorcet paradox* from voting theory. In a more general context, the Condorcet paradox states that ranking data need not be consistent. If A, B, C are among the alternatives we wish to rank or rate, the data may imply that $A < B$ and $B < C$ and yet $C < A$, which yields the cyclic ordinal relation $A < B < C < A$. In a pairwise comparison graph, these *ordinal inconsistencies* manifest themselves as triangles around which edge directions all point clockwise or counterclockwise.

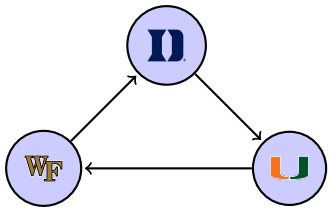

Figure 4.2: An ordinal cyclic relation between three basketball teams.

We also consider a more subtle form of inconsistency, which we refer to as a *cardinal inconsistency*. These cardinal inconsistencies manifest themselves as loops around which the net flow is non-vanishing, and paths between fixed pairs of vertices with different net flow along them.

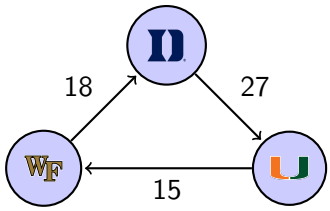

Figure 4.3: A cardinal cyclic relation between three basketball teams.

The idea behind a cardinal inconsistency is that a path between two pairs of alternatives i and j form indirect comparisons. For these indirect comparisons to be meaningful, we expect that any two indirect comparisons of i and j should be roughly the same. In terms of edge flows, we would like the net flow along any path between i and j to be approximately the same. Ideally, each alternative i and j would have been compared directly, but our data may be very incomplete, and so we cannot expect direct comparisons of each alternative. Shorter paths represent more direct comparisons, so we may assign more value to indirect comparisons of shorter length. For instance, a path of length two between two alternatives i and j represents a once-removed comparison of i and j, a path of length three represents a twice-removed comparison, etc.

A similar argument applies to loops and closed paths in our graph G. If we have a closed path in G beginning and ending at the alternative i, then the flow around this path can be thought of as an indirect self-comparison. It is reasonable to expect that alternative i should not be rated lower or higher than itself, and in the ideal case, flows resulting from pairwise comparisons should vanish around loops. As before, we would like these self-comparisons to be as direct as possible, so triangular loops, as in the previous two figures, will be of particular interest.

4.2.1 Interpretation of the Hodge Theorem

We recall from the Hodge Decomposition Theorem that the space of edge flows decomposes orthogonally into three subspaces.

Theorem 4.1. ***(The Combinatorial Hodge Decomposition Theorem)*** *The space of edge flows $C_1(G)$ on a graph $G = (V, E)$ admits an orthogonal decomposition*

$$C_1(G) = \text{im}(\text{grad }) \oplus \ker(\Delta_1) \oplus \text{im}(\text{curl}^*),$$

with

$$\ker(\Delta_1) = \ker(\text{curl}) \cap \ker(\text{div}).$$

Informally, we can state this decomposition as saying that the set of all edge flows on a graph G decomposes as

$$\text{Edge Flows} = \text{Gradient Flows} \oplus \text{Harmonic Flows} \oplus \text{Curly Flows}.$$

We recall that gradient flows satisfy the path-independence property and have vanishing flow around closed loops, so these represent the *consistent* pairwise rankings. That is, in a gradient flow, all indirect comparisons of the same two alternatives are the same, and consequently there are no problematic self-comparisons, we say that gradient flows are *globally consistent*. Furthermore, each gradient flow is determined by a score function $s \in \mathbb{R}^n$. In a gradient flow, a direct comparison between i and j is given by the difference $s_j - s_i$, so the score function s may be interpreted as a rating of the alternatives.

A harmonic flow has a less-obvious interpretation. We recall that harmonic flows are both curl-free and divergence-free. Since they are divergence-free, each alternative will have been compared favorably as much as it has been compared unfavorably.

However, the curl-free condition implies that self-comparisons along triangular loops are consistent, and consequently, the indirect comparisons of three alternatives which together form a triangle are consistent with one another. We say that harmonic flows are *locally consistent*, in that we can order (rank) alternatives within every triangle, but we may not be able to assign a meaningful global ranking.

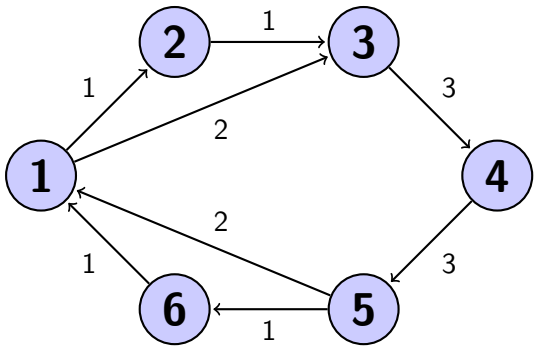

Figure 4.4: A harmonic flow on G.

As with harmonic flows, curl flows are also divergence-free, so each alternative will have been compared favorably as much as it has been compared unfavorably. As the name suggests, these are the only flows with non-zero curls. What this means is that, for a curl flow we cannot expect meaningful (consistent) comparisons or self-comparisons even within a triangle.

4.2.2 Least Squares via HodgeRank

In the context of ranking, gradient flows seem to be the most desirable. The potential function s which induces a gradient flow can be thought of as a rating of the alternatives. The other two subspaces, the harmonic and curl flows represent the locally consistent and globally inconsistent components of our data, respectively. If we are given a pairwise comparison matrix Y, we can project the corresponding edge flow

onto the space of gradient flows by formulating an appropriate least squares problem. The space of edge flows on our pairwise comparison graph is isomorphic to the space of skew-symmetric matrices $X(G)$ given by:

$$X(G) = \{X \in \mathbb{R}^{n\times n} : X = -X^T \text{ and } X_{ij} = 0 \text{ if } \{i,j\} \notin E\}. \tag{4.1}$$

We can use the standard matrix inner product on $X(G)$:

$$\langle X, Y\rangle_2 = \sum_{i,j} X_{ij}Y_{ij} = \mathrm{tr}(X^TY).$$

Alternatively, if $W = [w_{ij}]$ is our weight matrix, we can use the weighted inner product:

$$\langle X, Y\rangle_2^w = \sum_{i,j} w_{ij}X_{ij}Y_{ij}.$$

It should be noted that this is not exactly the same inner product as in (3.3), because there will be some double counting ($X_{ij}Y_{ij} = X_{ji}Y_{ji}$). We could normalize our inner product to account for this by simply dividing by two. However, it will not affect our results, so we will merely note this difference. Now that we have an inner product on our space of edge flows, we can orthogonally project our pairwise comparison matrix Y onto the subspace of gradient flows by solving the least squares problem:

$$\min_{s\in\mathbb{R}^n} = \|\text{grad } s - Y\|_2^w. \tag{4.2}$$

The normal equations for the least squares problem 4.2 are then:

$$\text{div} \circ \text{grad } s = \text{div } Y$$

But $\text{div} \circ \text{grad } = \Delta_0$, so that our least squares projection is given by:

$$s = \Delta_0^\dagger \,\text{div}\, Y,$$

where † indicates the Moore-Penrose psuedoinverse[1]. We may rightly worry that this least squares solution is not unique, indeed for each solution s and each $c \in \ker(\text{grad }\ )$, the score vector $s + c$ is also a solution. However, from our corollary to the combinatorial Hodge Decomposition Theorem, we know that

$$\ker(\text{grad }\ ) = \ker(\Delta_0) \cong H_0(K),$$

where K is the 3-clique complex of G. In particular, H_0 is just the number of connected components of G. Hence, for a connected graph, there a unique rating vector s up to adding a constant multiple of $[1, \cdots, 1]^T$ (See Proposition 3.1.2). In fact, the vector $[1, \cdots, 1]^T$ spans $\ker(\text{grad }\ )$.

The least squares residual R is then given by subtracting off the gradient component: $R = Y - \text{grad }(s)$. The residual R lies completely in the subspace $\ker(\Delta_1) \oplus \text{im}(\text{curl}^*)$, so we can decompose R further into its harmonic and curly component using orthogonal projection. The orthogonal projection onto $\ker(\Delta_1)$ may be done via the operator $1 - \Delta_1^\dagger \Delta_1$, so the harmonic component of Y is

$$H = R(1 - \Delta_1^\dagger \Delta_1).$$

Finally, $C = R - H$ is the projection of Y onto the curl flows.

In our numerical examples, we will not consider the decomposition of R into H and C. However, if we decompose our data into $Y = \text{grad }(s) + R$, we can define the *cyclicity ratio* our data to be the ratio:

$$C_r = \frac{\|R\|}{\|\text{grad }(s)\|},$$

where s is our rating determined by the least squares problem above. The projection grad (s) captures the component of our original data that contains no cardinal or

[1]See A.2 for a description of the Moore-Penrose Psuedoinverse and least squares problems.

ordinal inconsistencies. Hence, C_r is a measure of the consistent vs inconsistent parts of our ranking data. If $C_r \ll 1$, the consistent (gradient) component account for more of our data than the inconsistent component, so we should expect that our ratings reflect our original data well. On the other hand, if $C_r > 1$ there may be too many inconsistencies in our data to assign a good rating. In §4.3.1 we will consider two small examples with $C_r \approx 0.2$ and $C_r \approx 2$, that illustrate how these inconsistencies may affect our ratings.

4.2.3 Sports Data

In this section, we setup the standard sports ratings problem which we wish to solve. According to Langville and Meyer in [15], "the goal of any [sports] rating system is to associate a rating for each team in a league of n teams, where m total league games have been played thus far."

We will first describe how *HodgeRank* may be used as a sports rating method. We will also describe *Massey's Method*, a well-known and relatively simple sports rating method, since we wish to highlight several similarities between *HodgeRank* and *Massey's method*, and show how *HodgeRank* generalizes *Massey's method* using graph theory.

Massey's method is relatively simple compared to other rating methods in that we only need to know which teams played in each game, who won the game, and what the margin of victory was. We should note that what we refer to as *Massey's method* is actually one of several methods by Kenneth Massey, see [15] for a description of some of Massey's other rating methods.

