Sheaves on Graphs, Their Homological Invariants, and a Proof of the Hanna Neumann Conjecture

Joel Friedman

Author address:

Department of Computer Science, University of British Columbia, Vancouver, BC V6T 1Z4, CANADA, and Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2, CANADA.

E-mail address: jf@cs.ubc.ca or jf@math.ubc.ca

Contents

Abstract

In this paper we establish some foundations regarding sheaves of vector spaces on graphs and their invariants, such as homology groups and their limits. We then use these ideas to prove the Hanna Neumann Conjecture of the 1950's; in fact, we prove a strengthened form of the conjecture.

We introduce a notion of a sheaf of vector spaces on a graph, and develop the foundations of homology theories for such sheaves. One sheaf invariant, its "maximum excess," has a number of remarkable properties. It has a simple definition, with no reference to homology theory, that resembles graph expansion. Yet it is a "limit" of Betti numbers, and hence has a short/long exact sequence theory and resembles the L^2 Betti numbers of Atiyah. Also, the maximum excess is defined via a supermodular function, which gives the maximum excess much stronger properties than one has of a typical Betti number.

Our sheaf theory can be viewed as a vast generalization of algebraic graph theory: each sheaf has invariants associated to it—such as Betti numbers and Laplacian matrices—that generalize those in classical graph theory.

We shall use "Galois graph theory" to reduce the Strengthened Hanna Neumann Conjecture to showing that certain sheaves, that we call ρ-kernels, have zero maximum excess. We use the symmetry in Galois theory to argue that if the Strengthened Hanna Neumann Conjecture is false, then the maximum excess of "most of" these ρ-kernels must be large. We then give an inductive argument to show that this is impossible.

Received by the editor May 30, 2011.

2000 *Mathematics Subject Classification.* Primary 05C10, 55N30, 18F20; Secondary 05C50, 18F10, 14F20.

Key words and phrases. Graphs, sheaves, Hanna Neumann Conjecture, homology, Galois theory.

Research supported in part by an NSERC grant. Research done in part at the Centre Bernoulli, funded by the Swiss National Science Foundation.

Introduction

This memoir has two main goals. First, we develop some foundations on what we call "sheaves on graphs" and their invariants. Second, using these foundations, we resolve the Hanna Neumann Conjecture of the 1950's.

Although our foundations of sheaves on graphs seem likely to impact a number of areas of graph theory, the theme that is common to most of this memoir is the Hanna Neumann Conjecture (or HNC). Both this conjecture and a strengthening of it, known as the Strengthened Hanna Neumann Conjecture (or SHNC) have been extensively studied (see [**Bur71, Imr77b, Imr77a, Ser83, Ger83, Sta83, Neu90, Tar92, Dic94, Tar96, Iva99, Arz00, DF01, Iva01, Kha02, MW02, JKM03, Neu07, Eve08, Min10**]). These conjectures are usually stated as an inequality involving free groups, although both conjectures have well-known reformulations in term of finite graphs. In this memoir we prove both conjectures, using the finite graph reformulations, reducing both to the vanishing of a homology group of certain sheaves on graphs.

This work was originally written and posted to `arxiv.com` as two separate articles. The first aritcle, [**Fri11b**], contains the foundational material on sheaves of graphs, and comprises Chapter 1 of this manuscript. The second article, [**Fri11a**], resolves the SHNC (and HNC), and Chapter 2 of this manuscript consists of this material. This manuscript is easier to read than both articles separately, in that redundant definitions have been discarded, and references in Chapter 2 to material in Chapter 1 are now more specific. Yet, as we now explain, Chapters 1 and 2 are largely independently of one another, and Chapter 1, the foundations of sheaves on graphs, is of interest beyond the HNC and SHNC. To explain this interest, let us recall a bit about sheaf theory and its connection to discrete mathematics.

Among many (co)homology theories of topological spaces, the sheaf approach has many advantages. For one, it works with non-Hausdorff spaces, as done first by Serre in algebraic geometry with the Zariski topology (see [**Har77**]). Grothendieck's sheaf theory of [**sga72a, sga72b, sga73, sga77**] defined a notion of a sheaf on very general spaces now called "Grothendieck topologies." While Grothendieck's work has had remarkable success to cohomology theories in algebraic geometry, we believe that graph theory and combinatorics may greatly benefit by studying very special Grothendieck topologies formulated from finite, discrete structures. In particular, we will resolve the SHNC using a simple, finite Grothendieck topology associated to any finite graph.

Another aspect of sheaf cohomology is that it vastly generalizes the cohomology of a space. Each sheaf has injective resolutions that give cohomology groups. When the sheaf is take to be the "structure sheaf" of the space, we recover the cohomology groups of the space. However, there are many sheaves apart from the structure sheaf, and the resulting cohomology groups can represent a variety of aspects of the space. In particular, each open subset of a space, X, has an associated sheaf on X that reflects many properties of X; we will use such sheaves in our proof of the SHNC.

One fundamental aspect of any (co)homology theory is that it expresses relations between related (co)homology groups in terms of exact sequences. Furthermore, any exact sequence yields a triangle inequality between the dimensions or ranks of any three consecutive elements. In [**Fri05, Fri06, Fri07**] we began an investigation into applying such inequalities to complexity theory, in particular to construct formal complexity measures to obtain lower bounds for formula size. Similarly, in this manuscript, we prove the SHNC from such an inequality.

In Chapter 1 we define a sheaf on a graph with no reference to sheaf theory, rather as a collection of vector spaces indexed on the vertices and edges of the graphs along with certain "restriction" maps. We add that one can view such a sheaf as a simple genelization of an incidence matrix of a graph; it follows that sheaves on graphs can be viewed as a vast generalization of classical algebraic graph theory (of adjacency matrices, Laplacians, etc.). However, our sheaves on graphs can also be viewed as the very special case of sheaves of finite dimensional vector spaces on a simple Grothendieck topology that we associate to a finite graph. In the case where the graph has no self loops, the Grothendieck topology is equivalent to a simple topological space.

Chapter 1 begins with a simple definition of sheaves on graphs and some examples. However, quickly we begin to study "limits" of Betti numbers of these sheaves. The most remarkable invariant that we study in Chapter 1 is the *maximum excess* of a sheaf. We give a number of strong results regarding the maximum excess, and the related "twisted" homology. These are related to the L^2 Betti numbers first studied by Atiyah (see [**Ati76, Lüc02**]); however, the results we obtain in the case of finite graphs, especially regarding the maximum excess, seem especially strong.

To summarize the above few paragraphs, here are some reasons that Chapter 1 is of interest independent of the HNC:

(1) sheaf theory on graphs generalizes algebraic graph theory and, therefore, may strengthen its applications;
(2) our results on maximum excess give tools to study certain graph invariants such as the "reduced cyclicity" and number of "acyclic components;"
(3) our results on the maximum excess of sheaves may indicate what one can expect of limits of Betti numbers on more general structures;
(4) any results on sheaves on graphs may give new results and examples of what to expect on more general finite Grothendieck topologies, such as those of possible interest to complexity theory;
(5) any results on Betti numbers of sheaves may yield new inequalities on other integers that can be viewed as akin to Betti numbers on some discrete Grothendieck topology.

Of course, despite the above reasons for interest in sheaves on graphs, the reader will see that Chapter 1 is largely developed with an eye toward the reduced cyclicity and the HNC.

