Abelian Girth and Gapped Sheaves

by

Alice Izsak

B.Sc., Harvey Mudd College, 2007
M.Sc., The University of British Columbia, 2009

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Computer Science)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

July 2015

Abstract

In this work we study the abelian girth and sheaves on graphs. The girth of a graph is the length of the shortest cycle in a graph, and the abelian girth of a graph is the girth of the graph's universal abelian covering graph. We denote the abelian girth of a graph G as $\mathrm{Abl}(G)$ and show that for d-regular graphs on n vertices with d constant and n growing we have

$$\mathrm{Abl}(G) \leq 6 \log_{d-1} n + o_n(1).$$

This can be seen as an version of the Moore bound for abelian girth. We also prove $\mathrm{Girth}(G) \leq \mathrm{Abl}(G)/3$, which implies that any multiplicative important to the abelian girth Moore bound would also improve the standard Moore bound.

Sheaves on graphs and two of their homological invariants, the maximum excess and the first twisted Betti number, were used in the proof of the Hanna Neumann Conjecture from algebra and may be of use in proving several related conjectures. These conjectures can be proven by showing that certain sheaves called ρ-kernels have vanishing maximum excess. Gapped sheaves have maximum excess equal to the first twisted Betti number, and it is easy to compute the maximum excess of a given sheaf in the case that the sheaf is not gapped. For general sheaves though, there is no known way of computing the maximum excess in polynomial time. We give several conditions that a sheaf must satisfy if it is gapped. These conditions include that a sheaf must have edge dimension at least as large as the abelian girth of the underlying graph. The ρ-kernels are subsheaves of constant sheaves. We prove that gapped subsheaves of constant sheaves exist, implying that finding maximum excess of some ρ-kernels may be computationally difficult.

Preface

This dissertation is unpublished research done in collaboration with Joel Friedman, Lior Silberman and myself with the bulk of the research done in regular meetings between Joel Friedman and myself. The scope of the project was planned by Joel Freidman and myself. Though the vast majority of the research originates from discussions between myself and Friedman with Silberman joining occasionally, some results I first discovered on my own. These include the chain decomposition described in Chapter 9, the minimal gapped sheaves on the figure-eight graph and the theta graph described in Chapter 11 and the existence of the gapped subconstant sheaf also described in Chapter 11, though computations verifying the sheaf was gapped were done by myself and Friedman.

I wrote the first drafts of every chapter except Chapters 7 and 8 which were written by Joel Friedman and subsequently edited by the author.

Table of Contents

Acknowledgements

Joel Friedman's guidance, patience and valuable insights have propelled me through my graduate career. I could not have hoped for a more supportive supervisor, and his breadth of knowledge and eagerness to attack mathematical problems in unique ways have been constant inspirations. I also must thank Lior Silberman for his help answering many questions throughout this research and his extensive help in the editing process. I have learned much about how to write mathematical documents well from Lior. I also need to acknowledge Omer Angel, the other member of my supervisory committee, for helping oversee the research presented here. I should also mention that this research was supported by a UBC Four Year Fellowship.

I am beyond grateful for my family's love and support over the last few years. Thank you Nina, Benjamin, Ariela, Leon and Celine Izsak. I also send my heartfelt gratitude to colleagues and friends who have helped me in innumerable ways throughout my graduate career, particularly Craig Weidert, Jaimie Veale, Julieta Martinez, Herbie Carvajal, Samuel Rudolph, Stephen Jones, Julia Fornaca, and Jess Latimer. Finally, my love goes out to my partner Rachael Gillese who has stood with me through so much and brought so much happiness in my life. Thank you, thank you, thank you.

Chapter 1

Introduction

This dissertation is divided into two main parts. In the first, we discuss a graph invariant called the *abelian girth* of a graph, which Friedman introduced in his proof of the the Hanna Neumann Conjecture (or HNC) [17]. In the second part we discuss sheaves on graphs and their homological invariants. Though these two parts are fairly independent from one another, they are tied together by a main result of this dissertation which shows that what we call a minimally gapped sheaf on a graph has total dimension equal to the abelian girth of the underlying graph.

In the first part of this dissertation, we show several relations between the abelian girth of the graph and the standard girth of a graph as well as between the abelian girth and the volume bound on girth known as the Moore bound. Specifically, we give some evidence to show that bounding the abelian girth from above is a possible approach to improving the Moore bound for regular graphs of fixed degree with a large number of vertices. Let us make this more precise.

The *girth* of a graph is the length of the shortest non-trivial cycle in the graph. If $G_{n,d}$ is any d-regular graph on n vertices, it is known that for fixed d and large n we have

$$\mathrm{Girth}(G_{n,d}) \leq 2\log_{d-1} n + o_n(1), \tag{1.1}$$

and we are interested to know if the factor of 2 can be improved upon. This bound follows from the *Moore bound* (see [11]); although the Moore bound for regular graphs is very easy to prove, there has been only slight improvements to it in the last 50 years—only an additive constant of one or two—and, in particular, to the factor of 2 in (1.1) is the best factor known to date. Graphs of large girth have numerous applications ([2, 28, 33]) and much has been written on the girth and the Moore bound (see [3]).

The *abelian girth* of a graph, G, denoted $\mathrm{Abl}(G)$, is the girth of its *universal abelian cover*, or equivalently the shortest length of a non-trivial, closed, non-backtracking walk that traverses each edge the same number of times in each direction. The abelian girth is important in sheaf theory on graphs and the first proof of the Hanna Neumann Conjecture [17]; furthermore, in the second portion of this dissertation we show the abelian girth is strongly related to what we call minimal gapped sheaves on a graph. First though, we show that there is an analogue of the Moore bound for the abelian girth, and that improving this analogue would improve the Moore bound. We also provide evidence that such an improvement may be possible.

Specifically, we show that

$$\mathrm{girth}(G) \leq 3\,\mathrm{Abl}(G),$$

for any graph, G. Second, we show that there is an argument analogous to the Moore bound that shows that for fixed d we have

$$\mathrm{Abl}(G_{n,d}) \leq 6\log_{d-1} n + o_n(1)$$

for any d-regular graph on n vertices, $G_{n,d}$; we remark that the proof of this theorem is not as immediate as that of the Moore bound. It follows that any improvement to the factor of 6 in (1) would give an improvement to the factor of 2 in (1.1).

Ideally we would show that all known explicit constructions of families of d-regular graphs have abelian girth at most $c\log_{d-1} n$ for some $c < 6$. In this paper, we will focus on the only family of d-regular graphs with d fixed which has girth greater than $\log_{d-1} n$. It is known that the factor of 2 in (1.1) cannot be less than 4/3, at least for certain d: indeed, [27] constructs graphs, $X^{p,q}$, for primes $p, q \equiv 1 \pmod 4$, that are $d = p + 1$ regular on $n = q(q^2 + 1)$ vertices for which

$$\mathrm{girth}(X^{p,q}) = (4/3)\log_{d-1} n + o_n(1)$$

for fixed $d = p+1$ and large $n = q(q^2+1)$. Furthermore there are no known families of graphs which improve on the above 4/3; in fact, for general d, the best girth lower bound is 4/3 replaced with 1, by choosing a random graph and slightly modifying it [10], [12]. To show that it is plausible to improve on the factor of 6 in (1), we will show that

$$\mathrm{Abl}(X^{p,q}) \leq (16/3)\log_{d-1} n + o_n(1)$$

which suggests that there is room for the factor of 6 to decrease. We conjecture that the 16/3 can be replaced with 4, for reasons we shall explain later. We are unaware of any other graph constructions in the literature for which one can improve upon the 16/3 above.

To prove (1), we prove a fundamental lemma that suggests many possible generalizations of the above discussion of abelian girth and the Moore bound.

First we remark that the girth of a graph is smallest positive length of a cycle that embeds in the graph. Our fundamental lemma states that the abelian girth is the smallest number of edges in a graph of Euler characteristic -1 that embeds in the graph, provided that we weight the edges appropriately: namely, we count each edge twice, except that in "barbell graphs" we count each edge in the bar four times.

We remark that both girth and abelian girth can be viewed as linear algebraic invariants of graphs; indeed, girth is often studied as a property of the adjacency matrix of the graph (see, for example, [3]), and the abelian girth relates to sheaves of vector spaces over the graph. So our lemma fundamental to proving (1) also suggests that there could be other such girth-type invariants, both (1) arising as linear algebraically from the graphs, and (2) satisfying inequalities analogous to (1).

The rest of this paper is organized as follows. In Section 2 we give some precise definitions and state our main theorems. In Section 3 we prove Equation (1). In Section 4 we prove the fundamental lemma as well as Equation (1) which follows quickly from the fundamental lemma. In Section 5 we describe the graphs $X^{p,q}$ from [27] and show that they obey Equation (1).

The next part of this dissertation studies sheaves on graphs and their homological invariants. One invariant we study is the *maximum excess* of a sheaf. The maximum excess has a definition with no reference to homology theory, but it is also arises as a limit of Betti numbers akin to the L^2 Betti numbers studied by Atiyah [4] and has a long/short exact sequence theory. The maximum excess originally was discussed in Friedman's proof of the HNC [17]. The reason maximum excess is necessary for Friedman's proof and a reason we study it in this paper is that it generalizes the notion of reduced cyclicity, a graph invariant from the graph theoretic reformulation of the HNC (see [21], [22], [30], [20], [6], [32], [29], [13].) In particular, the maximum excess of a sheaf on a graph is equivalent to the reduced cyclicity of that graph if the sheaf is what we call a structure sheaf.

Our study of the maximum excess is partially motivated by our desire to prove certain conjectures related to the HNC (see [23], [24], [8], [9] [34].) These are conjectures on the rank of the intersection of two subgroups from a free product of groups. In Friedman's proof, he shows that the HNC is true if certain sheaves called ρ-kernels have vanishing maximum excess, and similarly these other conjectures can be verified if one can prove that certain other ρ-kernels also have vanishing maximum excess. Unfortunately, the techniques from Friedman's proof of the HNC are alone not sufficient in proving the other conjectures.

A drawback of using the maximum excess is that there is no known way to compute the maximum excess of a sheaf on a graph in time polynomial in the total dimension of the sheaf. Friedman's proof also defines the twisted homology groups and twisted Betti numbers of a sheaf, and the first twisted Betti number is easy to compute, bounds the maximum excess from above, and is in many cases equal to the maximum excess. We refer to sheaves as *gapped* if the maximum excess and first twisted Betti number are not equal. In this dissertation, we give several results describing the properties of gapped sheaves.

In our paper, we describe sheaves on graphs as a collection of vector spaces indexed along the vertices and edges of the graph, along some linear maps from the edge spaces to the vertex spaces that we call restriction maps. A *constant sheaf* is sheaf where all of those vector spaces are identical and the restriction maps are all the identity, and a *subconstant sheaf* is simply any subsheaf of a constant sheaf. One of the more surprising results of my dissertation is that there exists gapped subconstant sheaves. Since the ρ-kernels from the conjectures related to the HNC are subconstant sheaves, any attempted proof may need to take into account that the first twisted Betti number on these sheaves could be nonzero even though the maximum excess vanishes on that sheaf. In related work (to appear) we do find ρ-kernels that correspond to a conjecture in [8] that are gapped.

Not only do gapped subconstant sheaves exist, but there are gapped subconstant sheaves on a graph that is only two vertices and multiple edges but no self loops. This implies an interesting, purely linear algebraic result. The usual notion of linear independence of vectors in a vector space can be stated in many equivalent ways. In quiver representation theory, there arises a notion of what we call *linear k-indepedence*, which is defined for every integer k, and which reduces to the usual linear indepedence for $k = 1$. This notion seems most interesting for $k = 2$ in quiver theory (see [25]); the notion of linear 2-independence defines a remarkable type of sheaf on a graph which we call a *pseudobundle*. Unfortunately, linear k-independence, for $k \geq 2$, is difficult to check directly. However, there is a related notion, which we call *tensorial k-independence*, which is easier to verify and which immediately implies linear k-independence. We prove though, using the existence of gapped subconstant sheaves on two vertices and no self loops, that tensorial k-dependence does not imply linear k-independence.

We also find a connection between gapped sheaves on a graph and the abelian girth of the underlying graph. We prove that the abelian girth on a graph is equal to the minimum of the total dimensions of the gapped sheaves on the graph. This result can also be viewed as an improvement on

Theorem 1.10 from [17], which states that if a lift of a sheaf on a graph has a sufficiently large degree dependent on the abelian girth of the graph , then the pullback sheaf is not gapped. Our result implies that in the case that the lift is just the identity map, then we can remove the factor of 2 from Friedman's Theorem 1.10.

Though our research is mainly motivated by our desire to prove conjectures related to the HNC, there are other reasons we are interested in answering fundamental questions about sheaves on graphs. A sheaf on a graph can be viewed as a generalization of the incidence matrix of the graph; indeed the structure sheaf of a graph has the same incident matrix equal as the underlying graph. It follows that sheaves on graphs allow a generalization of standard algebraic graph theory. Tools such as long/short exact sequences can allow for new graph theoretic inequalities. Also, any morphism between graphs can be represented by a morphism on sheaves on graphs, but there are morphisms between sheaves on graphs that do not correspond to any morphism between graphs. These "new morphisms" are necessary in Friedman's proof of the HNC but may also be useful in studying other aspects of graph theory.

Section 6 introduces linear and tensorial 2-independence in terms of only linear algebra and gives a counterexample to the two being the same. Section 7 gives an introduction to sheaf theory on graphs, twisted Betti numbers and maximum excess. At the end of the chapter, we also explicitly state the main results of the paper. In Section 8 we define minimal gapped sheaves and establish several properties of minimal gapped sheaves. In Section 9 we prove the fundamental lemma of this paper, what we call the twist trick. This allows us to give lower bounds on the edge dimension of minimally gapped sheaves. Section 10 defines homotopy preserving operations that will allow us to transform a sheaf on a graph while not changing the gap of the sheaf. Finally, in Chapter 11 we introduce three minimally gapped sheaves which then allows us to prove that the abelian girth of a graph is the minimal edge dimension of a gapped sheaf on that graph. We also use one of those gapped sheaves to construct the counterexample to linear and tensorial 2-independence being equivalent.

Chapter 2

Main Results

In this section we fix some terminology regarding graphs and formally state the main results of the first part of this dissertation.

2.1 Graph Terminology

This entire subsection consists of definitions used throughout this paper; they are more or less standard.

By a *graph* we shall mean a quadruple $G = (V_G, E_G, t_G, h_G)$ where V_G and E_G*section* are sets (the "*vertices*" and "*edges*", respectively) and $t_G, h_G \colon E_G \to V_G$ are maps (the "tail" and "head" of each edge, respectively). The *Euler characteristic* of G is

$$\chi(G) = |V_G| - |E_G|,$$

where $|\ \cdot\ |$ denotes the cardinality, provided that G is finite, i.e., that $|V_G|$ and $|E_G|$ are finite. We define the *unoriented edge set* of G to be

$$U_G = E_G \times \{+, -\},$$

and we extend h_G and t_G to be functions on U_G such that for $e \in E_G$

$$h_G(e,+) = t_G(e,-) = h_G(e), \quad t_G(e,+) = h_G(e,-) = t_G(e);$$

for $e \in E_G$ we say that $(e,+)$ and $(e,-)$ are *inverses* of each other. We denote the inverse of $u \in U_G$ by u^{-1}. A *walk* of *length* m in G is a sequence of unoriented edges

$$w = (u_1, \ldots, u_m)$$

such that the head of u_i is the tail of u_{i+1} for $i = 1, \ldots, m-1$; we define the vertices $t_G(u_1)$ and $h_G(u_m)$ to be the *starting* and *terminating* vertices of w, jointly the *endpoints* of w, and the *length* of w to be m, denoted $l(w)$; we say that w is *closed* if its two endpoints are the same vertex; we refer to the *edges* of w as the $e \in E_G$ such that at least one of $(e,+), (e,-)$ is an unoriented edge of w; we say that w is *non-backtracking* if there is no

$i = 1, \ldots, m-1$ for which $u_i^{-1} = u_{i+1}$; we say that w is *strongly closed, non-backtracking* if w is closed and non-backtracking, and u_1, u_m are not inverses of each other; we say that w is a *path* (respectively, *cycle*) if w is non-backtracking with distinct endpoints (respectively strongly closed, non-backtracking), and each edge $e \in E_G$ appears at most once in w (i.e. $(e,+)$ and $(e,-)$ appears at most once in $u_1, \ldots, u_m$). The *inverse* of w is defined as the walk

$$w^{-1} = (u_m^{-1}, \ldots, u_1^{-1}).$$

Given a walk $w = (u_1, \ldots, u_m)$ and a walk $k = (t_1, \ldots, t_n)$ such that the terminating vertex of w is the starting vertex of k, we define the *product* wk to be the walk $(u_1, \ldots, u_m, t_1, \ldots, t_n)$. We say that a walk w as above *joins* its two endpoints. We say that G is *connected* if any two of its vertices are joined by some walk in G.

As usual, a *morphism* of graphs, $f \colon G \to H$, is a pair $f = (f_V, f_E)$ of maps $f_V \colon V_G \to V_H$ and $f_E \colon E_G \to E_H$ such that $t_H \circ f_E = f_V \circ t_G$ and $h_H \circ f_E = f_V \circ h_G$. We often drop the subscripts from f_V and f_E. The set of morphisms will be denoted $\mathrm{Mor}(G,H)$. Thus morphisms are homomorphisms of the underlying undirected graphs which preserve the orientation.

As usual, the *girth* of a graph is the length of its shortest closed, non-backtracking walk; this length is necessarily positive by our conventions above. (In the literature one often allows for walks of length zero, which we do not consider here.)

2.2 Our Fundamental Lemma

In this subsection we discuss our main results.

Definition 2.1. The *abelian girth* of a graph, G, denoted $\mathrm{Abl}(G)$, is the minimum $m \geq 1$ such that there is a closed, non-backtracking walk

$$w = (u_1, \ldots, u_m)$$

such that each edge is traversed the same number of times in w in both directions, i.e., for each $e \in E_G$ the edge $(e,+)$ appears the same number of times among $u_1, \ldots, u_m$ as does $(e,-)$.

The abelian girth is the same as the girth of the *universal abelian cover* of G; see [17]. Given a walk w and an $e \in E_G$, if $(e,+)$ appears i_+ times and $(e,-)$ appears i_- times we say e appears a net $i_+ - i_-$ times in w. We refer to a walk as *edge neutral* if every edge appears a net 0 times in it.

We begin with a fundamental lemma. We remind the reader that our conventions insists that cycles and paths are of positive length.

Definition 2.2. Let G be a connected graph of Euler characteristic -1 without leaves, i.e., without vertices of degree one. We say that G is

1. a *figure-eight* graph if G consists of two cycles, mutually edge disjoint, sharing the same endpoint;
2. a *barbell* graph if G consists of two cycles, w_1, w_2, and one path b, all mutually edge disjoint, such that b joins the endpoints of w_1 and w_2; we refer to b as the *bar* of G;
3. a *theta* graph if G consists of two vertices joined by three paths mutually edge disjoint.

It is well known that the above three cases classifies all connected graphs of Euler characteristic -1 without leaves (see, for example, [26].)

Definition 2.3. Let G be a connected graph of Euler characteristic -1 without leaves. We define the *abelian length* of G, denoted $l_{\mathrm{Abl}}(G)$, to be twice its number of edges except that each edge of its bar (if G is a figure-eight graph) is counted four times.

We call the following lemma the Fundamental Lemma, as it gives an alternative definition of the abelian girth that we use in several proofs.

Lemma 2.4. *For any graph, G, we have*

$$\mathrm{Abl}(G) = \min_{G' \subset G} l_{\mathrm{Abl}}(G')$$

taken over all subgraphs, G', that are connected, of Euler characteristic -1, and without leaves. If no such G' exist, then the above minimum is taken to be infinity, as is the abelian girth of G, and there are no closed non-backtracking walks that traverse each edge the same number of times in both directions.

2.3 Main Results on Abelian Girth

Now we can easily state the other main results of the first part of this dissertation.

Theorem 2.5. *For any graph we have*

$$\mathrm{Girth}(G) \le \mathrm{Abl}(G)/3.$$

This theorem is an easy corollary of Lemma 2.4.

Theorem 2.6. *For any fixed d, let $G_{n,d}$ be any d-regular graph on n vertices. Then for large n we have*

$$\mathrm{Abl}(G) \le 6 \log_{d-1} n + o_n(1).$$

section

Our last theorem uses the LPS graphs $X^{p,q}$ [27] whose definition we save for Section 5

Theorem 2.7. *Let $p, q \equiv 1 \pmod 4$ be prime with $(q/p) = 1$. The graph $X^{p,q}$ of [27] of our Definition 5.1 has degree $d = p + 1$ and $n = q(q^2 + 1)$, and for fixed p and large q (i.e., fixed d and large n) we have*

$$\mathrm{Abl}(X^{p,q}) \le (16/3) \log_{d-1} n + o_n(1).$$

Chapter 3

Proof of Theorem 2.6

In this section we prove Theorem 2.6. First we state a few well known facts regarding free groups and graphs.

Lemma 3.1. *Let G be a finite graph. Then the fundamental group, $\pi_1(G, u)$, of homotopy classes of closed walks in G about u for any $u \in V_G$ is isomorphic a free group on a finite set of generators.*

The contiguous appearance of an edge and its inverse in a walk is called a *reversal*. Each walk, w, can be *reduced* by successively discarding its reversals; the walk obtained is the *reduction* of w and is independent of the choice of pairs to reduce; the reduction therefore contains no reversals and is non-backtracking. We use the notation $\mathrm{red}(w)$ for the reduction of w. Note that $\mathrm{red}(w^{-1}) = \mathrm{red}(w)^{-1}$ for any walk w since the portion of a walk that are removed by reduction remain the same when we take the inverse of that walk.

Definition 3.2. In each element ω of $\pi_1(G, u)$, there exists one unique non-backtracking walk in G; any walk in ω reduces to that unique non-backtracking walk. By the *length* of ω, denoted $|\omega|$, we mean the length of the unique non-backtracking walk in ω. Note this is not length with respect to a set of generators of $\pi_1(G, u)$. Similarly, even though closed walks are members of homotopy classes in the fundamental group, when we refer to the length of a walk or the reduction of a walk w we mean length and reduction as defined earlier and not in with respect to a set of generators for a free group.

The following fact is well known. For ease of reading we provide a proof, for which we thank Seirius on Stackexchange[1].

Lemma 3.3. *Let F be a free group and let α and β be two elements of F that commute with each other. Then there exists a $\omega \in F$ such that $\alpha = \omega^m$ and $\beta = \omega^{m'}$ for $m, m' \in \mathbb{Z}$.*

[1]See `http://math.stackexchange.com/questions/213576/in-a-free-group-two-elements-commute-if-and-only-if-they-are-powers-of-a-common`

Proof. By the Nielsen-Schreier theorem, the subgroup that α, β generate is a free group. And yet, any two elements of this subgroup commute, and hence this subgroup cannot be free and rank more than one. Hence this subgroup is generated by some $\omega \in F$, and the conclusion follows. □

For any closed walk w, we use the notation $[w]$ for the homotopy class in $\pi_1(G, u)$ that has w as a representative. Suppose a and b are closed walks and [a] and [b] commute in $\pi_1(G, u)$. Note that equality in the previous lemma is with regards to $\pi_1(G, u)$, meaning an equality of homotopy classes but not necessarily walks. Reducing all the element of a homotopy class gives one unique walk though. So since $[a]$ and $[b]$ commute, we have $\mathrm{red}(a) = \mathrm{red}(w^m)$ and $\mathrm{red}(b) = \mathrm{red}(w^{m'})$ for $m, m' \in \mathbb{Z}$ and w a closed walk. These equalities may not exist without the reductions. We may assume w is reduced, since $\mathrm{red}(w^m) = \mathrm{red}(\mathrm{red}(w)^m)$.

The following lemma is also well known and easy.

Lemma 3.4. *Let w be a closed nontrivial non-backtracking walk in a graph. Then if for integers m, m' we have that the length of $\mathrm{red}(w^m)$ equals that of $\mathrm{red}(w^{m'})$, then $m = \pm m'$.*

Proof. Let y be a maximal length walk such that $w = yxy^{-1}$ for some walk x. Then it is easy to check that for $m \neq 0$ the length of $\mathrm{red}(w^m)$ with respect to S is precisely

$$2|y| + |m||x|.$$

□

Proof of Theorem 2.6. Fix any vertex, $v \in V_G$. Then for any integer $h \geq 1$, there are $d(d-1)^{h-1}$ non-backtracking walks of length h from v. So let h be the smallest integer for which

$$d(d-1)^{h-1} \geq 2n + 1;$$

then, by the pigeon hole principle, there are three distinct non-backtracking walks, a, b, c, in G beginning in v and terminating in the same vertex u (which may or may not equal v). We remark that

$$h = \big(1 + o_n(1)\big) \log_{d-1} n.$$

Hence, to prove the theorem it suffices to show that we have

$$\mathrm{Abl}(G) \leq 6h. \tag{3.1}$$

The rough idea is simple. We break the analysis into two cases: $u = v$ and $u \neq v$. In either case we define new walk, d, based on a, b, c such that (1) the length of d is $4h$ or $6h$ (in the respective two cases), (2) each edge of E_G appears a net 0 times in d, but (3) d is not necessarily non-backtracking. The main work is to show that d does not reduce to the empty word. Discarding a consecutive pair of an (unoriented) edge of d and its opposite retains the property that each edge is traversed the same number of times in each direction. Hence if d is reduced to a non-empty, non-backtracking walk, d', then d' is edge neutral, and so

$$\mathrm{Abl}(G) \leq 6h.$$

Let us describe d as above. If $u = v$, then at least two of a, b, c are not inverses of each other; if a, b are such walks, then we set $d = aba^{-1}b^{-1}$. This walk is of length $4h$. If $u \neq v$, then we set d to $d = ab^{-1}ca^{-1}bc^{-1}$, which is s walk of length $6h$. For any $e \in E_G$, if e appears a net k times in a walk w then e appears a net $-k$ times in the walk w^{-1}. Thus d in either case is edge neutral.

It remains to show that d reduces to a non-empty non-backtracking walk.

The case $u = v$ and $d = aba^{-1}b^{-1}$ where a and b (are distinct and) are not inverses of each other is relatively easy. If d reduces to the empty word e and if $\alpha = [a]$, $\beta = [b]$ and $\delta = [d]$ then $\delta = [e] = [\alpha, \beta]_{\mathrm{comm}}$ where $[\ ,\ ]_{\mathrm{comm}}$ denotes the commutator. So α and β commute as elements of $\pi_1(G, u)$. Hence, by Lemma 3.3, and since $\pi_1(G, u)$ is a free group, it follows that $a = \mathrm{red}(w^m)$ and $b = \mathrm{red}(w^{m'})$ for some closed walk w with $w \neq 1$; but then by Lemma 3.4, since a and b have the same length, we have $m = \pm m'$, which contradicts the fact that a and b are distinct and not inverses of each other.