4.2.4 Weights and Pairwise Comparisons

In this subsection, we describe how one can form a weight matrix W and pairwise comparison matrix Y using sports data. Suppose that we wish to rank basketball teams based on game results for one season, or even for a partial season. If there are n teams, we can assign each team a number so that our set of alternatives is given by $V = \{1, \cdots, n\}$. Our set of "voters" is the set of rematches, that is, if the most any two teams in our league have played each other is m times, we have a set $\Lambda = \{\alpha_1, \cdots, \alpha_m\}$ of m voters. For each $\alpha \in \Lambda$, we can define the obvious weight matrix:

$$[W_{ij}^{\alpha}] = \begin{cases} 1 & \text{if } i \text{ has played } j \text{ at least } \alpha \text{ times.} \\ 0 & \text{otherwise.} \end{cases}.$$

If team i has played team j at least α times, we define S_{ij}^{α} to be the margin of victory of team i over team j in the α^{th} game played between the two teams:

$$S_{ij}^{\alpha} = (\text{team } i \text{ score}) - (\text{team } j \text{ score}).$$

For each voter/rematch $\alpha \in \Lambda$, we can form a pairwise comparison matrix Y^{α} containing the game results for α^{th} game played between each pair of teams:

$$[Y_{ij}^{\alpha}] = \begin{cases} S_{ij}^{\alpha} & \text{if } W_{ij}^{\alpha} = 1 \\ 0 & \text{otherwise.} \end{cases}$$

We aggregate the margin of victory over each rematch in the obvious way. We use the total number of games played between i and j for our weight matrix:

$$[W_{ij}] = \sum_{\alpha \in \Lambda} W_{ij}. \tag{4.3}$$

We define a pairwise comparison matrix Y consisting of average margin of victory over all games played:

$$[Y_{ij}] = \frac{1}{W_{ij}} \sum_{\alpha \in \Lambda} Y_{ij}^{\alpha} = \frac{1}{W_{ij}} \sum_{\alpha \in \Lambda} S_{ij}^{\alpha}. \tag{4.4}$$

Our least squares problem is given by:

$$\min_{s \in \mathbb{R}^n} = \|\text{grad } s - Y\|_2^w.$$

We will also consider another type of pairwise comparison. Using the same weight matrices as above, define $B_{ij}^{\alpha} = \text{sgn } (S_{ij}^{\alpha})$. That is, $B_{ij}^{\alpha} = 1$ if i beat j in game α, and $B_{ij}^{\alpha} = -1$ if j beat i in game α, and $B_{ij}^{\alpha} = 0$ if i and j have played less than α games. From this, we can define another pairwise comparison matrix B^{α} which we will call the **binary comparison matrix**. Unlike with margin of victory, these binary comparison matrices B^{α} treat all wins and losses equally. With Y^{α} each win or loss carries with it a margin of victory, but with B^{α} each win or loss is given by the same cardinal pairwise comparisons ± 1. Just as with Y^{α}, we can aggregate these binary comparison matrices in a similar fashion:

$$[B_{ij}] = \frac{1}{W_{ij}} \sum_{\alpha \in \Lambda} B_{ij}^{\alpha}.$$

In general, B_{ij} will not take on values ± 1 (unless i or j won each rematch between i, j), but we will still refer to B as a binary comparison matrix. In the case of a binary comparison matrix, our least squares problem is

$$\min_{s \in \mathbb{R}^n} = \|\text{grad } s - B\|_2^w.$$

4.2.5 Massey's Method

The discussion in this section is largely based on the description of *Massey's method* in [15]. Suppose we wish to rate n basketball teams in a league, where a total of m

league games have been played. We will denote each team by a number $1, \cdots, n$. The goal of *Massey's method* is to associate ratings $r_1, \cdots, r_n$ to each team such that if y_k is the margin of victory of i over j in the k^{th} game, the ratings r_i and r_j satisfy:

$$y_k = r_i - r_j. \tag{4.5}$$

There is an equation of this form for each game, so we have a system of m linear equations in n unknowns that can be written as the matrix equation

$$Xr = y.$$

Here X is a $m \times n$ matrix with $X_{ki} = 1$ and $X_{kj} = -1$ if team i beats team j in the k^{th} game, the vector r stores the unknown team ratings r_i for each of the n teams, and the vector y stores the margin of victory for each of the m games.

In most cases there will be more league games played than teams in the league. Hence, $Xr = y$ is an overdetermined system of equations, meaning there are more equations than unknowns ($m > n$). In general, overdetermined solutions do not have solutions, but we can find a least squares solution[2] to $Xr = y$ by solving the normal equations:

$$X^T X r = X^T y.$$

Let $M = X^T X$ and $p = X^T y$. The matrix n-by-n matrix M is given by:

$$M_{ij} = \sum_{k=1}^{m} X_{ki} X_{kj}.$$

Let us first consider the off-diagonal elements ($i \neq j$). The only time both X_{ki} and X_{kj} are nonzero is if i and j both played in the k^{th} game. In this case, one of is -1 and the other is 1, so that $X_{ki} X_{kj} = -1$. Hence,

[2]see Appendix (A.2) for an overview of least squares.

$$M_{ij} = \sum_{k=1}^{m} X_{ki}X_{kj} = -\ (\#\text{ games played between } i \text{ and } j).$$

A similar argument applies to the diagonal elements M_{ii}. In this case, $X_{ki}X_{ki} = 1$ if i played in the k^{th} game and $X_{ki}X_{ki} = 0$ otherwise. Hence,

$$M_{ii} = \sum_{k=1}^{m} X_{ki}X_{ki} = \#\text{ total games played by team } i.$$

The n entries of the vector p are given by:

$$p_i = \sum_{k=1}^{m} X_{ki}y_i.$$

If team i did not play in the k^{th} game, then $X_{ki}y_i = 0$, if team i won the k^{th} game then $X_{ki}y_i = y_i$, otherwise $X_{ki}y_i = -y_i$. Therefore p_i is the cumulative point differential for team i.

There is still one obstacle keeping us from directly solving $Mr = p$, namely, the matrix M is singular. An off diagonal element M_{ij} is the negative of games played between i and j, while a diagonal element is the total # of games played by i. If one adds the elements in each row you get total games played minus total games played. Thus, the columns of M are linearly dependent. Massey's workaround is to replace the last row in M with all 1's and the last element of p with 0. This forces the ratings to sum to 0 and also results in M being nonsingular (see [15]). If M and p have been altered in this way, then our ratings r are given by:

$$r = M^{-1}p.$$

If we use *HodgeRank* as described in the previous section. We can represent the same data as a weight matrix W and pairwise comparison data Y where W_{ij} is the total number of games played between i and j and Y_{ij} is the average margin of victory

of i over j. To determine our ratings, we solve the least squares problem:

$$\min_{s \in \mathbb{R}^n} = \|\text{grad } s - Y\|_2^w.$$

The normal equations for the problem are given by:

$$\text{div} \circ \text{grad } (s) = \text{div } Y \quad \Longrightarrow \quad \Delta_0(s) = \text{div } Y.$$

We recall that the graph Laplacian on the space of edge flows with a weighted inner product given by weight $[W_{ij}]$ is given by the matrix

$$[\Delta_0]_{ij} = \begin{cases} -w_{ij} & \text{if } \{v_i, v_j\} \in E \\ \sum_k w_{ik} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}.$$

If we compare the Laplacian Δ_0 with M, we see that $M = \Delta_0$. Similarly, the vector $\text{div } Y$ is precisely the same as Massey's vector p. Hence, Massey's normal equation $Mr = p$ is exactly the normal equations we get from solving the least squares problem in *HodgeRank*. This implies that that the ratings determined via *Massey's method* and *HodgeRank* (using the matrices Y and W) should be the same. However, as we have seen *HodgeRank* allows us to use topology and graph theory to further analyze our ranking data by means of the residual R.

4.3 Numerical Examples

In this section, we discuss the actual implementation of *HodgeRank*. In particular, we have used *HodgeRank* on two different datasets coming from college sports. The first dataset we considered were game results from the 2012-13 NCAA Division I Football season, and the second dataset we considered was the 2012-13 NCAA Div. I Basketball season.

4.3.1 Small Examples

Before we apply *HodgeRank* to larger datasets, we give two smaller examples that make up a subset of one of our larger examples. In these examples, the vertex sets are small enough that we can explicitly show the pairwise comparison graphs corresponding to the least squares projection and residual.

Example 1: Mostly Consistent Data

Suppose we wish to rate the six NCAA basketball teams: Mercer, Wake Forest, Furman, Davidson, New Mexico, George Mason. We can assign each team a number 1 through 6 as below, then we can form our weight matrix W and pairwise comparison matrix Y as described in section (4.2.4) but only using games between these six teams. Note that most pairs of teams have not played, and only one pair of teams have played twice, namely Davidson and Furman.

1	Mercer
2	Wake Forest
3	Furman
4	Davidson
5	New Mexico
6	George Mason

$$W = \begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Figure 4.5: List of teams along with weight matrix.

Below is the corresponding pairwise comparison matrix. By inspection, we see that our data seems to be ordinally consistent. That is, on each of the two triangles in our graph, we do not get any ordinal cyclic relations. However, each of these two triangles are cardinally inconsistent; the absolute value of the curl around the triangles $\{1, 2, 3\}$ and $\{1, 5, 6\}$ are 6 and 14, respectively.

$$\begin{bmatrix} 0 & -3 & 27 & 0 & -18 & -3 \\ 3 & 0 & 24 & 0 & 0 & 0 \\ -27 & -24 & 0 & -27.5 & 0 & 0 \\ 0 & 0 & 27.5 & 0 & -5 & 0 \\ 18 & 0 & 0 & 5 & 0 & 1 \\ 3 & 0 & 0 & 0 & -1 & 0 \end{bmatrix}$$

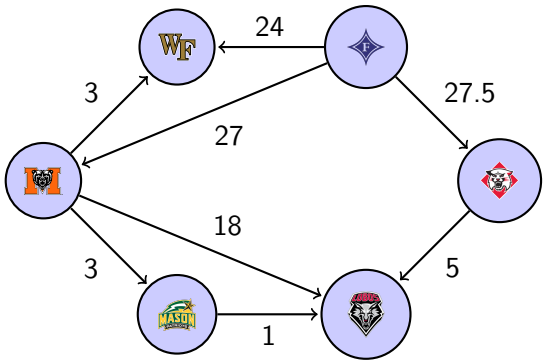

Figure 4.6: Pairwise comparison matrix and corresponding graph

By solving the least squares problem (4.2), we can find the l_2-projection of our pairwise comparison data onto the space of gradient flows. In Figure (4.7), we see the least squares projection of our game results on the space of gradient flows and the residual. By comparison with our original data in Figure 4.6, we see that our gradient projection appears to reflect our data well.

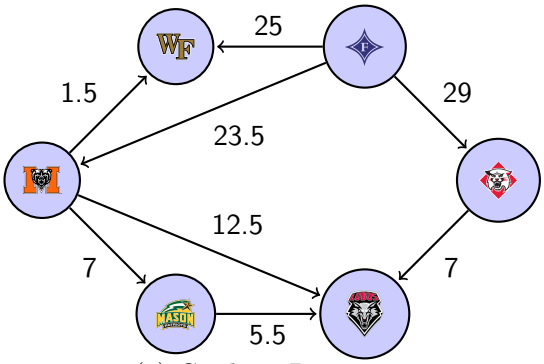

(a) Gradient Projection

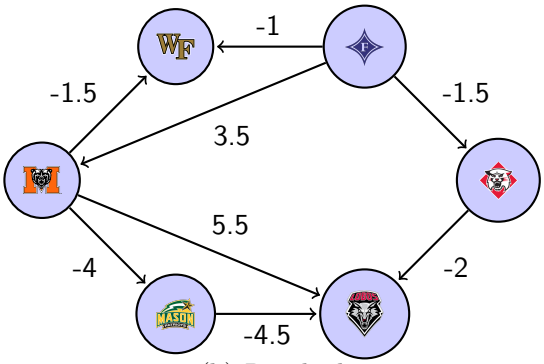

(b) Residual

Figure 4.7: Gradient component and Residual

We note that the gradient component is much larger than the residual. This implies that our original data was relatively consistent, and that the inconsistent and locally consistent components are small. Hence, we can expect that the rating obtained by solving grad (s) for s reflects our original data well.

Team	Score
N.Mexico	82
G. Mason	76.5
Davidson	75
Wake	71
Mercer	69.5
Furman	46

Figure 4.8: Ratings via *HodgeRank*

You may recall from an earlier section, that we described a crude comparison of the gradient and residual called the *cyclicity ratio*, which is given by

$$C_r = \frac{\|R\|}{\|\text{grad }(s)\|}.$$

Ideally, the gradient should make up a much larger component of our data, so we should hope $C_r \ll 1$. In the above example, $C_r \approx 0.2$, meaning the gradient component is about five times as large as the residual. So we can be somewhat confident in our ratings.