Let us summarize aspects of Chapter 2, our proof of the SHNC, in general terms. This will serve to highlight our approach to this problem via sheaves on graph, which is very different than previous approaches. We use a graph theoretic formulation of the SHNC that involves the reduced cyclicity of three graphs. However, using what we call "Galois graph theory" (of [**Fri93, ST96**], but also [**Gro77**]), the SHNC amounts to showing that the reduced cyclicity of one graph is less than that of another graph, and both of these graphs admit a natural map to the same Cayley graph.

We wish to emphasize that, to the best of our knowledge, our manuscript represents the first application of Galois graph theory to other parts of graph theory. That is, Galois graph theory occurs for its own interest (in [**Fri93**]) and for its connection to number theory (in [**ST96**]). However, in this manuscript we make essential use of Galois graph theory to two independent questions not obviously related to Galois graph theory. First, in Chapter 1 we use Galois graph theory to show that maximum excess scales under pulling back by a covering map; first we prove this for Galois morphisms, making essential use of the symmetry in Galois theory, and then we deduce the general case by the Normal Extension Theorem of Galois graph theory. Second, Galois graph theory is the basis of our construction of ρ-kernels, upon which our approach to the HNC and SHNC is based, and the symmetry of these ρ-kernels is used constantly in Chapter 2.

Let us return to the SHNC, and recall that exact sequences give triangle inequalities on the dimensions of consecutive terms. The reduced cyclicity is a type of limiting first Betti number. Hence, one graph has smaller reduced cyclicity than a second graph provided that there is a surjection from the first graph to the second, such that the kernel of this surjection has vanishing limiting first Betti number. Unfortunately there is no such graph surjection in the graphs that arise from the SHNC. However, both graphs admit a natural map to the same Cayley graph, and hence can be viewed as sheaves on this Cayley graph (much as open subsets of a topological space have associated sheaves). Remarkably, there is a surjection from the first graph to the second when viewed as sheaves. The kernel of such a sujection (generally a sheaf) will be called a ρ-kernel, and the SHNC turns out to be implied by the vanishing limiting first Betti number, or maximum excess, of an appropriate collection of ρ-kernels.

We emphasize that the ρ-kernels that we build seem almost forced upon us, once we look for the surjections described above. However, it does not seem to be an easy question, essentially of linear algebra, to determine whether or not these ρ-kernels have vanishing maximum excess. In fact, if we define a ρ-kernel as the kernel of any surjection of the two graphs of interest, then there are ρ-kernels whose maximum excess does not vanish.

To complete the proof of the SHNC, we shall show that the maximum excess of a "generic" ρ-kernel vanishes. This main idea is that there is a symmetry property of the "excess maximizer," which implies that maximum excess of a generic ρ-kernel must be a multiple of the order of an associated Galois group (the group associated to the Cayley graph mentioned above). From this point one knows that if the generic maximum excess doesn't vanish, it would be large; one can then use two different inductive arguments to show that this is impossible.

For the reader interested only in a proof of the HNC, we mention that can read its proof in Chapter 2 while skipping most of the material in Chapter 1. Indeed, Chapter 2 is based on the "stand alone" paper, [**Fri11a**], written without explicit reference to homology theory, using only sheaves and maximum excess. So to read Chapter 2, one needs the definitions of sheaves and maximum excess, of Section 1.2, the Galois graph theory of Section 1.3, and the submodularity of the excess in Section 1.6. Aside from these results, the proof in [**Fri11a**] needed the fact that the maximum excess is a "first quasi-Betti number," which relies on the main (and most difficult) theorem of Chapter 1. However, we have recently found a variant of the proof in [**Fri11a**] which does not require this fact. Hence one can read a complete proof of the HNC and SNHC in this manuscript, without most of Chapter 1 and any reference to homology. However, as explained in Chapter 2, homology still gives valuable insight into the proof.

We mention that as of writing [**Fri11b, Fri11a**], Mineyev has informed us of his independent proof of the HNC and SHNC, first using Hilbert modules ([**Min11b**], based on [**Min10**]), and then using only combinatorial group theory ([**Min11a**]). His approaches seem very different from ours.

We wish to thank Laurant Bartholdi, for conversations and introducing us to the SHNC, and Avner Friedman, for comments on a draft of this manuscript. We thank Luc Illusie, for an inspiring discussion regarding our ideas involving sheaf theory, homology, and the SHNC; this discussion was a turning point in our research. We wish to thank for following people for conversations: Goulnara Arjantseva, Warren Dicks, Bernt Everitt, Sadok Kallel, Richard Kent, Igor Mineyev, Pierre Pansu, and Daniel Wise. Finally, we thank Alain Valette and the Centre Bernoulli at the EPFL for hosting us during a programme on limits of graphs, where we met Bartholdi and Pansu and began this work.

CHAPTER 1

Foundations of Sheaves on Graphs and Their Homological Invariants

1.1. Introduction

The main goal of this chapter is to introduce a notion of a sheaf on a graph and to establish some foundational results regarding the homology groups of such sheaves and related invariants. After developing some general points we shall focus on a remarkable invariant of a sheaf that we call the *maximum excess*.

The *maximum excess* of a sheaf arises naturally as a "limit" of Betti numbers, akin to L^2 Betti number defined by Atiyah. Although such limits have been studied in many contexts, we are able to show some compellingly strong results about these limits in the case of sheaves on graphs. First, the maximum excess can be defined, with no reference to homology theory, in a manner that makes it resemble quantities seen in matching theory or expander graphs. Second, this definition amounts to the maximum of an "excess" function that is supermodular; this gives additional structure to the maximum excess that is not apparent from homology theory. Third, for any given sheaf, the limit is attained from "twisted Betti numbers" by passing to a finite cover (as opposed to an infinite limit of covers).

Our motivation for studying the maximum excess and certain Betti numbers arose from studying an important graph invariant that we call the *reduced cyclicity* of a graph. This invariant arises in one formulation of the much studied Hanna Neumann Conjecture of the 1950's (see [**Bur71, Imr77b, Imr77a, Ser83, Ger83, Sta83, Neu90, Tar92, Dic94, Tar96, Iva99, Arz00, DF01, Iva01, Kha02, MW02, JKM03, Neu07, Eve08, Min10**]); in Chapter 2 we shall use the results of this chapter to prove this conjecture. Moreover, our methods will prove what is known as the Strengthened Hanna Neumann Conjecture (or SHNC) of [**Neu90**].

Our sheaf theory on graphs is based on the sheaf theory of Grothendieck (see [**sga72a, sga72b, sga73, sga77**]), built upon what are now known as Grothendieck topologies. In the special case when the graph has no self-loops, the sheaf theory we describe is equivalent to the sheaf theory on certain topological spaces (see [**Har77**]). The basic definition of sheaves on graphs and their homology groups are special cases of theory developed in [**Fri05, Fri06, Fri07**] and are probably special cases of situations arising in the fields of toric varieties and quivers. However, in this chapter we study a special case of this general notion of sheaf theory, proving especially strong theorems particular to sheaves on graphs and obtaining new theorems in graph theory. In this process we also introduce new invariants in sheaf theory—such as "maximum excess" and "twisted homology"—and establish theorems about these invariants that may become useful to sheaf theories in other settings.