Now suppose $u \neq v$. We remark that as elements of $\pi_1(G, u)$ we have

$$d = ab^{-1}cb^{-1}ba^{-1}bc^{-1} = [ab^{-1}, cb^{-1}].$$

Let $\delta = [d]$, $\alpha = [ab^{-1}]$ and $\gamma = [cb^{-1}]$ and let e be the trivial walk. So if $d \in [e]$, we have that $1 = \delta = [\alpha, \gamma]_{\mathrm{comm}}$ and so $[ab^{-1}]$ and $[cb^{-1}]$ commute in the fundamental group. Hence $\mathrm{red}(ab^{-1}) = \mathrm{red}(w^m)$ and $\mathrm{red}(cb^{-1}) = \mathrm{red}(w^{m'})$ for some closed non-backtracking walk w. Let $w = yxy^{-1}$ with x a closed non-backtracking walk and y a walk of maximal length; then (1) $x \neq 1$ as that would imply $a = b$, (2) the first and last edges of x are not inverses of each other, and (3) $\mathrm{red}(w^m) = yx^m y^{-1}$.

Next we want to make some remarks on ab^{-1} and w^m based on the fact that $\mathrm{red}(ab^{-1}) = \mathrm{red}(w^m)$; we will later repeat similar remarks for cb^{-1} and

$w^{m'}$ based on the fact that $\mathrm{red}(cb^{-1}) = \mathrm{red}(w^{m'})$. Note

$$a = \mathrm{red}(ab^{-1}b) = \mathrm{red}(\mathrm{red}(ab^{-1})b) = \mathrm{red}(\mathrm{red}(w^m)b) = \mathrm{red}(w^m b)$$

as reductions are independent of the order in which we reduce portions of the word.

Let p be the maximal prefix of b that is a suffix of $\mathrm{red}(w^m) = yx^m y^{-1}$ (a priori p could be as large as the shorter of yxy^{-1} and b); since

$$a = \mathrm{red}(w^m b),$$

we have

$$|a| = 2|y| + |m|\,|x| + |b| - 2|p|.$$

But by assumption $|a| = |b|$, and hence

$$2|y| + |m|\,|x| = 2|p|.$$

It follows that

$$|p| = |y| + |m|\,|x|/2 > |y|, \tag{3.2}$$

since $x \neq 1$ and $m \neq 0$. Since $|p| > |y|$ and p is a suffix of $w^m = yx^m y^{-1}$, it follows that p^{-1} begins with y and contains at least one more edge; hence $p = yx_1$ where $|x_1| > 0$ and the first edge of x_1 is the inverse of the last edge of x^m.

However, since $\mathrm{red}(cb^{-1}) = \mathrm{red}(w^{m'})$, the very same arguments show that if p' is defined analogously, i.e., as the maximal prefix of b that is a suffix of $\mathrm{red}(w^{m'})$, then $p' = yx_1'$ where $|x_1'| > 0$ and the first letter of x_1' is the inverse of the last letter of $x^{m'}$. But since both p and p' are both prefixes of b, we must have that m and m' have the same sign: for if, say, $m > 0$ and $m' < 0$, then the last edge of x^m or it's orientation is different than the last edge of $x^{m'}$ by the maximality of y in the equation $w = yxy^{-1}$.

Hence

$$a = \mathrm{red}(w^m b), \quad c = \mathrm{red}(w^{m'} b),$$

a, b, c have the same length, and without loss of generality we may assume $m > m' > 0$. But then Equation (3.2) and its analog for $cb^{-1} = w^{m'}$ show that

$$m = (|p| - |y|)/|x|, \quad \text{and} \quad m' = (|p'| - |y|)/|x|.$$

It follows that $|p| > |p'|$. Since both p and p' are prefixes of b, we have that p' is a prefix of p. But p' is the maximal suffix of $w^{m'}$ that is also a prefix of b; since $w^{m'}$ is also a suffix of w^m it follows that $p' = w^{m'}$ (otherwise p could not be larger than p'). But in this case

$$|c| = |b| - |p'| < |b|,$$

which contradicts that fact that $|c| = |b|$. $\square$

Chapter 4

Proofs of the Fundamental Lemma and Theorem 2.5

Proof of the Fundamental Lemma. Let w be an edge neutral, closed, non-backtracking walk of minimal (positive) length, and let B be the subgraph of G of vertices and edges that occur in w. Then B is a connected subgraph of G such that each vertex has degree at least two. If every vertex in B has degree exactly two, then B would be a cycle, which is impossible any non-backtracing walk in a cycle traverses edges in at most one direction. Hence at least some vertex of B is strictly greater than two; hence the formula

$$\chi(B) = \sum_{v \in V_B} \bigl(2 - \deg(v)\bigr)$$

shows that $\chi(B) \geq -1$.

Now w traverses each edge of B at least once in each direction, and hence

$$l(w) \geq 2\,|E_B|.$$

We claim that if w is a figure-eight graph or a theta graph, then there is a non-backtracking walk on B traversing each edge exactly twice, and hence

$$l(w) = l_{\mathrm{Abl}}(B).$$

We claim that this formula also holds if B is a barbell graph; indeed, it suffices to show that the edges of the bar have to be traversed at least four times in w, since there is a closed, non-backtracking walk on a barbell graph that traverses each edge of the bar four times and the other edges twice. Since each edge is traversed an even number of times in w, it suffices to show that w cannot traverse an edge of the bar twice. Let e be an edge of the bar; by cyclically shifting w, we may assume that w beings by traversing $(e,+)$. Then w wraps some number of times around the cycle in this direction of e, and eventually traverses $(e,-)$, wrapping some number of times around the other cycle and returning to the tail of e. At this point each edge in the cycle has been traversed in only one direction, so that w must traverse $(e,+)$ again. Hence e cannot be traversed only twice. This establishes

$$l(w) = l_{\mathrm{Abl}}(B)$$

in the case where B is a barbell graph.

It suffices to show that for, w, as above, there is another walk, w' which is edge neutral and closed non-backtracking for which $l(w) \geq l(w')$ and the graph corresponding to w' a subgraph of B with Euler characteristic -1. So assume, for the sake of contradiction, that this fails to hold.

First we claim that B cannot contain a theta or figure-eight graph B'; indeed

$$l(w) \geq 2|E_B|;$$

but walking along the theta or figure-eight graph, there is a balanced, closed non-backtracking walk of length $2|E_{B'}|$; but $\chi(B') = -1$, so B' is a strict subgraph of the connected graph, B, and hence $2|E_{B'}| \leq 2|E_B|$, contradicting the minimality of w.

Second, we claim that B cannot contain an edge, e such that when we remove e, B is disconnected; indeed, assume otherwise, and extend e in both directions until it meets vertices v_1 and v_2 whose degrees are greater than two; let p denote the walk from v_1 to v_2 through e. The same argument used for the case where B is a barbell shows that (1) the path p is traversed in W four times, and (2) each endpoint of p is traversed in two cycles about that endpoint that do not meet p. If c_1 is the shortest cycle about v_1 that does not meet p, and similarly for c_2 and v_2, then

$$|W| \geq l_{\mathrm{Abl}}(B'),$$

where B' is the subgraph of B consisting of p, c_1, and c_2. But B' is just a barbell graph.

Third, we claim that either the first or second claim must be contradicted; this will complete the proof of the lemma. Indeed, B is not a tree, and therefore must contain a cycle, C. The cycle must contain a vertex, v, of degree at least three, or else B would consist entirely of C, which would then have Euler characteristic zero. Let e be any edge not in C that is incident upon v, and let u be the other endpoint of e. Since removing e does not disconnect B, there must be a walk from u to a vertex of C that does not contain e; let p be such a walk of minimal length. Then p begins in u and is a beaded path to a vertex of c. But then the path e followed by p is disjoint from C and joins v to another vertex of C, which yields either a figure-eight graph (if p terminates in v) or a theta graph (if p terminates in a vertex of C that is not v).

Hence the first claim is violated. $\square$

We now prove Theorem 2.5 as a consequence of the Fundamental Lemma.

Proof of Theorem 2.5. Let $G' \in \mathcal{C}$ be the subgraph of G of minimal abelian length. If G' is homeomorphic to the barbell graph, then $\mathrm{Abl}(G) \geq 4 + 4\,\mathrm{Girth}(G)$ since each of the two simple cycles in G' is at least as long as the girth. Similarly, $\mathrm{Abl}(G) \geq 4\,\mathrm{Girth}(G)$ if G' is homeomorphic to a figure-eight graph. In the case of a theta graph, we have two vertices u and v in G' of degree 3 and three simple paths from u to v. Let the lengths of the simple path be l, m and n. Then $l+m, m+n, n+l \geq \mathrm{Girth}(G)$. Adding these inequalities gives

$$3\,\mathrm{Girth}(G) \leq 2(l+m+n) = 2|E_{G'}| = \mathrm{Abl}(G).$$

$\square$

Chapter 5

Abelian Girth of the LPS Expanders

The Ramanujan graphs of Lubotzky, Phillips and Sarnak [27] are an infinite family of graphs with the largest known asymptotic girth. As mentioned earlier, theseare graphs $X^{p,q}$, for primes $p, q \equiv 1 \pmod 4$, that are $d = p+1$ regular on $n = q(q^2+1)$ vertices for which

$$\mathrm{girth}(X^{p,q}) = (4/3)\log_{d-1} n + o_n(1).$$

This immediately leads to the bound on abelian girth

$$\mathrm{Abl}(X^{p,q}) \geq 4\log_{d-1} n + o_n(1).$$

We will describe these Ramanujan graphs and obtain an upper bound on their abelian girth in order to show that their abelian girth isn't so large that it would be impossible to improve Theorem 2.6.

Definition 5.1. Let p and q be unequal primes congruent to 1 mod 4 with $(p/q) = -1$ and $q > \sqrt{p}$. The integral quaternions, denoted $H(\mathbb{Z})$, are given by

$$H(\mathbb{Z}) = \{\alpha = \alpha_0 + \alpha_1\mathbf{i} + \alpha_2\mathbf{j} + \alpha_3\mathbf{k} | a_j \in \mathbb{Z}\}.$$

We denote the conjugate of α by $\overline{\alpha}$ and define $N(\alpha) = \alpha\overline{\alpha}$. Let S be the set of all α in $H(\mathbb{Z})$ satisfying $N(\alpha) = p, \alpha \equiv 1(\mathrm{mod}2)$ and $\alpha_0 \geq 0$. It can be shown that $|S| = p+1$. Define $\Lambda'(2)$ as the set of $\alpha \in H(\mathbb{Z})$ such that $N(\alpha) = p^v$ for some non-negative integer v and $\alpha \equiv 1(\mathrm{mod}2)$. Define $\Lambda(2)$ as equivalence classes that identify α and $\beta \in H(\mathbb{Z})$ if $\pm p_1^v\alpha = p_2^v\beta$ for some $v_1, v_2 \in \mathbb{Z}$. It is known that the Cayley graph of $\Lambda(2)$ is the $(p+1)$-regular tree. Define $\Lambda(2q)$ by

$$\Lambda(2q) = \{[\alpha] \in \Lambda(2) | 2q \text{ divides } \alpha_j, j = 1, 2, 3\}.$$

This is a normal subgroup of $\Lambda(2)$ and $X^{p,q}$, the *LPS graph*, is defined as the Cayley graph of $\Lambda(2)/\Lambda(2q)$ with generators $S/\Lambda(2q)$. This graph is known to be a $(p+1)$-regular bipartite graph on $q(q^2+1)$ vertices.

Since $X^{p,q}$ is a Cayley graph it is vertex transitive which allows us to assume its smallest cycle goes from the identity element of $\Lambda(2)$ to some nontrivial element of $\Lambda(2q)$ along the infinite Cayley graph of $\Lambda(2)$. We now define a vertex' depth as its distance from the identity in the Cayley graph of $\Lambda(2)$. Biggs and Boshier [5] show that if $[b] \in \Lambda(2)$ is at depth $2r$ and $r > 0$ then there is some b in the equivalence class $[b]$ such that

$$b_0 = \pm(p^r - mq^2)$$

with $m > 0$ and even.

Definition 5.2. Positive integers are called *good* if they are not of the form $4^\alpha(8\beta + 7)$ for integers $\alpha, \beta \geq 0$.

The following is Lemma 2 of [5].

Lemma 5.3. *There exists a $[b] \in \Lambda(2q)$ at level $2r$ with $b_0 = p^r - mq^2$ with $m > 0$ and b_0 positive if and only if $2mp^r - m^2q^2$ is good.*

In the paragraph before this Lemma, Biggs and Boshier prove at least one of the integers $2mp^r - m^2q^2$ is good for the cases $m = 2$ and $m = 4$ so long as they are both positive.

The following lemma is original.

Lemma 5.4. *If $m = 4 + 8c$ for nonnegative integers c, then $2mp^r - m^2q^2$ is good if it is positive.*

Proof. Note

$$\frac{2mp^r - m^2q^2}{4} = (2 + 4c)p^r - (4 + 16c + 16c^2)q^2 \equiv 2 \mod 4$$

implying that $2mp^r - m^2q^2$ is good. □

So for $m = 4, 12, 20$, $2mp^r - m^2q^2$ is good if we can show it is positive. Let r_0 be the smallest positive integer such that $p^{r_0} > 10q^2$, which makes $2mp^{r_0} - m^2q^2$ positive for those values of m. Then $p^{r_0 - 1} < 10q^2$ and so

$$r_0 < 2\log_p q + \log_p 10 + 1.$$

Thus there exists three distinct $[b] \in \Lambda(2q)$, since each value of m produces a different b_0, which means there are three distinct closed walks of length

r_0 from the identity vertex to itself of length $2r_0$. From the proof of Theorem 2.6 we showed that this situation would imply the abelian girth of $X^{p,q}$ is $8r_0$. Since $n = q(q^2+1)$ and $d-1 = p$ we have

$$\mathrm{Abl}(X^{p,q}) \leq \frac{16}{3}\log_{d-1} n(1+o(1)).$$

The 16/3 constant above may be improved. In our arguments, we found three closed walks of length $2r_0$. If instead we found three walks share one starting vertex and another terminating vertex, all walks of length r_0, then our arguments from we showed that this situation would imply would mean we could instead have a constant of 4 in the previous equation. Any closed walk from a vertex v to itself of length $2r_0$ already implies there exists two distinct walks of length r_0 from v to the m, the middle vertex of the cycle. Identifying one more walk from v to m of length r_0 would be sufficient to improve the 16/3 coefficient.

Chapter 6

k-Independence

In this section, we introduce one of the main results from our paper in terms of only linear algebra. The main concepts discussed here are linear and tensorial k-independence. In the section following this one, we show a connection between our two forms of 2-independence and abelian girth using sheaf theory.

6.1 Oriented Graphs

In this subsection we fix some terminology on graphs and morphisms of graphs, and define some invariants of graphs that will be guiding examples in graph homology. We may state easy results without proof; in some cases more details may be found in [17].

Even though the graphs we consider are undirected, it is more natural to state the theory in terms of directed graphs (and arbitrarily orient the edges of every undirected graph). There is an equivalent formulation without the choice of orientation, but at the cost of more cumbersome notation; for more on this see [17]. Our choice of formulation motivates the following choice of nomenclature:

By an *oriented graph* (henceforce simply "graph") we shall mean a quadruple $G = (V_G, E_G, t_G, h_G)$ where V_G and E_G are sets (the "*vertices*" and "*edges*", respectively) and $t_G, h_G \colon E_G \to V_G$ are maps (the "tail" and "head" of each edge, respectively). We also sometimes use the notations $e_G^+ = h_G$ and $e_G^- = t_G$

Note that we allow multiple edges and self-loops (but not half-edges). Unless otherwise indicated all graphs are assumed finite (that is, the sets V_G, E_G are finite).

While our graphs are oriented, we shall treat them as undirected for graph-theoretic purposes, that is allow paths in the graphs to traverse each edge in either direction. Formally, a *walk* of *length* $m \geq 1$ in a graph G shall be a sequence of pairs $(e_1, s_1), \dots, (e_m, s_m)$ where $e_i \in E_G$ for each i, $s_i \in \{\pm\}$ and $e_G^{s_i}(e_i) = e_G^{-s_{i+1}}(e_{i+1})$ for $1 \leq i < m$. The vertices $e_G^{-s_1}(e_1)$,

$e_G^{s_m}(e_m)$ will be called the *starting* and *terminating* vertices of the walk, and jointly the *endpoints*.

This level of formality is mainly necessary for graphs with multiple edges or self-loops, since a self-loop may be traversed along or against its orientation and this cannot be determined from knowing the order at which the endpoint of the edge are met.

A walk is *closed* if its starting vertex is the same as its terminating vertex. A walk in which all the edges are distinct is called *simple*, and a closed, simple walk is also known as a *cycle*. Recall that a graph is *acyclic* if it has no cycles, *connected* if every two vertices in G are the endpoints of a walk, and a *tree* if it is connected and acyclic.

A walk is a *path* if all the edges and vertices of the walk are distinct. If G_1 and G_2 are subgraphs of a graph G, by a *walk (or path) from G_1 to G_2* we mean a walk (or path) in G with starting vertex in G_1, terminating vertex in G_2 and no other vertex of the walk in G_1 or G_2. We say a path or cycle is *beaded* if every vertex is of degree 2 besides possibly the starting and terminating vertex.

A *morphism* of graphs, $f\colon G \to H$, is a pair $f = (f_V, f_E)$ of maps $f_V\colon V_G \to V_H$ and $f_E\colon E_G \to E_H$ such that $t_H \circ f_E = f_V \circ t_G$ and $h_H \circ f_E = f_V \circ h_G$. We shall usually drop the subscripts from f_V and f_E, but may include them in the interests of clarity. The set of morphisms will be denoted $\mathrm{Mor}(G, H)$. Thus morphisms are homomorphisms of the underlying undirected graphs which preserve the orientation.

We say that $\pi \in \mathrm{Mor}(G, B)$ is a *covering map* (respectively, is étale[2]) if for each $v \in V_G$, π_E gives a bijection (respectively, injection) of incoming edges of v (i.e. those edges whose head is v) with those of $f_V(v)$, and a bijection (respectively, injection) of outgoing edges of v and $\pi_V(v)$.

If $\pi\colon G \to B$ is a covering map and B is connected, then the *fibres* $\pi^{-1}(v)$ ($v \in V_H$) and $\pi^{-1}(e)$ ($e \in E_H$) are all in bijection, and we call their joint cardinality the *degree* of f and denote it $[G : H]$. Even if H is not connected, one can still write $[G : H]$ when π is of *constant degree*, that is when the number of preimages of a vertex or edge in H is the same for all vertices and edges.

6.2 k-Independence

In this subsection we describe two notions of k-independence and give an example showing they are not equivalent. In later sections, we will show

[2]Stallings, in [31], uses the term "immersion."

how the two notions of 2-independence correspond to certain homological invariants of certain sheaves on graphs.

Definition 6.1. Let $A_1, \ldots, A_r$ be subspaces of a finite-dimensional vector space W over a field $\mathbb{F}$. We say that $A_1, \ldots, A_r$ are *linearly k-independent* with k a positive integer if for all subspaces $B \subset W$ we have

$$\sum_{i=1}^{r} \dim(A_i \cap B) \leq k \dim(B). \tag{6.1}$$

We define the *total dimension* of the subspaces to be $|A_1| + \ldots + |A_r|$

This notion arises implicitly in quiver representation theory in [25], where the case $k = 2$ is of special interest. We shall explain in this paper that the $k = 2$ case also arises in what we call *pseudobundles* on a graph. Notice that k-independence could easily be defined for any real number k.

It is easy to see that r vectors in a vector space, W, are linearly independent (in the usual sense) iff the one-dimensional spaces that the vectors span are linearly 1-independent.

Definition 6.2. Let $A_1, \ldots, A_r$ be subspaces of a vector space, W, over a field, $\mathbb{F}$, and let $\overline{\mathbb{F}}$ be an algebraic closure of $\mathbb{F}$. Let k be a positive integer. We say that $A_1, \ldots, A_r$ are *tensorially k-independent* if for any $f_1, \ldots, f_r \in \overline{\mathbb{F}}^k$ and any $a_i \in \overline{A}_i$ for $1 \leq i \leq r$ we have that

$$f_1 \otimes a_1 + \ldots + f_r \otimes a_r = 0$$

only if

$$a_1 = \ldots = a_r = 0$$

or the jth components of the $f_1, \ldots, f_r$ are all 0 for some $1 \leq j \leq k$.

It is easy to see that r vectors in a vector space, W, are linearly independent (in the usual sense) iff the one-dimensional spaces that the vectors span are tensorially 1-independent. It is also easy to see that if the A_i are not tensorially k-independent, they also are not tensorially $k+1$-independent

We claim that tensorial k-independence implies linear k-independence (this is also shown in[18].) For this, consider $A_1, \ldots, A_r \subset W$ that are not linearly k-independent; then there exists a $B \subset W$ such that

$$\dim(A_1 \cap B) + \cdots + \dim(A_r \cap B) > k \dim(B).$$

For a vector space, U, over $\mathbb{F}$, let $\overline{U}$ denote the corresponding vector space over $\overline{\mathbb{F}}$, namely $U \otimes_{\mathbb{F}} \overline{\mathbb{F}}$. Consider the map

$$\overline{A_1 \cap B} \oplus \cdots \oplus \overline{A_r \cap B} \to \overline{B}$$

given by

$$f_1 \otimes a_1 + \ldots + f_r a_r$$

for some $f_1, \ldots, f_r \in \overline{\mathbb{F}}^k$ and $a_i \in \overline{A}_i \cap \overline{B}$. Since the dimension of the domain is greater than that of the range, there must exist a nontrivial element in the kernel.

Theorem 6.3. *Linear* 2*-independence does not imply tensorial* 2*-independence.*

Here is an example of a collection of subspaces that are linearly but not tensorially 2-independent. If $W = \mathbb{F}^6$ has basis $\alpha, \beta, \gamma, \delta, \epsilon, \zeta$ then the subspaces are

$$A_1 = \text{Span}(\alpha, \delta), \quad A_2 = \text{Span}(\beta, \epsilon), \quad A_3 = Span(\gamma, \zeta),$$

$$A_4 = \text{Span}(\beta - \gamma, \delta - \epsilon, \zeta - \alpha), \quad A_5 = \text{Span}(\alpha - \beta, \gamma - \delta, \epsilon - \zeta).$$

We will provide sheaf theoretic tools in later sections in order to prove these subspaces are linearly but not tensorially 2-independent. We also conjecture that this is the smallest such collection of subspaces in terms of total dimension of the subspaces.

Chapter 7

Sheaves on Graphs

In this section we describe the notion of "sheaf on a graph" and show the 2-independence problem is a special case of a general problem concerning two homological invariants of certain sheaves on graphs with two vertices.

7.1 Sheaves and Homology

Definition 7.1. Let $G = (V, E, t, h) = (V_G, E_G, t_G, h_G)$ be a directed graph, and $\mathbb{F}$ a field. By a *sheaf of finite dimensional* $\mathbb{F}$*-vector spaces on* G, or simply *an* $\mathbb{F}$*-sheaf on* G, we mean the data, $\mathcal{F}$, consisting of

1. a finite dimensional $\mathbb{F}$-vector space, $\mathcal{F}(v)$, for each $v \in V$,
2. a finite dimensional $\mathbb{F}$-vector space, $\mathcal{F}(e)$, for each $e \in E$,
3. a linear map, $\mathcal{F}(t,e)\colon \mathcal{F}(e) \to \mathcal{F}(te)$ for each $e \in E$,
4. a linear map, $\mathcal{F}(h,e)\colon \mathcal{F}(e) \to \mathcal{F}(he)$ for each $e \in E$,

We will often write just refer to the sheaf, $\mathcal{F}$, with the graph, G, being implicit. The vector spaces $\mathcal{F}(P)$, ranging over all $P \in V_G \amalg E_G$ ($\amalg$ denoting the disjoint union), are called the *values* of $\mathcal{F}$. The morphisms $\mathcal{F}(t,e)$ and $\mathcal{F}(h,e)$ are called the *restriction maps.* If U is a finite dimensional vector space over $\mathbb{F}$, the *constant sheaf associated to* U, denoted $\underline{U}$, is the sheaf comprised of the value U at each vertex and edge, with all restriction maps being the identity map. The constant sheaf $\underline{\mathbb{F}}$ will be called the *structure sheaf* of G (with respect to the field, $\mathbb{F}$), for reasons to be explained later. We say that $\mathcal{F}_1$ is a *subsheaf* of $\mathcal{F}$ if $\mathcal{F}_1$ is a sheaf on G whose value at any vertex or edge is a subspace of the value of $\mathcal{F}$ at that vertex or edge, and if the restriction maps of $\mathcal{F}_1$ are induced by those of $\mathcal{F}$. Subsheaves of constant sheaves will be called *subconstant sheaves.*

To a sheaf, $\mathcal{F}$, on a digraph, G, we set

$$\mathcal{F}(E) = \bigoplus_{e \in E} \mathcal{F}(e), \quad \mathcal{F}(V) = \bigoplus_{v \in V} \mathcal{F}(v).$$

We associate a transformation

$$d_h = d_{h,\mathcal{F}} \colon \mathcal{F}(E) \to \mathcal{F}(V)$$

defined by taking $\mathcal{F}(e)$ (viewed as a component of $\mathcal{F}(E)$) to $\mathcal{F}(he)$ (a component of $\mathcal{F}(V)$) via the map $\mathcal{F}(h,e)$. Similarly we define d_t. We define the *differential of* $\mathcal{F}$ to be

$$d = d_{\mathcal{F}} = d_h - d_t,$$

and define the *Euler characteristic of* $\mathcal{F}$ to be

$$\chi(\mathcal{F}) = \dim\bigl(\mathcal{F}(V)\bigr) - \dim\bigl(\mathcal{F}(E)\bigr).$$

We call the sheaf $\mathcal{F}$ a *constant sheaf* if $\mathcal{F}(e) = \mathcal{F}(v)$ for all vertices v and edges e in G and the difference maps are all the identity.

Definition 7.2. We define the *zeroth* and *first homology groups of* $\mathcal{F}$ to be, respectively,

$$H_0(G,\mathcal{F}) = \text{cokernel}(d), \quad H_1(G,\mathcal{F}) = \text{kernel}(d).$$

We denote by $h_i(G,\mathcal{F})$ the dimension of $H_i(G,\mathcal{F})$ as an $\mathbb{F}$-vector space, and call it the *i-th Betti number of* $\mathcal{F}$. This definition agrees with sheaf homology theory. We often just write $h_i(\mathcal{F})$ and $H_i(\mathcal{F})$ if G is clear from the context (when no confusion will arise between $h_i(\mathcal{F})$, the dimension, and h the head map of a graph). We call $H_i(\underline{\mathbb{F}})$ the i-th homology group of G with coefficients in $\mathbb{F}$, denoted $H_i(G)$ or, for clarity, $H_i(G,\underline{\mathbb{F}})$. This agrees with the standard definition of $H_i(G)$.