Example 2: Inconsistent Data

In the last subsection, we gave an example of a small ranking problem with relatively consistent data. In this example, we show how inconsistent data (with large cyclicity ratios), does not lend itself to good ratings. Suppose that we would like to rate the six ACC teams: Miami, NC State, Wake Forest, Maryland, UNC, and Duke. In this example, only the earliest results were used, and some later games had been omitted to reduce clutter in our graphs. We may imagine that these results represent a time early on in the season before each team has played each other.

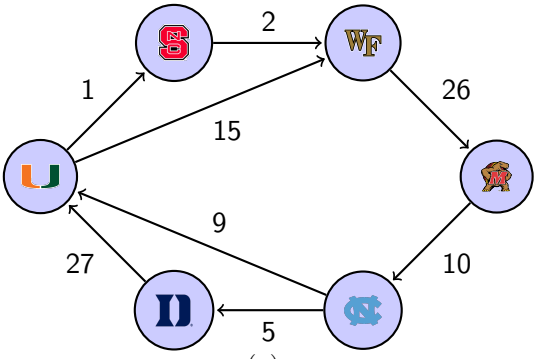

(a)

1	Miami
2	NC State
3	Wake Forest
4	Maryland
5	North Carolina
6	Duke

(b)

Figure 4.9: Pairwise Comparison Graph for Example 2.

By inspection, we see that both triangles are ordinally consistent, however there is a longer ordinal cyclic relation given by

$$1 < 2 < 3 < 4 < 5 < 6 < 1.$$

The absolute curl on $\{1, 2, 3\}$ and $\{4, 5, 6\}$ is 13 and 23 respectively, however along the closed loop $1 - 2 - 3 - 4 - 5 - 6 - 1$ the net flow is 71. By inspection, there does not seem to be a clear choice of which teams should be rated higher than which. However, we can still apply the same methods as before to compute the l_2-projection our data on the gradient flows. The gradient component and the residual are given below.

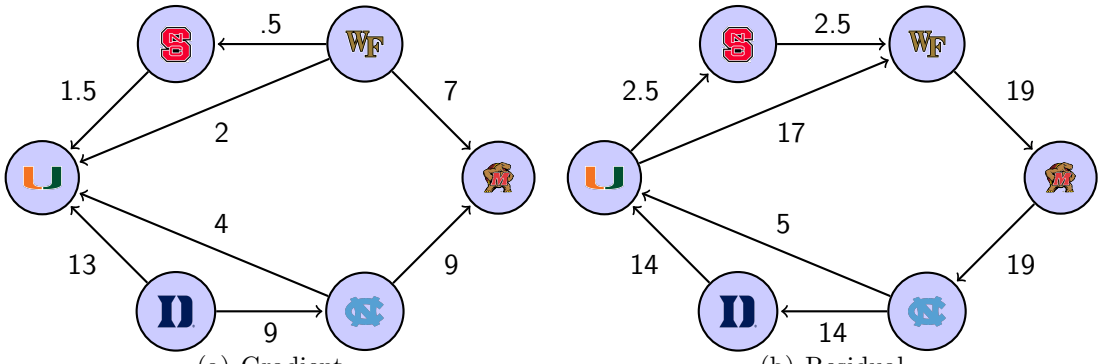

(a) Gradient (b) Residual

Figure 4.10: Cyclic and Noncyclic Components

We see that in this case, there appears to be more flow in the inconsistent component (residual) than in the consistent component (gradient). The cyclicity ratio for this decomposition is $C_r \approx 1.9$, so we should not expect a very reliable rating given our score vector s. The individual ratings for these teams are given below, juxtaposed with our original data:

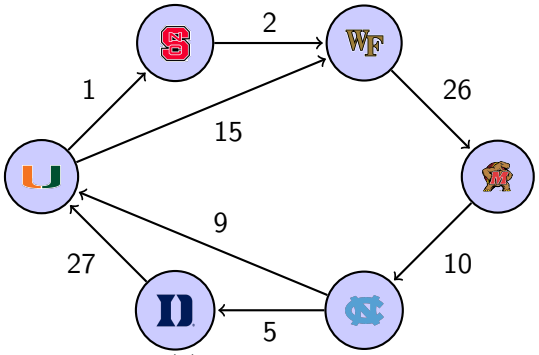

(a) Original Data

Team	Score
Maryland	77.5
Miami	72.5
NSCU	71
Wake	70.5
UNC	68.5
Duke	59.5

(b) Ratings

Figure 4.11: Comparison of original data with *HodgeRank* ratings.

By comparing the ratings with what our data may suggest, we can see that there are many discrepancies. For example, in the original data the highest rated team (Maryland) lost to the second lowest rated team (UNC) who in turn lost to the lowest rated team (Duke). From this one could argue that Duke should be rated higher than Maryland. However, a similar argument can be made for just about every one of these six teams. In the larger datasets, we cannot visualize inconsistencies like we can in these two examples, but we can still use the cyclicity ratio.

4.3.2 NCAA Football

The first large dataset we will consider is the game results for the 2012-13 NCAA Division I Football season. In this example, there are a total of 246 colleges with Div. I football teams and a total of 2944 league games. In the regular season, each team plays a total of 12 games against 12 different opponents, except for some teams which

qualify for the bowl games and championships. In this case, a team plays a 13^{th} game which could potentially be against a team they have already played. In the 2012-13 season, only 12 of the 246 teams played the same team twice (a total of 6 rematch games).

Suppose we have assigned to each team a number between 1 and 246, then we may form the weight matrix W in the usual way by letting W_{ij} be the number of games played between teams i and j. In this example, W_{ij} is a very sparse matrix, as there are only 12 or 13 nonzero elements in each row and column. Furthermore, almost all of the nonzero entries of W_{ij} are 1 since there were so few rematches.

For our pairwise comparison matrices, we use both the average margin of victory Y and the binary comparisons B as described in section (4.2.4). In both cases, the weight matrix W is the same as above. To determine our rankings, we solve the l_2-minimization problems:

$$\min_{s\in\mathbb{R}^n} = \|\text{grad } s - Y\|_2^w \quad \text{and} \quad \min_{s\in\mathbb{R}^n} = \|\text{grad } s - B\|_2^w.$$

Below we list the top 10 NCAA Div. I football teams as determined by *HodgeRank* using the binary comparison matrix B. In each case, all games through the bowl series were included.

Rank	Team	Rating
1	Stanford	1.6294
2	Notre Dame	1.6038
3	Florida	1.5542
4	Alabama	1.5432
5	Oregon	1.4959
6	Ohio St.	1.4857
7	Texas A&M	1.4642
8	South Carolina	1.4438
9	Georgia	1.4085
10	Kansas St.	1.3716

Figure 4.12: Top 10 NCAA Football Teams (Binary)

For the complete rankings of all 246 teams we refer the reader to Appendix C. In the next table, we list the top 10 football teams as determined by average margin of victory.

Rank	Team	Rating
1	Alabama	49.0962
2	Oregon	47.8779
3	Texas A&M	44.9811
4	Kansas St.	40.9600
5	Georgia	39.1104
6	Oklahoma St.	37.4924
7	Oklahoma	37.1732
8	Notre Dame	36.2345
9	Florida	35.2030
10	Florida St.	35.1155

Figure 4.13: Top 10 NCAA Football Teams (Avg. MOV)

For comparison, the Associated Press Poll had Alabama, Oregon, Ohio State, Notre Dame (Tied), Texas A&M (Tied) as the top five teams at the end of the 2012-13 season. In this example, the cyclicity ratio was approximately 1, which suggests that there is a good amount of inconsistency in our data. However, since our rankings are comparable to other established rankings, there may not be enough inconsistency

in this data to negatively affect our ratings on a large scale.

4.3.3 NCAA Basketball

The second dataset we considered were the game results for the 2012-13 NCAA Division I basketball season. There are a total of 347 colleges with Div. I basketball teams and a total of 5320 league games. Each team begins the season with approximately 12 games against non-conference opponents or occasionally non-divisional opponents (in which case, we omit the game result). After these non-conference games, the remaining regular season games are played between conference opponents. By this point in the season, each team will have played around 30 games. Each conference then holds a tournament in which some teams may play up to 4 or 5 games, depending on how successful they are in the tournament.

In general, there will be no rematches between non-conference opponents in the regular season. However, within each conference, most opponents will have played each other during conference play, sometimes twice. If we include games from the conference tournaments, it is possible that some opponents would have played three times. In football, so few games are played that even teams within a conference will not necessarily play each other. However, in college basketball, individual conferences will often form cliques in our pairwise comparison graph. That is, within most conferences every pair of teams will play each other at least once.

Suppose we have assigned to each team a number between 1 and 347, then we may form the weight matrix W in the usual way by letting W_{ij} be the number of games played between teams i and j. In this example, W_{ij} is also very sparse, between 29 and 35 nonzero elements in each row and column. The nonzero entries of W_{ij} are

either $1, 2$ or 3, since no two teams had played each other more than three times in the regular season.

For our pairwise comparison data, we will again use both the average margin of victory Y and the binary comparison matrix B as described in (4.2.4). Below are the top 10 NCAA Div. I basketball teams as determined by *HodgeRank* using the binary comparison matrix B.

Rank	Team	Rating
1	New Mexico	1.1796
2	Duke	1.1565
3	Louisville	1.1540
4	Kansas	1.1058
5	Miami FL	1.0839
6	Gonzaga	1.0820
7	Indiana	1.0730
8	Ohio St	1.0206
9	Georgetown	1.0200
10	Michigan St	1.0046

Figure 4.14: Top 10 NCAA Basketball Teams (Binary)

In the next table, we list the top 10 basketball teams rated using the average margin of victory matrix Y.

Rank	Team	Rating
1	Indiana	26.3030
2	Florida	25.0507
3	Louisville	24.6464
4	Duke	22.2694
5	Gonzaga	21.9128
6	Kansas	21.5680
7	Ohio St	20.8467
8	Michigan	20.5705
9	Pittsburgh	20.0057
10	Syracuse	19.7535

Figure 4.15: Top 10 NCAA Basketball Teams (Avg. MOV)

For comparison, the Associated Press had Gonzaga, Duke, Indiana, Louisville, Georgetown as the top five teams at the conclusion of the regular season. In this example, the cyclicity ratio of the residual R to our least squares projection grad (s) was approximately 1. Again, we found that our ratings were comparable to other established ranking methods.

Chapter 5: Conclusion

In their original paper, Jiang, Lim, Yao and Ye suggested that *HodgeRank* may be well-suited to ranking datasets which are incomplete or imbalanced, meaning that most values are missing or that some alternatives are rated more frequently than others. We have shown that *HodgeRank* can be applied very naturally to sports rating problems. In particular, *HodgeRank* generalizes the popular sports rating method known as *Massey's method*. Massey's method imposes extra conditions on these ratings so that a unique rating is guaranteed to exist. In *HodgeRank*, we can classify all of our potential ratings by looking at the homology of the underlying graph. If the graph is connected, our ratings are unique up to an additive constant. If the graph has vanishing homology, then local consistency in our data implies global consistency.