In this chapter we explore primarily those aspects of sheaf theory directly related to our future study of the SHNC, namely general properties of the maximum excess. However, we believe sheaf theory is a concept fundamental to graph theory, and that there will probably emerge other applications of these ideas. One reason for this belief is that many areas in graph theory, such as expanding graphs, work with the adjacency matrix of a graph. Any sheaf on a graph, G, has an adjacency matrix (and incidence matrix, Laplacian, etc.) with many of the properties that graph adjacency matrices have. A graph has a particularly simple sheaf that we call its "structure sheaf." The adjacency matrix of the structure sheaf turns out to be the adjacency matrix of G. In this way the adjacency matrix of a graph, and all of traditional algebraic graph theory, can be generalized to sheaf theory; the sheaf theory, given its more general nature and expressiveness, may shed new light on traditional algebraic graph theory and its applications.

New graph theoretic inequalities arise in our sheaf theory out of "long exact sequences," analogous to long exact sequences that appear in virtually any homology theory. Indeed, relations between different homology groups are often expressed in exact sequences, and in any exact sequence of vector spaces, the dimensions of three consecutive elements satisfy a triangle inequality. It is such triangle inequalities that inspire and form the basis of our approach to the SHNC.

One remarkable aspect of our sheaf theory is that it adds "new morphisms" between graphs. In other words, consider two graphs, G_1 and G_2 that each admit a morphism to another graph, G. It is possible to associate with each G_i a sheaf, $\mathcal{S}(G_i)$, over G, that contains all the information present in G_i. Any G-morphism from G_1 to G_2 gives rise to a morphism of sheaves, from $\mathcal{S}(G_1)$ to $\mathcal{S}(G_2)$; however, there are sheaf morphisms from $\mathcal{S}(G_1)$ to $\mathcal{S}(G_2)$ that do not arise from any graph morphism. For example, there may be a surjection from $\mathcal{S}(G_1)$ to $\mathcal{S}(G_2)$ when there is no graph theoretic surjection $G_1 \to G_2$. Some such "new surjections" are crucial to our proof of the SHNC; the kernel of such "new surjections" give a type of sheaf that we call a *ρ-kernel*, which is the basis of our approach to the SHNC. Said otherwise, for any graph, G, there is a faithful functor from the category of "graphs over G" to the category of "sheaves over G;" however this functor is not full, and some of the "new morphisms" between graphs over G, viewed as sheaves over G, ultimately yield new concepts in graph theory needed in our proof of the SHNC.

This chapter will focus on four types of invariants of sheaves: (1) homology groups and resulting Betti numbers, (2) twisted homology groups and resulting twisted Betti numbers, (3) the maximum excess, and (4) limiting twisted Betti numbers. Let us briefly motivate our interest in these invariants and describe the main theorems in this chapter. This discussion will be made more precise, with more background, in Section 1.2.

Our first type of invariant, homology groups of sheaves and resulting Betti numbers, will not involve any difficult theorems. The main novelty of this type of invariant is in its definition; it is chosen in a way that it has appropriate properties for our needs and can express some traditional invariants of a graph; these invariants include its Euler characteristic and the traditional zeroth and first Betti numbers. In sheaf theory, usually sheaf cohomology based on the global section functor is a central object of study; however, these cohomology groups do not yield the invariants of interest to us in this chapter. Instead, our homology groups are based on global cosections; i.e., our homology groups are essentially Ext groups in the first variable, where the second variable is fixed to be the structure sheaf.

The SHNC conjecture can be reformulated in graph theoretic terms, involving a more troubling graph invariant, $\rho(G)$, of a graph, G, which we call the *reduced cyclicity* of G. The reason this graph invariant is troubling is that its usual definition seems to require that we know how many connected components of G are acyclic, i.e., are isolated vertices or trees. Prior to this paper, all non-trivial techniques we know to bound $\rho(G)$ either presuppose something about the number of acyclic components of G, or else they overlook such components; as such, previous results on the reduced cyclicity usually either require special assumptions or give results that are not sharp. Our second set of invariants, the twisted homology groups and their dimensions, i.e., the twisted Betti numbers, give $\rho(G_1)$ as the first twisted Betti number of a certain sheaf on G, for any graph, G_1, with a graph morphism to G. As such, the long exact sequences arising in twisted homology give the first sharp relations between values of ρ; however, these relations usually involve sheaves and not just graphs alone.

Let us sketch the idea of why reduced cyclicity is a special case of a twisted Betti number. In this chapter we observe that $\rho(G)$ is the limit of $h_1(K)/[K\colon G]$ over "generic Abelian coverings maps," $K \to G$, where the degree, $[K\colon G]$, of the covering map tends to infinity. It is well known that for Abelian covering maps $K \to G$, we can recover spectral properties of the adjacency matrix of K by working with that of G and "twisting its entries," i.e., multiplying certain entries by roots of unity that appear in the characters of the underlying Abelian group. So we form "twisted" homology groups by "generically twisting" a sheaf, with twists that are parameters or indeterminates, and compute that the reduced connectivity, $\rho(G)$, equals the first "twisted" Betti number of the structure sheaf of G. This gives a generalization of the definition of ρ from graphs to sheaves, and the resulting twisted Betti numbers satisfy triangle inequalities coming from the long exact sequences in twisted homology.

Another promising fact about twisted Betti numbers is that, via the theory of long exact sequences, one can reduce the SHNC to the vanishing of the first twisted Betti number of a collection of sheaves that we call ρ-kernels.

The problem is that the twisted homology approach often seems to be the "wrong" way to view the reduced cyclicity, mainly for the following reason. The Euler characteristic and reduced cyclicity have a remarkable scaling property under covering maps, $\phi\colon K \to G$, i.e.,

$$\chi(K) = \chi(G)\deg(\phi), \quad \rho(K) = \rho(G)\deg(\phi).$$

Twisted Betti numbers do not always scale in this way; this makes us suspect that the twisted Betti number is not always a good generalization of the reduced cyclicity.

The remedy comes in our third type of invariant, a single invariant of a sheaf that we shall define and call its *maximum excess*. This is an integer that one can define simply and with no reference to homology theory. Its definition resembles combinatorial invariants arising in matching theory or expander graphs. The maximum excess of any sheaf is at most the first twisted Betti number, and the two are equal on many types of sheaves, including all constant sheaves. Hence the two concepts are related but not identical. Furthermore, the SHNC is implied by the (*a priori* weaker) vanishing maximum excess of ρ-kernels, and the maximum excess satisfies stronger properties that yield better bounds than what one would get for the first twisted Betti number. So for the SHNC, we largely abandon the idea of using twisted Betti numbers to generalize ρ from graphs to sheaves, and instead use the maximum excess. The problem is that to proof the SHNC we require inequalities involving the maximum excess akin to those holding of Betti numbers of homology theories via long exact sequences; there is no *a priori* reason that such inequalities should hold.

The main theorem of this chapter, Theorem 1.10, says that for any fixed sheaf on a graph, G, there is an integer, q, with the following property: the maximum excess and first twisted Betti number agree when the sheaf is "pulled back" along a covering map $G' \to G$, provided that the girth of G' is at least q.

The main theorem implies that the maximum excess is a *first quasi-Betti number*, meaning that the maximum excess satisfies certain triangular inequalities that we use to prove the SHNC. However, in Chapter 2 we see that a variation of our proof avoids these triangular inequalities.