In particular, we have

$$\chi(\mathcal{F}) = h_0(\mathcal{F}) - h_1(\mathcal{F}).$$

Example 7.3. For $\mathcal{F} = \underline{\mathbb{F}}$, d is just the usual incidence matrix; thus, if $\mathbb{F}$ is of characteristic zero, then the $h_i(G)$, i.e., the dimension of the $H_i(G)$, are the usual Betti numbers of G.

Example 7.4. For any morphism $\phi \colon G' \to G$ of digraphs, and any sheaf $\mathcal{F}$ on G', there is a sheaf $\phi_!\mathcal{F}$ such that $H_i(G',\mathcal{F})$ is naturally isomorphic to $H_i(G,\phi_!\mathcal{F})$. See [16, 18] for more details on this and other examples of sheaves.

One of the main tools in sheaf homology is the ability to get information about them by the long exact sequence in homology associated to a short exact sequence of sheaves. For this we need to describe what is meant by a morphism of sheaves.

Definition 7.5. A morphism of sheaves $\alpha\colon \mathcal{F} \to \mathcal{G}$ on G is a collection of linear maps $\alpha_v\colon \mathcal{F}(v) \to \mathcal{G}(v)$ for each $v \in V$ and $\alpha_e\colon \mathcal{F}(e) \to \mathcal{G}(e)$ for each $e \in E$ such that for each $e \in E$ we have $\mathcal{G}(t,e)\alpha_e = \alpha_{te}\mathcal{F}(t,e)$ and $\mathcal{G}(h,e)\alpha_e = \alpha_{he}\mathcal{F}(h,e)$.

It is not hard to check that all Abelian operations on sheaves, e.g., taking kernels, taking direct sums, checking exactness, can be done "vertexwise and edgewise," i.e., $\mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3$ is exact iff for all $P \in V_G \amalg E_G$, we have $\mathcal{F}_1(P) \to \mathcal{F}_2(P) \to \mathcal{F}_3(P)$ is exact. This is actually well known, since our sheaves are presheaves of vector spaces on a category (see [14] or Proposition I.3.1 of [1]).

The following theorem results from a straightforward application of classical homological algebra.

Theorem 7.6. *To each "short exact sequence" of sheaves, i.e.,*

$$0 \to \mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3 \to 0$$

(in which the kernel of each arrow is the image of the preceding arrow), there is a natural long exact sequence of homology groups

$$0 \to H_1(\mathcal{F}_1) \to H_1(\mathcal{F}_2) \to H_1(\mathcal{F}_3) \xrightarrow{\delta}$$

$$H_0(\mathcal{F}_1) \to H_0(\mathcal{F}_2) \to H_0(\mathcal{F}_3) \to 0.$$

7.2 Twisted Homology

Here we define twisted homology. We refer to [16, 17] for its motivation; in brief, it represents a "scaled Abelian limit" of ordinary homology, and has "reduced cyclicity" as a special case. Here we give a self-contained description of these invariants, and the reader who does not consult [16, 18] may regard twisted homology as simply an alternate homology theory.

Definition 7.7. Let $\mathcal{F}$ be a sheaf of $\mathbb{F}$-vector spaces on a digraph, G, and let $\mathbb{F}'$ be a field containing $\mathbb{F}$. A *twist* or $\mathbb{F}'$*-twist*, ϕ, on G is a map

$$\phi\colon E_G \to \mathbb{F}'.$$

By the *twisting* of $\mathcal{F}$ by ϕ, denoted $\mathcal{F}^\phi$, we mean the sheaf of $\mathbb{F}'$-vector spaces given via

$$\mathcal{F}^\phi(P) = \big(\mathcal{F}(P)\big) \otimes_{\mathbb{F}} \mathbb{F}'$$

for all $P \in V_G \amalg E_G$, and

$$\mathcal{F}^\phi(h,e) = \mathcal{F}(h,e), \quad \mathcal{F}^\phi(t,e) = \phi(e)\mathcal{F}(t,e),$$

where $\mathcal{F}(h,e)$ and $\mathcal{F}(t,e)$ are viewed as $\mathbb{F}'$-linear maps arising from their original $\mathbb{F}$-linear maps.

In other words, $\mathcal{F}^\phi$ is the sheaf obtained by extending scalars to $\mathbb{F}'$ and twisting the tail maps. The map, $d_{\mathcal{F}^\phi}$, viewed as a matrix, has entries in the field $\mathbb{F}'$ and the groups $H_i(\mathcal{F}^\phi)$ are $\mathbb{F}'$-vector spaces.

Definition 7.8. Let $\mathcal{F}$ be a sheaf of $\mathbb{F}$-vector spaces on a digraph, G. By the *full twist of* $\mathcal{F}$ we mean the twist $E_G \to \mathbb{F}(\psi)$, where $\psi = \{\psi(e)\}_{e\in E_G}$ is a collection of $|E_G|$ independent indeterminates, and $\mathbb{F}(\psi)$ is the field of rational functions in the $\{\psi(e)\}$ over $\mathbb{F}$, and where the map $E_G \to \mathbb{F}(\psi)$ is given by $e \mapsto \psi(e)$. We shall refer to this twist as ψ when no confusion will arise. The twist ψ can be defined in the same way given a graph G and a field $\mathbb{F}$ even if no sheaf is specified. In this case we refer to ψ as *the full twist of* G.

In the above situation, $d = d_{\mathcal{F}^\psi}$ can be viewed as a morphism of finite dimensional vector spaces over $\mathbb{F}(\psi)$, given by a matrix with entries in $\mathbb{F}(\psi)$.

Definition 7.9. We define the *i-th twisted homology group of* $\mathcal{F}$, denoted by

$$H_i^{\mathrm{tw}}(\mathcal{F}) = H_i^{\mathrm{tw}}(\mathcal{F},\psi),$$

for $i = 0, 1$, respectively, to be the cokernel and kernel, respectively, of $d_{\mathcal{F}^\psi}$ described above as a morphism of $\mathbb{F}(\psi)$ vector spaces. We define the *i-th twisted Betti number of* $\mathcal{F}$, denoted $h_i^{\mathrm{twist}}(\mathcal{F})$, to be dimension of $H_i^{\mathrm{tw}}(\mathcal{F})$.

Notice that twisted homology as above is just the homology of certain sheaves over $\mathbb{F}(\psi)$. In particular, the following results from applying Theorem 7.6.

Theorem 7.10. *Let $\mathbb{F}$ be a field, and let ψ be a full twist on G. For any short exact sequence*

$$0 \to \mathcal{F}_1 \to F_2 \to \mathcal{F}_3 \to 0$$

of sheaves of $\mathbb{F}$-vector spaces on G, there is a long exact sequence of $\mathbb{F}(\psi)$ vector spaces,

$$0 \to H_1^{\mathrm{tw}}(\mathcal{F}_1) \to H_1^{\mathrm{tw}}(\mathcal{F}_2) \to H_1^{\mathrm{tw}}(\mathcal{F}_3) \xrightarrow{\delta}$$
$$H_0^{\mathrm{tw}}(\mathcal{F}_1) \to H_0^{\mathrm{tw}}(\mathcal{F}_2) \to H_0^{\mathrm{tw}}(\mathcal{F}_3) \to 0.$$

7.3 Maximum Excess of a Sheaf

Definition 7.11. Let $\mathcal{F}$ be a sheaf on a graph, G. We define the *excess* of $\mathcal{F}$ to be

$$\mathrm{excess}(\mathcal{F}) = -\chi(\mathcal{F}).$$

We define the *maximum excess* of $\mathcal{F}$ to be the maximum excess over all subsheaves of $\mathcal{F}$, i.e.,

$$\mathrm{m.e.}(\mathcal{F}) = \max_{\mathcal{F}'\subset\mathcal{F}} \mathrm{excess}(\mathcal{F}').$$

The excess and maximum excess has a number of remarkable properties that are established in [16, 18]; let us mention a few here without proof, referring the reader to [16, 18] for more details and the proofs.

Let $\mathcal{F}$ be a sheaf on a graph, G. The dimension of the kernel of any linear map from $\mathcal{F}(E)$ to $\mathcal{F}(V)$ is at least

$$\dim\bigl(\mathcal{F}(E)\bigr) - \dim\bigl(\mathcal{F}(V)\bigr) = -\chi(\mathcal{F}) = \mathrm{excess}(\mathcal{F}).$$

If ψ is a full twist on G, then the map $d_{\mathcal{F}\psi}$ resticts to a map from $\mathcal{F}'(E)$ to $\mathcal{F}'(V)$ for any subsheaf, $\mathcal{F}'$, of $\mathcal{F}$; hence

$$h_1^{\mathrm{tw}}(\mathcal{F}) \geq \mathrm{m.e.}(\mathcal{F}).$$

Definition 7.12. The *gap* of a sheaf, $\mathcal{F}$, is the non-negative integer

$$\mathrm{gap}(\mathcal{F}) = h_1^{\mathrm{tw}}(\mathcal{F}) - \mathrm{m.e.}(\mathcal{F}).$$

We say that a sheaf *is gapped* or *has a gap* if its gap is positive; otherwise we say that the sheaf *has no gap*.

This paragraph is written to give the reader some intuition for the gap, h_1^{tw}, and the maximum excess, though the results mentioned here will also be useful later on. In [16, 17] Friedman defines the pullback, $\mu^*\mathcal{F}$, for any graph morphism $\mu\colon G'\to G$ and any sheaf on G (in a natural and standard fashion); he shows that if μ is a covering map then

$$\mathrm{m.e.}(\mu^*\mathcal{F}) = \deg(\mu)\mathrm{m.e.}(\mathcal{F}), \tag{7.1}$$

and he shows that for any fixed sheaf, $\mathcal{F}$, on G, we have

$$\mathrm{m.e.}(\mu^*\mathcal{F}) = h_1^{\mathrm{tw}}(G',\mu^*\mathcal{F}), \tag{7.2}$$

provided that $\mu\colon G' \to G$ is a covering map and the girth of G' is sufficiently large. Therefore, if we order covering maps under refinement, we can say that any sheaf has no gap if pulled-back via a covering map to a graph of sufficiently large girth. We can also write

$$\mathrm{m.e.}(\mathcal{F}) = \lim_{\mu} \frac{h_1^{\mathrm{tw}}(\mu^*\mathcal{F})}{\deg(\mu)}, \tag{7.3}$$

where the limit is taken over the directed set of covering maps, u, under refinement.

As a consequence of (7.3), we have that the maximum excess is additive, i.e., if $\mathcal{F}_1, \mathcal{F}_2$ are sheaves on a graph, then

$$\mathrm{m.e.}(\mathcal{F}_1 \oplus \mathcal{F}_2) = \mathrm{m.e.}(\mathcal{F}_1) + \mathrm{m.e.}(\mathcal{F}_2);$$

this follows since h_1^{tw} is additive and (arbitrary) pullbacks are additive. Similarly, we get inequalities from short exact sequences, which we now describe.

Definition 7.13. For a sheaf, $\mathcal{F}$, on a graph, we define its *dual maximum excess*, denoted d.m.e.$(\mathcal{F})$ to be

$$\mathrm{d.m.e.}(\mathcal{F}) = \chi(\mathcal{F}) + \mathrm{m.e.}(\mathcal{F}).$$

We easily see that

$$\mathrm{d.m.e.}(\mathcal{F}) = \max_{\mathcal{F} \twoheadrightarrow \mathcal{F}'} \chi(\mathcal{F}'),$$

since

$$\mathrm{d.m.e.}(\mathcal{F}) = \chi(\mathcal{F}) + \max_{\mathcal{F}'' \subset \mathcal{F}} -\chi(\mathcal{F}'') = \max_{\mathcal{F}'' \subset \mathcal{F}} \chi(\mathcal{F}/\mathcal{F}'').$$

It follows that the dual maximum excess is non-negative and

$$\mathrm{d.m.e.}(\mathcal{F}) - \mathrm{m.e.}(\mathcal{F}) = \chi(\mathcal{F}) = h_0^{\mathrm{tw}}(\mathcal{F}) - h_1^{\mathrm{tw}}(\mathcal{F}),$$

for any sheaf, $\mathcal{F}$; in particular

$$h_1^{\mathrm{tw}}(\mathcal{F}) - \mathrm{m.e.}(\mathcal{F}) = \mathrm{gap}(\mathcal{F}) = h_0^{\mathrm{tw}}(\mathcal{F}) - \mathrm{d.m.e.}(\mathcal{F}),$$

which gives an alternate expression for the gap of a sheaf.

Theorem 7.14. *To each short exact sequence of sheaves,*

$$0 \to \mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3 \to 0,$$

the sequence of non-negative integers

$$\ldots, 0, 0, \mathrm{m.e.}(\mathcal{F}_1), \mathrm{m.e.}(\mathcal{F}_2), \mathrm{m.e.}(\mathcal{F}_3),$$

$$\mathrm{d.m.e.}(\mathcal{F}_1), \mathrm{d.m.e.}(\mathcal{F}_2), \mathrm{d.m.e.}(\mathcal{F}_3), 0, 0, \ldots$$

is "triangular," meaning any element of the sequence is at most the sum of its successor and its predecessor.

This follows from (7.3) and Theorem 7.6. Alternatively, this can be proven from scratch (see Warren Dicks' appendix in [18].)

The following definition is a crucial observation to the proof of the Hanna Neumann Conjecture in [15, 18] and to this paper.

Definition 7.15. A *maximizer* of the sheaf $\mathcal{F}$ is any subsheaf whose excess is the maximum excess of $\mathcal{F}$.

If $\mathcal{F}_1, \mathcal{F}_2$ are subsheaves of the sheaf $\mathcal{F}$ one can easily verify that

$$\chi(\mathcal{F}_1) + \chi(\mathcal{F}_2) = \chi(\mathcal{F}_1 \cap \mathcal{F}_2) + \chi(\mathcal{F}_1 + \mathcal{F}_2).$$

It follows that the set of maximizers of $\mathcal{F}$ is closed under intersection and addition of subsheaves. In particular, each sheaf $\mathcal{F}$ has a unique maximum (or maximal) maximizer (namely the sum of all the maximizers) and unique minimum (or minimal) maximizer (namely the intersection of all the maximizers).

It is not hard to see that the structure sheaf, $\underline{\mathbb{F}}$, of a connected graph G has

$$\mathrm{m.e.}(G, \underline{\mathbb{F}}) = \max\bigl(0, -\chi(G)\bigr),$$

and that

$$\mathrm{d.m.e.}(G, \underline{\mathbb{F}}) = \begin{cases} 1 & \text{if } G \text{ is cyclic, and} \\ 0 & \text{if } G \text{ is acyclic,} \end{cases}$$

where G is acyclic if G is a vertex or a tree, and otherwise G is cyclic; if G is not connected, then the maximum excess of $\underline{\mathbb{F}}$ is the sum composed of the maximum excess of each of G's connected components, and similarly with "maximum excess" replaced with "dual maximum excess."

7.4 Main Theorems

In this section we fix some terminology regarding sheaves on graphs and formally describe the main theorems of the second part of this dissertation.

Definition 7.16. Let $\mathcal{F}$ be a sheaf on a graph. By a *subquotient* of $\mathcal{F}$ we mean any sheaf of the form $\mathcal{F}_2/\mathcal{F}_1$, where $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}$.

Theorem 7.17. *If G is a graph on two vertices, no self loops, and at least five edges, then there exists a gapped subconstant sheaf on G with maximum excess* 0*.*

Theorem 6.3 follows quickly from this theorem in the following way. Let $\mathcal{F}$ be the gapped subquotient sheaf on the graph G with vertices v_1 and v_2 produced by the previous theorem. Let the vector space $W = \mathcal{F}(v_1) + \mathcal{F}(v_2)$. If G has edges $e_1, ..., e_r$, then let $A_i = \mathcal{F}(e_i)$ for $i = 1, \ldots, r$. For any subconstant sheaf the head and tail maps are inclusions since they are induced by identity maps, and so the A_i are subspaces of W. Since $\text{m.e.}(\mathcal{F}) = 0$, for any $B \subset W$

$$\sum_{i=1}^{r} \dim(A_i \cap B) \leq 2\dim(B),$$

showing that the A_i are linearly 2-independent.

Given any nonzero $\alpha \in H_1^{\text{tw}}(\mathcal{F})$, let a_i be α restricted to the edge space on e_i. Then

$$\sum_{i=1}^{r} f_i \otimes a_i = 0$$

where $f_i = (1, -\psi_{e_i})$. This shows that the A_i are not tensorially 2-independent.

The proof for Theorem 7.17 is constructive, using a sheaf that is modified from a minimally gapped sheaf necessary in the proof of the following theorem.

Theorem 7.18. *For any graph G we have*

$$\text{Abl}(G) = \min\{\dim(\mathcal{F}(E)) | \mathcal{F} \textit{ is a gapped sheaf on } G\}.$$

The proof of this theorem includes explicit constructions of gapped sheaves with minimal edge dimension for any graph. This theorem implies that sheaf theory could potentially be useful for finding better upper bounds on the girth of a regular graph.

Corollary 7.19. *Let G be a graph with $\chi(G) < 0$. Then there exists a subconstant sheaf on G with positive gap. If $\mathcal{F}$ is a subquotient of a constant sheaf on G with positive gap then* $\dim(\mathcal{F}(E)) \geq 6$.

The smallest edge dimension known to us for a a gapped subconstant sheaf on a graph is 12. Before this research, it was not known whether gapped subconstant sheaves even existed.

Proving the results in this subsection will occupy the rest of this paper.

Chapter 8

Remarks on Gapped Sheaves

Recall that that gap of a sheaf, $\mathcal{F}$, on a graph is defined to be

$$\mathrm{gap}(\mathcal{F}) = h_1^{\mathrm{tw}}(\mathcal{F}) - \mathrm{m.e.}(\mathcal{F}) = h_0^{\mathrm{tw}}(\mathcal{F}) - \mathrm{d.m.e.}(\mathcal{F})$$

and that we say that $\mathcal{F}$ *has a gap* (or *is gapped*) if its gap is positive.

The goal of the next few sections is to prove that certain sheaves have no gap. Furthermore, when a collection of sheaves contains gapped sheaves, we wish to describe the "simplest" example of a gapped sheaf in the collection. This section contains some important definitions and some easy observations regarding these matters.

As an example, Theorems 7.6 and 7.14 imply that $h_1^{\mathrm{tw}}, h_0^{\mathrm{tw}}$ and the maximum excess and dual maximum excess are additive invariants, and hence so is the gap; i.e.,

$$\mathrm{gap}(\mathcal{F}_1 \oplus \mathcal{F}_2) = \mathrm{gap}(\mathcal{F}_1) + \mathrm{gap}(\mathcal{F}_2)$$

for any sheaves, $\mathcal{F}_1, \mathcal{F}_2$ on a graph. Hence if $\mathcal{F}_1 \oplus \mathcal{F}_2$ is gapped, then so is $\mathcal{F}_1$ or $\mathcal{F}_2$. On certain collections of sheaves, with a certain partial order, this will mean that a minimal gapped sheaf (if it exists) is not a proper direct sum. However, we need to specify the collections of sheaves and partial orders we use.

8.1 Minimality

In this subsection we define what is a "simplest" or "minimal" sheaf in a collection, in a sense useful to our theory. This section consists only of definitions and very simple remarks.

Definition 8.1. Let $\mathcal{F}$ be a sheaf on a graph, G. The *total dimension (or T-dim) of* $\mathcal{F}$, denoted $\mathrm{T} - \dim(\mathcal{F})$ shall mean the quantity

$$\mathrm{T} - \dim(\mathcal{F}) = \sum_{P \in V_G \amalg E_G} \dim\bigl(\mathcal{F}(P)\bigr).$$

Definition 8.2. By a *sheaf collection* we mean a set of pairs, $(G, \mathcal{F})$, where G is a graph and $\mathcal{F}$ is a sheaf on G.

If $\mathcal{C}$ is a sheaf collection, we may write $\mathcal{F} \in \mathcal{C}$ instead of the more cumbersome $(G, \mathcal{F}) \in \mathcal{C}$ when G is implicit.

Definition 8.3. Given a sheaf collection, $\mathcal{C}$, a *T-minimal* element of $\mathcal{C}$ is an $\mathcal{F} \in \mathcal{C}$ of minimal total dimension.

Clearly any nonempty sheaf collection has T-minimal element.

Definition 8.4. Given a sheaf collection, $\mathcal{C}$, a *T-minimal gapped* element is a T-minimal element among the subcollection of sheaves in $\mathcal{C}$ which have a positive gap (if this subcollection is nonempty).

8.2 Minimal Elements and Stability

We now wish to state the main result in this section. It will be proved over the next few subsections. First though, we include some addition definitions related to the ones we need.

Definition 8.5. The sheaf $\mathcal{F}$ is called

1. *cyclic* if $\text{m.e.}(\mathcal{F}) = -\chi(\mathcal{F})$, i.e., $\chi(\mathcal{F}') \geq \chi(\mathcal{F})$ for each $\mathcal{F}' \subset \mathcal{F}$;
2. *semistable* if it is cyclic, and $\chi(\mathcal{F}) = 0$;
3. *stable* if it is semistable, and for all $\mathcal{F}' \subset \mathcal{F}$ we have $\chi(\mathcal{F}') > \chi(\mathcal{F})$ unless $\mathcal{F}'$ is 0 or $\mathcal{F}$;
4. *superstable* if for all $\mathcal{F}' \subset \mathcal{F}$ we have $\chi(\mathcal{F}') > \chi(\mathcal{F})$ unless $\mathcal{F}'$ is $\mathcal{F}$.

The minimum maximizer of $\mathcal{F}$ is called its *supercore*.

A supercore is therefore necessarily superstable. Our terminology is borrowed from [25] and [7]. It is easy to see that for $\mathcal{F}' \subset \mathcal{F}$, if $\mathcal{F}$ is semistable, stable, or superstable, then so is $\mathcal{F}/\mathcal{F}'$. Furthermore, if $\mathcal{F}$ is stable and $\mathcal{F}' \neq \mathcal{F}$, then $\mathcal{F}/\mathcal{F}'$ is superstable.

A graph is cyclic if it contains no components that are trees; it is semistable if its connected components all have Euler characteristic zero (i.e., when its leaves are repeatedly "pruned," the graph becomes a union of cycles); it is stable if it consists of a single cycle; it is superstable if all its connected components have no leaves and have negative Euler characteristic.

One can also define dual or "co" notions of stability.

Definition 8.6. A sheaf, $\mathcal{F}$, on a graph is called *co-acyclic*, *co-semistable*, *co-stable*, *co-superstable*, by taking the respective notion in Definition 8.5 (omitting the prefix "co-") and replacing "maximum excess" with "dual maximum excess," "χ" with "$-\chi$," and "$\mathcal{F}' \subset \mathcal{F}$" with "$\mathcal{F}'$ is a quotient of $\mathcal{F}$."

Generally speaking, to most of the notions and lemmas in this section, there is a "co" or "dual" notion; typically a lemma and its dual lemma are different and both important to proving Theorem 8.9 below. We mention, however, that stability, one of the cornerstones of this article, is equivalent to its dual notion, co-stability. Indeed, to every subsheaf, $\mathcal{F}'$, of a sheaf, $\mathcal{F}$, there is a short exact sequence

$$0 \to \mathcal{F}' \to \mathcal{F} \to \mathcal{F}'' \to 0,$$

(with $\mathcal{F}'' = \mathcal{F}/\mathcal{F}'$), and to every quotient, $\mathcal{F}''$, of $\mathcal{F}$ there is a short exact sequence (with $\mathcal{F}'$ the kernel of the map $\mathcal{F} \to \mathcal{F}''$). So whenever $\chi(\mathcal{F}) = 0$, the condition of having a subsheaf with non-negative Euler characteristic is equivalent to having a quotient sheaf with non-positive Euler characteristic. Hence stability and co-stability amount to the same thing.

Definition 8.7. We say that a sheaf collection, $\mathcal{C}$, is *closed under subsheaves and quotients* provided that for all $\mathcal{F} \in \mathcal{C}$, any subsheaf or quotient sheaf of $\mathcal{F}$ lies in $\mathcal{C}$.

Definition 8.8. We say that a sheaf is *faithful* if every restriction map is an injection.

Here is the main theorem of this section. Parts of this theorem are valid under weaker hypothesis, as will be evident in the subsections that follow.

Theorem 8.9. *Let $\mathcal{C}$ be a sheaf collection that is closed under subsheaves and quotients. Assume that $\mathcal{C}$ contains at least one gapped sheaf, and let $\mathcal{F}$ be any T-minimal gapped sheaf. Then $\mathcal{F}$ is stable, and $h_i^{\mathrm{tw}}(\mathcal{F}) = 1$ for $i = 0, 1$.*

The proof will take up the rest of this section and will be completed in Subsection 8.6; the proof is straightforward and independent of the rest of this paper. The same theorem holds for any T-minimal gapped sheaf, as well, although we will not need this result.

8.3 Edge Subsheaf Closed Collections

Definition 8.10. An *edge subsheaf* of a sheaf, $\mathcal{F}$, is a subsheaf $\mathcal{F}' \subset \mathcal{F}$ such that $\mathcal{F}'$ and $\mathcal{F}$ agree on their values of each vertex. A sheaf collection, $\mathcal{C}$, is *edge subsheaf closed* if $\mathcal{F} \in \mathcal{C}$ implies that $\mathcal{F}' \in \mathcal{C}$ whenever $\mathcal{F}'$ is an edge subsheaf of $\mathcal{F}$.