As this discussion suggests, the novelty of *HodgeRank* is in its use of topology and graph theory to determine information about the ranking data. However, in the examples considered in this thesis, we have considered only the simplest pairwise comparisons, namely margin of victory. A future direction for research would be to apply *HodgeRank* to more advanced statistics and pairwise comparisons. The more established *Massey's method* is already widely used and adapted to other contexts, so it seems likely that *HodgeRank* could be used in these same instances.

Another aspect of this work that may merit further research would be analyzing various subgraphs of our larger datasets. In both of our larger examples, it is impossible to visualize inconsistency as we could in the smaller six-team examples. However, as we saw when we calculated cyclicity ratios, both the consistent and inconsistent

components of our data were roughly the same size. At first this is rather alarming, however as we saw, we can still determine a reasonable global ranking of teams. This suggests that maybe the cyclicity ratio is not very meaningful on larger datasets. We may be able to get more meaningful information about cyclicity by examining smaller subgraphs such as individual conferences.

If our data has a relatively large inconsistent component, does this say anything about the prospects of good predictive rankings? For instance, should we expect more upsets and variability in a post-season tournament if our regular season data had many inconsistencies? Can we meaningfully assign an inconsistency rating to a single team or game?

Another point of interest that could direct future research would be to study ranking interpretations of higher dimensional chains/cochains and boundary/coboundary maps. In *HodgeRank*, we interpret the first coboundary map as taking ratings to pairwise comparisons and the second coboundary map as giving a measure of local consistency of our ranking data. Do the higher dimensional boundary/coboundary maps have any meaningful ranking theoretic interpretation?

We have shown how *HodgeRank* use graph theory and topology to both determine a global rating and to perform a meta-analysis of sports ranking data. The latter function of *HodgeRank* is particularly intriguing and provides many directions for further research.

Bibliography

[1] Arfken, George B. and Weber, Hans J. *Mathematical Methods for Physicists*. Elsevier Academic Press (2005).

[2] Axler, Sheldon. *Linear Algebra Done Right*. Springer (1997)

[3] Bau, David and Trefethen, Lloyd N. *Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics (1997)

[4] Canterella, J., DeTurck, D. and Gluck H. *Vector Calculus and the Topology of Domains in 3-Space*. The American Mathematical Monthly, Vol. 109, No. 5, (May, 2002), pp. 409-442

[5] Demmel, James W. *Applied Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics (1997)

[6] Evans, Lawrence C. *Partial Differential Equations*. American Mathematical Society. (1998)

[7] Flanders, Harley. *Differential Forms with Applications to the Physical Sciences*. Dover Publications. (1989)

[8] Folland, Gerald B. *Introduction to Partial Differential Equations. 2nd Ed*. Princeton University Press. (1976)

[9] Garabedian, Paul R. *Partial Differential Equations*. John Wiley & Sons, Inc. (1964)

[10] Gilbarg, David and Trudinger, Neil S. *Elliptic Partial Differential Equations of the Second-Order*. Springer-Verlag (1977)

[11] Goldberg, Timothy E. *Combinatorial Laplacians of Simplicial Complexes.* (Senior Thesis) Bard College (2002).

[12] Griffiths, David J. *Introduction to Electrodynamics, 3rd Ed.* Prentice-Hall. (1999)

[13] Hatcher, Allen. *Algebraic Topology.* Cambridge University Press. (2001)

[14] Jiang, X., Lim. L-H., Yao, Y. and Ye, Y. *Statistical Ranking and combinatorial Hodge theory.* Mathematical Programming Series B. Vol. 127 pp. 203-244. (2008)

[15] Langville, Amy and Meyer, Carl. *Who's #1?: The Science of Rating and Ranking.* Princeton University Press. (2012)

[16] Lee, John M. *Introduction to Smooth Manifolds.* Springer-Verlag New York, Inc. (2003)

[17] Lee, John M. *Introduction to Topological Manifolds.* Springer-Verlag New York, Inc. (2007)

[18] Munkres, James R. *Elements of Algebraic Topology.* Addison-Wesley Publishing Company. (1984)

[19] Munkres, James R. *Topology, 2nd Ed.* Prentice-Hall. (2000)

[20] Spivak, Michael. *Calculus on Manifolds.* Westview Press. (1998)

[21] Strauss, Walter A. *Partial Differential Equations: An Introduction.* John Wiley & Sons. (2008)

[22] Tu, Loring W. *An Introduction to Manifolds. 2nd edition* Springer. (2010)

[23] Warner, Frank W. *Foundations of Differentiable Manifolds and Lie Groups.* Scott, Foresman and Company. (1971)

Appendix A: Linear Algebra

In this appendix, we outline several results from linear algebra. It will be assumed throughout that all vector spaces are finite-dimensional vector spaces over $\mathbb{R}$ that have been endowed with an inner-product. We begin by introducing linear functionals, adjoints and dual spaces, which will be useful in our discussion of (co)chain groups. Next, we will introduce the general least squares problem, normal equations and the Moore-Penrose psuedoinverse. It is beyond the scope of this thesis to work through all of these concepts in detail, however we hope this appendix serves as a useful reference. See [2], [3], or [5] for a full treatment.

A.1 Dual Spaces and the Adjoint

In this section, we may adopt the notation $\mathcal{L}(V, W)$ for the vector space of linear operators from V to W. If V is a real vector space, then a **linear functional** is a linear map $\varphi : V \to \mathbb{R}$. The vector space of linear functionals $\mathcal{L}(V, \mathbb{R})$ is called the **dual space** of V and is sometimes denoted V^*. The first result we consider states that every linear functional on a finite-dimensional inner product space can be represented using the inner product.

Proposition A.1.1. *If $\varphi \in V^*$ is a linear functional on V, then there is a unique vector $v \in V$ such that*

$$\varphi(u) = \langle u, v \rangle ,$$

for every $u \in V$

With this proposition, we can can more easily describe the dual space V^*. In particular, for a given basis for V we can find a dual basis for V^* using this characterization of linear functionals.

Proposition A.1.2. ***(Dual Space)*** *Let V be a vector space and $V^* = \mathcal{L}(V, \mathbb{R})$ be the* ***dual space*** *of V. If $\{e_1, \cdots, e_n\}$ is a basis for V, then a basis for V^* is given by the set of linear functionals $\varphi_1, \cdots, \varphi_n$ such that:*

$$\varphi_i(v) = \langle v, e_i \rangle ,$$

for each $v \in V$. So, in particular, $V \cong V^*$.

The idea behind this proposition is that a linear map $L : V \to W$ is determined by its action on a basis for V. In particular, a linear functional φ takes each basis vector e_i to some real number c_i, we can then write φ as the linear combination:

$$\varphi = c_1\varphi_1 + \cdots + c_n\varphi_n.$$

We leave it to reader to verify that every linear functional can be written this way and that this expression is unique. The next concept we will introduce is that of an adjoint operator, if a linear operator between two vector spaces is represented by a matrix A, then the adjoint is given by the matrix transpose.

Definition 19. (Adjoint) Let $L \in \mathcal{L}(V, W)$. Then the **adjoint** L^* of L, is defined to be the unique operator such that for all $v \in V$ and $w \in W$:

$$\langle Lv, w \rangle = \langle v, L^*w \rangle .$$

The next proposition provides a very useful orthogonal decomposition, which will play an integral role in the proof of the combinatorial Hodge Decomposition Theorem.

Theorem A.1. *Let $T \in \mathcal{L}(V, W)$ and $T^* \in \mathcal{L}(W, V)$ denote the adjoint of T. Then V and W decompose orthogonally as:*

$$V = \ker(T) \oplus \mathrm{im}(T^*) \quad \textit{and} \quad W = \mathrm{im}(T) \oplus \ker(T^*)$$

Proof. Let $v \in V$. We want to show that $v \in \ker(T) \Leftrightarrow v \in (\mathrm{im}(T^*))^\perp$ and therefore $V = \ker(T) \oplus \mathrm{im}(T^*)$. The following chain of equivalent statements shows the result:

$$\begin{aligned} v \in \ker(T) &\Leftrightarrow Tv = 0 \\ &\Leftrightarrow \langle Tv, w \rangle_W = 0 \text{ for all } w \in W \\ &\Leftrightarrow \langle v, T^* w \rangle_V = 0 \text{ for all } w \in W \\ &\Leftrightarrow v \in (\mathrm{im}(T^*))^\perp \end{aligned}$$

Since $(T^*)^* = T$, the same argument with the roles of T^* and T switched shows us that $W = \mathrm{im}(T) \oplus \ker(T^*)$. □

The final proposition we state in this section is a simple lemma that will be used in the corollary to the Hodge Decomposition Theorem.

Lemma A.1.1. *Let V be a vector space, and U, W be subspaces of V such that*

$$V = U \oplus W.$$

Then,

$$W \cong V/U.$$

Proof. Let $V = U \oplus W$, and $A : V \to U$ be the projection map onto the subspace U. Clearly, A is surjective since it restricts to the identity on U, and by construction $W = \ker A$. Therefore, by the First Isomorphism Theorem:

$$\ker A \cong V/\mathrm{im}A \implies W \cong V/U$$

□

A.2 Least Squares

In this section, we describe a basic least squares problem and its solution using normal equations. It is beyond the scope of this thesis to describe the theory of least squares in its full detail, see [3] or [5] for a more detailed exposition on least squares.

Suppose we are given a matrix equation $Ax = b$ where A is an m-by-n matrix with $m > n$. This gives a systems of m linear equations in n unknowns. Since $m > n$, it may not be the case that for a given $b \in \mathbb{R}^m$ that $Ax = b$ has a solution. So, instead of seeking an explicit solution, we seek a nearest match, namely we wish to find an $\hat{x}$ such that

$$\|A\hat{x} - b\| \leq \|Ax - b\|,$$

for all $x \in \mathbb{R}^n$. Alternatively, we seek a $\hat{x}$ that is a solution to the minimization problem:

$$\min_{x \in \mathbb{R}^n} \|Ax - b\|. \tag{A.1}$$

In the following proposition, we introduce the normal equations. The solutions of the normal equations are precisely the solutions of our least squares problem. For a derivation of the normal equations, see [5].

Proposition A.2.1. *The vector $\hat{x}$ is a solution of (A.1) iff $\hat{x}$ is a solution of the matrix equation:*

$$A^T A\hat{x} = A^T b,$$

*this equation is known as the **normal equation** of the least squares problem (A.1). If A is full-rank, then $A^T A$ is nonsingular, and so our solution is given by:*

$$\hat{x} = (A^T A)^{-1} A^T b. \tag{A.2}$$

If A is full-rank, there is a unique solution x given by (A.2). Otherwise, as we will see in the next proposition, we get another solution for each element of $\ker A$. Unfortunately, the matrices and linear maps we will consider in this thesis are generally rank-deficient.

Proposition A.2.2. *Suppose A is m-by-n with $m > n$ and that* rank $A = r < n$. *Then the solution set of*

$$\min_{x \in \mathbb{R}^n} \|Ax - b\|,$$

is a $(n - r)$-dimensional set.

Proof. If rank $A = r < n$, then $\ker A$ is a $(n - r)$- dimensional set. Let $z \in \ker A$, then for any solution x that is a solution of the minimization problem:

$$\min_{x \in \mathbb{R}^n} \|Ax - b\|.$$

The vector $x + z$ is as well, since

$$\|A(x + z) - b\| = \|Ax + Az - b\| = \|Ax - b\|.$$

Hence, if $\ker A = \{z_1, \cdots, z_{n-r}\}$ the set of solutions is given by $\{x + z_1, \cdots, x + z_{n-r}\}$. □

The best we can do is find a minimum norm solution, that is, a solution x to (A.1) such that $\|x\|$ is minimized over all solutions. To describe this solution, we need to introduce the Moore-Penrose psuedoinverse. In the next definition, we will assume familiarity with the singular value decomposition (SVD). See [3] for a full discussion of the SVD and its properties.