Another view of our main theorem is that there exists a "limit" to the ratio of a twisted Betti number of a pullback of a fixed sheaf along a graph covering to the degree of the covering. We shall call this limiting ratio a "limiting twisted Betti number," which is our fourth type of invariant. Our main theorem can be rephrased as saying that the first limiting twisted Betti number is just the maximum excess. It is easy to see that limiting twisted Betti numbers satisfy the triangular inequalities we desire for the maximum excess; hence proving the main theorem proves the desired inequalities for the maximum excess. However, as a limiting Betti number, the maximum excess actually has associated homology groups whose dimensions divided by the covering degree approximate the maximum excess. And it may turn out that the homology groups themselves may contain useful information beyond knowing merely their dimension; however, for our proof of the SHNC, all that we need is the dimensions of these homology groups, i.e., their Betti numbers.

Lior Silbermann has pointed out to us that our notion of limiting twisting Betti numbers is a discrete analogue of "L^2 Betti numbers" introduced by Atiyah on manifolds ([**Ati76**]); the theory involved in the study of L^2 Betti numbers (see[**Lüc02**]), especially the von Neumann dimension of certain "matrices" of this theory, may already imply that our limiting twisting Betti numbers do have a limit and that it is an integer (because the fundamental group of a graph is a free group). So part of our results can be viewed as a very explicit type of L^2 or limiting Betti number calculation (for the very special case sheaves on graphs), that includes stronger information; indeed, we give a simple interpretation of this number (the maximum excess) and a finite procedure for computing it (pulling back to a graph of sufficiently large girth and computing a twisted Betti number).

We note that for the purpose of proving the SHNC, the main results needed from this chapter are the definitions of a sheaf and its maximum excess, and a few properties we prove regarding the maximum excess. If we could prove such properties without using homology theory, we could study the SHNC without homology theory. Nonetheless, we find that twisted homology gives important intuition for the maximum excess; for example, we first proved the SHNC using twisted homology, and only discovered during the writing of [**Fri11a**] that the proof could be written entirely in terms of the maximum excess. As we remark at the end of

Chapter 2, there is a way to prove the SHNC with no reference to homology theory, but this requires some extra combinatorial analysis (namely Appendix A).

The rest of this chapter is organized as follows. In Section 1.2 we give precise definitions and statements of the theorems in this chapter. In Section 1.3 we review part of what might be called "Galois theory of graphs" that we will use in this chapter. In Section 1.4 we give the basic properties of sheaves and homology, pullbacks and their adjoints; then we explain everything in terms of cohomology of Grothendieck topologies (this explanation will help the reader to understand the context of our definitions, but this explanation is not necessary to read the rest of this paper). In Section 1.5 we define the twisted homology and compute the twisted homology of the constant sheaf of a graph; we also interpret twisted homology in terms of Abelian covers. In Section 1.6 we establish the basic properties of the maximum excess, including its bound on the twisted homology. The next two sections establish our main theorem. In Section 1.7 we show how to interpret elements of the first twisted homology group of a graph in terms of the first homology group of the maximum Abelian covering of the graph. In Section 1.8 we prove Theorem 1.10, that says that the first twisted Betti number and the maximum excess agree after an appropriate pullback. In Section 1.9 we make some concluding remarks.

1.2. Basic Definitions and Main Results

In this section we will define sheaves and all the main invariants of sheaves that we use in this paper. We will state the main theorem in this chapter, and state or describe other results in this chapter. In most of this paper we work with directed graphs (digraphs), which makes things notationally simpler; as we remark in Section 1.9, all this sheaf and homology theory works just as well with undirected graphs, although it is slightly more cumbersome if one wants to avoid orienting the edges.

1.2.1. Definition of Sheaves and Homology. We will allow directed graphs to have multiple edges and self-loops; so in this paper a directed graph (or digraph) consists of tuple $G = (V_G, E_G, t_G, h_G)$ where V_G and E_G are sets—the vertex and edge sets—and $t_G \colon E_G \to V_G$ is the "tail" map and $h_G \colon E_G \to V_G$ the "head" map. Throughout this paper, unless otherwise indicated, a digraph is assumed to be finite, i.e., the vertex and edge sets are finite.

Recall that a morphism of digraphs, $\mu \colon K \to G$, is a pair $\mu = (\mu_V, \mu_E)$ of maps $\mu_V \colon V_K \to V_G$ and $\mu_E \colon E_K \to E_G$ such that $t_G\mu_E = \mu_V t_K$ and $h_G\mu_E = \mu_V h_K$. We can usually drop the subscripts from μ_V and μ_E, although for clarity we shall sometimes include them.

Recall that fibre products exist for directed graphs (see, for example, [**Fri93**], or [**Sta83**], where fibre products are called "pullbacks") and the fibre product, $K = G_1 \times_G G_2$, of morphisms $\mu_1 \colon G_1 \to G$ and $\mu_2 \colon G_2 \to G$ has

$$V_K = \{(v_1, v_2) \mid v_i \in V_{G_i},\ \mu_1 v_1 = \mu_2 v_2\},$$

$$E_K = \{(e_1, e_2) \mid e_i \in E_{G_i},\ \mu_1 e_1 = \mu_2 e_2\},$$

$$t_K = (t_{G_1}, t_{G_2}), \quad \text{and} \quad h_K = (h_{G_1}, h_{G_2}).$$

For $i = 1, 2$, respectively, there are natural digraph morphisms, $\pi_i \colon G_1 \times_G G_2 \to G_i$ called projection onto the first and second component, respectively, given by the respective set theoretic projections on V_K and E_K.

We say that $\nu\colon K \to G$ is a covering map (respectively, étale[1]) if for each $v \in V_K$, ν gives a bijection (respectively, injection) of incoming edges of v (i.e., those edges whose head is v) with those of $\nu(v)$, and a bijection (respectively, injection) of outgoing edges of v and $\nu(v)$. If $\nu\colon K \to G$ is a covering map and G is connected, then the *degree* of ν, denoted $[K\colon G]$, is the number of preimages of a vertex or edge in G under ν (which does not depend on the vertex or edge); if G is not connected, one can still write $[K\colon G]$ when the number of preimages of a vertex or edge in G is the same for all vertices and edges.

Given a digraph, G, we view G as an undirected graph (by forgetting the directions along the edges), and let $h_i(G)$ denote the i-th Betti number of G, and $\chi(G)$ its Euler characteristic; hence $h_0(G)$ is the number of connected components of G, $h_1(G)$ is the minimum number of edges needed to be removed from G to leave it free of cycles, and

$$h_0(G) - h_1(G) = \chi(G) = |V_G| - |E_G|.$$

Let $\mathrm{conn}(G)$ denote the connected components of G, and let

$$\rho(G) = \sum_{X \in \mathrm{conn}(G)} \max(0, h_1(X) - 1), \tag{1.1}$$

which we call the *reduced cyclicity of G*.