Lemma 8.11. *Let $\mathcal{F}$ be a sheaf on a graph G with* $\mathrm{m.e.}(\mathcal{F}) \geq 1$*. Then there is an edge subsheaf, $\mathcal{G}$, of $\mathcal{F}$ such that $\mathcal{F}/\mathcal{G}$ is of total dimension one with* $\mathrm{m.e.}(\mathcal{G}) = \mathrm{m.e.}(\mathcal{F}) - 1$*.*

Proof. (Given in [17].) Let $\mathcal{F}'$ be the minimal subsheaf of $\mathcal{F}$ with $-\chi(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F})$. Since $\mathrm{m.e.}(\mathcal{F}) \geq 1$, there exists an $e \in E_G$ such that $\mathcal{F}'(e) \neq 0$; let $A \subset \mathcal{F}(e)$ be a subspace of codimension one such that $\mathcal{F}'(e) \cap A$ is a proper subspace of $\mathcal{F}'(e)$. Let $\mathcal{G}$ be the (A, e) edge subsheaf of $\mathcal{F}$. Then $\mathcal{F}/\mathcal{G}$ is of total dimension one. We claim $\mathrm{m.e.}(\mathcal{G}) = \mathrm{m.e.}(\mathcal{F}) - 1$; indeed, any subsheaf, $\mathcal{G}'$, of $\mathcal{G}$ is also a subsheaf of $\mathcal{F}$, but $\mathcal{G}'$ does not contain the miniminal excess maximizer of $\mathcal{F}$; hence $-\chi(\mathcal{G}') \leq \mathrm{m.e.}(\mathcal{F}) - 1$. $\square$

Lemma 8.12. *Let $\mathcal{F}$ be a sheaf on a graph, G with $h_1^{\mathrm{tw}}(\mathcal{F}) \geq 1$. Then there is an edge subsheaf, $\mathcal{G}$, of $\mathcal{F}$ such that $\mathcal{F}/\mathcal{G}$ is of total dimension one with $h_1^{\mathrm{tw}}(\mathcal{G}) = h_1^{\mathrm{tw}}(\mathcal{F}) - 1$.*

Proof. Let $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F}, \psi)$ with $\alpha \neq 0$. Since α is nonzero, there is an $e \in E_G$ with $\alpha(e) \neq 0$. For some $A \subset \mathcal{F}(e)$ we have $\alpha(e) \notin A(\psi)$. Fix some such A, and let $\mathcal{G}$ be the (A, e) edge subsheaf of $\mathcal{F}$. Then $H_1^{\mathrm{tw}}(\mathcal{G}, \psi)$ is a subspace of $H_1^{\mathrm{tw}}(\mathcal{F}, \psi)$, but since the former does not contain α, the former is a strict subspace of the latter. But since $\mathcal{F}/\mathcal{G}$ is of total dimension one, the former cannot be of codimension greater than one in the latter. Hence $h_1^{\mathrm{tw}}(\mathcal{G}) = h_1^{\mathrm{tw}}(\mathcal{F}) - 1$. $\square$

Lemma 8.13. *Let $\mathcal{F}$ be any gapped sheaf on a graph. Then there exists an edge subsheaf, $\mathcal{F}'$, of $\mathcal{F}$, such that* $\mathrm{m.e.}(\mathcal{F}') = 0$ *and* $h_1^{\mathrm{tw}}(\mathcal{F}') = 1$*.*

Proof. We claim that it suffices to consider the case where $\mathrm{m.e.}(\mathcal{F}) = 0$. Indeed, assume that $\mathrm{m.e.}(\mathcal{F}) > 0$. By Lemma 8.11, $\mathcal{F}$ has an edge subsheaf, $\mathcal{G}$, with $\mathrm{m.e.}(\mathcal{G}) = \mathrm{m.e.}(\mathcal{F}) - 1$ with $\mathcal{F}/\mathcal{G}$ of total dimension one. But $h_1^{\mathrm{tw}}(\mathcal{G}) \geq h_1^{\mathrm{tw}}(F) - 1$, since $\mathcal{F}cG$ is of total dimension one. Hence $\mathcal{G}$ is a gapped edge subsheaf of $\mathcal{F}$ with maximum excess one less than $\mathcal{F}$. Repeating this argument gives a gapped edge subsheaf, $\mathcal{G}'$, of $\mathcal{F}$ with maximum excess zero, and it suffices to prove the lemma for $\mathcal{G}'$.

Similarly, if $\mathrm{m.e.}(\mathcal{F}) = 0$ but $h_1^{\mathrm{tw}}(\mathcal{F}) \geq 2$, we may repeatedly apply Lemma 8.12 to find an edge subsheaf, $\mathcal{F}'$, of $\mathcal{F}$, with $h_1^{\mathrm{tw}}(\mathcal{F}') = 1$.

But passing to edge subsheaves cannot increase the maximum excess, so m.e.$(\mathcal{F}') = 1$. $\square$

Corollary 8.14. *Let $\mathcal{C}$ be a sheaf collection that is closed under edge subsheaves and that contains a gapped element. Then any T-minimal gapped element, $\mathcal{F}$, of $\mathcal{C}$ has* m.e.$(\mathcal{F}) = 0$ *and* $h_1^{\rm tw}(\mathcal{F}) = 1$.

Proof. We know that $\mathcal{F}$ has an edge subsheaf, $\mathcal{F}'$, with m.e.$(\mathcal{F}') = 0$ and $h_1^{\rm tw}(\mathcal{F}') = 1$. But $\mathcal{F}'$ is also gapped, and by minimality must equal $\mathcal{F}$. $\square$

8.4 Vertex Quotient Closed Collections

At this point we give "dual" versions of all the statements in the previous subsection.

We remark that some readers can infer all the dual statements from the following principle (and therefor need not read any proofs in this section.) Our sheaves on graphs can be viewed as presheaves of vector spaces on a 1-dimensional category (in the evident sense.) For any presheaf of vector spaces of finite dimension on a category, the dual spaces form a presheaf on the opposite category. Since the statements in the previous subsection generalize to this wider setting, the dual statements follow immediately.

Definition 8.15. A *vertex quotient* of a sheaf, $\mathcal{F}$, is a quotient of $\mathcal{F}$ whose values agree with those of $\mathcal{F}$ on all the edges. A sheaf collection, $\mathcal{C}$, is *vertex quotient closed* if $\mathcal{F} \in \mathcal{C}$ implies that $\mathcal{F}' \in \mathcal{C}$ whenever $\mathcal{F}'$ is a vertex quotient of $\mathcal{F}$.

Lemma 8.16. *Let $\mathcal{F}$ be a sheaf on a graph, G with* d.m.e.$(\mathcal{F}) \geq 1$. *Then there exists a vertex quotient, $\mathcal{G}$, of $\mathcal{F}$, such that* d.m.e.$(\mathcal{G}) =$ d.m.e.$(\mathcal{F}) - 1$ *and the total dimension of $\mathcal{G}$ is one less than the total dimension of $\mathcal{F}$. Furthermore, $\mathcal{G}$ can be taken equal to $\mathcal{F}$ except at any vertex, v, for which the maximum excess maximizer's value at v is not all of $\mathcal{F}(v)$.*

Proof. Let $\mathcal{F}'$ be the maximum subsheaf of $\mathcal{F}$ that maximizes the excess. Since

$$1 \leq \text{d.m.e.}(\mathcal{F}) = \text{m.e.}(\mathcal{F}) + \chi(\mathcal{F}),$$

we have $\mathcal{F}' \neq \mathcal{F}$ and $-\chi(\mathcal{F}') > -\chi(\mathcal{F})$. It follows that for some v we have $\mathcal{F}'(v) \neq \mathcal{F}(v)$; let A be a one dimensional subspace of $\mathcal{F}(v)$ that does not lie in $\mathcal{F}'(v)$, for such a v, and let $\mathcal{G}$ be the (A, v) quotient of $\mathcal{F}$.

We claim that m.e.$(\mathcal{G})$ = m.e.$(\mathcal{F})$; this will finish the proof of the lemma, since then

$$\text{d.m.e.}(\mathcal{G}) = \text{m.e.}(\mathcal{G}) + \chi(\mathcal{G}) = \text{m.e.}(\mathcal{F}) + \chi(\mathcal{F}) - 1 = \text{d.m.e.}(\mathcal{F}) - 1.$$

To see that m.e.$(\mathcal{G})$ = m.e.$(\mathcal{F})$, consider the image of $\mathcal{F}'$ under the inclusion to $\mathcal{F}$ followed by the quotient map to $\mathcal{G}$; this gives a subsheaf, $\mathcal{G}'$, of $\mathcal{G}$ that is isomorphic to $\mathcal{F}'$. Hence

$$\text{m.e.}(\mathcal{G}) \geq \text{m.e.}(\mathcal{F}). \tag{8.1}$$

But for any subsheaf, $\mathcal{G}''$, of $\mathcal{G}$, consider its inverse image, $\mathcal{F}''$, in $\mathcal{F}$: we have $\chi(\mathcal{F}'') = \chi(\mathcal{G}'') + 1$, because of the quotienting at v, but $\mathcal{F}''$ contains A and therefore cannot maximize the excess of $\mathcal{F}$. Hence

$$-\chi(\mathcal{F}'') \leq \text{m.e.}(\mathcal{F}) - 1,$$

and hence

$$-\chi(\mathcal{G}'') = -\chi(\mathcal{F}'') + 1 \leq \text{m.e.}(\mathcal{F}).$$

Hence, maximizing the above over all $\mathcal{G}''$ in $\mathcal{G}$, we have

$$\text{m.e.}(\mathcal{G}) \leq \text{m.e.}(\mathcal{F}).$$

Combing this with (8.1) shows that m.e.$(\mathcal{G})$ = m.e.$(\mathcal{F})$, completing the proof of the lemma. $\square$

Lemma 8.17. *Let $\mathcal{F}$ be a sheaf on a graph, G with $h_0^{\mathrm{tw}}(\mathcal{F}) \geq 1$. Then there exists a vertex quotient, $\mathcal{G}$, of $\mathcal{F}$, such that $h_0^{\mathrm{tw}}(\mathcal{G}) = h_0^{\mathrm{tw}}(\mathcal{F}) - 1$ and the total dimension of $\mathcal{G}$ is one less than the total dimension of $\mathcal{F}$.*

Proof. For $v \in V_G$ and $w \in \mathcal{F}(v)$, consider the element $\delta_{v,w}$ of $\mathcal{F}(V)$ defined as 0 outside of v and w at $\mathcal{F}(v)$. Clearly the $\delta_{v,w}$ span $\mathcal{F}(V)(\psi)$. Hence if $H_0^{\mathrm{tw}}(\mathcal{F}) \neq 0$, there is some $\delta_{v,w}$ with $w \neq 0$ that is not in the image of $d_{\mathcal{F}\psi}$. Fix such a v and w. Let $\mathcal{G}$ be the same as $\mathcal{F}$ except that its value at v is $\mathcal{F}(v)$ modulo the span of w. Then α is an equivalence class modulo the image of $\mathcal{F}(E)$ in $\mathcal{F}(V)$ under $d_{cf.\ \psi}$. $\square$

Lemma 8.18. *Let $\mathcal{F}$ be any gapped sheaf on a graph. Then there exists a vertex quotient, $\mathcal{F}'$, of $\mathcal{F}$, such that* d.m.e.$(\mathcal{F}') = 0$ *and* $h_0^{\mathrm{tw}}(\mathcal{F}') = 1$.

Proof. We claim that it suffices to consider the case where d.m.e.$(\mathcal{F}) = 0$. Indeed, assume that d.m.e.$(\mathcal{F}) > 0$. By Lemma 8.16, $\mathcal{F}$ has a vertex quotient, $\mathcal{G}$, with d.m.e.$(\mathcal{G}) =$ d.m.e.$(\mathcal{F}) - 1$ with $\mathcal{F}/\mathcal{G}$ of total dimension one. But $h_0^{\mathrm{tw}}(\mathcal{G}) \geq h_0^{\mathrm{tw}}(F) - 1$, since $\mathcal{F}cG$ is of total dimension one. Hence $\mathcal{G}$ is a gapped vertex quotient of $\mathcal{F}$ with dual maximum excess one less than $\mathcal{F}$. Repeating this argument gives a gapped vertex quotient, $\mathcal{G}'$, of $\mathcal{F}$ with dual maximum excess zero, and it suffices to prove the lemma for $\mathcal{G}'$.

Similarly, if d.m.e.$(\mathcal{F}) = 0$ but $h_0^{\mathrm{tw}}(\mathcal{F}) \geq 2$, we may repeatedly apply Lemma 8.17 to find a vertex quotient, $\mathcal{F}'$, of $\mathcal{F}$, with $h_0^{\mathrm{tw}}(\mathcal{F}') = 1$. But passing to vertex quotients cannot increase the dual maximum excess, so d.m.e.$(\mathcal{F}') = 1$. ☐

Corollary 8.19. *Let $\mathcal{C}$ be a sheaf collection that is vertex quotient closed and that contains a gapped element. Then any T-minimal gapped element, $\mathcal{F}$, has* d.m.e.$(\mathcal{F}) = 0$ *and* $h_0^{\mathrm{tw}}(\mathcal{F}) = 1$.

Proof. Let $\mathcal{F}$ be T-minimal gapped. Then $\mathcal{F}'$ has a vertex quotient with d.m.e.$(\mathcal{F}') = 0$ and $h_0^{\mathrm{tw}}(\mathcal{F}') = 1$. But its T-dimension is strictly less unless $\mathcal{F}' = \mathcal{F}$. ☐

8.5 Edge Quotient Closed Collections and Injectivity

Definition 8.20. Let $\mathcal{F}$ be a sheaf on a graph, G. Given $e \in E_G$, and a subspace $A \subset \mathcal{F}(e)$, we define the *A skyscraper (of $\mathcal{F}$) at e*, denoted $A|_e$ (with $\mathcal{F}$ understood), to be the subsheaf whose values for $P \in V_G \amalg E_G$ are

$$A|_e(P) = \begin{cases} A & \text{if } P = e, \\ \mathcal{F}(e,h)A & \text{if } P = he \text{ and } P \neq te, \\ \mathcal{F}(e,t)A & \text{if } P = te \text{ and } P \neq he, \\ \mathcal{F}(e,h)A + \mathcal{F}(e,t)A & \text{if } P = he = te, \text{ and} \\ 0 & \text{otherwise} \end{cases}$$

and where the restriction maps are given by restricting $\mathcal{F}(e,h)$ and $\mathcal{F}(e,t)$. By the *A edge quotient of $\mathcal{F}$ at e* we mean the quotient $\mathcal{F}/A|_e$. We say that a sheaf collection, $\mathcal{C}$, is *edge quotient closed* if for any $\mathcal{F} \in \mathcal{C}$, any edge quotient of $\mathcal{F}$ also lies in $\mathcal{C}$.

Lemma 8.21. *Let $\mathcal{C}$ be a sheaf collection that is edge subsheaf, vertex quotient, and edge quotient closed. Let $\mathcal{F} \in \mathcal{C}$ be T-minimal gapped. Then for any edge, e we have that both restriction maps at $\mathcal{F}(e)$ are injective.*

Proof. Let e be any edge with two distinct endpoints. Assume that $\mathcal{F}(e) \neq 0$. For any nonzero $a \in \mathcal{F}(e)$ we claim that we must have $\mathcal{F}(e,h)a \neq 0$ and $\mathcal{F}(e,t)a \neq 0$. Indeed, let $\mathcal{F}'$ be the A edge quotient of $\mathcal{F}$ at e, where A is the span of a; then $\mathcal{F}' \in \mathcal{C}$ and $\mathcal{F}' < \mathcal{F}$ in the T-order. If $\mathcal{F}(e,h)a = 0$ and $\mathcal{F}(e,t)a = 0$, then $\mathcal{F}$ is the direct sum of $\mathcal{F}'$ plus $A|_e$. But in this case $\mathrm{m.e.}(A|_e) = h_1^{\mathrm{tw}}(A|_e)$ and both the maximum excess and the first twisted Betti of a direct sum is the sum of the individual invariants. Hence $\mathrm{m.e.}(\mathcal{F}') > h_1^{\mathrm{tw}}(\mathcal{F}')$, contradicting the minimality of $\mathcal{F}$. The other case to consider is $\mathcal{F}(e,h)a = 0$ and $\mathcal{F}(e,t)a \neq 0$ (the case with h and t reversed is argued identically). In this case we have an exact sequence

$$0 \to A|_e \to \mathcal{F} \to \mathcal{F}' \to 0,$$

and $h_0^{\mathrm{tw}}(A|_e) = h_1^{\mathrm{tw}}(A|_e) = 0$, so $\mathcal{F}$ and $\mathcal{F}'$ have the same twisted Betti numbers. Similary for any covering map $\phi\colon G' \to G$ we have that $h_0^{\mathrm{tw}}(\phi^* A|_e) = h_1^{\mathrm{tw}}(\phi^* A|_e) = 0$, so we have that $\mathcal{F}$ and $\mathcal{F}'$ have the same maximum excess and dual maximum excess. This again contradicts the minimality of $\mathcal{F}$. $\square$

8.6 The Proof of Theorem 8.9

Proof of Theorem 8.9. Let $\mathcal{C}$ be a sheaf collection that is closed under subsheaves and quotients. Assume that $\mathcal{C}$ contains at least one gapped sheaf, and let $\mathcal{F}$ be any T-minimal gapped sheaf. According to Corollaries 8.14 and 8.19 we have that $\mathrm{m.e.}(\mathcal{F}) = \mathrm{d.m.e.}(\mathcal{F}) = 0$ and $h_i^{\mathrm{tw}}(\mathcal{F}) = 1$ for $i = 0, 1$. Furthermore, by Lemma 8.21, we have that each restriction is an injection.

Consider any short exact sequence

$$0 \to \mathcal{F}' \to \mathcal{F} \to \mathcal{F}'' \to 0 \tag{8.2}$$

in which $\mathcal{F}'$ and $\mathcal{F}''$ are both nonzero. Then $\mathcal{F}'$ and $\mathcal{F}''$ are both less than than $\mathcal{F}$ in T-order, and hence $\mathcal{F}'$ and $\mathcal{F}''$ have no gap.

By Theorem 7.14, we have

$$\mathrm{m.e.}(\mathcal{F}') = \mathrm{d.m.e.}(\mathcal{F}'') = 0.$$

Since $\mathcal{F}'$ and $\mathcal{F}''$ have no gap, we have

$$h_1^{\mathrm{tw}}(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F}') = 0, \quad h_0^{\mathrm{tw}}(\mathcal{F}'') = \mathrm{d.m.e.}(\mathcal{F}') = 0. \tag{8.3}$$

We also have

$$\chi(\mathcal{F}') = \mathrm{d.m.e.}(\mathcal{F}') - \mathrm{m.e.}(\mathcal{F}') = \mathrm{d.m.e.}(\mathcal{F}') \geq 0. \tag{8.4}$$

Furthermore,

$$\chi(\mathcal{F}'') = \chi(\mathcal{F}) - \chi(\mathcal{F}') = -\chi(\mathcal{F}'). \tag{8.5}$$

We claim that $\chi(\mathcal{F}') > 0$. Indeed, otherwise, by (8.4), $\chi(\mathcal{F}') = 0$, and, by (8.5), $\chi(\mathcal{F}'') = 0$. But $\chi(\mathcal{F}') = 0$ implies that $h_i^{\mathrm{tw}}(\mathcal{F}')$ is independent of i, and similarly with $\mathcal{F}''$ replacing $\mathcal{F}'$. Then (8.3) implies that $h_i^{\mathrm{tw}}(\mathcal{F}') = 0$ for $i = 0, 1$, and the same with $\mathcal{F}''$ replacing $\mathcal{F}'$. But Theorem 7.10 shows that if $\mathcal{F}'$ and $\mathcal{F}''$ have all their twisted homology groups vanishing, then so does $\mathcal{F}$. But we know $h_i^{\mathrm{tw}}(\mathcal{F}) = 1$ for $i = 0, 1$.

But if $\mathcal{F}'$ is any subsheaf of $\mathcal{F}$, then it forms a subsequence of the form in (8.2), with $\mathcal{F}'' = \mathcal{F}/\mathcal{F}'$. Hence for any subsheaf, $\mathcal{F}'$, we have $\chi(\mathcal{F}') > 0$. Hence $\mathcal{F}$ is stable. $\square$

Chapter 9

The Twist Trick and Edge Detachment

In this section we describe what we call the "twist trick." It amounts to modifying a sheaf by setting the head or tail map to zero along an edge, e, and noting that the new sheaf's twisted homology groups do not depend on the twist at e; this seemingly trivial observation yields one of our main tools. Intuitively, these theorems seem to specialize the value of an edge twist to be zero or infinity or some other value, and draw interesting conclusions (this "specialization" has to be explained and justified before it makes sense).

We then define an operation on graphs and sheaves called "edge loopification" and interpret our twist trick results in those terms. This leads us to a conjecture on sheaves on graphs which will occupy us for most of the rest of this paper.

9.1 The Fundamental "Twist Trick" Lemma

The basic "twist trick" is to take a sheaf, $\mathcal{F}$, on a graph and modify it by setting its head or tail map at some edge to zero, obtaining $\mathcal{F}'$. If S is a vector space or a linear map, we use $S^\vee$ to denote the dual of S. Clearly $h_0^{\rm twist}$ cannot decrease, and the elements of $H_0^{\rm twist}(\mathcal{F})^\vee$ can be written in terms of $H_0^{\rm twist}(\mathcal{F}')^\vee$; however, $H_0^{\rm twist}(\mathcal{F}')^\vee$ elements are independent of the twist at the discarded edge, and this gives interesting results.

First we give a lemma that describes the basic technique. Then we give its consequences and some variants. The simplest consequence roughly says that "one can always set a twist to zero or infinity.'

Lemma 9.1. *Let $\mathcal{F}$ be a sheaf on a graph, G, and let $e_0 \in E_G$. Let G' be G with e_0 deleted, and let $\mathcal{F}'$ be $\mathcal{F}$ restricted to G'. If $h_0^{\rm twist}(G;\mathcal{F}) \neq 0$, then there exists an element of $H_0^{\rm twist}(G';\mathcal{F}')^\vee$ whose value at he_0 (respectively te_0) is a linear functional that vanishes at the image of $\mathcal{F}(e_0,h)$ (respectively $\mathcal{F}(e_0,t)$).*

Our primary application is the case where $te_0 \neq he_0$, but the lemma holds even if e_0 is a self-loop.

Proof. Let ψ be a full twist of G, and which we will write as $\psi = (\psi', \psi_{e_0})$, where ψ' is a full twist of G' and ψ_{e_0} is the twist at e_0. There is a natural injection:

$$H_0^{\mathrm{twist}}(G;\mathcal{F})^\vee \to H_0^{\mathrm{twist}}(G';\mathcal{F}')^\vee(\psi_{e_0}) = H_0^{\mathrm{twist}}(G';\mathcal{F}')^\vee \otimes_{\mathbb{F}(\psi')} \mathbb{F}(\psi),$$

simply because an element, α, of $H_0^{\mathrm{twist}}(G;\mathcal{F})^\vee$ is just a collection of linear functionals, $\alpha(v) \in \mathcal{F}(v)(\psi)^\vee$ that satisfy consistency conditions along the edges, namely for each $e \in E$:

$$\alpha(he)\mathcal{F}(e,h)^\vee = \psi_e \alpha(te)\mathcal{F}(e,t)^\vee.$$

The elements of $H_0^{\mathrm{twist}}(G';\mathcal{F}')^\vee(\psi_{e_0})$ are the same, except that the condition of consistency is not required at $e = e_0$; since G' has the same vertices as G, and $\mathcal{F}'$ has the same vertex spaces as $\mathcal{F}$, the above map is an injection. Let $\alpha_1, \ldots, \alpha_r$ be a basis for $H_0^{\mathrm{twist}}(G';\mathcal{F}')^\vee$. Any non-zero element, β, of $H_0^{\mathrm{twist}}(G';\mathcal{F}')^\vee$ may therefore be written as

$$\beta = \sum_{i=1}^{r} c_i(\psi)\alpha_i(\psi'),$$

where the $c_i(\psi) \in \mathbb{F}(\psi)$ and not all c_i are zero, and where we have written α_i as $\alpha_i(\psi')$ to emphasize the fact that the α_i depend only on the twists of ψ' are are independent of ψ_{e_0}. By clearing denominators we may assume that $c_i(\psi) \in \mathbb{F}[\psi]$, i.e., are polynomials in the ψ; we may therefore write, for some integer, s,

$$c_i(\psi) = \sum_{j=0}^{s} \psi_{e_0}^j c_{ij}(\psi'),$$

where the $c_{ij} \in \mathbb{F}[\phi']$; furthermore, by dividing common factors of ψ_{e_0}, we may assume that $c_{i0} \neq 0$ for at least one value of i. But then the consistency condition at e_0 implies that

$$\sum_{i=1}^{r}\sum_{j=0}^{s} \psi_{e_0}^j c_{ij}(\psi')\alpha_i(\psi')(he)\mathcal{F}(e,h)^\vee$$
$$= \psi_{e_0} \sum_{i=1}^{r}\sum_{j=0}^{s} \psi_{e_0}^j c_{ij}(\psi')\alpha_i(\psi')(te)\mathcal{F}(t,h)^\vee. \quad (9.1)$$

Setting ψ_{e_0} to zero yields

$$\sum_{i=1}^{r} c_{i0}(\psi')\alpha_i(\psi')(he)\mathcal{F}(e,h)^\vee = 0.$$

In other words, if

$$\gamma = \sum_{i=1}^{r} c_{i0}(\psi')\alpha_i(\psi') \in H_0^{\text{twist}}(G';\mathcal{F}')^\vee,$$

then $\gamma(he)\mathcal{F}(e,h)^\vee = 0$, or, in other words, the entire image of $\mathcal{F}(e)$ via the map $\mathcal{F}(e,h)$ in $\mathcal{F}(he)$ is taken to zero by the linear functional $\gamma(he)$. The coefficient of $\psi_{e_0}^{s+1}$ must be equal on either side of Equation 9.1 and so

$$\psi_{e_0}^{s+1} \sum_{i=1}^{r} c_{is}(\psi')\alpha_i(\psi')(te)\mathcal{F}(e,t)^\vee = 0.$$

Thus we also have a linear functional in $H_0^{\text{twist}}(G';\mathcal{F}')$ that is zero over the image of $\mathcal{F}(e_0,t)$. $\square$

9.2 Edge Detachment

In this section we define "edge detachment" and explain that conjectures about this process imply that certain sheaf collections have no gap.

Definition 9.2. Let $\mathcal{F}$ be a sheaf on a graph, G, and e an edge. By the *e-head detachment of $\mathcal{F}$* we mean the sheaf on G that is identical to $\mathcal{F}$ except that $\mathcal{F}'(h,e) = 0$. We define *e-tail detachment* analogously.

Corollary 9.3. *Let $\mathcal{F}$ be a sheaf on a graph, G, such that $h_0^{\text{twist}}(\mathcal{F}) \neq 0$. Then $h_0^{\text{twist}}(\mathcal{F}') \neq 0$, where $\mathcal{F}'$ is any head or tail detachment of $\mathcal{F}$*

Proof. Let G_e be the graph G with an edge e deleted, and let $\mathcal{F}_e$ be $\mathcal{F}$ restricted to G_e. Any element of $H_0^{\text{twist}}(G_e;\mathcal{F}_e)^\vee$ that satisfies the condition $\alpha(he)\mathcal{F}(e,h)^\vee = 0$ in the case $\mathcal{F}'$ is the tail detachment is also an element of $H_0^{\text{twist}}(G';\mathcal{F}')^\vee$. Such an element is guaranteed to exist by Lemma 9.1 and so $H_0^{\text{twist}}(G';\mathcal{F}')^\vee$ and it's dual space are nonempty.