Definition 20. (Psuedoinverse) Suppose that A is m-by-n with $m > n$, and that A is rank-deficient with rank $A = r < n$. Consider the singular value decomposition

$A = U\Sigma V^T$, where U, V are unitary. Since A is rank-deficient we can write this in block form as:

$$A = [U_1, U_2] \begin{bmatrix} \Sigma_1 & 0 \\ 0 & 0 \end{bmatrix} [V_1, V_2]^T = U_1 \Sigma_1 V_1^T.$$

Where Σ_1 is r-by-r and U_1, V_1 have r columns. Then the **Moore-Penrose pseudoinverse** of A is given by:

$$A^\dagger = U_1 \Sigma_1^{-1} V_1^T.$$

Although this definition seems appears complicated, the singular value decomposition is a standard matrix decomposition. For our purposes, the Moore-Penrose pseudoinverse will be most useful for the following fact.

Proposition A.2.3. *Suppose that A is m-by-n with $m > n$ and that A is rank-deficient. Then the minimum-norm solution to the least squares problem*

$$\min_{x \in \mathbb{R}^n} \|Ax - b\|,$$

with normal equations $A^T Ax = A^T b$ is given by:

$$\hat{x} = (A^T A)^\dagger A^T b.$$

Hence, by the previous proposition, for any $z \in \ker A$, $x + z$ is also a solution.

Appendix B: MATLAB & Ruby Code

The implementation of *HodgeRank* is relatively simple. To find a rating, we need to form the matrices W and Y containing our weights and pairwise comparisons. The Laplacian matrix Δ_0 and the vector $\text{div}(Y)$ are easily formed from W and Y. The hardest part is not the actual implementation, but rather obtaining the data and preprocessing it. Assuming we can find the data in a reasonably convenient format, we would still have to do some processing in order to determine W and Y. For all of the preprocessing, we used the Ruby .CSV libraries. The actual implementation of *HodgeRank*, which was done in Octave, assumes that Y and W have already been formed correctly and stored as comma-separated files (.CSV). For completeness, we have included all code used in our basketball ratings, however it would need to be altered in order to be used.

The following Ruby code was used to form our weight matrix, pairwise comparison matrix, and binary comparison matrix. It assumes we have raw data in a particular format stored in the file 'Mar18.csv' (indicating game results through this date), and a file 'masseydiv1.csv' which matches each team to a number $1 - 347$.

```
#
# Create Comparison Array
#

require 'csv'
require 'matrix'

data = CSV.read('Mar18.csv')
teams = CSV.read('masseydiv1.csv')

temp = Array.new
temp = teams.transpose.first
```

```
mat = Matrix.zero(teams.length).to_a
data.each do |arr|
  if mat[temp.index(arr[2])][temp.index(arr[5])] == 0
    a = arr[4].to_i - arr[7].to_i
    mat[temp.index(arr[2])][temp.index(arr[5])] = -a
    mat[temp.index(arr[5])][temp.index(arr[2])] = a
  elsif mat[temp.index(arr[2])][temp.index(arr[5])]
  .kind_of?(Array)
    a = arr[4].to_i - arr[7].to_i
    mat[temp.index(arr[5])][temp.index(arr[2])] =
    mat[temp.index(arr[5])][temp.index(arr[2])] << a
    mat[temp.index(arr[2])][temp.index(arr[5])] =
    mat[temp.index(arr[2])][temp.index(arr[5])] << -a
  else
    a = arr[4].to_i - arr[7].to_i
    mat[temp.index(arr[5])][temp.index(arr[2])] =
    [mat[temp.index(arr[5])][temp.index(arr[2])]] << a
    mat[temp.index(arr[2])][temp.index(arr[5])] =
    [mat[temp.index(arr[2])][temp.index(arr[5])]] << -a
  end
end

CSV.open('comparisons.csv', 'w') do |csv|
  mat.each { |i| csv << i }
end

#
# Create Weight Matrix
#

weight = Array.new
mat.each do |arr|
  temp = Array.new
  arr.each do |j|
    if j.kind_of?(Array)
      temp << j.length
    elsif j.to_i == 0
      temp << j
    else temp << 1
    end
  end
```

```
  weight << temp
end

CSV.open('weights.csv', 'w') do |csv|
  weight.each { |i| csv << i }
end


#
# Create Binary Comparison Matrix
#

sgn = Array.new
mat.each do |arr|
  temp = Array.new
  arr.each do |j|
    if j.kind_of?(Array)
      tmp = Array.new
      j.each do |i|
        tmp << (i.to_i <=> 0)
      end
      temp << tmp.inject(:+).to_f/j.length.to_f
    else temp << (j.to_i <=> 0)
    end
  end
  sgn << temp
end

CSV.open('binary.csv', 'w') do |csv|
  sgn.each { |i| csv << i }
end

#
# Create Margin of Victory Matrix
#


avg = Array.new
mat.each do |arr|
  temp = Array.new
  arr.each do |j|
    if j.kind_of?(Array)
```

```
        temp << j.inject(:+).to_f/j.length.to_f
      else temp << j
      end
    end
    avg << temp
end

CSV.open('average.csv', 'w') do |csv|
  avg.each { |i| csv << i }
end
```

The formatted text that follows is the MATLAB/Octave source code for implementing *HodgeRank*. The files 'weights.csv', 'binary.csv' and 'average.csv' were created by the previous Ruby source code. The actual score vector determined is stored as 'scorebinary.csv' or 'scoreavg.csv'.

```
# Binary Comparisons

W = csvread('weights.csv');
L = -W;
L = L + diag(sum(W));
Y = csvread('binary.csv');
s = diag(Y*W');
D = pinv(L);
r = -D*s;
csvwrite('score_binary.csv',r);

# Average Score Differential

W = csvread('weights.csv');
L = -W;
L = L + diag(sum(W));
Y = csvread('average.csv');
s = diag(Y*W');
D = pinv(L);
r = -D*s;
csvwrite('score_avg.csv',r);
```

Given the rating determined by the previous code, we would run a final Ruby script (given below) which converts our rating vector into a more readable format.

The output is .CSV file with three columns, the first column is the number assigned to each team, the second column is the team name, and the third column is the *HodgeRank* rating. The file is sorted so that the teams are listed in descending order by rating.

```
#
# Rank via Margin of Victory
#

require 'csv'

score = CSV.read('score_avg.csv')
teams = CSV.read('masseydiv1.csv')

rank = Array.new

rank = teams.transpose + score.transpose
rank = rank.transpose
rank = rank.sort_by { |c| c[2].to_f }
rank = rank.reverse

CSV.open('rankings_avg.csv', 'w') do |csv|
  rank.each { |i| csv << i }
end

#
# Rank via Binary Comparisons
#

require 'csv'

score = CSV.read('score_bin.csv')
teams = CSV.read('masseydiv1.csv')

rank = Array.new

rank = teams.transpose + score.transpose
rank = rank.transpose
rank = rank.sort_by { |c| c[2].to_f }
rank = rank.reverse
```

```
CSV.open('rankings_bin.csv', 'w') do |csv|
  rank.each { |i| csv << i }
end
```

Appendix C: Complete Football and Basketball Rankings

In this appendix, we list the full NCAA football and basketball rankings. We used these basketball rankings to fill out a NCAA Tournament bracket, by taking the team with the higher rating in each match up. For our margin of victory pairwise comparison, we predicting a final four of Louisville, Gonzaga, Florida and Indiana with Indiana beating Louisville in the championship game. Our binary comparison bracket had Duke, New Mexico, Kansas and Miami in the final four with New Mexico beating Kansas.

Table C.1: Full Ranking NCAA Div. I Football (Margin of Victory)

Rank	Team	Rating	Rank	Team	Rating
1	Alabama	49.1	124	Central Mich.	0.82
2	Oregon	47.88	125	Fla. Atlantic	0.8
3	Texas A& M	44.98	126	Harvard	0.75
4	Kansas St.	40.96	127	Appalachian St.	0.59
5	Georgia	39.11	128	Villanova	0.14
6	Oklahoma St.	37.49	129	UAB	-0.1
7	Oklahoma	37.17	130	Buffalo	-0.39
8	Notre Dame	36.23	131	Eastern Ky.	-1.01
9	Florida	35.2	132	Chattanooga	-1.02
10	Florida St.	35.12	133	UNLV	-1.09
11	South Carolina	34.61	134	UTSA	-1.14
12	Oregon St.	34.42	135	Richmond	-1.17
13	Stanford	34.34	136	Missouri St.	-1.67
14	LSU	34.25	137	Army	-2.24
15	Baylor	32.71	138	James Madison	-2.35
16	Arizona St.	32.01	139	McNeese St.	-2.87
17	Clemson	31.6	140	Northern Ariz.	-3.26
18	Texas	31.35	141	Southern Utah	-3.28
19	Ohio St.	30.78	142	Samford	-3.38
20	Wisconsin	30.66	143	Miami (OH)	-3.51
21	Southern California	30.21	144	Maine	-3.78
22	Michigan	29.74	145	Montana	-3.8
23	UCLA	28.46	146	Tenn.-Martin	-3.95
24	Utah St.	27.81	147	Akron	-4.59
25	BYU	27.71	148	Citadel	-5.08
26	TCU	26.65	149	Tulane	-5.16
27	Ole Miss	26.56	150	Eastern Ill.	-5.24
28	Texas Tech	26.09	151	Hawaii	-5.57
29	Nebraska	26.06	152	Southern Miss.	-5.59
30	Penn St.	25.61	153	South Ala.	-5.67
31	Northwestern	25.57	154	New Hampshire	-5.97
32	Vanderbilt	25.56	155	Colorado	-6.2
33	Arizona	24.52	156	Eastern Mich.	-6.27
34	West Virginia	24.34	157	Jacksonville St.	-6.93
35	Iowa St.	23.96	158	Coastal Caro.	-6.97
36	Michigan St.	23.91	159	Albany (NY)	-7.15
37	North Carolina	23.78	160	Murray St.	-7.3
38	Fresno St.	23.71	161	Portland St.	-7.33
39	Missouri	23.56	162	Stephen F. Austin	-7.92
40	Boise St.	23.4	163	Bethune-Cookman	-8.06
41	Northern Ill.	23.39	164	North Dakota	-8.4
42	UCF	23.32	165	Tennessee St.	-8.41
43	Mississippi St.	23.29	166	UC Davis	-8.72
44	San Jose St.	22.24	167	Sacramento St.	-9.44
45	Syracuse	22.1	168	Delaware	-9.69
46	Cincinnati	21.58	169	Furman	-9.75
47	Tulsa	21.42	170	William& Mary	-9.96
48	Georgia Tech	21.37	171	Colgate	-10.42
49	Louisiana Tech	21.18	172	South Dakota	-11.24
50	Washington	21.15	173	New Mexico St.	-12.12