For each digraph, G, and field, $\mathbb{F}$, our sheaf theory is the theory of sheaves of finite dimensional $\mathbb{F}$-vector spaces on a certain finite Grothendieck topology (see [**sga72a, sga72b, sga73, sga77**], where a Grothendieck topology is called a "site") that we associate to G; this Grothendieck topology has many properties in common with topological spaces; in [**Fri05**] we have called these spaces *semitoplogical*, and have worked out the structure of their injective and projective modules, which allows us to compute derived functors (e.g., cohomology, Ext groups), used in [**Fri05, Fri06, Fri07**]. Here we define sheaves and describe a homology theory "from scratch," without appealing to projective or injective modules; later we explain how our homology theory fits into standard sheaf theory as the derived functors of global cosections.

Definition 1.1. *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* a 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 $e \in E$,*

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.*

[1] Stallings, in [**Sta83**], uses the term "immersion."

The field, $\mathbb{F}$, is arbitrary, although at times we insist that it not be finite, and at times that it have characteristic zero.

Now we define homology groups. 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

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

Definition 1.2. *We define the* zeroth *and* first homology groups of $\mathcal{F}$ *to be, respectively,*

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

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 $\mathcal{F}$. *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, $H_i(G,\underline{\mathbb{F}})$.*

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.

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

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

Since $d_{\mathcal{F}}$ has domain $\mathcal{F}(E)$ and codomain $\mathcal{F}(V)$, we have

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

If $j\colon G' \to G$ is a digraph morphism, there is a naturally defined sheaf $j_!\underline{\mathbb{F}}$ on G such that $H_i(j_!\underline{\mathbb{F}})$ is naturally isomorphic to $H_i(G')$ ($j_!$ will be defined as a functor from sheaves on G' to sheaves on G in Subsection 1.4.1); when j is an inclusion, then $j_!\underline{\mathbb{F}}$ is just the sheaf whose values are $\mathbb{F}$ on G' and 0 outside of G' (i.e., on vertices and edges not in G'); we will usually use $\underline{\mathbb{F}}_{G'}$ to denote $j_!\underline{\mathbb{F}}$ (which is somewhat abusive unless j is understood). If $\phi\colon G' \to G''$ is a morphism of digraphs over G, then ϕ gives rise to a natural morphism of sheaves $\underline{\mathbb{F}}_{G'} \to \underline{\mathbb{F}}_{G''}$. In this way the functor $G' \mapsto \underline{\mathbb{F}}_{G'}$ includes the category of digraphs over G as a subcategory of sheaves over G. As mentioned before, one key aspect of sheaf theory is that the functor is not full, i.e., there exist (very important) morphisms of sheaves $\underline{\mathbb{F}}_{G'} \to \underline{\mathbb{F}}_{G''}$ that do not arise from a morphism of digraphs $G' \to G''$; such morphisms will be needed to define sheaves (their kernels) that we call ρ-kernels, which will be crucial to our approach to the SHNC.

Next we give the long exact sequence in homology associated to a short exact sequence of sheaves.

Definition 1.3. *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 [**Fri05**] or Proposition I.3.1 of [**sga72a**]).

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

Theorem 1.4. *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

$$0 \to H_1(\mathcal{F}_1) \to H_1(\mathcal{F}_2) \to H_1(\mathcal{F}_3) \to H_0(\mathcal{F}_1) \to H_0(\mathcal{F}_2) \to H_0(\mathcal{F}_3) \to 0.$$

1.2.2. Quasi-Betti Numbers and Maximum Excess. For any digraph, G, we have that the pair h_0, h_1 assign non-negative integers to each sheaf over G, and these integers satisfy certain properties. In this chapter we introduce other pairs of invariants, essentially variations of h_0, h_1, that satisfy the same properties. Our proof of the SHNC will use the fact that the "maximum excess" is part of such a pair. Let us make these notions precise.

Definition 1.5. *A sequence of real numbers, $x_0, \ldots, x_r$ is a* triangular sequence *if for any $i = 1, \ldots, r-1$ we have*

$$x_i \le x_{i-1} + x_{i+1}.$$

Definition 1.6. *Given a digraph, G, and a field, $\mathbb{F}$, consider the category of sheaves of $\mathbb{F}$-vector spaces on G. Let α_0, α_1 be two functions from sheaves to the non-negative reals. We shall say that (α_0, α_1) is a* quasi-Betti number pair (for G and $\mathbb{F}$) *provided that:*

(1) *for each sheaf, $\mathcal{F}$, we have*

$$\alpha_0(\mathcal{F}) - \alpha_1(\mathcal{F}) = \chi(\mathcal{F}); \tag{1.2}$$

(2) *for any sheaves, $\mathcal{F}_1, \mathcal{F}_2$ on G we have*

$$\alpha_i(\mathcal{F}_1 \oplus \mathcal{F}_2) = \alpha_i(\mathcal{F}_1) + \alpha_i(\mathcal{F}_2) \quad \text{for } i = 0, 1;$$

(3) *for any short exact sequence of sheaves on G*

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

the sequence of integers

$$0, \alpha_1(\mathcal{F}_1), \alpha_1(\mathcal{F}_2), \alpha_1(\mathcal{F}_3), \alpha_0(\mathcal{F}_1), \alpha_0(\mathcal{F}_2), \alpha_0(\mathcal{F}_3), 0$$

is triangular.

Moreover, we say that a function, α, from sheaves to non-negative reals is a first quasi-Betti number *if the pair (α_0, α_1) with*

$$\alpha_1(\mathcal{F}) = \alpha(\mathcal{F}), \quad \alpha_0(\mathcal{F}) = \chi(\mathcal{F}) + \alpha(\mathcal{F})$$

are quasi-Betti number pair. The relationship between quasi-Betti numbers and a first quasi-Betti numbers is forced by equation (1.2).

Notice that (h_0, h_1) is a quai-Betti number pair; the only issue in establishing this is property (3) of the definition, and this follows from the long exact sequence given by Theorem 1.4.

Of course, if (α_0, α_1) is a quasi-Betti number pair, then clearly α_1 is a first quasi-Betti number.

Let us give other quasi-Betti number pairs, beginning with the one of main interest in this paper.

DEFINITION 1.7. *Let $\mathcal{F}$ be a sheaf on a digraph, G. For any $U \subset \mathcal{F}(V)$ we define the* head/tail neighbourhood of U*, denoted* $\Gamma_{\rm ht}(G, \mathcal{F}, U)$*, or simply* $\Gamma_{\rm ht}(U)$*, to be*

$$\Gamma_{\rm ht}(U) = \bigoplus_{e \in E_G} \{w \in \mathcal{F}(e) \mid d_h(w), d_t(w) \in U\}; \tag{1.3}$$

we define the excess of $\mathcal{F}$ at U *to be*

$$\operatorname{excess}(\mathcal{F}, U) = \dim\bigl(\Gamma_{\rm ht}(U)\bigr) - \dim(U).$$

Furthermore we define the maximum excess *of $\mathcal{F}$ to be*

$$\text{m.e.}(\mathcal{F}) = \max_{U \subset \mathcal{F}(V_G)} \operatorname{excess}(\mathcal{F}, U).$$

We shall see that the excess is a supermodular function, and hence the maximum excess occurs on a lattice of subsets of $\mathcal{F}(V)$. It is not hard to see that for the structure sheaf, $\underline{\mathbb{F}}$, we have

$$\text{m.e.}(\underline{\mathbb{F}}) = \rho(G).$$

It is instructive to determine which subsets of $\underline{\mathbb{F}}(V_G)$ obtain this maximum excess of $\rho(G)$. So let G' be obtained from G by discarding all components with positive Euler characteristic and, optionally, discarding some components of zero Euler characteristic, and then, optionally repeatedly pruning any of its leaves (i.e., removing a vertex of degree one and its incident edge); then the excess of $U(G') = \oplus_{v \in V_{G'}} \mathbb{F}(v)$ of $\underline{\mathbb{F}}$ on G is $\rho(G)$, and, conversely, any subspace $U \subset \underline{\mathbb{F}}(V_G)$ achieving the maximum excess of $\rho(G)$ is of the form $U(G')$ for a G' as above. The reader can easily see that such $U(G')$ form a lattice (i.e., are closed under intersection and sum).