$\square$

Corollary 9.4. *Let $\mathcal{F}$ be a sheaf on a graph, and let $\mathcal{F}'$ be any head or tail detachment of $\mathcal{F}$. Then for $i = 0, 1$ we have*

$$h_i^{\text{tw}}(\mathcal{F}') \geq h_i^{\text{tw}}(\mathcal{F}).$$

Proof. Since $\mathcal{F}, \mathcal{F}'$ have the same Euler characteristic, it suffices to prove the inequality for $i = 0$. We claim that for any positive integer n, if $n \leq h_0^{\mathrm{tw}}(\mathcal{F})$, then $n \leq h_0^{\mathrm{tw}}(\mathcal{F}')$; we will use induction on n. The case $n = 1$ is just Corollary 9.3. Assume that our claim is true for some value of $n \geq 1$, and assume that $n + 1 \leq h_0^{\mathrm{tw}}(\mathcal{F})$; then $h_0^{\mathrm{tw}}(\mathcal{F}') \neq 0$, and so, by Lemma 8.17 there exists $\mathcal{G}'$ a vertex quotient of $\mathcal{F}'$ of total dimension one less with $h_0^{\mathrm{tw}}(\mathcal{G}') = h_0^{\mathrm{tw}}(\mathcal{F}') - 1$. But the if $\mathcal{G}$ agrees with $\mathcal{F}$ except that it has the same vertex values as $\mathcal{F}'$ (i.e., it equals $\mathcal{F}$ except that its value at one vertex is a quotient of $\mathcal{F}$'s value there of dimension one less), then $h_0^{\mathrm{tw}}(\mathcal{G}) \geq h_0^{\mathrm{tw}}(\mathcal{F}) - 1 \geq n$. Hence, by the inductive hypothesis, $h_0^{\mathrm{tw}}(\mathcal{G}') \geq n$, and hence

$$h_0^{\mathrm{tw}}(\mathcal{F}') = h_0^{\mathrm{tw}}(\mathcal{G}') + 1 \geq n + 1.$$

This proves the inductive step. $\square$

Definition 9.5. Given a sheaf $\mathcal{F}$ on a graph G with some edge e, we define the *other space* to e, h to be

$$\mathrm{span}\{\mathrm{im}\mathcal{F}(e', h) | he = he', e' \neq e\} + \mathrm{span}\{\mathrm{im}\mathcal{F}(e', t) | te = te'\}$$

and define the other space to e, t to be

$$\mathrm{span}\{\mathrm{im}\mathcal{F}(e', t) | te = te', e' \neq e\} + \mathrm{span}\{\mathrm{im}\mathcal{F}(e', h) | he = he'\}$$

. These spaces are denoted by $\mathrm{other}(e, h)$ and $\mathrm{other}(e, t)$ respectively. We say that the edge e is *internal* if $\mathrm{im}\mathcal{F}(e, h) \in \mathrm{other}(e, h)$ and $\mathrm{im}\mathcal{F}(e, t) \in \mathrm{other}(e, t)$.

The technical lemma below immediately implies that if $\mathcal{F}$ is a minimally gapped sheaf from a sheaf collection closed under quotients, then $\mathcal{F}$ is supported on a graph with no leaves. It also shows that if e_1 and e_2 are two edges incident to a degree two vertex v by, without loss of generality, their tail maps then $\mathrm{im}\mathcal{F}(e_1, t) = \mathrm{im}\mathcal{F}(e_2, t)$.

Lemma 9.6. *Let $\mathcal{F}$ be a gapped sheaf on a graph G. If there is some $e \in E(G)$ such that e is not internal, then there is a gapped proper subsheaf of $\mathcal{F}$.*

Proof. Without loss of generality, assume $\mathrm{im}\mathcal{F}(e, h) \not\subset \mathrm{other}(e, h)$. Let $\mathcal{F}'$ be the same as $\mathcal{F}$ except $\mathcal{F}'(e) = \mathcal{F}(e, h)^{-1}(\mathrm{other}(e, h) \cap \mathrm{im}\mathcal{F}(e, h))$. This is a proper subsheaf. Any element of $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ can be restricted to an element of $H_1^{\mathrm{tw}}(\mathcal{F}')$, else d_ψ would map $\alpha|_e$ to something nonzero in $\mathrm{im}\mathcal{F}(e, h)/\mathrm{other}(e, h)$ and no other head or tail map could cancel that out .

Since it is a subsheaf, we have $h_1^{\mathrm{tw}}(\mathcal{F}) = h_1^{\mathrm{tw}}(\mathcal{F}')$. Any subsheaf of $\mathcal{F}$ must have at least as large an excess as the corresponding subsheaf in $\mathcal{F}'$, since their only difference is that we made the edge dimension of $\mathcal{F}'$ smaller. This directly implies that m.e.$(\mathcal{F}) \geq$ m.e.$(\mathcal{F}')$ and so $\mathcal{F}'$ is gapped.

$\square$

Corollary 9.7. *Let $\mathcal{C}$ be a sheaf collection that is closed under subsheaves and quotients and let $\mathcal{F}$ be a T-minimal gapped sheaf in $\mathcal{C}$ over a graph G. Then all $e \in E(G)$ must satisfy*

$$\dim(\mathcal{F}(e)) > 1.$$

Proof. Suppose to the contrary there exists $e \in E(G)$ with $\dim(\mathcal{F}(e)) = 1$. Let span(a) be the image of $\mathcal{F}(e,h)$ with v the vertex mapped to by the head of e and $a \in \mathcal{F}(v)$. Let $\mathcal{F}'$ be the the e-head detachment of $\mathcal{F}$. By the previous corollary and theorem 8.9 we know $h_i^{\mathrm{tw}}(\mathcal{F}') \geq 1$ and that $\mathcal{F}$ is stable with m.e.$(\mathcal{F}) = 0$. For any subsheaf $\mathcal{U}$ of $\mathcal{F}$ assign a corresponding subsheaf $\mathcal{U}'$ of $\mathcal{F}'$ that has the same values on each edge and vertex as $\mathcal{U}$, implying that we have excess$(\mathcal{U})$ = excess$(\mathcal{U}')$. A subsheaf $\mathcal{R}'$ of $\mathcal{F}'$ is not mapped to in this way if and only if span$(a) \not\subset \mathcal{R}'(v)$ and $\mathcal{R}'(e) \neq 0$. For any such $\mathcal{R}'$ let $\mathcal{R}$ be the subsheaf of $\mathcal{F}$ with $\mathcal{R}'(E) = \mathcal{R}(E)$ and the values on the vector spaces also agree except for $\mathcal{R}(v) = \mathcal{R}'(v) + \mathrm{span}(a)$. By the definition of excess, excess$(\mathcal{R}')$ = excess$(\mathcal{R}) + 1$. By the stability of $\mathcal{F}$, we have excess$(\mathcal{R}') \leq 0$ unless $\mathcal{R}$ is 0 or $\mathcal{F}$. Since $\mathcal{R}(e) \neq 0$, we know $\mathcal{R} \neq 0$. If $\mathcal{R} = \mathcal{F}$ on the other hand, then by Lemma 9.6 since $\mathcal{F}$ is a T-minimal gapped sheaf there is some element of $\mathcal{R}'(E)$ that maps to $\mathcal{F}(he)$ under a sum of the head and tail maps, which contradicts that $\mathcal{R}'(he) = 0$. Thus m.e.$(\mathcal{F}') = 0$.

Let span(b) be the image of $\mathcal{F}(e,t)$ for some $b \in \mathcal{F}(te)$. Let $\mathcal{S}_1$ be the subsheaf of $\mathcal{F}$ that agrees with $\mathcal{F}$ on all vertices and edges except $\mathcal{S}_1(e) = 0$ and let $\mathcal{S}_2$ be the subsheaf of $\mathcal{F}$ that is zero on all vertices and edges except $\mathcal{F}(te) = \mathrm{span}(b)$. The sheaf $\mathcal{S} = \mathcal{S}_1/\mathcal{S}_2$ is a subquotient of $\mathcal{F}$. For any subsheaf $\mathcal{S}'$ of $\mathcal{S}$ assign a corresponding subsheaf $\mathcal{G}$ of $\mathcal{F}'$ that is equal to $\mathcal{S}'$ on all vertices and edges except $\mathcal{G}(e) = \mathcal{F}(e)$ and $\mathcal{G}(te) = \mathcal{S}'(te) + \mathrm{span}(b)$. The excess of $\mathcal{S}'$ is equal to the excess of $\mathcal{G}$, as $\mathcal{G}$ is exactly one dimension larger in the vertex space and one dimension larger in the edge space than $\mathcal{S}'$. So for any subsheaf of $\mathcal{S}$ there is a subsheaf of $\mathcal{F}$ with equal excess, implying m.e.$(\mathcal{S}) \leq$ m.e.$(\mathcal{F}) = 0$.

Given any $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F}')$, let α' be the restriction of α to all edge spaces except the edge space on e. So $d_{\mathcal{F}'\psi}(alpha')$ is 0 on all vertex spaces except possibly on te where it is in span(b). This implies that $\alpha' \in H_1^{\mathrm{tw}}(\mathcal{S})$ and so $H_1^{\mathrm{tw}}(\mathcal{S}) \geq H_1^{\mathrm{tw}}(\mathcal{F}')$. Thus $\mathcal{S}$ is a gapped subquotient of $\mathcal{F}$

$\square$

9.3 Alternate Proof of the Twist Trick

One can also give a proof of Corollary 9.4 by arguing about h_1^{tw} directly rather than h_0^{tw}. The nice thing about such a proof is that one can state a very general linear algebra lemma that illustrates the main technique in broad terms.

To explain the lemma, recall that if $\mathbb{F}$ is a field and ψ an indeterminate, then to every vector space, W, over $\mathbb{F}$, there is a naturally associated vector space, $W(\psi) = W \times_{\mathbb{F}} \mathbb{F}(\psi)$ over $\mathbb{F}(\psi)$. Any vector of W can be identified with one in $W(\psi)$ (via $w \mapsto w \otimes 1$), and any basis of W as such gives a basis of $W(\psi)$; similarly for any subspace of W. One can reverse this process, to a certain extent. Any vector in $W(\psi)$ is a finite sum of terms $w_i \times f_i(\psi)$, where $w_i \in W$ and $f_i(\psi) \in \mathbb{F}(\psi)$ is a rational function in ψ; scaling the vector by the product of the denominators of the $f_i(\psi)$ gives a proportional vector in $W[\psi] = W \times_{\mathbb{F}} \mathbb{F}[\psi]$, i.e., a finite sum of terms $w_i \times p_i(\psi)$, where the $p_i(\psi)$ are polynomials. For any vector in $W[\psi]$ and any $c \in \mathbb{F}$, we can substitute c for ψ to get a vector in W; if $w_1, \ldots, w_r$ are linearly independent vectors in $W[\psi]$, then $w_1 \wedge \ldots \wedge w_r$ is a nonzero element of $(\Lambda^r W)[\psi]$, and therefore for all but at most finitely many $c \in \mathbb{F}$, this element is nonzero when we substitute c for ψ. It follows that for any subspace of $W(\psi)$, there is a basis of elements of $W[\psi]$, for which we may substitute all but at most finitely many $c \in \mathbb{F}$, we get a subspace of W of the same dimension.

Lemma 9.8. *Let $\mathbb{F}$ be a field, and ψ an indeterminate. Let A, B be subspaces of $\mathbb{F}$, and let J, K be subspaces of $A \oplus B$. As usual, let $A(\psi) = A \times_{\mathbb{F}} \mathbb{F}(\psi)$ be A extended as an $\mathbb{F}(\psi)$ vector space, and similarly for $B(\psi), J(\psi), K(\psi)$. Let*

$$J \cap_\psi K = \{(j_1, j_2) \in J(\psi) \mid (j_1, \psi j_2) \in K(\psi)\},$$

and for $c \in \mathbb{F}$ let

$$J \cap_c K = \{(j_1, j_2) \in J \mid (j_1, c j_2) \in K\}.$$

Then for all $c \in \mathbb{F}$ we have

$$\dim_{\mathbb{F}(\psi)}(J \cap_\psi K) \leq \dim_{\mathbb{F}}(J \cap_c K) \ . \tag{9.2}$$

The paragraph above the lemma shows that, by considering a basis for $J\cap_\psi K$ in $(A\times B)[\psi]$, that (9.2) holds for all but at most finitely many $c\in\mathbb{F}$. The novelty in the above lemma is that it holds for every $c\in\mathbb{F}$, without a finite number of exceptions.

Proof. We may assume $J\cap_\psi K$ is nonzero. We start by proving that this implies that $J\cap_c K$ is nonzero. So let $(j_1,j_2)\in J\cap_\psi K$ be nonzero; we may assume that $(j_1,j_2)\in J[\psi]$, and write

$$(j_1,j_2)=\sum_{i=0}^{r}(j_1^i,j_2^i)(\psi-c)^i$$

for some j_1^i,j_2^i and r; by dividing by an appropriate power of $(\psi-c)$ we may assume that at least one of j_1^0,j_2^0 is not zero. It follows that

$$(j_1,\psi j_2)=\sum_{i=0}^{r}(j_1^i,\psi j_2^i)(\psi-c)^i$$

lies in $K[\psi]$. Hence

$$(j_1^0,\psi j_2^0)=(j_1,\psi j_2)+(\psi-c)w,$$

where $w\in(A\times B)[\psi]$. Hence

$$(j_1^0,cj_2^0)=(j_1^0,\psi j_2^0)+(\psi-c)(0,-j_2^0)=(j_1,\psi j_2)+(\psi-c)w',$$

where $w'\in(A\times B)[\psi]$. In other words,

$$(j_1^0,cj_2^0)-(\psi-c)w'\in K[\psi],$$

and so

$$(j_1^0,cj_2^0)-(\psi-c)w'=\sum_{i=0}^{s}(\psi-c)^i k_i,$$

for some $k_i\in K$ and s. Hence

$$(j_1^0,cj_2^0)=k_0\in K.$$

Hence $J\cap_c K$ contains the nonzero vector (j_1^0,j_2^0).

We have shown

$$1\le\dim_{\mathbb{F}(\psi)}(J\cap_\psi K)\quad\Rightarrow\quad 1\le\dim_{\mathbb{F}}(J\cap_\psi K),$$

and we can use this to establish that

$$n \le \dim_{\mathbb{F}(\psi)}(J\cap_\psi K) \quad \Rightarrow \quad n \le \dim_{\mathbb{F}}(J\cap_\psi K)$$

for any positive integer n. There are a number of ways of doing this, at least one of which is similar to our inductive proof of Corollary 9.4.

For example, we can use induction on n. The case $n = 1$ has been established. Assuming the claim for a particular value of $n \ge 1$, assume that

$$n+1 \le \dim_{\mathbb{F}(\psi)}(J\cap_\psi K).$$

Then we know there is a nonzero element, $j = (j_1^0, j_2^0) \in J\cap_c K$; let J' be J modulo the span of j. Then

$$\dim(J'\cap_c K) = \dim(J\cap_c K) - 1;$$

but since $\dim(J') = \dim(J) - 1$, we have

$$\dim_{\mathbb{F}(\psi)}(J'\cap_\psi K) \ge \dim_{\mathbb{F}(\psi)}(J\cap_\psi K) - 1 \ge n.$$

Hence

$$\dim(J\cap_c K) = \dim(J'\cap_c K) + 1 \ge n + 1.$$

$\square$

9.4 The Chain Decomposition

The main point of this subsection is to show that if $\mathcal{F}$ is a stable sheaf on a graph, G, and $\mathcal{F}(e) \ne 0$ for a self-loop e about a vertex, then $\mathcal{F}(e)$ and its restriction maps to $\mathcal{F}(v)$ have a fairly simple description in terms of "chains," provided that $\mathcal{F}$ is not supported entirely on e and v; this is Theorem 9.12 below. This will be useful for what we call "the second twist trick." However, we will make a number of additional remarks regarding chains.

Let us give some definitions that will allow us to state the main theorem of this section.

Definition 9.9. Let $H\colon \mathcal{E} \to \mathcal{V}$ and $T\colon \mathcal{E} \to \mathcal{V}$ be two linear maps of vector spaces over some field, and let $i \ge 0$ be an integer. We say that (H,T) is *superstable* if for any finite dimensional subspace $\mathcal{E}' \subset \mathcal{E}$ such that $\mathcal{E}' \ne 0$ we have

$$\dim\bigl(H(\mathcal{E}') + T(\mathcal{E}')\bigr) > \dim(\mathcal{E}'). \tag{9.3}$$

By an (H,T)-*chain of order* i we mean an alternating sequence of elements of $\mathcal{V}$ and $\mathcal{E}$
$$c = (v_1, e_1, \dots, v_i, e_i, v_{i+1})$$
(i.e., $v_j \in \mathcal{V}$ for $1 \le j \le i+1$ and $e_j \in \mathcal{E}$ for $1 \le j \le i+1$) such that
$$Te_j = v_j \quad \text{and} \quad He_j = v_{j+1}.$$
for all $1 \le j \le i$. By the *edge sequence* of the chain, c, we mean the sequence $(e_1, \dots, e_i)$, and by the *edge space* of c, denoted $c_{\mathcal{E}}$, we mean
$$c_{\mathcal{E}} \stackrel{\text{def}}{=} \operatorname{span}(e_1, \dots, e_i);$$
we similarly define the *vertex sequence* and *vertex space*, $c_{\mathcal{V}}$, of the chain, c, based on the v_j. If $i \ge 0$ and $j \ge 1$ are integers with $j \le 2i+1$, we refer to the j-th element of a chain as its j*-th component* (which lies in $\mathcal{V}$ or $\mathcal{E}$ according to whether j is odd or even). We often drop the "(H,T)" or "order i" from the term *chain* when H,T or i are understood (or when i does not matter).

A chain superficially resembles a walk in a graph. Also, a chain of order 0 is simply an element of $\mathcal{V}$. The last definition we need to state the main theorem is the notion of a *canonical form* in the above setting.

Definition 9.10. Let $H : \mathcal{E} \to \mathcal{V}$ and $T : \mathcal{E} \to \mathcal{V}$ be two linear maps of vector spaces over some field. By a *chain decomposition* for (H,T) we mean a collection of (H,T)-chains such that their edge sequences are mutually disjoint and their union comprises a basis for $\mathcal{E}$, and similarly for vertex spaces and $\mathcal{V}$. In analogy with Jordan canonical form, we call also call a chain decomposition a *canonical form* for (H,T).

Example 9.11. Let U be the unhappy 4-bundle over the graph with vertex, v, and two self-loops, e_1 and e_2. Let H and T be the respective restrictions of d_h and d_t to $U(e_1)$. Then if α, β is a basis for $U(e_1)$, then $H\alpha, H\beta, T\alpha, T\beta$ are linearly independent, spanning $U(v)$. In this case $T\alpha, \alpha, H\alpha$ and $T\beta, \beta, H\beta$ are two chains that give a canonical form for H,T; the same is true if α and β are replaced by any two linear combinations of α and β that are linearly independent.

Here is the main theorem of this section.

Theorem 9.12. *Let (H,T) be a superstable pair of linear maps $\mathcal{E} \to \mathcal{V}$ of finite dimensional vector spaces over some field. Then there exists a* canonical form *for (H,T).*

We remark that if $\mathcal{V}$ is larger than the span of the images of H and T then the canonical form for (H,T) will necessarily contain a number of chains of order zero, whose number is precisely the codimension of $H(\mathcal{E})+T(\mathcal{E})$ in $\mathcal{V}$.

The rough idea of the proof of the above theorem is to build a canonical form by starting with any chain of maximal order, and successively adding new chains that are maximal subject to being linearly independent of the previous chains.

The setting of our notion of canonical form shares a number of other similarities with the setting of Jordan canonial form. First, in either setting there is a polynomial time algorithm to compute a maximal chain, and hence a polynomial time algorithm to compute a canonical form. Second, in either setting chains of any order form a vector space. Third, if the chains of maximal order form a vector space of dimension d, then exactly d chains of maximal order appear in any canonical form. Fouth, the number of chains of a given order in any canonical form is independent of the particular choice of canonical form. Fifth, the dimension of chains of a given order and a basis for this space of chains can be easily calculated once we find a canonial form. Sixth, there is a way to view our canonical form as related to the Jordan canonical form of a certain nilpotent matrix, which we now state.

Theorem 9.13. *Let (H,T) be a superstable pair of linear maps $\mathcal{E}\to\mathcal{V}$ of finite dimensional vector spaces over some field. Let $\mathcal{T}\subset\mathcal{V}$ be the span of the first elements of the chains in any canonical form. Then*

1. *the first element of any maximal chain lies in $\mathcal{T}$ regardless of the particular choice canonical form;*
2. *the spaces $H(\mathcal{E})$ and $\mathcal{T}$ are complementary spaces of $\mathcal{V}$, i.e.,*

$$H(\mathcal{E})\cap\mathcal{T}=0, \quad \textit{and} \quad H(\mathcal{E})+\mathcal{T}=\mathcal{V}.$$

We remark this implies there is a unique linear map $A\colon\mathcal{V}\to\mathcal{V}$ such that A restricts to TH^{-1} on $H(\mathcal{E})$ and A maps $\mathcal{T}$ to 0; we have that A is nilpotent, and there is a one-to-one correspondence between canonical forms for (H,T) given $\mathcal{T}$, the span of the first elements of the chains, and Jordan canonical forms for A.

Below we will state some lemmas on (H,T)-chains which will not only prove the above theorems, but give more insight into chains.

Let us record some simple remarks whose proofs are immediate.

Proposition 9.14. *Let $H\colon\mathcal{E}\to\mathcal{V}$ and $T\colon\mathcal{E}\to\mathcal{V}$ be two linear maps of vector spaces. Then*

1. *for any integer $i \geq 0$, and any integer $j \geq 1$ with $j \leq 2i+1$, the map that takes a chain to its j-th component is a linear map; the chains of order i are a vector space under term-by-term addition and term-by-term scalar multiplication;*

2. *any chain is determined by its edge part;*

3. *if H and T are injections, then any chain of a given order is determined by any of its components.*

Let us begin with some remarks about superstable pairs.

Proposition 9.15. *Let (H, T) be a superstable pair of linear maps $\mathcal{E} \to \mathcal{V}$ of vector spaces over some field. Then*

1. *H and T are injections, and hence any chain is uniquely determined from any of its components;*

2. *if*
$$c = (v_1, e_1, \ldots, e_i, v_{i+1})$$
is any nonzero chain of order i, then $v_1, \ldots, v_i, v_{i+1}$ are linearly independent in $\mathcal{V}$, and $e_1, \ldots, e_i$ are linearly independent in $\mathcal{E}$.

In particular, if $\mathcal{E}$ and $\mathcal{V}$ are finite dimensional, then there is no nonzero chain of order greater than $\dim(\mathcal{E})$.

Proof. If H were not an injection, then $Hu = 0$ for some $u \neq 0$, and then (9.3) would be violated for $\mathcal{E}'$ being the space spanned by u. Hence H is an injection. Similarly T is an injection.

Let us prove the second part by induction on i. If $i = 0$, then c consists of a single element of $\mathcal{V}$, and the claim is clear. Now say that the second part holds for a value of $i \geq 0$, and let us prove that the same is true for i replaced with $i+1$. So consider a chain

$$(v_1, e_1, \ldots, e_i, v_{i+1}, e_{i+1}, v_{i+2}).$$

If this chain is not the zero chain, then since the value of v_1 determines the values of the entire chain, we have $v_1 \neq 0$. By induction, we know that $v_1, \ldots, v_{i+1}$ are linearly independent; hence, since T is an injection, and $e_j = T^{-1} v_j$ for $j = 1, \ldots, i+1$, we have that $e_1, \ldots, e_{i+1}$ are linearly independent. But then

$$\mathcal{E}' = \operatorname{span}(e_1, \ldots, e_{i+1})$$

is $i+1$ dimensional, and hence

$$T(\mathcal{E}') + H(\mathcal{E}') = \mathrm{span}(v_1, \ldots, v_{i+2})$$

must be of dimension at least $i+2$, by (9.3); hence $v_1, \ldots, v_{i+2}$ are linearly independent. This proves the inductive argument, and hence the second part of the proposition holds. $\square$

We remark that the inductive proof of independence can be viewed as starting from v_1, then proceeding to u_1, then v_2, then u_2, etc. As such the proof can be viewed as zig-zagging up and down, or moving along the components of the chain. This type of induction will also be used in Lemma 9.18.

We now define a canonical form for superstable pairs.

Definition 9.16. Let (H, T) be a superstable pair of linear maps $\mathcal{E} \to \mathcal{V}$. We see that a sequence, $c_1, \ldots, c_m$, of chains is *successively maximal* if

1. c^1 is a maximum order chain; and
2. for each $j \geq 2$,
$$c^j = (v_1^j, e_1^j, \ldots, v_{i_j+1}^j)$$
which is of maximum order among all chains for which v_1^j is not in the span of the $v_\ell^{j'}$ ranging over all $j' \leq j-1$ and any ℓ (i.e., $1 \leq \ell \leq i_{j'}+1$).

We say that the a successively maximum sequence, $c_1, \ldots, c_m$ is *complete* if the span of the vertex spaces of $c_1, \ldots, c_m$ is all of $\mathcal{V}$.

Clearly a successively maximum sequence of chains exists if $\mathcal{V}$ is finite dimensional, since vertex spaces of some chains is not all of $\mathcal{V}$, then we may take any element of $\mathcal{V}$ that is not in this chain as the first element of a new chain.

Now we come the some lemmas that culminate in showing that any complete sequence of successively maximum chains actually comprise a canonical form.

Lemma 9.17. *Let (H, T) be a superstable pair of maps of vector spaces over a field, $\mathbb{F}$. Let $c^1, \ldots, c^m$ be a sequence of successive maximum (H, T)-chains, with*

$$c^j = (v_1^j, u_1^j, \ldots, v_{i_j+1}).$$

Then we have

$$i_1 \geq i_2 \geq \ldots \geq i_m.$$

For $j = 1, \ldots, m-1$ let $\tilde{c}^j$ be the chain c^j that is truncated to a chain of order i_m taking its first $2i_m + 1$ components, i.e.,

$$\tilde{c}^j = (v_1^j, u_1^j, \ldots, v_{i_m+1}).$$

Then for any $\alpha_1, \ldots, \alpha_{m-1} \in \mathbb{F}$, we have that $c^1, \ldots, c^{m-1}, c$ is also a sequence of successive maximum chains, for

$$c = c_m - \alpha_1 \tilde{c}^1 - \cdots - \alpha_{m-1} \tilde{c}^{m-1}.$$

Proof. If $i_j < i_{j+1}$ for some $j = 1, \ldots, m-1$, then c^{j+1} is of order greater than c^j and yet v_1^{j+1} is linearly independent of the $v_k^{j'}$ with $j' < j$ and any k; hence c^j is not of maximal order subject to v_1^j being linearly independent of $v_k^{j'}$ for $j' < j$.