51	North Dakota St.	20.94	174	Weber St.	-12.32
52	Utah	20.72	175	Princeton	-12.53
53	Louisville	20.11	176	Lehigh	-12.77
54	Tennessee	19.85	177	Idaho	-13.04
55	Miami (FL)	19.22	178	Northern Colo.	-13.43
56	SMU	18.42	179	Wagner	-13.43
57	Arkansas St.	18	180	Liberty	-13.47
58	Pittsburgh	17.77	181	Elon	-13.66
59	San Diego St.	17.7	182	Northwestern St.	-16.09
60	Virginia Tech	17.03	183	Brown	-16.19
61	Rutgers	16.79	184	Penn	-16.56
62	North Carolina St.	16.29	185	Dartmouth	-16.77
63	Arkansas	16.27	186	Southeast Mo. St.	-17.27
64	Iowa	15.31	187	Fordham	-17.3
65	Purdue	15.15	188	San Diego	-17.68
66	La.-Lafayette	14.94	189	Tennessee Tech	-17.69
67	California	14.65	190	N.C. A& T	-17.7
68	Kent St.	14.5	191	Western Caro.	-18.23
69	Sam Houston St.	14.1	192	Massachusetts	-18.43
70	Minnesota	13.82	193	Delaware St.	-18.72
71	La.-Monroe	13.27	194	Alabama St.	-18.84
72	Auburn	12.48	195	Drake	-19.02
73	Nevada	11.75	196	Southeastern La.	-19.55
74	Toledo	11.33	197	Florida A& M	-20.16
75	Ball St.	11.11	198	Western Ill.	-20.42
76	Duke	10.81	199	Monmouth	-20.52
77	Ga. Southern	10.36	200	Duquesne	-20.8
78	Rice	10.08	201	St. Francis (PA)	-20.94
79	Indiana	10.01	202	Ark.-Pine Bluff	-21.69
80	Virginia	9.74	203	N.C. Central	-21.81
81	Western Ky.	9.44	204	South Carolina St.	-21.83
82	South Fla.	8.97	205	Lafayette	-22.36
83	Kansas	8.95	206	Cornell	-22.37
84	UNI	8.73	207	Howard	-22.75
85	East Carolina	8.48	208	Holy Cross	-22.85
86	Kentucky	8.43	209	Jackson St.	-23.62
87	Washington St.	8.35	210	Dayton	-25
88	Troy	8.33	211	Georgetown	-25.69
89	Connecticut	8.21	212	Jacksonville	-25.9
90	Bowling Green	8.06	213	Alabama A& M	-26.06
91	Navy	7.99	214	Lamar	-26.18
92	Maryland	7.68	215	Robert Morris	-26.37
93	Boston College	7.37	216	Morehead St.	-26.63
94	Ohio	7.27	217	Georgia St.	-26.65
95	South Dakota St.	7.2	218	Bucknell	-27.22
96	Temple	6.88	219	Gardner-Webb	-27.4
97	Middle Tenn.	6.21	220	Norfolk St.	-28.41
98	Houston	6.17	221	Marist	-29.19
99	Montana St.	6.12	222	Sacred Heart	-29.24
100	Marshall	5.64	223	Butler	-29.33
101	Eastern Wash.	5.12	224	Nicholls St.	-29.37
102	Wofford	4.92	225	Bryant	-29.52
103	UTEP	4.7	226	Morgan St.	-29.62
104	Indiana St.	4.63	227	Hampton	-29.8
105	Cal Poly	4.62	228	Columbia	-30.02
106	Western Mich.	4.23	229	Yale	-30.32
107	Wyoming	3.03	230	Austin Peay	-30.83
108	FIU	2.91	231	Central Conn. St.	-31.06
109	Southern Ill.	2.82	232	Charleston So.	-31.27
110	Towson	2.75	233	Mississippi Val.	-32.57
111	Air Force	2.68	234	Southern U.	-33.51
112	Texas St.	2.6	235	VMI	-34.69
113	Old Dominion	2.6	236	Prairie View	-35.04
114	Youngstown St.	2.4	237	Rhode Island	-36.58
115	Wake Forest	2.37	238	Presbyterian	-40.33
116	Illinois St.	2.26	239	Idaho St.	-42.19
117	Illinois	1.87	240	Grambling	-42.59
118	Colorado St.	1.74	241	Alcorn St.	-42.63
119	Central Ark.	1.59	242	Savannah St.	-46.38
120	Stony Brook	1.54	243	Davidson	-47.19
121	Memphis	1.35	244	Texas Southern	-47.65
122	North Texas	1.34	245	Valparaiso	-56.15
123	New Mexico	1.24	246	Campbell	-57.05

Table C.2: Full Ranking NCAA Div. I Football (Binary)

Rank	Team	Rating	Rank	Team	Rating
1	Stanford	1.63	124	Towson	0
2	Notre Dame	1.6	125	Marshall	-0.01
3	Florida	1.55	126	UNI	-0.01

4	Alabama	1.54	127	North Texas	-0.03
5	Oregon	1.5	128	Kansas	-0.03
6	Ohio St.	1.49	129	Chattanooga	-0.04
7	Texas A&M	1.46	130	Buffalo	-0.05
8	South Carolina	1.44	131	Colorado St.	-0.05
9	Georgia	1.41	132	Air Force	-0.06
10	Kansas St.	1.37	133	Western Mich.	-0.07
11	LSU	1.3	134	Coastal Caro.	-0.07
12	Oklahoma	1.25	135	Southeastern La.	-0.09
13	Clemson	1.19	136	Lehigh	-0.09
14	Florida St.	1.09	137	Villanova	-0.09
15	Oregon St.	1.06	138	Illinois	-0.1
16	Nebraska	1.06	139	Wyoming	-0.1
17	Texas	1.04	140	Texas St.	-0.1
18	Louisville	1.03	141	Memphis	-0.12
19	North Dakota St.	1.02	142	Sacramento St.	-0.12
20	Northwestern	1.01	143	Boston College	-0.13
21	San Jose St.	1	144	Richmond	-0.14
22	Michigan	0.96	145	Missouri St.	-0.15
23	Baylor	0.95	146	Southern Utah	-0.15
24	Arizona	0.91	147	Fla. Atlantic	-0.15
25	Vanderbilt	0.91	148	North Dakota	-0.16
26	Utah St.	0.9	149	Colorado	-0.16
27	Northern Ill.	0.89	150	Albany (NY)	-0.18
28	Tulsa	0.89	151	UTEP	-0.18
29	Penn St.	0.87	152	UAB	-0.19
30	Oklahoma St.	0.85	153	James Madison	-0.2
31	UCLA	0.84	154	Eastern Mich.	-0.21
32	Cincinnati	0.84	155	FIU	-0.22
33	Wisconsin	0.83	156	Alabama St.	-0.22
34	Arkansas St.	0.81	157	New Hampshire	-0.22
35	Boise St.	0.81	158	Stephen F. Austin	-0.24
36	Ole Miss	0.81	159	Howard	-0.26
37	Michigan St.	0.8	160	Montana	-0.29
38	Southern California	0.8	161	Army	-0.31
39	Mississippi St.	0.78	162	New Mexico	-0.33
40	Kent St.	0.78	163	Tulane	-0.33
41	Arizona St.	0.77	164	Harvard	-0.34
42	Syracuse	0.77	165	N.C. A&T	-0.34
43	Texas Tech	0.76	166	South Carolina St.	-0.35
44	Missouri	0.74	167	Northern Colo.	-0.36
45	Washington	0.74	168	UC Davis	-0.38
46	Miami (FL)	0.73	169	Alabama A&M	-0.38
47	Ball St.	0.73	170	Idaho	-0.38
48	Rutgers	0.71	171	Hawaii	-0.39
49	Louisiana Tech	0.69	172	Jackson St.	-0.39
50	West Virginia	0.69	173	Wagner	-0.39
51	Toledo	0.68	174	Colgate	-0.43
52	TCU	0.64	175	Delaware St.	-0.45
53	North Carolina	0.63	176	Northwestern St.	-0.46
54	Iowa St.	0.62	177	UNLV	-0.47
55	BYU	0.61	178	Furman	-0.49
56	Fresno St.	0.58	179	Tennessee Tech	-0.49
57	La.-Lafayette	0.57	180	Southeast Mo. St.	-0.5
58	San Diego St.	0.57	181	N.C. Central	-0.5
59	Virginia Tech	0.56	182	Drake	-0.5
60	Sam Houston St.	0.55	183	Liberty	-0.52
61	UCF	0.54	184	South Ala.	-0.52
62	Montana St.	0.52	185	Massachusetts	-0.54
63	Arkansas	0.51	186	Brown	-0.55
64	Purdue	0.5	187	Western Ill.	-0.56
65	Navy	0.49	188	Florida A&M	-0.57
66	Middle Tenn.	0.49	189	New Mexico St.	-0.58
67	Ga. Southern	0.48	190	Penn	-0.6
68	Georgia Tech	0.48	191	Elon	-0.6
69	Eastern Wash.	0.48	192	Portland St.	-0.6
70	Tenn.-Martin	0.48	193	Akron	-0.62
71	La.-Monroe	0.47	194	Mississippi Val.	-0.63
72	North Carolina St.	0.46	195	Fordham	-0.64
73	Tennessee	0.45	196	Austin Peay	-0.65
74	Duke	0.44	197	Southern Miss.	-0.65
75	Central Ark.	0.43	198	South Dakota	-0.66
76	East Carolina	0.43	199	Maine	-0.66
77	Ohio	0.42	200	San Diego	-0.67
78	Illinois St.	0.4	201	Weber St.	-0.68
79	Minnesota	0.4	202	Delaware	-0.71
80	Pittsburgh	0.38	203	Dartmouth	-0.73
81	Utah	0.38	204	Alcorn St.	-0.74
82	Bowling Green	0.37	205	Lamar	-0.74
83	SMU	0.37	206	Princeton	-0.77
84	Wofford	0.35	207	Jacksonville	-0.79
85	Tennessee St.	0.35	208	Butler	-0.79
86	Eastern Ky.	0.34	209	Western Caro.	-0.84
87	South Dakota St.	0.31	210	Charleston So.	-0.85
88	Iowa	0.3	211	Southern U.	-0.86
89	Wake Forest	0.29	212	Monmouth	-0.86
90	Appalachian St.	0.29	213	Norfolk St.	-0.86

91	Auburn	0.29	214	Dayton	-0.89
92	Western Ky.	0.28	215	St. Francis (PA)	-0.9
93	California	0.28	216	Nicholls St.	-0.91
94	Jacksonville St.	0.26	217	Gardner-Webb	-0.92
95	Central Mich.	0.25	218	Hampton	-0.92
96	Cal Poly	0.25	219	Prairie View	-0.93
97	Indiana St.	0.25	220	Lafayette	-0.94
98	Eastern Ill.	0.23	221	Georgetown	-0.95
99	Indiana	0.22	222	Duquesne	-0.98
100	UTSA	0.21	223	Robert Morris	-1
101	Youngstown St.	0.21	224	Cornell	-1.03
102	Ark.-Pine Bluff	0.21	225	William & Mary	-1.04
103	Southern Ill.	0.21	226	Morgan St.	-1.06
104	Rice	0.21	227	Texas Southern	-1.08
105	Nevada	0.18	228	Idaho St.	-1.1
106	Kentucky	0.16	229	Georgia St.	-1.12
107	Old Dominion	0.15	230	Columbia	-1.15
108	Northern Ariz.	0.15	231	VMI	-1.16
109	Connecticut	0.14	232	Bryant	-1.22
110	Virginia	0.14	233	Holy Cross	-1.23
111	Washington St.	0.14	234	Presbyterian	-1.24
112	Temple	0.13	235	Central Conn. St.	-1.24
113	Bethune-Cookman	0.08	236	Savannah St.	-1.25
114	Troy	0.06	237	Bucknell	-1.28
115	Citadel	0.06	238	Yale	-1.35
116	Miami (OH)	0.06	239	Rhode Island	-1.38
117	McNeese St.	0.06	240	Morehead St.	-1.39
118	South Fla.	0.06	241	Grambling	-1.44
119	Houston	0.05	242	Marist	-1.45
120	Samford	0.05	243	Sacred Heart	-1.45
121	Stony Brook	0.04	244	Davidson	-1.8
122	Murray St.	0.04	245	Valparaiso	-1.83
123	Maryland	0.01	246	Campbell	-2.01