THEOREM 1.8. *The maximum excess is a first quasi-Betti number.*

Theorem 1.8 will be crucial to our proof of the SHNC (although, as mentioned before, in an alternate proof we avoid the need for this theorem). Somewhat surprisingly, the statement of this theorem and all the necessary definitions do not involve any homology theory.

We shall show Theorem 1.8 by identifying the maximum excess with a certain "limit" Betti number.

1.2.3. Twisted Homology. One graph theoretic reformulation of the SHNC involves the reduced cyclicity defined in equation (1.1). This definition seems difficult to deal with, because of the $\max(0, h_1(X) - 1)$ term, and of the possibility of $h_1(X) = 0$ for some components, X, of G. For a digraph, G, one can realize $\rho(G)$ as a "twisted first Betti number;" constructing this "twisted homology theory" is our first step towards showing that the maximum excess is a first quasi-Betti number.

Let us first briefly motivate our definitions of twisted homology. We begin by noticing that for G connected we have

$$\rho(G) = \lim_{n\to\infty} h_1(L_n)/n, \tag{1.4}$$

where for each positive integer n we choose a covering $L_n \to G$ of degree n such that L_n is connected (for then $h_0(L_1) = 1$ and $h_1(L_n) = h_0(L_n) - \chi(L_n) = 1 + n\rho(G)$).

One way of choosing n and $L_n \to G$ of degree n such that L_n is connected is to take $n = p$ a prime number, and take $L_p \to G$ to be a "generic" $\mathbb{Z}/p\mathbb{Z}$ covering of G (see Section 1.3). It is well known that for $\mathbb{Z}/p\mathbb{Z}$ coverings $G' \to G$, or for any Abelian covering, the eigenvalues of the adjacency matrix of G' can be computed from those of G after "twisting" appropriately; here "twisting" means multiplying the entries of G's adjacency matrix by appropriate roots of unity, according to the characters of the "Galois group" of G' over G (see Section 1.3). The same holds for homology groups.

This leads us to a new homology theory, as follows. 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*, ψ, on G is a map

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

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

$$\mathcal{F}^\psi(P) = \bigl(\mathcal{F}(P)\bigr) \otimes_{\mathbb{F}} \mathbb{F}'$$

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

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

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, $\mathcal{F}^\psi$ is the sheaf on the same vector spaces extended to $\mathbb{F}'$-vector spaces, but with the tail restriction maps twisted by ψ. The map, $d_{\mathcal{F}^\psi}$, viewed as a matrix, has entries in the field $\mathbb{F}'$. The groups $H_i(\mathcal{F}^\psi)$ are defined as $\mathbb{F}'$-vector spaces.

Now let $\psi = \{\psi(e)\}_{e\in E_G}$ be viewed as $|E_G|$ indeterminates, and let $\mathbb{F}(\psi)$ denote the field of rational functions in the $\psi(e)$ over $\mathbb{F}$. Then $d = d_{\mathcal{F}^\psi}$ 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 1.9. *We define the* i-th twisted homology group of $\mathcal{F}$, *denoted by*

$$H_i^{\rm twist}(\mathcal{F}) = H_i^{\rm twist}(\mathcal{F},\psi),$$

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* $h_i^{\rm twist}(\mathcal{F})$, *to be dimension of* $H_i^{\rm twist}(\mathcal{F})$.

We easily see, akin to equation (1.4), that

$$\rho(G) = h_1^{\rm twist}(\underline{\mathbb{F}}).$$

The analogous short/long exact sequences theorem holds in twisted homology, and this easily implies that $h_1^{\rm twist}$ is a quasi-Betti number. We wish to mention that we can interpret

$$h_0^{\rm twist}(\underline{\mathbb{F}}) = \chi(\underline{\mathbb{F}}) + h_1^{\rm twist}(\underline{\mathbb{F}}) = \chi(G) + \rho(G)$$

as the number of "acyclic components" of G, i.e., the number of connected components that are free of cycles.

1.2.4. Maximum Excess Versus Twisted Betti Numbers, and The Unhappy 4-Bundle. Note that for the constant sheaf, $\underline{\mathbb{F}}$, on a digraph, G, the values of $h_1^{\rm twist}$ and the maximum excess agree and equal $\rho(G)$. Notice also that it is immediate that $h_1^{\rm twist}$ is a first quasi-Betti number, but it seems to us more difficult to show that the maximum excess is a first quasi-Betti number. This indicates that it would be easier to work with $h_1^{\rm twist}$ rather than the maximum excess in studying the SHNC (and this can be done). We give two reasons why we nonetheless use the maximum excess.

First, the SHNC is more directly related to the vanishing maximum excess of a certain sheaves we call ρ-kernels; and this vanishing is weaker (at least *a priori*) than the vanishing of $h_1^{\rm twist}$ of the ρ-kernels. Second, the Euler characteristic, reduced cyclicity, and the maximum excess have a nice scaling property under "pullbacks" via covering maps, that $h_1^{\rm twist}$ does not share. This makes $h_1^{\rm twist}$ seem to be, at times, the "wrong" invariant for certain situations, like those arising in the SHNC.

Let us discuss the above remarks in more precise terms. It is easy to see that

$$h_1^{\rm twist}(\mathcal{F}) \ge {\rm m.e.}(\mathcal{F}),$$

and one can show that equality holds if for each $e \in E_G$, $\mathcal{F}(e)$ is either zero or one dimensional. In particular, this holds for $\mathcal{F} = \mathbb{C}_L$ for any subgraph, L of G. However, there are sheaves, such as the "unhappy 4-bundle," that we will soon describe, which have maximum excess zero but positive $h_1^{\rm twist}$. The above inequality does show that if $h_1^{\rm twist}$ vanishes then so does the maximum excess; in the case of the SHNC and ρ-kernels this means that vanishing $h_1^{\rm twist}$ of ρ-kernels is at least as strong a condition as the SHNC.