For the second part of the theorem, c has the same order as c^m, and since the first component of c is

$$v_1^m - \alpha_1 v_1^1 - \cdots \alpha_{m-1} v_1^{m-1},$$

this cannot be linearly dependent on the v_k^j with $j < m$ unless v_1^m is linearly dependent as well. $\square$

Lemma 9.18. *Let*

$$c^j = (v_1^j, e_1^j, \ldots, v_{i_j+1}^j)$$

for $j = 1, \ldots, m$ be a sequence of successively maximal chains. Then

1. *the subspace*
$$\mathcal{T} \stackrel{def}{=} \operatorname{span}(v_1^1, v_1^2, \ldots, v_1^m) \subset \mathcal{V}$$
satisfies $\mathcal{T} \cap H(\mathcal{E}) = 0$;

2. *the v_k^j, ranging over all k and j, are linearly independent, and similarly for the u_k^j; and*

3. *the subspace*
$$\mathcal{H} \stackrel{def}{=} \operatorname{span}(v_{i_1}^1, v_{i_2}^2, \ldots, v_{i_m}^m) \subset \mathcal{V}$$
satisfies $\mathcal{H} \cap T(\mathcal{E}) = 0$.

Proof. We will prove this by induction on m. The inductive step will be very similar to the proof of Proposition 9.15. Of course, this proposition also proves almost everything we need for the base case $m = 1$.

If $m = 1$, i.e., we have a single maximum chain, c_1, then by Proposition 9.15 we know that the v_k^1 are linearly independent in $\mathcal{V}$, as are the u_k^1 in $\mathcal{E}$. Since c is a maximum order chain, v_1^1 cannot lie in the image of $H(U)$, for otherwise we could lengthen the chain; similarly for $v_{i_1}^1$. This establishes the claims about a sequence of successively maximum chain in the lemma in the case where $m = 1$. Let us establish the situation $m \geq 2$ by induction.

So assume that for some $m \geq 2$ we have established the claims regarding any sequence of maximum chains $c_1, \ldots, c_{m-1}$, and let us add to this sequence c_m, a chain of maximum order subject to v_1^m is not in the span of the v_k^j for $j < m$ and all k.

We first show that $\mathcal{T} \cap H(U) = 0$ for $c^1, \ldots, c^m$. Assume, to the contrary, that

$$\alpha_1 v_1^1 + \cdots \alpha_m v_1^m \in H(U)$$

for some $\alpha_i \in \mathbb{F}$ that are not all zero; since

$$\operatorname{span}(v_1^1, \ldots, v_1^{m-1}) \cap H(U) = 0,$$

we have that $\alpha_m \neq 0$. By subtracting multiples of the $c^1, \ldots, c^{m-1}$ truncated appropriately, as in Lemma 9.17 from c^m, we may assume that α_j for $j < m$ are all zero. But then v_1^m is in the span of H, and there is a unique way to extent c^m "backwards" to a longer chain

$$(v_0^m, u_0^m, v_1^m, \ldots, v_{i_m+1}^m).$$

We wll now get a contradiction by showing that v_0^m is not in the span of the v_k^j for all k and j with $j \leq m-1$.

For the sake of contradiction, assume that

$$v_0^m = \sum_{j \leq m-1} \alpha_k^j v_k^j.$$

Since

$$\sum_{j \leq m-1} \alpha_{i_j+1} v_{i_j+1}^j = v_0^m - \sum_{j \leq m-1} \sum_{k \leq i_j} \alpha_k^j v_k^j,$$

and the right hand side is in the image of $T(U)$, we have

$$\sum_{j \leq m-1} \alpha_{i_j+1} v_{i_j+1}^j \in T(U);$$

but by the inductive assumption, this implies that

$$\sum_{j \leq m-1} \alpha_{i_j+1} v_{i_j+1}^j = 0.$$

Hence
$$v_0^m = \sum_{j\le m-1}\sum_{k\le i_j}\alpha_k^j v_k^j.$$
But v_0^m and all the v_k^j are in the image of T, so we can apply HT^{-1} to both sides and conclude that
$$v_1^m = HT^{-1}v_0^m = \sum_{j\le m-1}\sum_{k\le i_j} HT^{-1}\alpha_k^j v_k^j =$$
$$= \sum_{j\le m-1}\sum_{k\le i_j}\alpha_k^j v_{k+1}^j.$$
But this contradicts the fact that v_1^m is not in the span of the v_k^j for $j \le m-1$. Hence v_0^m would also not be in this span, and we must have $\mathcal{T} \cap H(U) = 0$ unless c^m is not successively maximum when added to $c^1, \ldots, c^m$.

Next we claim that u_1^m (assuming that c^m is of order at least one) is linearly independent of the u_k^j with $j \le m-1$: indeed, if not, then by applying T to u_1^m and its expression as a linear combination of the u_k^j with $j \le m-1$, we would conclude that $Tu_1^m = v_1^m$ is a linear combination of some v_k^j with $j \le m-1$, which is impossible.

Now, similar to the proof of Proposition 9.15, we show that for $\ell = 2, 3, \ldots, i_m$ we have that v_ℓ^m is not a linear combination of the "previous v_k^j," meaning $v_{\ell'}^m$ with $\ell' < k$ and the v_k^j with $j \le m-1$ (and k arbitrary), and also that the same holds for u replacing the v's everywhere. Again we use induction on ℓ, having established $\ell = 1$ as a base case. Indeed, since v_ℓ^m is in the image of H, and so are all the v_k^j with $k \ge 2$ and $j \le m-1$ and with $2 \le k < \ell$, the same argument as in the previous paragraph shows that any way for writing v_ℓ^m in terms of the previous v_k^j cannot involve any of the v_1^j, with $j \le m$, since this would mean that some linear combination of the v_1^j would lie in H. Hence we need show that v_ℓ^m cannot be written as a sum of the "previous v_k^j" where $k \ge 2$; but then we can apply H^{-1}, and conclude that expressing v_ℓ^m as such would imply that $u_{\ell-1}^m$ could be written in term of "previous u_k^j" (meaning with $j \le m-1$ or $j = m$ and $k < \ell - 1$), contradicting the inductive claim. And then the claim for u_ℓ^m instead of v_ℓ^m follows by applying T.

To finish the lemma, we see, by the same argument, that $v_{i_m+1}^m$ cannot be written in terms of the previous v_k^j (meaning either $j \le m-1$ or $j = m$ and $k \le i_m$). Now we finish by showing that
$$\mathcal{H} = \operatorname{span}(v_{i_1+1}^1, \ldots, v_{i_m+1}^m)$$

has $\mathcal{H} \cap T(U) = 0$. But here we can use the same argument as used for $\mathcal{T} \cap H(U) = 0$, by truncating $c^1, \ldots, c^{m-1}$ from their beginning, rather than their end, and using the inductively known fact that

$$\operatorname{span}(v^1_{i_1+1}, \ldots, v^{m-1}_{i_{m-1}+1}) \cap T(U) = 0.$$

□

Now we get to the whole point of successively maximum chains.

Lemma 9.19. *Let $c^1, \ldots, c^m$ be a sequence of successively maximum (H,T) chains that is complete, where (H,T) is a superstable pair. Then this sequence is a canonical form for (H,T).*

As remarked earlier, complete sequences exist since we can always augment a sequence by adding chains of order zero, consisting of a single element of $\mathcal{V}$.

Proof. By definition, the vertex spaces of the chains spans all of $\mathcal{V}$. Similarly, if the edge spaces of the chains does not span all of $\mathcal{E}$, then choose some nonzero ϵ in $\mathcal{E}$ that is not in this span. Then $(T\epsilon, \epsilon, H\epsilon)$ is a new chain. But now we claim that $T\epsilon$ cannot be a linear combination of elements of the vertex spaces of the chains: indeed, otherwise this linear combination cannot involve any of the last components of the chains (the $v^j_{i^j+1}$ in the previous notation), for otherwise $\mathcal{H}$ would intersect T. But then we could apply T^{-1} to the equation representing $T\epsilon$ as a linear combination of vertex space elements and conclude that ϵ is a linear combination of the edge spaces of $c^1, \ldots, c^m$.

Hence the sum of the vertex and edge spaces of the chains are, respectively, the entirety of $\mathcal{V}$ and $\mathcal{E}$, and by Lemma 9.18, these chains give bases for $\mathcal{V}$ and $\mathcal{E}$ and are therefore a canonical form for (H,T). □

Proof of Theorem 9.12. Any complete sequence of successively maximal chains gives a canonical form for (H,T). □

Proof of Theorem 9.13. Let $c^1, \ldots, c^m$ be a complete sequence of successively maximum chains of (H,T), with

$$c^j = (v^j_1, e^j_1, \ldots, v^j_{i_j+1}).$$

Let ℓ be the order of the chain of maximum order, and let the dimension of the vector space of maximum order chains be d. Since any chain is determined by its first component, if we write down any sequence of at most

$d-1$ chains, there is always a chain of maximal order whose first component is linearly independent of the edge spaces of the chains in this sequence. It follows that $c^1, \ldots, c^d$ must all be of maximum order. However, then the $v_1^1, \ldots, v_1^d$ span the space of first components of chains of maximal order. It follows that $c^{d+1}, \ldots, c^m$ are all of order strictly less than ℓ.

□

9.5 The Second Twist Trick

This section shows that the twist trick can give a better result when the edge we are detaching is a self loop. This will give us stronger conditions on minimally gapped sheaves with self loops.

Definition 9.20. Given a sheaf $\mathcal{F}$ on a graph G and a basis for $\mathcal{F}(E)$, $b_1, \ldots, b_r$, we can express any $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ as $d_1 b_1 + \ldots + d_r b_r$ with $d_i \in \mathbb{F}(\psi)$. If all of the b_i are elements of $\bigcup_{e \in E} \mathcal{F}(e)$, we call the sequence $(d_1, \ldots, d_r)$ a *representation* of α with respect to the sequence $(b_1, \ldots, b_r)$. We say $\mathcal{F}$ is *full* if it satisfies $h_1^{\mathrm{tw}}(\mathcal{F}) = 1$ and for any nontrivial $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ and given any basis of $\mathcal{F}(E)$, the representation of α does not contain 0. Note that any minimally gapped sheaf $\mathcal{H}$ from a sheaf collection closed under taking subsheaves and quotients is also a full sheaf, since if there exists a $\alpha \in H_1^{\mathrm{tw}}(\mathcal{H})$ with a representation $(d_1, \ldots, d_r)$ with respect to $(b_1, \ldots, b_r)$ and some $d_i = 0$ then we may quotient out $\mathrm{span}(b_i)$ from the edge space of $\mathcal{H}$ to create a new quotient sheaf $\mathcal{H}'$ with maximal excess still 0 and $H_1^{\mathrm{tw}}(\mathcal{H})$ containing

$$d_1 b_1 + \ldots + d_{i-1} b_{i-1} + d_{i+1} b_{i+1} + \ldots + d_r b_r,$$

contradicting that $\mathcal{H}$ is minimal.

Lemma 9.21. *Let $\mathcal{C}$ be a sheaf collection closed under taking subsheaves and quotients and let $\mathcal{F}$ be a full, stable, gapped sheaf in $\mathcal{C}$ on a graph G. Suppose G has a self loop l on a vertex v and $\mathcal{F}$ is supported on l. Let $\tilde{\mathcal{F}}$ be $\mathcal{F}$ except $\tilde{\mathcal{F}}(l) = 0$, $\tilde{\mathcal{F}}(v) = 0$ and any restriction map into $\mathcal{F}(v)$ is now the zero map. Then $h_1^{\mathrm{tw}}(\tilde{\mathcal{F}}) \geq 2$.*

Proof. Let $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ be nonzero on l and let ψ and $\tilde{\psi}$ be the full twists for $\mathcal{F}$ and $\tilde{\mathcal{F}}$ respectively. By definition α is a direct sum consisting of a vector $\alpha|_e \in \mathcal{F}(e)(\psi)$ for each $e \in E_G$. For each vertex $u \in V_G$, α must satisfy the condition that

$$\sum_{e \in h_G^{-1}(u)} \mathcal{F}(h,e)(\alpha|_e) + \sum_{e \in t_G^{-1}(u)} \psi_e \mathcal{F}(t,e)(\alpha|_e) = 0.$$

Fix a basis $b_1^e, \ldots, b_{r_e}^e$ for $\mathcal{F}(e)$ for each edge e. We can write $\alpha|_e$ as

$$b_1^e d_1^e + \ldots + b_{r_e}^e d_{r_e}^e$$

with $d_1^e, \ldots, d_{r_e}^e \in \mathbb{F}(\psi)$. By multiplying out denominators for every edge though, we may assume that $d_i^e \in \mathbb{F}[\psi]$ for any edge e and $1 \le i \le r_e$. In other words, the d_i^e are polynomials in ψ.

Given a $f \in \mathbb{F}$, we may define $\alpha_f \in \tilde{\mathcal{F}}(e)(\tilde{\psi})$ to be α restricted to all edges except l and for each d_i^e we substitute every instance of ψ_l with f. Then $\alpha_f \in H_1^{\mathrm{tw}}(\tilde{\mathcal{F}})$ since for every $u \in V_G$ we have

$$\sum_{e \in h_G^{-1}(u)} \tilde{\mathcal{F}}(h, e)(\tilde{a}|_e) - \sum_{e \in t_G^{-1}(u)} \psi_e \tilde{\mathcal{F}}(t, e)(\tilde{a}|_e) = 0.$$

Let $C = (v_1, e_1, \ldots, e_k, v_{k+1})$ be a chain from a canonical form on $\mathcal{F}(l)$. Let S be the span of the e_i. We can express $\alpha|S$ as

$$e_1\gamma_1 + \ldots + e_k\gamma_k$$

for $\gamma \in \mathcal{F}[\psi]$. Since C is from a canonical form, $(d_h - \psi_l d_t)\alpha|_l$ restricted to the span of the v_i is the same as $(d_h - \psi_l d_t)\alpha|_S$. If $\Theta = \{\alpha_f | f \in \mathbb{F}\} \in H_1^{\mathrm{tw}}(\tilde{\mathcal{F}})$ must have dimension at most 1, then so must $(d_h - \psi_l d_t)\Theta|_E$. This implies

$$\begin{bmatrix} \psi_l\gamma_1 \\ \psi_l\gamma_2 - \gamma_1 \\ \vdots \\ \psi_l\gamma_k - \gamma_{k-1} \\ -\gamma_k \end{bmatrix} = c_0 \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_k \\ c_{k+1} \end{bmatrix} \tag{9.4}$$

where $c_0 \in \mathcal{F}(\psi)$ and for $i = 1, \ldots, k+1$ we have $c_i \in \mathcal{F}(\tilde{\psi})$. We also have c_1 and c_{k+1} are nonzero since otherwise $\mathcal{F}$ wouldn't be a full sheaf. Also c_0 is nonzero or else stability is violated on the chain. Dividing by c_0 gives

$$\begin{bmatrix} \psi_l\gamma_1' \\ \psi_l\gamma_2' - \gamma_1' \\ \vdots \\ \psi_l\gamma_k' - \gamma_{k-1}' \\ -\gamma_k' \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_k \\ c_{k+1} \end{bmatrix} \in \left(\mathcal{F}(\tilde{\psi})\right)^{k+1}. \tag{9.5}$$

So

$$\begin{aligned}
\gamma'_k &= \pm c_{k+1} \\
\gamma'_{k-1} &= \pm \psi_l c_{k+1} \pm c_l \\
\gamma'_{k-2} &= \pm \psi_l^2 c_{k+1} \pm \psi_l c_l \pm c_{k-1} \\
&\vdots \\
\gamma'_1 &= \pm \psi_l^l c_{k+1} \pm \psi_l^{k-1} c_l \pm \ldots \pm c_2.
\end{aligned}$$

But $c_1 = \psi_l \gamma'_1 \notin \mathcal{F}(\tilde{\psi})$ since $c_{k+1} \neq 0$.

□

Corollary 9.22. *Let G be a graph with a self loop e on a vertex v and let $\mathcal{F}$ be a minimally gapped sheaf supported on e. Then* $\dim(\mathcal{F}(v)) > \dim(\mathcal{F}(e)) + 1$.

Proof. Note $\dim(\mathcal{F}(v)) \neq \dim(\mathcal{F}(e))$ simply by stability. If $\dim(\mathcal{F}(v)) = \dim(\mathcal{F}(e)) + 1$ then $\tilde{\mathcal{F}}$ has maximum excess 1 and by the previous lemma $h_1^{\mathrm{tw}}(\tilde{\mathcal{F}}) = 2$ contradicting that $\mathcal{F}$ is minimally gapped. □

Chapter 10

Homotopy Simplifications and Vector Bundles

In Lemma 2.4 we gave an alternate definition of the abelian girth using certain subgraphs with no leaves. Pruning a graph, or more generally taking subdivisions can be seen as examples of "homotopy preserving operations" and this view will allow us to define many different ways of transforming a sheaf on a graph while preserving maximum excess and twisted Betti numbers. This will be useful in our proof of Theorem 7.18.

Homotopy operations are also intrinsically interesting, as each sheaf in a homotopy equivalence class has similar properties. In this section we will not be extremely formal in discussing homotopy; rather, we shall define some operations on sheaves and show that they preserve certain sheaf invariants. Later we will describe in rough terms what one might mean by homotopy equivalence and briefly quote theorems in [18] to justify certain theorems.

10.1 Some Homotopy Operations

The following operations are examples of homotopy operations; we will justify the terms later.

Definition 10.1. Let $\mathcal{F}$ be a sheaf on a graph, G, and let $e \in E_G$. If $u \subset \mathcal{F}(e)$ with $\mathcal{F}(e,h)u = 0$ and $\mathcal{F}(e,t)u \neq 0$, we say that u is *tail retractible*. By the *tail retract (at e) of $\mathcal{F}$ along u* we mean the sheaf $\mathcal{F}'$ on G such that

1. the values of $\mathcal{F}'$ equal those of $\mathcal{F}$ except that $\mathcal{F}'(e) = \mathcal{F}(e)/U$, and $\mathcal{F}'(te) = \mathcal{F}(te)/W$, where U is the span of u, and W is the span of $\mathcal{F}(e,t)u$;

2. the restriction maps of $\mathcal{F}'$ are inherited from those of $\mathcal{F}$ in the natural way.

We similarly define a head retract.

Note that in the above definition, the edge e may be a self-loop.

Here we introduce particular notion for the standard notion of a graph contraction.

Definition 10.2. Let G be a graph and let e be an edge. We define the *subdivision of* e *(in* G*)* to be the graph, $s_e(G)$, where e is discarded and replaced with a new vertex and two new edges incident to it; formally $s_e(G)$ has

1. edge set $\{e_1, e_2\} \amalg E_G \setminus \{e\}$;

2. vertex set $V_G \amalg \{e\}$, i.e., V_G with a new vertex with the label e;

3. the head/tail maps of G along with new maps $te_1 = te_2 = e$, $he_1 = he$ and $he_2 = te$.

We caution the reader that while e is an edge in G, e becomes a vertex in $s_e(G)$; this (perhaps unusual) convention for the "new vertex" makes edge contraction in a sheaf quite simple.

Definition 10.3. Let $\mathcal{F}$ be a sheaf on a graph, G, and let $e \in E_G$ be an edge for. We define the *subdivision of* e *(in* $\mathcal{F}$*)*, also called $s_e(\mathcal{F})$ to be the sheaf of $s_e(G)$ (with notation as in Definition 10.2) given by

1. the edge and vertex spaces are the same as in $\mathcal{F}$ except $s_e(\mathcal{F})(e) = \mathcal{F}(e)$ (note e is an edge in G and a vertex in $s_e(G)$) and $s_e(\mathcal{F})(e_1) = s_e(\mathcal{F})(e_2) = s_e(\mathcal{F})(e)$,

2. the restriction maps of $s_e(\mathcal{F})$ are the same as those in $\mathcal{F}$, except tail maps of the new edges e_1, e_2 are both the identity, $s_e(\mathcal{F})(e_1, h) = \mathcal{F}(e, h)$ and $s_e(\mathcal{F})(e_2, h) = \mathcal{F}(e, t)$.

Now we note that the above defined homotopy operations do not change many of the fundamental invariants of a sheaf.

Definition 10.4. We say that two sheaves have *isomorphic homology* if they have isomorphic homology groups and twisted homology groups, and the same maximum excess.

Any such two sheaves also, therefore, have the same Euler characteristic and dual maximum excess. In Subsection 10.4 we show that it is enough to check that the homology groups are the same, when the two sheaves are related by a "functorial procedure," in a certain sense.

Theorem 10.5. *Two sheaves have isomorphic homology if one is obtained from the other by a head or tail retract. (Definition 10.1).*

Proof. Let $\mathcal{F}$ be a sheaf on a graph G with $\mathcal{F}'$ being a tail or head retract on some edge e. The natural map from $\mathcal{F}$ to $\mathcal{F}'$ also maps subsheaves of $\mathcal{F}$ to subsheaves of $\mathcal{F}'$ with the same excess. Any member of $H_1(\mathcal{F})$ or $H_1^{\mathrm{tw}}(\mathcal{F})$ under the same map would also be a member of $H_1(\mathcal{F}')$ or $H_1^{\mathrm{tw}}(\mathcal{F}')$ respectively. Let $v \in G$ be the vertex affected the by head or tail retract, and let W be defined as they are in Definition 10.1. Any member of $H_1(\mathcal{F}')$ or $H_1^{\mathrm{tw}}(\mathcal{F}')$ can be mapped to some subspace S of $\mathcal{F}(E)$ in a natural way. Then d or $d_{\mathcal{F}\psi}$ on that subspace must be zero on all vertices except $\mathcal{F}(v)$ where it is some element of $W(\psi)$. Thus adding some element of V to the edge space over e in S creates an element of $H_1(\mathcal{F})$ or $H_1^{\mathrm{tw}}(\mathcal{F})$ respectively. Similarly, given any $\mathcal{S}$ subsheaf of $\mathcal{F}'$ we can create a subsheaf of $cf.$ with the same excess by adding W to $\mathcal{S}(v)$, adding U to $\mathcal{S}(e)$ and letting all other vertex and edge spaces remain the same. □

Theorem 10.6. *If $\mathcal{F}$ is a sheaf on a graph G and $e \in E_G$ then $\mathcal{F}$ and $s_e(\mathcal{F})$ have isomorphic homology.*

Proof. First we note that

$$\dim\left(s_e(\mathcal{F})(V)\right) = \dim\left(\mathcal{F}(V)\right) + \dim\left(\mathcal{F}(e)\right)$$

and

$$\dim\left(s_e(\mathcal{F})(E)\right) = \dim\left(\mathcal{F}(E)\right) + \dim\left(\mathcal{F}(e)\right)$$

and so subdividing a sheaf preserves the excess of the sheaf. Since subdivision maps subsheaves of $\mathcal{F}$ to subsheaves of $s_e(\mathcal{F})$, we have m.e.$(\mathcal{F}) \leq$ m.e.$(s_e(\mathcal{F}))$. Let $\mathcal{K}$ be a maximizer of $s_e(\mathcal{F})$. Then $\mathcal{K}(e) = \mathcal{K}(e_1) + \mathcal{K}(e_2)$ since if not we could restrict $\mathcal{K}(e)$ to that subspace and have a sheaf with larger excess. Thus

$$\dim(\mathcal{K}(e)) = \dim(\mathcal{K}(e_1)) + \dim(\mathcal{K}(e_e)) - \dim(\mathcal{K}(e_1) \cap \mathcal{K}(e_2)).$$

Define $\mathcal{K}'$ to be the same as $\mathcal{K}$ except

$$\mathcal{K}'(e_1) = \mathcal{K}'(e_2) = \mathcal{K}'(v) = \mathcal{K}(e_1) \cap \mathcal{K}(e_2).$$

The excess of $\mathcal{K}'$ is the same as $\mathcal{K}$ and $\mathcal{K}'$ is a subdivision of some sheaf on G, implying m.e.$(\mathcal{F}) =$ m.e.$(s_e(\mathcal{F}))$.

We now define a map $m : H_1^{\mathrm{tw}}(\mathcal{F}) \to H_1^{\mathrm{tw}}(s_e(\mathcal{F}))$. Given $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ let $(c_1, \ldots, c_t)$ be a representation of α with respect to $(b_1, \ldots b_t)$ where the b_i

are a basis for $\mathcal{F}(E)$ with $b_1, \ldots, b_r$ a basis for $\mathcal{F}(e)$ for some $r < t$. Here we may assume that $c_i \in \mathbb{F}[\psi_G]$ by multiplying out denominators. We define c_i' for $1 \leq i \leq t$ to be the same as c_i except we substitute every occurrence of ψ_e with $\psi_{e_1}\psi_{e_2}^{-1}$. We also define $m(\alpha)$ restricted to the sum of all the edge spaces in $E_{s_e(G)} \setminus \{e_1, e_2\}$ has the representation $c_{r+1}', c_{r+2}', \ldots c_t'$. Let $g_1, \ldots, g_t$ and $h_1, \ldots, h_t$ be the bases for e_1 and e_2 respectively with $g_i = h_i = b_i$. The representation of $m(\alpha)$ restricted to e_1 is $(c_1', \ldots, c_t')$ with respect to $(g_1, \ldots, g_t)$ and the representation of $m(\alpha)|_{e_2}$ is $(\psi_{e_1}\psi_{e_2}^{-1}c_1', \ldots, \psi_{e_1}\psi_{e_2}^{-1}c_t')$ with respect to $(h_1, \ldots h_r)$. All of this gives a representation for $m(\alpha)$. For any $\beta \in H_1^{\mathrm{tw}}(s_e(\mathcal{F}))$, we must have

$$\beta|_{e_1} = \psi_{e_1}\psi_{e_2}^{-1}\beta|_{e_2}$$

or else the twisted differential map would sent β to something nonzero on the vertex e. Thus m is invertible and a bijection between $H_1^{\mathrm{tw}}(\mathcal{F})$ and $H_1^{\mathrm{tw}}(s_e(\mathcal{F}))$.

□

10.2 A Twisted Homotopy Operation

Here we define an example of what we will call a "twisted homotopy operation;" this will be formally explained in Subsection 10.4. For now we just define the opearation.

Definition 10.7. Let $\mathcal{F}$ be a sheaf of $\mathbb{F}$-vector spaces on a graph, G. Let $e \in E_G$ be a self-loop such that if I is the sum of the images of $\mathcal{F}(e, h)$ and $\mathcal{F}(e, t)$, then for some $\psi \in \mathbb{F}$ we have that $\mathcal{F}(e, h) + \psi\mathcal{F}(e, t)$ has no kernel. We define the *contraction of e (in $\mathcal{F}$)*, denoted $\mathcal{F}//e$, to be the sheaf in G given by

1. the values of $\mathcal{F}//e$ agree with those of $\mathcal{F}$, except that $(\mathcal{F}//e)(he) = \mathcal{F}(e)/I$ and $(\mathcal{F}//e)(e) = 0$;
2. the restriction maps of $\mathcal{F}//e$ are the same as those in $\mathcal{F}$, except that the restrictions taken to $\mathcal{F}(he)$ are set to zero.