Table C.3: Full Ranking NCAA Div. I Basketball (Margin of Victory)

Rank	Team	Rating	Rank	Team	Rating
1	Indiana	26.30	174	Northeastern	-0.66
2	Florida	25.05	175	Albany NY	-0.68
3	Louisville	24.65	176	Army	-0.86
4	Duke	22.27	177	W Illinois	-1.07
5	Gonzaga	21.91	178	IL Chicago	-1.13
6	Kansas	21.57	179	Buffalo	-1.15
7	Ohio St	20.85	180	Idaho	-1.32
8	Michigan	20.57	181	Tennessee St	-1.38
9	Pittsburgh	20.01	182	Manhattan	-1.43
10	Syracuse	19.75	183	TX Southern	-1.46
11	Wisconsin	19.74	184	Towson	-1.49
12	Michigan St	18.97	185	Florida Intl	-1.58
13	Miami FL	17.54	186	Georgia St	-1.63
14	Arizona	17.38	187	CS Fullerton	-1.66
15	Creighton	17.19	188	Bryant	-1.68
16	Georgetown	17.17	189	James Madison	-1.74
17	Minnesota	17.16	190	Charleston So	-1.75
18	VA Commonwealth	16.85	191	NC Central	-1.76
19	Oklahoma St	16.58	192	Toledo	-1.79
20	Missouri	16.49	193	Youngstown St	-1.98
21	North Carolina	15.26	194	South Alabama	-2.07
22	NC State	15.21	195	UC Davis	-2.23
23	Cincinnati	15.11	196	Loy Marymount	-2.33
24	St Mary's CA	15.09	197	Long Island	-2.34
25	St Louis	15.08	198	Missouri St	-2.35
26	Kentucky	15.06	199	Lafayette	-2.37
27	Mississippi	14.90	200	Bowling Green	-2.62
28	Marquette	14.83	201	Columbia	-2.66
29	Colorado St	14.81	202	Southern Univ	-2.66
30	New Mexico	14.80	203	Elon	-2.67
31	Notre Dame	14.80	204	Gardner Webb	-2.68
32	Iowa	14.42	205	Texas Tech	-2.68
33	Iowa St	14.34	206	CS Northridge	-2.72
34	San Diego St	14.34	207	Wagner	-2.77
35	Kansas St	14.18	208	Santa Barbara	-2.86
36	UNLV	14.15	209	UNC Asheville	-3.04
37	UCLA	14.09	210	Pepperdine	-3.05
38	Baylor	13.99	211	Marshall	-3.14
39	Virginia	13.74	212	Yale	-3.35
40	Illinois	13.60	213	Oakland	-3.36
41	Wichita St	13.55	214	SC Upstate	-3.40
42	Oregon	13.47	215	William & Mary	-3.59
43	Memphis	12.97	216	Seattle	-3.68

44	Colorado	12.73	217	Ark Little Rock	-3.74
45	Stanford	12.65	218	TCU	-3.75
46	Connecticut	12.50	219	Jacksonville St	-3.82
47	Oklahoma	12.31	220	Mississippi St	-3.83
48	Middle Tenn St	12.09	221	High Point	-3.96
49	Maryland	11.89	222	Duquesne	-4.02
50	Belmont	11.36	223	Quinnipiac	-4.26
51	Villanova	11.35	224	ULL	-4.28
52	Butler	11.34	225	Morehead St	-4.64
53	California	11.07	226	SE Missouri St	-4.80
54	Boise St	10.88	227	Morgan St	-4.87
55	Arkansas	10.43	228	Coastal Car	-5.00
56	La Salle	10.34	229	Savannah St	-5.10
57	BYU	10.24	230	Portland	-5.12
58	Denver	10.23	231	Hartford	-5.24
59	Tennessee	10.15	232	FL Atlantic	-5.26
60	Akron	10.08	233	Mt St Mary's	-5.28
61	Illinois St	10.07	234	St Francis NY	-5.43
62	Providence	9.94	235	Norfolk St	-5.53
63	Dayton	9.83	236	C Michigan	-5.54
64	Alabama	9.80	237	Holy Cross	-5.55
65	Southern Miss	9.74	238	Texas St	-5.56
66	Purdue	9.62	239	E Michigan	-5.59
67	Temple	9.50	240	CS Bakersfield	-5.63
68	Northern Iowa	9.20	241	UT San Antonio	-5.76
69	Arizona St	9.03	242	North Texas	-5.76
70	Washington	8.80	243	Marist	-5.77
71	Xavier	8.29	244	Brown	-5.80
72	St Joseph's PA	7.96	245	Sam Houston St	-5.86
73	Richmond	7.83	246	Miami OH	-5.89
74	Ohio	7.81	247	W Carolina	-5.90
75	Bucknell	7.80	248	Ball St	-5.98
76	Valparaiso	7.76	249	Wofford	-6.16
77	Davidson	7.74	250	Stetson	-6.29
78	Texas	7.68	251	IPFW	-6.31
79	Detroit	7.61	252	Fordham	-6.36
80	Washington St	7.33	253	NC A&T	-6.45
81	Georgia Tech	7.23	254	Penn	-6.47
82	Texas A&M	7.16	255	Cornell	-6.71
83	Santa Clara	7.09	256	Cleveland St	-6.75
84	USC	7.05	257	South Dakota	-6.77
85	Wyoming	6.96	258	North Florida	-7.11
86	Vanderbilt	6.91	259	Old Dominion	-7.14
87	Massachusetts	6.71	260	N Colorado	-7.16
88	Boston College	6.66	261	Appalachian St	-7.39
89	Oregon St	6.63	262	N Kentucky	-7.58
90	Rutgers	6.61	263	NJIT	-7.74
91	LSU	6.60	264	Central Conn	-7.76
92	St John's	6.57	265	New Hampshire	-7.78
93	Clemson	6.48	266	American Univ	-7.95
94	Florida St	6.45	267	Winthrop	-8.09
95	Louisiana Tech	6.34	268	Troy	-8.16
96	New Mexico St	6.28	269	Colgate	-8.19
97	Evansville	6.15	270	E Illinois	-8.29
98	Seton Hall	6.08	271	North Dakota	-8.37
99	N Dakota St	6.08	272	St Peter's	-8.47
100	Georgia	6.00	273	Liberty	-8.54
101	Air Force	5.99	274	Dartmouth	-8.66
102	Indiana St	5.86	275	San Jose St	-8.68
103	Northwestern	5.78	276	Hampton	-8.89
104	Stony Brook	5.76	277	UNC Greensboro	-8.90
105	Utah	5.25	278	Tennessee Tech	-8.97
106	UTEP	5.16	279	CS Sacramento	-9.00
107	West Virginia	5.10	280	Bethune-Cookman	-9.03
108	Nebraska	4.91	281	Ga Southern	-9.15
109	Iona	4.88	282	Maine	-9.15
110	South Florida	4.81	283	Sacred Heart	-9.29
111	G Washington	4.73	284	Samford	-9.50
112	Fresno St	4.71	285	UNC Wilmington	-9.63
113	SF Austin	4.66	286	Siena	-9.67
114	Utah St	4.49	287	WI Milwaukee	-9.73
115	Lehigh	4.45	288	Campbell	-9.76
116	Weber St	4.44	289	Chattanooga	-9.88
117	Harvard	4.32	290	Nicholls St	-9.96
118	Wake Forest	4.05	291	SE Louisiana	-10.05
119	St Bonaventure	4.02	292	UC Riverside	-10.18
120	Drake	3.91	293	Jacksonville	-10.20
121	UCF	3.88	294	Hofstra	-10.32
122	Princeton	3.79	295	Cent Arkansas	-10.32
123	Pacific	3.68	296	Radford	-10.32
124	Kent	3.62	297	E Washington	-10.32
125	Penn St	3.37	298	Edwardsville	-10.47
126	Charlotte	3.21	299	Montana St	-10.48
127	San Francisco	3.12	300	Portland St	-10.48
128	S Dakota St	2.95	301	Rice	-10.64
129	Canisius	2.63	302	Delaware St	-10.68
130	E Kentucky	2.56	303	Navy	-10.76

131	WI Green Bay	2.51	304	Utah Valley	-10.84
132	Virginia Tech	2.34	305	VMI	-10.86
133	Tulane	2.33	306	McNeese St	-10.93
134	DePaul	2.30	307	Chicago St	-11.07
135	Wright St	2.25	308	TAM C. Christi	-11.17
136	Cal Poly SLO	2.04	309	N Illinois	-11.27
137	George Mason	1.94	310	Lipscomb	-11.42
138	Murray St	1.89	311	Monmouth NJ	-11.51
139	UC Irvine	1.75	312	Ark Pine Bluff	-11.56
140	UT Arlington	1.72	313	Southern Utah	-11.57
141	Nevada	1.68	314	Northern Arizona	-11.70
142	Loyola MD	1.57	315	Austin Peay	-11.71
143	Bradley	1.50	316	MD Baltimore Co	-12.05
144	FL Gulf Coast	1.40	317	Coppin St	-12.11
145	Arkansas St	1.38	318	Idaho St	-12.62
146	Niagara	1.37	319	TX Pan American	-12.72
147	Northwestern LA	1.27	320	ETSU	-12.94
148	East Carolina	1.08	321	Missouri KC	-13.00
149	Long Beach St	1.02	322	ULM	-13.37
150	Auburn	0.86	323	Alcorn St	-13.83
151	W Michigan	0.81	324	Howard	-13.96
152	UAB	0.73	325	Jackson St	-14.29
153	Mercer	0.68	326	NE Omaha	-14.51
154	Col Charleston	0.64	327	Houston Bap	-14.55
155	Vermont	0.43	328	St Francis PA	-14.62
156	Robert Morris	0.31	329	Kennesaw	-15.09
157	San Diego	0.06	330	Prairie View	-15.15
158	Boston Univ	0.03	331	IUPUI	-15.32
159	Houston	0.02	332	Florida A&M	-15.69
160	Montana	0.01	333	TN Martin	-15.93
161	Fairfield	0.01	334	Citadel	-16.39
162	South Carolina	-0.01	335	Presbyterian	-16.87
163	SMU	-0.07	336	Furman	-16.91
164	Rhode Island	-0.10	337	Alabama St	-17.33
165	Tulsa	-0.16	338	F Dickinson	-17.51
166	S Illinois	-0.29	339	S Carolina St	-17.71
167	Oral Roberts	-0.31	340	Alabama A&M	-17.95
168	Rider	-0.32	341	Binghamton	-18.11
169	Delaware	-0.33	342	MS Valley St	-18.63
170	Loyola-Chicago	-0.46	343	Longwood	-18.81
171	W Kentucky	-0.49	344	MD E Shore	-19.09
172	Hawaii	-0.59	345	Lamar	-20.04
173	Drexel	-0.62	346	New Orleans	-20.56
			347	Grambling	-35.36

Table C.4: Full Ranking NCAA Div. I Basketball (Binary)