We now describe a sheaf we call the *unhappy* 4-*bundle*. It is a highly instructive example that illustrates a number of points on maximum excess and twisted homology. Let B_2 be the bouquet of two self-loops, i.e., the digraph with one vertex, v, and two self-loops, e_1, e_2. Let $\mathcal{U}$ be defined as

$$\mathcal{U}(v) = \mathbb{F}^4, \quad \mathcal{U}(e_i) = \mathbb{F}^2 \quad \text{for } i = 1, 2, \tag{1.5}$$

and

$$d_h = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad d_t = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \tag{1.6}$$

where these matrices multiply the coordinates of $\mathcal{U}(E)$ arranged as a column vector (the column vector to the right of the matrix), where $\mathcal{U}(E)$'s coordinates are ordered as $\mathcal{U}(e_1) \oplus \mathcal{U}(e_2)$. The twisted incidence matrix of $\mathcal{U}$ (which characterizes $\mathcal{U}$) is given by

$$d_{\mathcal{U}^\psi} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & -\psi(e_2) & 0 \\ -\psi(e_1) & 0 & 0 & 1 \\ 0 & -\psi(e_1) & 0 & -\psi(e_2) \end{bmatrix}. \tag{1.7}$$

This matrix has a kernel of dimension one in $\mathbb{F}(\psi)$, however its maximum excess is zero. Equivalently, if $\mathcal{F}(v) = \mathbb{F}^4$ has $\alpha, \beta, \gamma, \delta$ as its standard basis (i.e., $\alpha = (1, 0, 0, 0)$, $\beta = (0, 1, 0, 0)$, etc.), then the image of the four standard coordinates on

$\mathcal{F}(E)$ via $d_{\mathcal{U}^\psi}$ is

$$\nu_1 = \alpha - \psi(e_1)\gamma, \quad \nu_2 = \beta - \psi(e_1)\delta, \quad \nu_3 = \alpha - \psi(e_2)\beta, \quad \nu_4 = \gamma - \psi(e_2)\delta. \tag{1.8}$$

The fact that $h_1^{\rm twist}(\mathcal{U}) \neq 0$ follows from the simple computation that

$$\nu_1 \wedge \nu_2 \wedge \nu_3 \wedge \nu_4 = 0$$

or the linear dependence relation

$$\nu_1 - \psi(e_2)\nu_2 - \nu_3 + \psi(e_1)\nu_4 = 0$$

The reason we call $\mathcal{U}$ a 4-bundle is that is four dimensional at the vertex of B_2, and it is has properties akin to a vector bundle; this will be explained more fully in a sequel to this paper.

For any sheaf, $\mathcal{F}$, on a digraph, G, and any morphism $\phi\colon K \to G$ of directed graphs, we define the *pullback of $\mathcal{F}$ via ϕ* to be the sheaf $\phi^*\mathcal{F}$ on K given via

$$(\phi^*\mathcal{F})(P) = \mathcal{F}(\phi(P)) \qquad \text{for all } P \in V_K \amalg E_K,$$

and for all $e \in E_K$,

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

It is easy to see that if μ is a covering map of degree $\deg(\mu)$ then

$$\chi(\mu^*\mathcal{F}) = \deg(\mu)\chi(\mathcal{F}),$$

and, with a little more work (and using Galois graph theory, oddly enough), that

$$\text{m.e.}(\mu^*\mathcal{F}) = \deg(\mu)\text{m.e.}(\mathcal{F}). \tag{1.9}$$

The "unhappy 4-bundle" also shows that $h_1^{\rm twist}$ does not enjoy this "scaling by $\deg(\mu)$ under pullback" property. Indeed, $h_1^{\rm twist}(\mathcal{U}) = 1$; however if $\phi\colon G' \to B_2$ (recall $\mathcal{U}$ is defined on the graph B_2) is the degree two cover of B_2 in which the G' edges mapping to e_1 are self-loops, and the edges mapping to e_2 are not, then $h_1^{\rm twist}(\phi^*\mathcal{U}) = 0$. In other words, via taking wedge products or solving for a linear relation, it is straightforward to verify the linear independence of the eight vectors

$$\nu_1^1 = \alpha^1 - \psi(e_1^1)\gamma^1,\ \nu_2^1 = \beta^1 - \psi(e_1^1)\delta^1,\ \nu_3^1 = \alpha^1 - \psi(e_2^1)\beta^2,\ \nu_4^1 = \gamma^1 - \psi(e_2^1)\delta^2.$$

$$\nu_1^2 = \alpha^2 - \psi(e_1^2)\gamma^2,\ \nu_2^2 = \beta^2 - \psi(e_1^2)\delta^2,\ \nu_3^2 = \alpha^2 - \psi(e_2^2)\beta^1,\ \nu_4^2 = \gamma^2 - \psi(e_2^2)\delta^1.$$

1.2.5. The Fundamental Lemma and Limit Homology. The following is the main and most difficult theorem in this chapter; it allows us to connect twisted homology and maximum excess. For any digraph we shall define the notion of its *Abelian girth*, which is always at least as large as its girth.

THEOREM 1.10. *For any sheaf, $\mathcal{F}$, on a digraph, G, let $\mu\colon G' \to G$ be a covering map where the Abelian girth of G' is at least*

$$2\Big(\dim\big(\mathcal{F}(V)\big) + \dim\big(\mathcal{F}(E)\big)\Big) + 1.$$

Then

$$h_1^{\rm twist}(\mu^*\mathcal{F}) = \text{m.e.}(\mu^*\mathcal{F}).$$

From this lemma it is easy to see that the maximum excess is a first quasi-Betti number.

1.2.6. Limits and Limiting Betti Numbers. In this subsection we give a new interpretation to our main theorem, Theorem 1.10. For any two covering maps,

$$\phi_1 \colon G_1 \to G \quad \text{and} \quad \phi_2 \colon G_2 \to G,$$

their fibre product

$$\phi \colon G_1 \times_G G_2 \to G$$

factors through both ϕ_1 and ϕ_2, i.e., ϕ is a "common cover." It follows that the set, $\mathrm{cov}(G)$, of covering maps of a fixed digraph, G, is a directed set, under the partial order $\phi_1 \le \phi_2$ if ϕ_2 factors through ϕ_1. As such we may speak of limits in the usual sense of limits of a directed sets; i.e., if f is, say, a real-valued function on covering maps, then we write

$$\lim_{\phi \in \mathrm{cov}(G)} f(\phi) = L$$

if for any $\epsilon > 0$ there is a $\phi_\epsilon \in \mathrm{cov(G)}$ such that $|f(\phi') - L| \le \epsilon$ provided that ϕ' factors through ϕ_ϵ (such a limit, L, is necessarily unique).

Theorem 1.10 implies that for any sheaf, $\mathcal{F}$, on G, we have

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

Of course, Theorem 1.10 amounts to saying that this limiting value is exactly attained at any $\phi \colon G' \to G$ with G' of sufficiently large girth or Abelian girth.

For a sheaf, $\mathcal{F}$, on a digraph, G, we define

$$\lim_{\phi \in \mathrm{cov}(G)} \frac{h_i^{\mathrm{twist}}(\mathcal{F})}{\deg(\phi)}$$

to be the i-th limiting Betti number, which we denote $h_i^{\lim}(\mathcal{F})$. Evidently,

$$h_1^{\lim}(\mathcal{F}) = \mathrm{m.e.}(\mathcal{F}), \quad h_0^{\lim}(\mathcal{F}) = \chi(\mathcal{F}) + \mathrm{m.e.}(\mathcal{F}).$$

It is easy to see that the limit of quasi-Betti pairs is also a quasi-Betti pair, and that for any fixed covering map $\phi \colon G' \to G$, the functions for $i = 0, 1$ given by

$$h_i^{\mathrm{twist}}(\phi^* \mathcal{F}) / \deg(\phi)$$

form a quasi-Betti pair. This is another way of saying that Theorem 1.10 implies Theorem 1.8.