Example 10.8. Let G be the sheaf with a single vertex and a single self loop. Let $\mathcal{F}_1$ be the sturcture sheaf of G, and let $\mathcal{F}_2$ be the sheaf which is the same as the structure sheaf except that the tail restriction (at the single edge) is minus the identity. Then the Betti numbers of $\mathcal{F}_1$ both equal one; those of $\mathcal{F}_2$ are both zero, provided that the underlying field, $\mathbb{F}$, has characteristic different than two. (Morally, $\mathcal{F}_2$ is essentially the Möbius strip.) Hence self-loop contraction does not preserve Betti numbers. (The twisted Betti numbers of $\mathcal{F}_1, \mathcal{F}_2$ are all zero.)

Definition 10.9. We say that two sheaves have *isomorphic twisted invariants* if they have isomorphic twisted homology groups, and the same maximum excess.

Any such two sheaves also, therefore, have the same Euler characteristic and dual maximum excess. Similarly to Definition 10.4, in Subsection 10.4 we will explain that it suffices to check equality of twisted homology groups when one sheaf is obtained from another by a "sufficiently functorial operation."

For completeness we mention the following theorem.

Theorem 10.10. *Let $\mathcal{F}$ be a sheaf on a graph, G. Assume that for some self-loop $e \in E_G$ we have that $\mathcal{F}(e,h)$ and $\mathcal{F}(e,t)$ are isomophisms. Then $\mathcal{F}$ and the contraction of $\mathcal{F}$ at e have isomorphic twisted invariants.*

Proof. Easy, using the fact that any square matrix has finitely many eigenvalues. ☐

10.3 Example: Vector Bundles

Vector bundles, which we now define, is an illustrative application of edge contraction.

Definition 10.11. A sheaf, $\mathcal{F}$, on a connected graph, G, is a *vector bundle* if all its restriction maps are isomorphisms. In this case, all $\mathcal{F}$'s values have the dimension, we which call the *dimension* of the vector bundle.

Theorem 10.12. *Let $\mathcal{F}$ be a vector bundle of dimension d on a connected graph, G. Then*

$$h_1^{\mathrm{tw}}(G,\mathcal{F}) = \mathrm{m.e.}(\mathcal{F}) = d\,\mathrm{m.e.}(G) = d\,h_1^{\mathrm{tw}}(G),$$

$$h_0^{\mathrm{tw}}(G,\mathcal{F}) = \mathrm{d.m.e.}(\mathcal{F}) = d\,\mathrm{d.m.e.}(G) = d\,h_0^{\mathrm{tw}}(G),$$

and

$$\chi(\mathcal{F}) = d\,\chi(G).$$

Proof. Any vector bundle remains so after an edge contraction, and a vector bundle can be contracted along every edge. Hence it suffices to prove the theorem when G has exactly one vertex. If G has no edges, then the theorem is immediate, and if G has more than one edge we contract along any self-loop. ☐

10.4 A High Road to Homotopy

The goal of this subsection is to describe homotopy in simple and fairly general terms. However, some of the theory we will use (the L^2 Betti number computation implicit in [17], for example), is not self-contained here. This subsection is not essential to the rest of this paper; rather, it unifies and simplifies the ad hoc constructions of the previous subsections in this section.

In topology one defines a notion of homotopic spaces; among other things, homotopy preserves certain invariants. Regarding homology groups, topological homotopies usually give "chain homotopies" of the chains from which the homology groups are computed.

Ideally we would define a topological notion of homotopy for sheaves, say based on the Zariski topology or étale topology, and then show that they result in the appropriate chain homotopies. However, we will content ourselves here to view the notion of a homotopy in terms of chains alone.

Definition 10.13. Let $\mathbb{F}$ be a field, and let $u\colon \mathcal{F} \to \mathcal{F}$ be an morphism of a sheaf of $\mathbb{F}$-vector spaces, $\mathcal{F}$, on a graph, G, to itself. We say that u is *null homotopic* if there is a linear map $k\colon \mathcal{F}(V) \to \mathcal{F}(E)$ for which

$$u_E = kd_{\mathcal{F}}, \quad u_V = d_{\mathcal{F}}k,$$

where, as usual, $d_{\mathcal{F}} = d_{\mathcal{F},h} - d_{\mathcal{F},t}$ is the differential map of $\mathcal{F}$. We say that u is *weakly null homotopic* if for a full twist, ψ, on G there is a linear map $K\colon \mathcal{F}(V)(\psi) \to \mathcal{F}(E)(\psi)$ for which

$$u_E = Kd_{\mathcal{F},\psi}, \quad u_V = d_{\mathcal{F},\psi}K,$$

where, as usual, $d_{\mathcal{F},\psi} = d_{\mathcal{F},h} - \psi d_{\mathcal{F},t}$ is the twisted differential map of $\mathcal{F}$ (with respect to the fill twist ψ).

Definition 10.14. Let $\mathcal{F}, \mathcal{G}$ be sheaves on a graph, G. Given maps $u\colon \mathcal{F} \to \mathcal{G}$ and $w\colon \mathcal{G} \to \mathcal{F}$, we say that (u, w) is a *homotopy pair* (respectively, *weak homotopy pair*) for $(\mathcal{F}, \mathcal{G})$ if $uw - 1$ and $wu - 1$ are null homotopic (respectively, weakly null homotopic), where 1 denotes the identity map.

We shall use (as is typical) the term "homotopy" and "weak homotopy" in diverse ways; for example, we refer to a morphism $u\colon \mathcal{F} \to \mathcal{G}$ alone as a *(weak) homotopy*, provided there exists there exists a w for which (u, w) is a (weak) homotopy pair; we say $\mathcal{F}$ and $\mathcal{G}$ are *(weakly) homotopic* if there exists a weakly homotopy pair for $(\mathcal{F}, \mathcal{G})$.

As an example we show that if $\mathcal{F}$ is a sheaf on a graph G then $s_e(\mathcal{F})$ are homotopic. Let e_1 and e_2 be the new edges of $s_e(\mathcal{F})$ as described in

Definition 10.3) and let v_1 and v_2 be the head and tail respectively of e in G. We define the map $u\colon \mathcal{F} \to s_e(\mathcal{F})$ as the following collection of linear maps

1. for any $f \in E_G$ other than e $u_f\colon \mathcal{F}(f) \to s_e(\mathcal{F})(f)$ is the identity map,
2. $u_e\colon \mathcal{F}(e) \to s_e(\mathcal{F})(e_1) \oplus s_e(\mathcal{F})(e_2)$ is given by $u_e(x) = (x, -x)$ for all $x \in \mathcal{F}(e)$,
3. $u_v\colon \mathcal{F}(v) \to s_e(\mathcal{F})(v)$ is the identity map for any $v \in V_G$ other than v_1 or v_2,
4. $u_{v_i}\colon \mathcal{F}(v_i) \to s_e(\mathcal{F})(v_i)$ is the zero map for $i \in \{1, 2\}$.

We now define another map $w\colon s_e(\mathcal{F}) \to \mathcal{F}$ as the following collection of linear maps

1. $w_f\colon s_e(\mathcal{F})(f) \to \mathcal{F}(f)$ is the identity map for any $f \in E_{s_e(G)}$ other than e_1 or e_2,
2. $w_{e_1}\colon s_e(\mathcal{F})(e_1) \to \mathcal{F}(e)$ is also the identity map,
3. $w_{e_2}\colon s_e(\mathcal{F})(e_2) \to \mathcal{F}(e)$ is the zero map,
4. $w_v\colon s_e(\mathcal{F})(v) \to \mathcal{F}(v)$ is the identity map for any $v \in V_{s_e(G)}$ other than e ,
5. $w_e\colon s_e(\mathcal{F})(e) \to \mathcal{F}(v_2)$ is $\mathcal{F}(e, t)$.

Checking each of the vertex and edge spaces shows wu is the identity map on $\mathcal{F}$, implying $wu - 1$ is null homotopic by setting k from Definition 10.13 to be the zero map. The map $(uw-1)_E$ is the zero map on each of the edge spaces except $(uw-1)_{e_1 \oplus e_2}(x, y) = (0, -x-y)$. For $(uw-1)_V$ we find the zero map on all vertex spaces except $(uw-1)_{e \oplus v_2}(x, y) = (-x, \mathcal{F}(e, t)(x))$. If we let $k_e\colon e \to e_2$ be the identity and define $k\colon s_e(\mathcal{F})(V) \to s_e(\mathcal{F})(E)$ be the zero map elsewhere, then k satisfies the conditions of null homotopy.

Theorem 10.15. *Homotopic sheaves have naturally isomorphic homology groups. Weakly homotopic sheaves have isomorphic twisted homology groups.*

Recall from [17] that if $\mathcal{F}$ is a sheaf on a graph, G, and $u\colon G' \to G$ is a morphism of graphs, then the *pullback of $\mathcal{F}$ via u*, denoted $u^*\mathcal{F}$, is the naturally arising sheaf on G', i.e., the sheaf whose value at any $P \in V_{G'} \amalg E_{G'}$ equal is just $\mathcal{F}(u(P))$, and whose restriction maps are given by $(u^*\mathcal{F})(e, h) = \mathcal{F}(u(e), h)$, and similarly with h replacing t.

Definition 10.16. We say that sheaves $\mathcal{F}, \mathcal{G}$ on a graph, G, are *universally homotopic* (respectively, *weakly universally homotopic*) if for every covering map $u\colon G' \to G$ we have that $u^*\mathcal{F}$ and $u^*\mathcal{G}$ are homotopic (respectively, weakly homotopic).

(The notion of a covering map is the usual one, also given in [17].)

The operations of retracting and edge contraction (and other such "general" operations) have a sort of "functoriality" or "commuting with pullbacks" in a way that makes the universality almost immediate from the basic notion. For example, contracting a sheaf along an edge, e, and pulling it back via a covering map is the same as first pulling back and then contracting along all the edges in the preimage of e.

Theorem 10.17. *Universally homotopic sheaves have isomorphic twisted homology groups and the same maximum excess.*

Proof. The twisted homology groups arise from the the homology groups of abelian covers of the graph. Let $\mathcal{F}$ be a sheaf on a graph G. Take a random Galois cover $\mu : G' \to G$ with group $\mathbb{Z}/q\mathbb{Z}$ with q prime. Friedman shows as an immediate result of Lemma 1.17 in [17] that $h_1(\mu^*\mathcal{F})/p$ tends to $h_1^{\mathrm{tw}}(\mathcal{F})$ in probability as p goes to infinity. Thus if $\mathcal{F}$ is universally homotopic to a sheaf $\mathcal{G}$ on a graph, then they have isomorphic twisted homology groups.

In [17] Friedman shows that $\mathrm{cov}(G)$, the set of covering maps over a graph G, is a directed set using fibre products to create a partial ordering of the covers. Friedman then shows for a sheaf $\mathcal{F}$ over G that

$$\mathrm{m.e.}(\mathcal{F}) = \frac{\lim_{\phi\in\mathrm{cov}(G)} h_1^{\mathrm{tw}}(\phi^*\mathcal{F})}{\deg(\phi)}.$$

Hence universally homotopic sheaves have equal maximum excess.

□

10.5 Contracting and Subdividing Gapped Sheaves

The inverse operation of subdivision is called *contraction*. Instead of being defined for a given edge like subdivision, contraction is defined for a given degree 2 vertex with two distinct incident edges. Contraction simply replaces the degree 2 vertex v and the incident edges, e_1 and e_2, with a single edge. Contraction of a vertex v on a sheaf $\mathcal{F}$ is well-defined when v has degree 2, and if e_1 and e_2 are incident upon v, then they are distinct and the restriction maps from $\mathcal{F}(e_1)$ and $\mathcal{F}(e_2)$ are both the identity map.

Corollary 10.18. *Let G be a graph and let W be a beaded path or beaded cycle in G. If $\mathcal{F}$ is a minimally gapped sheaf on G from a sheaf collection closed under quotients then for each vertex v of W that is not the starting or terminating vertex, $s_v^{-1}(\mathcal{F})$ is a well defined sheaf and if $\mathcal{F}'$ is the sheaf produced by contracting every vertex of W besides the starting and terminating vertices then $\mathcal{F}'$ is gapped and stable.*

Proof. Let e_1 and e_2 be the edges incident upon a degree 2 vertex v in G and without loss of generality, assume the tail map sends both edges to v. By Lemma 9.6, $\operatorname{im}(\mathcal{F}(e_1, t)) = \operatorname{im}(\mathcal{F}(e_2, t))$. We may assume $\mathcal{F}(v) = \operatorname{im}(\mathcal{F}(e_1, t))$ since if not we could restrict $\mathcal{F}(v)$ to $\operatorname{im}(\mathcal{F}(e_1, t))$ and have a gapped sheaf with smaller total dimension. By a change of basis, and since the tail maps are all injections, we may also assume $\mathcal{F}(e_1, t)$ and $\mathcal{F}(e_2, t)$ are both the identity map. Thus $s_v^{-1}(\mathcal{F})$ is a well-defined sheaf for every degree 2 vertex and is gapped since subdivision produces universally homotopic sheaves. This remains true after repeated contractions of the vertices in W besides the starting and terminating vertex. Contraction preserves stability since it preserves the Euler characteristic on sheaves for which it is well-defined, and contraction is a bijection from the set of subsheaves of $\mathcal{F}$ for which contraction is well-defined to set subsheaves of $s_v^{-1}(\mathcal{F})$. $\square$

We now show a condition under which subdivision of a minimally gapped sheaf produces another minimally gapped sheaf. We say a gapped sheaf $\mathcal{F}$ from $\mathcal{C}$ is *minimal on an edge e* if, for any gapped sheaf $\mathcal{H}$ from $\mathcal{C}$ we have that $\dim(\mathcal{F}(e)) \leq (\mathcal{H}(e))$.

Theorem 10.19. *Let G be a graph and e an edge of G. If $\mathcal{F}$ is a minimally gapped sheaf on G from $\mathcal{C}$, the collection of all sheaves on G, and $\mathcal{F}$ is minimal on an edge e then $s_e(\mathcal{F})$ is minimally gapped over the sheaf collection of all sheaves on $s_e(G)$. If e_1 and e_2 are the new edges in $s_e(G)$, then $s_e(\mathcal{F})$ is minimal on the edges e_1 and e_2.*

Proof. Let $\mathcal{K}$ be a minimally gapped sheaf on $s_e(G)$ over the collection of all sheaves on $s_e(G)$ and suppose $\dim_T(\mathcal{K}) < \dim_T(s_e(\mathcal{F}))$. Since e is a vertex of degree 2 in $s_e(G)$ and $\mathcal{K}$ is minimally gapped, by Corollary 10.18 $s_e^{-1}(\mathcal{K})$ is a well-defined sheaf on G. We call this sheaf $\mathcal{K}'$.

Since subdivision on an edge e simply adds a new edge and vertex with the same dimension as e we have that

$$\dim_T(\mathcal{K}') + 2\dim(\mathcal{K}'(e)) = \dim_T(\mathcal{K})$$

and

$$\dim_T(\mathcal{F}) + 2\dim(\mathcal{F}(e)) = \dim_T(s_e(\mathcal{F})).$$

Since $\dim_T(\mathcal{K}) < \dim_T(s_e(\mathcal{F}))$ and $\dim_T(\mathcal{F}) \leq \dim_T(\mathcal{K}')$ since $\mathcal{F}$ is minimally gapped, we have that $\dim(\mathcal{K}(e)) < \dim(\mathcal{F}(e))$, contradicting our assumption.

Let $\mathcal{H}$ be a gapped sheaf on $s_e(G)$. By Lemma 9.6 and arguments from Corollary 10.18, there exists a sheaf $\mathcal{I}$ that has the same dimension as $\mathcal{H}$ or less on every edge and $\dim(\mathcal{I}(e_1)) = \dim(\mathcal{I}(e_2))$ and $\mathcal{I}(e_1, t)$ and $\mathcal{I}(e_2, t)$ are both the identity map. Thus $s_e^{-1}(I)$ is a well-defined sheaf on G which we will call $\mathcal{I}'$. Then

$$\dim(\mathcal{I}'(e)) = \dim(\mathcal{I}'(e_1)) = \dim(\mathcal{I}'(e_2))$$

and since $\dim(cF(e)) \leq \dim(\mathcal{I}'(e))$ we have that $\dim(s_e(\mathcal{F})(e_1)) \leq \dim(\mathcal{I}(e_1))$ and $\dim(s_e(\mathcal{F})(e_2)) \leq \dim(\mathcal{I}(e_2))$.

□

In other words, if a minimally gapped sheaf from the collection of all sheaves results is minimal on all edges, any sequence of subdivisions on that sheaf also results in minimally gapped sheaf from the collection of all sheaves results is minimal on all edges.

Theorem 10.20. *Subdivision and contraction map full sheaves to full sheaves and stable sheaves to stable sheaves.*

Proof. First we show subdivision can't map a full sheaf to a sheaf that isn't full. Let $\mathcal{F}$ be a full sheaf on a graph G with an edge e. Let $b_1, \ldots, b_r$ be a basis for $\mathcal{F}(e)$ and $b_1, \ldots, b_r, \ldots, b_s$ a basis for $\mathcal{F}(E)$. Let $\mathcal{H} = s_e(\mathcal{F})$ be a sheaf on the graph $H = s_e(G)$ and let e_1 and e_2 be the new edges produced by the subdivision. Since $\mathcal{F}(e_1)$ and $\mathcal{F}(e_2)$ are identical to $\mathcal{F}(e)$, let $b_1^i, \ldots, b_r^i$ be the basis the for $\mathcal{H}(e_i)$ for $i = 1, 2$ where we identify $b_1^i = b_i$. If $\alpha \in H_1^{\mathrm{tw}}(\mathcal{F})$ has representation $(d_1, \ldots, d_s)$ with respect to $(b_1, \ldots, b_s)$, then

$$(\psi_{e_1}\psi_{e_2}^{-1}d_1, \ldots, \psi_{e_1}\psi_{e_2}^{-1}d_r, d_1, \ldots, d_r, d_{r+1}, \ldots d_s)$$

is a representation of an $\beta \in H_1^{\mathrm{tw}}(\mathcal{H})$ with respect to

$$(b_1^1, \ldots, b_r^1, b_1^2, \ldots, b_r^2, b_{r+1}, \ldots, b_s).$$

All elements of the representation are nonzero since the d_i are all nonzero, and that is true no matter our original choice of basis. If contraction maps a full sheaf to a sheaf that isn't full, then subdivision maps a sheaf that isn't full to a sheaf that is. But from our previous arguments, if any of the d_i are zero in our representation of α then there is also a zero in our representation of β.

Subdivision on any subsheaf of $\mathcal{F}$ does not change it's Euler characteristic but it does not map to all subsheaves of $\mathcal{H}$. Let v_1 be the vertex in H incident to e_1 but not e_2 if e is not a self loop in G, and define v_2 similarly. Let $\mathcal{G}$ be a subsheaf $\mathcal{H}$ and let $\mathcal{G}'$ be defined the same of $\mathcal{G}$ except $\mathcal{G}'(e_1) = \mathcal{G}'(e_2) = \mathcal{G}(e_1) + \mathcal{G}(e_2)$ and, if e is not a self loop in G, $\mathcal{G}'(v_i) = \mathcal{G}(v_i) + \mathrm{im}\mathcal{G}'(e_i, h)$. We also set $\mathcal{G}'(e) = \mathcal{G}(e_1) + \mathcal{G}(e_2)$, and we note $\mathcal{G}'(e)$ must be a subspace of $\mathcal{G}(e)$. For e a self loop in G, we do not need to change the vertex spaces for the subsheaf $\mathcal{G}'$ to be well defined. The excess of $\mathcal{G}'$ is at least as large as the excess of $\mathcal{G}$, since the amount we add to the total edge dimension is at least as large as the amount we add to the total vertex dimension. $\mathcal{G}'$ must have excess at most zero though if $\mathcal{F}$ is stable, since it can be mapped to by a subdivision of a subsheaf of $\mathcal{F}$.

Contraction also doesn't change the Euler characteristic on any subsheaf of a sheaf, and contraction on the subsheaves of $\mathcal{H}$ for which contraction is well-defined maps to all the subsheaves of $\mathcal{F}$. □

We finish this section by remarking that if $\mathbb{F}$ is the reals or complex numbers, then there is a natural way of viewing a sheaf on a graph as a CW-complex, much in the way that a graph can be viewed as a CW-complex; in this case, retracting and contracting are clearly topological homotopies. However, we imagine that one could do this, say for any algebraically closed field, say for the Zariski topology (or, perhaps, the étale topology if need be). We shall leave this for future work.

Chapter 11

Proofs of Theorems 7.17 and 7.18

In this section we give the proofs of the main theorems of our paper.

11.1 Three Minimal Gapped Sheaves

We will now describe three gapped sheaves that will be necessary for the main proofs in this paper. The first sheaf known as the *unhappy bundle* is given by Friedman in [18]. This sheaf is over the minimal figure-eight graph on vertex v with two directed self-loops, e_1 and e_2. The sheaf $\mathcal{U}$ is defined as

$$\mathcal{U}(v) = \mathbb{F}^4, \quad \mathcal{U}(e_i) = \mathbb{F}^2 \text{for } i = 1, 2$$

with

$$d_h = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad d_t = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}$$

where these matrices multiply the coordinates of $\mathcal{U}(E)$ as a column vector to the right of the matrix with $\mathcal{U}(E)$'s coordinates ordered as $\mathcal{U}(e_1) \oplus \mathcal{U}(e_2)$. Friedman also gives $H_1^{\mathrm{tw}}(\mathcal{U})$ and shows that maximum excess is zero. Corollary 9.7 clearly shows that $\mathcal{U}$ is a T-minimal gapped sheaf over the figure-eight graph, B_2.

We now introduce a gapped T-minimal sheaf over the minimal barbell graph, though we will not prove its minimality until our proof of Theorem 7.18. Consider B, the barbell graph labelled in the following way. Vertex v_1 has a self loop e_1 and vertex v_2 has a self loop e_2. The edge e_0 is directed from v_1 to v_2. The sheaf $\mathcal{V}$ which we will call the *small barbell* has

$$\mathcal{V}(v_i) = \mathbb{F}^4, \quad \mathcal{V}(e_i) = \mathbb{F}^2 \text{ for } i = 1, 2, \quad \mathcal{V}(e_0) = \mathbb{F}^4.$$

Both restriction maps of $\mathcal{V}(e_0)$ are the identity map. For $i = 1, 2$, the restriction maps for $\mathcal{V}(e_i)$ are the same as the restriction maps for $\mathcal{U}(e_i)$ except they map to the $\mathbb{F}^4$ over v_i. A quick computation shows that $H_1^{\mathrm{tw}}(\mathcal{V}) = 1$.

Let $\mathcal{V}'$ be a subsheaf of $\mathcal{V}$ with excess equal to the maximum excess of $\mathcal{V}$. The excess of $\mathcal{V}'$ is at most that of the subsheaf $\mathcal{V}''$ of $\mathcal{V}$ where $\mathcal{V}''$ is defined by $\mathcal{V}''(e_i) = \mathcal{V}'(e_i)$ for $i = 1, 2$ and

$$\mathcal{V}''(v_1) = \mathcal{V}''(v_2) = \mathcal{V}''(e_0) = \mathcal{V}'(v_1) + \mathcal{V}'(v_2).$$

This can be seen by noting that any increase in vertex dimension is compensated for by an increase in the dimension of the edge space of e_0. Let the subsheaf $\mathcal{U}'$ of $\mathcal{U}$ be the same as $\mathcal{V}''$ over e_1 and e_2 and set $\mathcal{U}'(v) = \mathcal{V}''(v_1)$. This subsheaf has the same excess as $\mathcal{V}''$. Thus

$$\text{m.e.}(\mathcal{U}) = \text{m.e.}(\mathcal{V}) = 0.$$

Given an element of α of $H_i^{\mathrm{tw}}(\mathcal{U})$ we create an element of $H_i^{\mathrm{tw}}(\mathcal{V})$ in order to show that $\mathcal{V}$ is a gapped sheaf. The element α is a collection of vectors, $\alpha(e_i) \in \mathcal{U}(e_i)$, such that

$$d_{\mathcal{U}^\psi}(\alpha(e_1)) + d_{\mathcal{U}^\psi}(\alpha(e_2)) = 0.$$

Define β as the collection of vectors from the edge spaces of $\mathcal{V}$ with

$$\beta(e_1) = \psi_{e_0}\alpha(e_1), \quad \beta(e_2) = \alpha(e_2) \text{and} \beta(e_0) = d_{\mathcal{U}^\psi}(\alpha(e_1)).$$

It is now easy to verify that β is then a member of $H_i^{\mathrm{tw}}(\mathcal{V})$.

Now we describe $\mathcal{W}$, the T-minimal sheaf over the minimal figure eight graph with vertices v_1, v_2 and edges e_1, e_2 and e_3. Let $W = \mathbb{F}^6$ have basis $\alpha, \beta, \gamma, \delta, \epsilon, \zeta$. Then

$$\mathcal{W}(e_1) = \text{Span} < \alpha, \delta >, \quad \mathcal{W}(e_2) = \text{Span} < \beta, \epsilon >, \quad \mathcal{W}(e_3) = \text{Span} < \gamma, \zeta >,$$

and

$$\mathcal{W}(v_1) = W/\,\text{Span} < \beta{-}\gamma, \delta{-}\epsilon, \zeta{-}\alpha >, \quad \mathcal{W}(v_2) = W/\,\text{Span} < \alpha{-}\beta, \gamma{-}\delta, \epsilon{-}\zeta > .$$

The restriction maps for any edge are the canonical quotient maps. Computing the kernel of the twisted difference map shows that $h_1^{\mathrm{tw}}(\mathcal{V}) = 1$. Let G' be the figure eight graph without e_3 and let u be a connected degree 2 covering map of the figure eight graph for which $u^{-1}(G')$ is a disconnected graph. By computation, we have that $h_1^{\mathrm{tw}}(u^*\mathcal{V}) = 0$. Since $h_1^{\mathrm{tw}}(\mathcal{F}) \geq \text{m.e.}(\mathcal{F})$ for any sheaf $\mathcal{F}$, $\text{m.e.}(u^*\mathcal{V}) = 0$. Equation 7.1 implies $\text{m.e.}(\mathcal{V}) = 0$ as well. Thus the sheaf is gapped and, since every edge has dimension 2, Corollary 9.7 shows it is in fact a gapped T-minimal sheaf.