Rank	Team	Rating	Rank	Team	Rating
1	New Mexico	1.18	175	Mt St Mary's	0.00
2	Duke	1.16	176	S Illinois	0.00
3	Louisville	1.15	177	Fairfield	0.00
4	Kansas	1.11	178	Hawaii	-0.01
5	Miami FL	1.08	179	Long Island	-0.01
6	Gonzaga	1.08	180	Texas Tech	-0.01
7	Indiana	1.07	181	James Madison	-0.02
8	Ohio St	1.02	182	Lafayette	-0.02
9	Georgetown	1.02	183	South Carolina	-0.03
10	Michigan St	1.00	184	Towson	-0.03
11	Arizona	0.97	185	Boston Univ	-0.03
12	Michigan	0.97	186	DePaul	-0.04
13	Florida	0.94	187	Norfolk St	-0.04
14	St Louis	0.93	188	Oakland	-0.06
15	Syracuse	0.93	189	Elon	-0.06
16	Memphis	0.92	190	Southern Univ	-0.06
17	Kansas St	0.91	191	Youngstown St	-0.07
18	Marquette	0.91	192	Toledo	-0.08
19	UCLA	0.90	193	Hartford	-0.08
20	UNLV	0.88	194	TCU	-0.10
21	Colorado St	0.87	195	SMU	-0.11
22	Butler	0.87	196	Marshall	-0.11
23	Oklahoma St	0.86	197	Savannah St	-0.12
24	Creighton	0.85	198	TX Southern	-0.13
25	North Carolina	0.84	199	Rhode Island	-0.14
26	Notre Dame	0.82	200	Missouri St	-0.14
27	Wisconsin	0.82	201	Morehead St	-0.14
28	San Diego St	0.80	202	Manhattan	-0.14
29	VA Commonwealth	0.80	203	Mississippi St	-0.15
30	Oregon	0.80	204	Loyola-Chicago	-0.16
31	Pittsburgh	0.79	205	Gardner Webb	-0.17
32	NC State	0.79	206	Ball St	-0.18

33	St Mary's CA	0.78	207	Pepperdine	-0.18
34	Minnesota	0.76	208	Idaho	-0.19
35	Illinois	0.74	209	Charleston So	-0.19
36	Wichita St	0.74	210	Quinnipiac	-0.19
37	Colorado	0.73	211	Georgia St	-0.20
38	Connecticut	0.73	212	UC Davis	-0.21
39	Cincinnati	0.72	213	Cleveland St	-0.21
40	Temple	0.71	214	Loy Marymount	-0.21
41	Mississippi	0.71	215	E Michigan	-0.21
42	Boise St	0.69	216	Auburn	-0.22
43	California	0.68	217	Army	-0.22
44	Missouri	0.68	218	Duquesne	-0.23
45	Middle Tenn St	0.68	219	FL Atlantic	-0.23
46	Oklahoma	0.67	220	Portland	-0.23
47	Iowa St	0.66	221	Drexel	-0.24
48	La Salle	0.65	222	SE Missouri St	-0.24
49	Belmont	0.64	223	Yale	-0.24
50	Villanova	0.61	224	Buffalo	-0.24
51	Kentucky	0.60	225	CS Northridge	-0.24
52	Southern Miss	0.59	226	Stetson	-0.24
53	Massachusetts	0.59	227	Ark Pine Bluff	-0.25
54	Akron	0.59	228	ULL	-0.25
55	Iowa	0.56	229	Central Conn	-0.25
56	Tennessee	0.56	230	Fordham	-0.27
57	Bucknell	0.54	231	Sam Houston St	-0.28
58	Maryland	0.54	232	CS Bakersfield	-0.29
59	Alabama	0.54	233	Brown	-0.29
60	Valparaiso	0.52	234	NC A&T	-0.29
61	Arizona St	0.52	235	Santa Barbara	-0.30
62	Virginia	0.52	236	IPFW	-0.30
63	Stanford	0.51	237	SC Upstate	-0.30
64	Louisiana Tech	0.51	238	North Texas	-0.31
65	Baylor	0.50	239	UNC Asheville	-0.32
66	Wyoming	0.49	240	Bowling Green	-0.32
67	Charlotte	0.49	241	St Francis NY	-0.33
68	Washington	0.48	242	CS Fullerton	-0.33
69	BYU	0.47	243	North Florida	-0.34
70	Florida St	0.46	244	Morgan St	-0.34
71	New Mexico St	0.46	245	SE Louisiana	-0.35
72	SF Austin	0.46	246	Cornell	-0.35
73	Providence	0.45	247	Tennessee Tech	-0.35
74	LSU	0.43	248	McNeese St	-0.36
75	Arkansas	0.42	249	Texas St	-0.37
76	Air Force	0.42	250	High Point	-0.37
77	Denver	0.42	251	Holy Cross	-0.37
78	Ohio	0.41	252	C Michigan	-0.37
79	Xavier	0.41	253	Lipscomb	-0.37
80	Indiana St	0.40	254	North Dakota	-0.38
81	Northern Iowa	0.40	255	Troy	-0.40
82	St Joseph's PA	0.40	256	William & Mary	-0.41
83	St John's	0.40	257	Jacksonville	-0.42
84	Texas	0.39	258	N Kentucky	-0.42
85	Nebraska	0.39	259	UT San Antonio	-0.42
86	Davidson	0.36	260	Columbia	-0.43
87	Santa Clara	0.36	261	Marist	-0.43
88	Dayton	0.36	262	Cent Arkansas	-0.44
89	USC	0.34	263	Appalachian St	-0.44
90	Texas A&M	0.34	264	South Dakota	-0.45
91	Richmond	0.34	265	Miami OH	-0.45
92	UCF	0.34	266	NE Omaha	-0.45
93	UTEP	0.33	267	San Jose St	-0.45
94	Purdue	0.33	268	Colgate	-0.46
95	Boston College	0.32	269	Coastal Car	-0.46
96	S Dakota St	0.32	270	NJIT	-0.46
97	Detroit	0.31	271	E Illinois	-0.47
98	Rutgers	0.31	272	Delaware St	-0.47
99	N Dakota St	0.31	273	American Univ	-0.47
100	Evansville	0.30	274	W Carolina	-0.47
101	Pacific	0.30	275	Wofford	-0.48
102	Stony Brook	0.30	276	Sacred Heart	-0.48
103	Illinois St	0.29	277	Monmouth NJ	-0.48
104	Georgia Tech	0.28	278	Penn	-0.48
105	Utah St	0.28	279	CS Sacramento	-0.49
106	Harvard	0.28	280	Winthrop	-0.49
107	E Kentucky	0.28	281	Maine	-0.50
108	Montana	0.27	282	Ga Southern	-0.50
109	Vanderbilt	0.26	283	Hampton	-0.50
110	Northwestern LA	0.25	284	ETSU	-0.51
111	Georgia	0.24	285	Northern Arizona	-0.51
112	Seton Hall	0.24	286	N Colorado	-0.52
113	Loyola MD	0.23	287	TX Pan American	-0.52
114	Weber St	0.23	288	St Peter's	-0.52
115	Murray St	0.23	289	Bethune-Cookman	-0.52
116	FL Gulf Coast	0.22	290	Chattanooga	-0.55
117	Long Beach St	0.22	291	Seattle	-0.55
118	Lehigh	0.22	292	UNC Wilmington	-0.56
119	Iona	0.21	293	Nicholls St	-0.56

120	East Carolina	0.20	294	Siena	-0.57
121	Northwestern	0.20	295	Dartmouth	-0.58
122	W Illinois	0.20	296	Houston Bap	-0.58
123	Utah	0.20	297	Campbell	-0.59
124	South Florida	0.19	298	Edwardsville	-0.59
125	Col Charleston	0.18	299	WI Milwaukee	-0.59
126	Drake	0.18	300	Missouri KC	-0.59
127	Robert Morris	0.18	301	Samford	-0.60
128	West Virginia	0.18	302	New Hampshire	-0.60
129	UT Arlington	0.17	303	Liberty	-0.60
130	Wright St	0.17	304	Montana St	-0.60
131	St Bonaventure	0.16	305	Prairie View	-0.61
132	Albany NY	0.16	306	Utah Valley	-0.62
133	Kent	0.15	307	Jackson St	-0.62
134	Fresno St	0.15	308	TN Martin	-0.64
135	Wake Forest	0.15	309	Rice	-0.64
136	Tulsa	0.14	310	VMI	-0.65
137	Nevada	0.13	311	Southern Utah	-0.66
138	Niagara	0.13	312	ULM	-0.67
139	Oregon St	0.13	313	MD Baltimore Co	-0.67
140	W Michigan	0.13	314	E Washington	-0.69
141	Canisius	0.13	315	Chicago St	-0.69
142	Virginia Tech	0.12	316	Alcorn St	-0.69
143	G Washington	0.12	317	F Dickinson	-0.70
144	UC Irvine	0.12	318	Navy	-0.70
145	Mercer	0.11	319	Hofstra	-0.71
146	Washington St	0.11	320	Radford	-0.72
147	Northeastern	0.10	321	Coppin St	-0.74
148	Bradley	0.10	322	TAM C. Christi	-0.74
149	Princeton	0.10	323	Austin Peay	-0.75
150	Vermont	0.10	324	Alabama A&M	-0.75
151	Cal Poly SLO	0.09	325	Old Dominion	-0.76
152	Penn St	0.09	326	IUPUI	-0.77
153	Florida Intl	0.09	327	UC Riverside	-0.77
154	W Kentucky	0.09	328	St Francis PA	-0.77
155	UAB	0.09	329	Alabama St	-0.77
156	Tulane	0.08	330	UNC Greensboro	-0.80
157	Clemson	0.08	331	N Illinois	-0.80
158	Delaware	0.08	332	Portland St	-0.85
159	Tennessee St	0.07	333	Florida A&M	-0.86
160	WI Green Bay	0.07	334	Howard	-0.88
161	South Alabama	0.07	335	New Orleans	-0.89
162	Arkansas St	0.07	336	Citadel	-0.90
163	Bryant	0.07	337	Lamar	-0.90
164	San Francisco	0.07	338	Longwood	-0.91
165	George Mason	0.06	339	Idaho St	-0.93
166	Wagner	0.06	340	MS Valley St	-0.95
167	Rider	0.05	341	Presbyterian	-0.95
168	Houston	0.04	342	S Carolina St	-0.97
169	IL Chicago	0.04	343	Kennesaw	-0.98
170	Jacksonville St	0.03	344	Furman	-0.98
171	Oral Roberts	0.03	345	Binghamton	-1.02
172	NC Central	0.02	346	MD E Shore	-1.07
173	San Diego	0.02	347	Grambling	-1.37
174	Ark Little Rock	0.01			

Curriculum Vitae

Robert Kelly Sizemore

534 Camway Drive

Wilmington, NC 28403

Education

University of North Carolina at Wilmington August 2005 - December 2010

Bachelor of Science, Physics

Bachelor of Science, Mathematics

Wake Forest University August 2011 - May 2013

Master of Arts, Mathematics

Conferences/Talks

- Rating Sports Teams with Hodge Theory. *Carolina Sports Analytics Meeting.* Furman University. April 13, 2013.
- Removing Inconsistencies in Sports Ranking Data. *Wake Forest Graduate School's Research Day.* March 21, 2013.

Honors

- Dean's List (Fall 2006 - Spring 2008)
- Phi Mu Epsilon Physics Honor Society Inductee (Spring 2008)