11.2 Proof of Theorem 7.17

Let $\mathcal{W}'$ be the following sheaf on the graph G with vertices t and v and five edges $e_1, ..., e_5$ from one vertex to the other. The vertex spaces both are $\mathbb{F}^6$ and let $\mathcal{W}'(e_i) = \mathcal{W}(e_i)$ for $i \in 1, 2, 3$. Here $\mathcal{W}$ is the gapped sheaf over the figure eight graph we described in the previous subsection. Using the same notation for the basis of $\mathbb{F}^6$ as we did in defining $\mathcal{W}$, we define

$$\mathcal{W}'(e_4) = \text{Span} < \beta - \gamma, \delta - \epsilon, \zeta - \alpha >$$

and

$$\mathcal{W}(e_5) = \text{Span} < \alpha - \beta, \gamma - \delta, \epsilon - \zeta > .$$

Similarly to the case of $\mathcal{W}$, we show the maximum excess of $\mathcal{W}'$ is zero by examine a cover of G. Let G' have vertices four vertices, t_1, t_2, u_1 and u_2 and edges e'_i, e''_1 for $i = 1, ..., 5$. For $i = 1, 2$ and 3 let e'_i have tail t_1 and head v_1 and e''_i have tail t_2 and head v_2. For all other i, let e'_i have tail t_1 and head v_2 while e''_i has tail t_2 and head v_1. Then G' is a covering graph of G with covering map u that takes v_i to v, t_i to t, e'_i to e_i and e''_i to e_i as well. Computing the kernel of the twisted difference map for $u^*\mathcal{W}'$ show that $h_1^{\text{tw}}(u^*\mathcal{W}') = 0$ and so m.e.$(\mathcal{W}') = 0$

11.3 Proof of Theorem 7.18

First, we prove that there does not exist gapped sheaves supported on only a cycle or a tree. Then we show that if a graph has Euler characteristic -1 than the edge dimension of a minimally gapped sheaf is derived readily from the three examples of minimally gapped sheaves discussed earlier. Finally we consider the case of a minimally gapped sheaf supported on a graph G of Euler characteristic less than -1. Since the figure-eight graph and theta graph have dimension 2 on all edges, it will be easy to show that minimality implies the G cannot contain either as a subgraph. We then argue about the structure of G and use the twist trick on loops to show that there exists a smaller gapped sheaf on a barbell subgraph of G.

Let $\mathcal{F}$ be a sheaf on a graph G and $\pi : G[\mathbb{Z}] \to G$ be the universal abelian covering. Lemma 1.30 from [17] states that $H_1^{\text{tw}}(\mathcal{F})$ is non-trivial iff $H_1(\pi^*\mathcal{F})$ non-trivial. If G is a tree or cycle then $G[\mathbb{Z}]$ is a tree and so $H_1(\pi^*\mathcal{F})$ and $H_1^{\text{tw}}(\mathcal{F})$ must both be trivial. Thus any sheaf supported on only a cycle or graph is not gapped.

Now we argue that the previously defined gapped sheaf $\mathcal{V}$ on the minimal barbell graph is T-minimal. Corollary 9.7 implies this so long as $\mathcal{V}$ is minimal on the edge e_0. By Corollary 9.22, if $\mathcal{F}'$ is identical to $\mathcal{F}$ except $\mathcal{F}'(e_0) = 0$ then $\chi(\mathcal{F}') \geq 4$ as it is just two self loops. So by stability, $\dim(\mathcal{F}(e_0)) \geq 4$, showing $\mathcal{V}$ is minimally gapped.

As mentioned before, $\mathcal{U}$ and $\mathcal{W}$ are minimal gapped sheaves on the figure-eight graph and theta graph respectively since each edge on those sheaves has dimension 2, the minimal dimension possible. By Lemma 9.6, any minimal gapped sheaf is supported on a pruned graph. So Theorem 10.19 implies Theorem 7.18 in the case the graph has Euler characteristic -1, as any such graph is either a figure-eight, theta or barbell graph.

Let $\mathcal{F}$ be a minimally gapped sheaf from the collection of all sheaves on a graph G and now suppose $\mathcal{F}$ is supported on a graph G' with $\chi(G') < -1$. If G' contains a subgraph S that is homeomorphic to a figure-eight graph or theta graph then $2|E_S| < 2|E_{G'}| \leq |\mathcal{F}(E_G)|$, but $|2E_S|$ is the edge dimension of a gapped sheaf on S when S is a figure-eight graph or a theta graph contradicting the minimality of $\mathcal{F}$.

Thus we may assume G' contains no figure-eight graph or theta graph as a subgraph. So any pruned subgraph of G' with Euler characteristic -1 is a barbell graph. Let B be a barbell graph that is a subgraph of G' with minimal bar length out of all the barbell graphs that are subgraphs of G'. Let R be the bar of B and let r be the number of edges of R. Let C_1 and C_2 be the two cycles in B. We also set c_1 and c_2 to be the number of edges in C_1 and C_2 respectively. If $\mathcal{B}$ is the minimal gapped sheaf on B, then

$$\dim(\mathcal{B}(E)) = 2c_1 + 2c_2 + 4r.$$

Since

$$\dim(\mathcal{B}(E)0 > \dim(\mathcal{F}(E)) > 2|E_{G'}|$$

we have

$$|E_{G'}| < c_1 + c_2 + 2r. \tag{11.1}$$

In other words, there are at most r edges in G' that are not also in B.

There can be no walk from C_1 to C_2 nor a walk from either C_1 or C_2 to B in G' else G' would contain a theta graph. Any walk from either C_1 or C_2 to itself must begin and end with the same edge or else G' contains a theta graph or a figure-eight graph in the case the walk is closed. Suppose there exists a walk not in B but in G' from C_i to C_i for $i = 1, 2$. Since, G' contains no leaves and the only Euler characteristic -1 subgraphs are barbell graphs, walk contain a cycle as a subgraph that is not in B. This implies there is a barbell graph in G' that contains C_1 as a subgraph but no other part of B. The bar of this barbell graph has at least r edges, since the bar length of B is minimal, but then this violates Equation 11.1. Hence any walk from B to itself must begin and end in R, excluding the starting and terminating vertices of R as those in C_1 or C_2.

If a walk from R to R in G' has different first and last edges, than G' contains a cycle with at least one vertex in R. This implies there is barbell graph as a subgraph of G' whose bar is only a portion of R, contradicting the minimality of the length of R. Thus any non-backtracking walk from R to R contains a path P from R to a cycle C_3. Let v be the vertex in P and R and let R_i be the path in R from from C_i to v. Set r_i to be the number of edges in R_i, for $i = 1, 2$, let p be the number of edges in P and let c_3 be the number of edges in C_3. By Equation 11.1, $p \leq r$. Note P, C_3, R_i and C_i form a barbell graph for $i = 1, 2$. Thus by the minimality of R, $r_i + p <\geq r$ for $i = 1, 2$ and since $r_1 + r_2 = r$ this means $p \geq r/2$. Thus the only edge of G' not in B with one endpoint in B is in P, or else we'd have another path with at least $r/2$ edges and another cycle with at least 1 edge which would contradict Equation 11.1.

The same arguments can be repeated from earlier to show there is no walk in G' from P and C_3 to themselves or to B. Thus G' is the union of the vertex sets and edge sets of C_1, C_2, C_3, R and P. So the C_i are beaded cycles in G' starting and terminating at a degree 3 vertex. By Corollary 10.18 we may contract each C_i to a self loop on the degree 3 vertex, resulting in a new gapped sheaf $\mathcal{H}$ supported on a graph H. Note H also contains P and R as subgraphs and the sheaf $\mathcal{H}$ is identical to $\mathcal{F}$ on P and R. It is also a full, stable gapped sheaf by Theorem 10.20 and because contraction is universally homotopic. If s_i is the self loop that results from contracting C_i then we also have that $\dim(\mathcal{H}(s_i)) = \dim(\mathcal{F}(c_i))$ where c_i is any edge in C_i. For each s_i define S_i to be the cycle composed of s_i and the vertex v_i that s_i is incident to. Then by Corollary 9.22, $\dim(\mathcal{H}(s_i)) + 1 < \dim(\mathcal{H}(v_i)$. Any chain on $\mathcal{H}(s_i)$ has edge dimension exactly one less than it's vertex dimension. Thus a canonical form on $\mathcal{H}(s_i)$ contains at least two chains. Let F_i be the span of the first elements of the chains from a canonical form on $\mathcal{H}(s_i)$ and L_i be the span of the last elements of the chain from the canonical form. So $F_i \in \mathcal{H}(s_i, t)$ but $F_i \notin \mathcal{H}(s_i, h)$ and $L_i \in \mathcal{H}(s_i, h)$ but $L_i \notin \mathcal{H}(s_i, t)$. Let e_i be the only edge incident to s_i and without loss of generality assume the head map sends e_i to v_i. Then $F_i + L_i \in \mathcal{H}(e_i, h)$ since $\mathcal{H}$ is a full sheaf and so any nontrivial element of the first twist homology group is nonzero on the portion of the edge space that maps to F_i and L_i. The dimension of $F + L$ is at least 4 and so the d_i is also at least 4.

Chapter 12

Subquotients, Submodularity and Supermodularity

In this section, we give a few results that aren't necessary for the main theorems of this dissertation but still are of interest. Theorem 8.9 characterizes a minimal gapped sheaf in any collection of sheaves that is closed under taking subquotients. We now wish to show that there are many such collections; the most basic such collection is the collection of all subquotients of a given sheaf. We also prove a theorem that can simplify computations on such sheaves. Using that theorem, we show there are no gapped subquotient sheaves of the constant sheaf on any graph with only one vertex. After doing this we show that although excess is a supermodular function, as proven in Section 1.6 of [17], the gap is neither supermodular nor submodular.

12.1 Subquotients of a Sheaf

We wish to show that there are many interesting sheaf collections that are subquotient closed.

Definition 12.1. Let $\mathcal{F}$ be a sheaf on a graph, G. We say that a sheaf, $\mathcal{F}'$, on G is a *subquotient of* $\mathcal{F}$ if there are sheaves $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}$ such that $\mathcal{F} \simeq \mathcal{F}_2/\mathcal{F}_1$. We denote the set of subquotients of $\mathcal{F}$ by $\mathrm{SubQuo}(\mathcal{F})$.

Lemma 12.2. *The set of subquotients,* $\mathrm{SubQuo}(\mathcal{F})$*, of a sheaf* $\mathcal{F}$ *on a graph* G *is closed under taking subsheaves and quotient sheaves.*

Proof. It suffices to take $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}$ and consider subsheaves and quotients of $\mathcal{F}_2/\mathcal{F}_1$.

Let $\mathcal{F}'$ be a subsheaf of $\mathcal{F}_2/\mathcal{F}_1$. Then it follows that for each $P \in V_G \amalg E_G$ we have $\mathcal{F}'(P) \subset \mathcal{F}_2(P)/\mathcal{F}_1(P)$, i.e., $\mathcal{F}'(P)$ consist of a collection of $\mathcal{F}_1(P)$ equivalence classes in $\mathcal{F}_2(P)$; let $\mathcal{F}_2'(P)$ be all the vectors in these classes. We easily check that $\mathcal{F}_2'$, with the natural restriction maps obtained from $\mathcal{F}_2$, forms a subsheaf of $\mathcal{F}_2$ that contains $\mathcal{F}_1$. Hence $\mathcal{F}'$, which is clearly isomorphic to $\mathcal{F}_2'/\mathcal{F}_1$, is a subquotient of $\mathcal{F}$.

The case of a quotient of $\mathcal{F}_2/\mathcal{F}_1$ is handled analogously. □

12.2 Simplifications for Stable Sheaves

Theorem 7.17 is statement about subconstant sheaves. The reason that we consider a larger collection of sheaves, namely subquotients of constant sheaves, is that there are certain theorems, such as Theorem 8.9, regarding what one can say about a minimal gapped sheaf (if it exists). The problem with working with subquotients is the values are quotient spaces, and this can complicate calculations. The following theorem will simplify our calculation.

Theorem 12.3. *Let $\mathcal{F}$ be a sheaf on a graph, G, and $\mathcal{F}' \in \mathrm{SubQuot}(\mathcal{F})$. Then there exist $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}$ such that (1) $\mathcal{F}' \simeq \mathcal{F}_2/\mathcal{F}_1$, (2) $\mathcal{F}_1(e) = 0$ for all $e \in E_G$, and (3) $\mathcal{F}_2(v) = \mathcal{F}(v)$ for all $e \in E_G$.*

Proof. By definition, $\mathcal{F} = \mathcal{F}_2/\mathcal{F}_1$ for some $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}$.

For each $e \in E_G$, choose a $W_e \subset \mathcal{F}_2(e)$ that is a subspace of “representatives” of $\mathcal{F}_2(e)/\mathcal{F}_1(e)$, i.e., such that $W_e \cap \mathcal{F}_1(e) = 0$ and $W_e + \mathcal{F}_1(e) = \mathcal{F}_2(e)$. Then replace the value $\mathcal{F}_1(e)$ with 0 and the value $\mathcal{F}_2(e)$ with W_e, and replace the restriction maps from e to he or te with zero maps at $\mathcal{F}_1(e)$, and with the old restriction maps at $\mathcal{F}_2(e)$ restricted to W_e. We easily see that this replacement does not change the isomorphism class of $\mathcal{F}_2/\mathcal{F}_1$.

Similarly, for each $v \in V_G$, we choose a subspace $W_v \subset \mathcal{F}(v)$ such that $W_v + \mathcal{F}_2(v) = \mathcal{F}(v)$; then we can replace $\mathcal{F}_1(v)$ and $\mathcal{F}_2(v)$, respectively, with $\mathcal{F}_1(v) + W_v$ and $\mathcal{F}(v)$. □

12.3 Single Vertex Graphs

Now we use several of our results on gapped sheaves to establish that a gapped sheaf on one vertex cannot be a subquotient of a constant sheaf.

Lemma 12.4. *Let G be a graph with exactly one vertex. Let $\mathcal{C}$ be the sheaf collection of subquotient sheaves of a constant sheaf on G. Then no element of $\mathcal{C}$ has a gap.*

Proof. Assume, to the contrary, that $\mathcal{C}$ contains a sheaf with a gap, and let $\mathcal{F}$ be a T-minimal gapped sheaf. Then Theorem 8.9 implies $\mathcal{F}$ satisfies

$$\chi(\mathcal{F}) = \mathrm{m.e.}(\mathcal{F}) = \mathrm{d.m.e.}(\mathcal{F}) = 0, \quad h_0^{\mathrm{tw}}(\mathcal{F}) = h_1^{\mathrm{tw}}(\mathcal{F}) = 1.$$

Let $V_G = \{v\}$, and let $E_G = \{e_1, \ldots, e_r\}$. Theorem 12.3 implies that if $\mathcal{F}$ is a subquotient of the constant sheaf $\underline{L}$, for a $\mathbb{F}$-vector space, L, then we may assume, by passing to an isomorphic subquotient of L, that

$$\mathcal{F}(e_1) = A_1, \ \ldots, \ \mathcal{F}(e_r) = A_r, \quad \mathcal{F}(v) = B/C$$

for some subspaces, $A_1, \ldots, A_r, B, C$, of L. Now we gather a few facts about these subspaces.

First, $A_i \subset B$ for all i, since $\mathcal{F}$ is a sheaf, and, of course, $C \subset B$. Second, $A_1 + \cdots + A_r + C = B$, or else $\mathcal{F}$ would have a quotient sheaf whose value is zero at all the edges and nonzero at v (namely $B/(A_1 + \cdots + A_r + C)$), contradicting the fact that d.m.e.$(\mathcal{F}) = 0$.

Third, we claim that $A_1, \ldots, A_r, C$ are linearly independent, i.e., that if $a_i \in A_i$ for $i = 1, \ldots, r$ and $c \in C$, and

$$a_1 + \cdots + a_r + c = 0,$$

then $a_1 = \cdots = a_r = c = 0$. If not, then consider the subsheaf, $\mathcal{F}'$, of $\mathcal{F}$ whose values are $\mathcal{F}(e_i)$ being the (one-dimensional) span of a_i for all i, and $\mathcal{F}(v) = A'/(A' \cap C)$, where A' is the span of of the a_i. Since $a_1 + \cdots + a_r \in C$, we have that $\mathcal{F}(v)$ is at most $r-1$ dimensional. Hence

$$-\chi\bigl(\mathcal{F}(v)\bigr) = \dim\bigl(\mathcal{F}(E)\bigr) - \dim\bigl(\mathcal{F}(V)\bigr) \geq r - (r-1) = 1,$$

which contradicts the fact that m.e.$(\mathcal{F}) = 0$. Hence $A_1, \ldots, A_r, C$ are linearly independent.

But we will easily see that linear independence of $A_1, \ldots, A_r, C$ implies that $h_1^{\mathrm{tw}}(\mathcal{F}) = 0$. Indeed, fixing a full twist, ψ, on G, we have that elements of $H_1^{\mathrm{tw}}(\mathcal{F})$ are given by tuples $(a_1, \ldots, a_r)$ such that $a_i \in A_i(\psi)$ and

$$(1 - \psi_1)a_1 + \ldots + (1 - \psi_r)a_r \in C(\psi).$$

But the linear independence of $A_1, \ldots, A_r, C$ over $\mathbb{F}$ implies the linear independence of $A_1(\psi), \ldots, A_r(\psi), C(\psi)$ over $\mathbb{F}(\psi)$ (this is an easy exercise), and this implies that $a_1 = \cdots = a_r = 0$. Hence the only element of $H_1^{\mathrm{tw}}(\mathcal{F})$ is the zero element; i.e. $h_1^{\mathrm{tw}}(\mathcal{F}) = 0$. But this contradicts the fact that, by Theorem 8.9, $h_1^{\mathrm{tw}}(\mathcal{F}) = 1$. Hence our assumption that $\mathcal{C}$ contains a sheaf with a gap is impossible. $\square$

12.4 Submodularity and Supermodularity

We say that a function g over sheaves on graphs is *supermodular* if given any $\mathcal{A}$ and $\mathcal{B}$ that are subsheaves of any sheaf $\mathcal{F}$ over a graph G we have

that
$$g(\mathcal{A} \cap \mathcal{B}) + g(\mathcal{A} \cup \mathcal{B}) \geq g(\mathcal{A}) + g(\mathcal{B}).$$
Alternatively we say g is *submodular* if we have
$$g(\mathcal{A} \cap \mathcal{B}) + g(\mathcal{A} \cup \mathcal{B}) \leq g(\mathcal{A}) + g(\mathcal{B})$$
for any subsheaves $\mathcal{A}$ and $\mathcal{B}$. In this section we give two quick counterexamples to show the gap is neither submodular nor supermodular.

To disprove submodularity let $\mathcal{F}$ be the unhappy bundle on the figure-eight graph discussed in the previous section and let $\mathcal{A}$ be $\mathcal{F}$ restricted to v and e_1 and let $\mathcal{B}$ $\mathcal{F}$ restricted to v and e_2. Since $\mathcal{A} \cap \mathcal{B}$ has no edges, it isn't gapped. The sheaves $\mathcal{A}$ and $\mathcal{B}$ are only supported on single self loops, and so have no gap as well. But $\mathrm{gap}(\mathcal{A} \cup \mathcal{B}) = 1$ since the union is the unhappy bundle.

To disprove supermodularity, let G be a single vertex v with four self loops e_i for $i = 1, \ldots, 4$. Let $\mathcal{F}(v), \mathcal{F}(e_1)$ and $\mathcal{F}(e_2)$ be the same as in the unhappy bundle and also let $\mathcal{F}(v)$, $\mathcal{F}(e_3)$ and $\mathcal{F}(e_4)$ be another copy of the unhappy bundle. Then define $\mathcal{A}$ and as the copy of the unhappy bundle over v, e_1 and e_2 and $\mathcal{B}$ as a copy of the unhappy bundle over v the other two edges. Then $\mathrm{m.e.}(\mathcal{A} \cup \mathcal{B}) = 4$ since $\mathcal{A}$ is a stable sheaf and $\mathcal{A} \cup \mathcal{B}$ has the same vertex space as $\mathcal{A}$ and edge dimension that is larger by four. The sheaves $\mathcal{A}$ and $\mathcal{B}$ each have positive gap, but by computing the first twisted Betti number of the union we find that it isn't a gapped sheaf. Once again, the intersection isn't gapped since it is not supported on any edges.

Chapter 13

Conclusion

This dissertation has shown several links between the abelian girth of a graph and the standard girth as well as between gapped sheaves and the abelian girth. In Section 3 we proved Theorem 2.6, which can be seen as a version of the Moore bound for abelian girth. In Section 4 we proved Theorem 2.5, which implied that any multiplicative improvement on Theorem 2.6 would also improve the Moore bound. Then in Section 5, we show that the bipartite LPS graphs do not have abelian girth so large as to preclude any possibility of improving Theorem 2.6.We also showed in Section 11 that the abelian girth of a graph is the minimal dimension of any gapped sheaf on that graph. These results might allow for techniques that improve the Moore bound or answer other questions about the girth.

Conjectures related to the HNC (see [9]) may potentially be verified by showing that certain ρ-kernels have vanishing maximum excess, as is done in the proof of the HNC [17]. These ρ-kernels are subconstant sheaves on two vertices and in Section 11 we demonstrated the existence of gapped subconstant sheaves on two vertices. Thus any attempt to prove the conjectures related to the HNC using the first twisted Betti number of the ρ-kernels should take into account that the ρ-kernels may be gapped. Though we can still compute the maximum excess of a gapped sheaf by finding the first twisted Betti number of a cover of the sheaf with sufficient degree, this may be a computationally intensive task. It remains to be shown whether computing the maximum excess of a sheaf can be done in time polynomial in the degree of the sheaf.

Though it may be time-intensive to compute the maximum excess of a gapped sheaf, thoughout this dissertation we have proved several results that can be used to verify that a sheaf is not gapped, in which case computing the maximum excess can be done quickly. Theorem 7.18 implies that any sheaf with total dimension less than the abelian girth of the underlying graph is not gapped. Lemma 8.21 and Theorem 8.9 show that any gapped sheaf must have a gapped, stable, faithful sheaf as a subquotient. We have shown that any gapped sheaf must have dimension at least 2 on any edge and dimension at least 4 on self loops. Lemma 9.6 shows that any gapped sheaf contains a gapped subsheaf with every edge being internal.

Bibliography

[1] *Théorie des topos et cohomologie étale des schémas. Tome 1: Théorie des topos.* Springer-Verlag, Berlin, 1972. Séminaire de Géométrie Algébrique du Bois-Marie 1963–1964 (SGA 4), Dirigé par M. Artin, A. Grothendieck, et J. L. Verdier. Avec la collaboration de N. Bourbaki, P. Deligne et B. Saint-Donat, Lecture Notes in Mathematics, Vol. 269.

[2] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.

[3] Alon Amit, Shlomo Hoory, and Nathan Linial. A continuous analogue of the girth problem. *J. Combin. Theory Ser. B*, 84(2):340–363, 2002.

[4] M. F. Atiyah. Elliptic operators, discrete groups and von Neumann algebras. In *Colloque "Analyse et Topologie" en l'Honneur de Henri Cartan (Orsay, 1974)*, pages 43–72. Astérisque, No. 32–33. Soc. Math. France, Paris, 1976.

[5] A. L. Biggs and N. G. Boshier. Note on the girth of ramanujan graphs. *J. Combin. Theory Ser. B*, 49(2):190–194, 1990.

[6] Warren Dicks. Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture. *Invent. Math.*, 117(3):373–389, 1994.

[7] Warren Dicks and M. J. Dunwoody. *Groups acting on graphs.* Cambridge University Press, Cambridge, 1989.

[8] Warren Dicks and Sergei V Ivanov. On the intersection of free subgroups in free products of groups. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 144, pages 511–534. Cambridge Univ Press, 2008.

[9] Warren Dicks, SV Ivanov, et al. On the intersection of free subgroups in free products of groups with no 2-torsion. *Illinois Journal of Mathematics*, 54(1):223–248, 2010.

[10] David Ellis and Nathan Linial. On regular hypergraphs of high girth. *The Electronic Journal of Combinatorics*, 21(1):P1–54, 2014.

[11] P Erdős, A Rényi, and VT Sós. On a problem of graph theory. *Studia Sci. Math. Hungar*, 1:215–235, 1966.

[12] Paul Erdös and Horst Sachs. Reguläre graphen gegebener taillenweite mit minimaler knotenzahl. *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, 12:251–257, 1963.

[13] Brent Everitt. Graphs, free groups and the Hanna Neumann conjecture. *J. Group Theory*, 11(6):885–899, 2008.

[14] Joel Friedman. Cohomology of grothendieck topologies and lower bounds in boolean complexity. 2005. `http://www.math.ubc.ca/~jf`, also `http://arxiv.org/abs/cs/0512008`, to appear.

[15] Joel Friedman. Sheaves on graphs and a proof of the hanna neumann conjecture. April 2011. available at `http://arxiv.org/pdf/1105.0129v1` and at `http://www.math.ubc.ca/~jf`.

[16] Joel Friedman. Sheaves on graphs and their homological invariants. April 2011. available at `http://arxiv.org/pdf/1104.2665v1` and at `http://www.math.ubc.ca/~jf`.

[17] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture: with an appendix by Warren Dicks. *Mem. Amer. Math. Soc.*, 233(1100):xii+106, 2015. With an appendix by Warren Dicks.

[18] Joel Friedman. Sheaves on graphs, their homological invariants, and a proof of the hanna neumann conjecture: with an appendix by warren dicks. 2015.

[19] Joel Friedman, Alice Izsak, and Lior Silberman. Gapped sheaves on graphs. To appear.

[20] S. M. Gersten. Intersections of finitely generated subgroups of free groups and resolutions of graphs. *Invent. Math.*, 71(3):567–591, 1983.

[21] Wilfried Imrich. On finitely generated subgroups of free groups. *Arch. Math. (Basel)*, 28(1):21–24, 1977.

[22] Wilfried Imrich. Subgroup theorems and graphs. In *Combinatorial mathematics, V (Proc. Fifth Austral. Conf., Roy. Melbourne Inst. Tech., Melbourne, 1976)*, pages 1–27. Lecture Notes in Math., Vol. 622. Springer, Berlin, 1977.

[23] S. V. Ivanov. On the intersection of finitely generated subgroups in free products of groups. *Internat. J. Algebra Comput.*, 9(5):521–528, 1999.

[24] S. V. Ivanov. Intersecting free subgroups in free products of groups. *Internat. J. Algebra Comput.*, 11(3):281–290, 2001.

[25] Alastair D King. Moduli of representations of finite dimensional algebras. *The Quarterly Journal of Mathematics*, 45(4):515–530, 1994.

[26] Nati Linial and Doron Puder. Word maps and spectra of random graph lifts. *Random Structures Algorithms*, 37(1):100–135, 2010.

[27] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[28] David JC MacKay and Radford M Neal. Near shannon limit performance of low density parity check codes. *Electronics letters*, 32(18):1645–1646, 1996.

[29] J. Meakin and P. Weil. Subgroups of free groups: a contribution to the Hanna Neumann conjecture. In *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, volume 94, pages 33–43, 2002.

[30] Brigitte Servatius. A short proof of a theorem of Burns. *Math. Z.*, 184(1):133–137, 1983.

[31] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[32] Gábor Tardos. Towards the Hanna Neumann conjecture using Dicks' method. *Invent. Math.*, 123(1):95–104, 1996.

[33] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

[34] Alexander Zakharov. On the rank of the intersection of free subgroups in virtually free groups. *Journal of Algebra*, 418:29–43, 2014.