1.2.7. Sheaves, Adjacency Matrices, and Laplacians. We remark that from the incidence matrix, $d_{\mathcal{F}} = d_h - d_t$, of a sheaf, $\mathcal{F}$, one can define a Laplacians, adjacency matrices, and related matrices that are analogues of those used for graphs. This construction can also be viewed as a very special, discrete case of Hodge theory. We require that for each $P \in V_G \amalg E_G$, we have that each $\mathcal{F}(P)$ be endowed with an inner product. In that way $\mathcal{F}(V), \mathcal{F}(E)$ become inner product spaces, and we have adjoint operators d_h^*, d_t^* and $d^* = d_h^* - d_t^*$ from $\mathcal{F}(V)$ to $\mathcal{F}(E)$. We define

$$\Delta_0 = dd^*, \quad \Delta_1 = d^*d$$

to be the *Laplacians* of $\mathcal{F}$, which, of course, depend on the inner products chosen for the values, $\mathcal{F}(P)$, of $\mathcal{F}$; we easily see that Δ_i is an operator on $\mathcal{F}(V)$ and $\mathcal{F}(E)$ respectively for $i = 0$ and $i = 1$ respectively; if $\mathbb{F}$ is of characteristic zero, then the Δ_i are positive semi-definite operators, and the kernel of Δ_i is $H_i(\mathcal{F})$. In the special case $\mathcal{F} = \underline{\mathbb{F}}$, with the same, standard inner products on all $\mathcal{F}(P) = \mathbb{F}$, the Laplacians become the usual Laplacians of the graph.

Furthermore, given $\mathcal{F}$ and inner products on the values of $\mathcal{F}$, we get generalizations of the adjacency matrix and degree matrix. For example, if we set

$$D_0 = d_h d_h^* + d_t d_t^*, \quad A_0 = d_h d_t^* + d_t d_h^*,$$

we have that $\Delta_0 = D_0 - A_0$; in the case $\mathcal{F} = \underline{\mathbb{F}}$ and standard inner products, D_0, A_0, respectively amount to the usual degree and adjacency matrices, respectively. One can define D_1, A_1 analogously.

One could define a sheaf to be *regular* in the way that one would define a graph to be regular, i.e., if both D_0 and D_1 are both multiples of the identity. One could measure the *expansion* of a sheaf by the eigenvalues of A_0, A_1 or Δ_0, Δ_1.

We believe that the spectral theory of such matrices and related properties such as expansion could be quite interesting to pursue. However, we shall not pursue them further in this paper.

1.3. Galois and Covering Theory

In this section we establish a number of important definitions and facts concerning graph coverings, Abelian coverings, and Galois coverings.

There is a collection of facts about number fields that may be called Galois theory; this would include classical Galois theory, but also more recent statements such as if k' is a Galois extension field of k, then

$$k' \otimes_k k' \simeq \bigoplus_{\mathrm{Aut}(k'/k)} k'$$

(see [**Del77**], Section I.5.1). Such facts have analogues in graph theory, which one might call "graph Galois theory." Such facts were described in [**Fri93, ST96**]; at least some of these some of these facts were known much earlier, in [**Gro77**]; since these facts are fairly simple and quite powerful, we presume they may occur elsewhere in the literature (perhaps only implicitly).

1.3.1. Galois Theory of Graphs. We shall summarize some theorems of [**Fri93**]; the reader is referred to there and [**ST96**] for more discussion. In this article Galois group actions, when written multiplicatively (i.e., not viewed as functions or morphisms) will be written on the right, since our Cayley graphs are written with its generators acting on the left.

Let $\pi\colon K \to G$ be a covering map of digraphs. We write $\mathrm{Aut}(\pi)$, or somewhat abusively $\mathrm{Aut}(K/G)$ (when π is understood), for the automorphisms of K over G, i.e., the digraph automorphisms $\nu\colon K \to K$ such that $\pi = \pi\nu$.

Now assume that K and G are connected. Then it is easy to see ([**Fri93, ST96**]) that for every $v_1, v_2 \in V_K$ there is at most one $\nu \in \mathrm{Aut}(K/G)$ such that $\nu(v_1) = v_2$; the same holds with edges instead of vertices. It follows that $|\mathrm{Aut}(K/G)| \le [K : G]$, with equality iff $\mathrm{Aut}(K/G)$ acts transitively on each vertex and edge fibre of π. In this case we say that π is *Galois*.

If $\pi\colon K \to G$ is Galois but K is not connected, $|\mathrm{Aut}(K/G)|$ can be as large as $[K\colon G]$ factorial (if K is a number of copies of G). So when K is not connected, we say that a covering map $\pi\colon K \to G$ is *Galois* provided that we additionally specify a subgroup, $\mathcal{G}$, of $\mathrm{Aut}(K/G)$ of that acts simply (without fixed points) and transitively on each of the vertex and edge fibres of π; we declare $\mathcal{G}$ to be the Galois group. Again, this additional specification does not change any of the theorems here, although it does mean that certain $\pi\colon K \to G$ can be Galois on each component of G without being Galois in our sense (consider $G = G_1 \amalg G_2$, and $K_i = \pi^{-1}(G_i)$, where G_1, G_2 are connected and $\mathrm{Aut}(K_i/G_i)$ are non-isomorphic groups).

THEOREM 1.11 (Normal Extension Theorem). *If $\pi\colon G \to B$ is a covering map of digraphs, there is a covering map $\mu\colon K \to G$ such that $\pi\mu$ is Galois.*

In this situation we say that K is a *normal extension* of G (assuming the maps μ and π are understood). By convention, all graphs are finite in the paper unless otherwise specified. Generally speaking, we will not address infinite graphs in the context of Galois theory; however, if the $\pi\colon G \to B$ in this theorem is a morphism of finite degree, even if G and B are infinite digraphs, then the proof of the Normal Extension Theorem due to Gross is still valid.

Let us outline two proofs of the Normal Extension Theorem. The proof in [**Fri93**] uses the fact that G corresponds to a subgroup, S, of index $n = |V_G|$ of the group $\pi_1(B)$, the fundamental group of B (which is the free group on $h_1(B)$ elements). The intersection of xSx^{-1} over a set of coset representatives of $\pi_1(B)/S$ is a normal subgroup, N, of finite size (at worst n^n, since there are n cosets and each xSx^{-1} is of index n); $\pi(B)/N$ then naturally corresponds to a Galois cover $K \to B$ of at most n^n vertices.

There is a very pretty proof of the Normal Extension Theorem discovered earlier by Jonathan Gross in [**Gro77**], giving a better bound on the number of vertices of K. For any positive integer k at most $n = |V_G|$, let $\Omega^k(G)$ be the subgraph of $G \times_B G \times_B \cdots \times_B G$ (multiplied k times) induced on the set of vertices of the form $(v_1, \ldots, v_k)$ where $v_i \neq v_j$ for all i, j with $i \neq j$. Each $\Omega^k(G)$ admits a covering map to G by projecting onto any one of its components. But $\Omega^n(G)$, which has edge and vertex fibers of size $n!$, is Galois by the natural, transitive action of S_n (the symmetric group on n elements) on $\Omega^n(G)$. So $\Omega^n(G)$ is a Galois cover of degree at most $n!$ over B.

1.3.2. Galois Coordinates. Given a graph, G, and a group, $\mathcal{G}$, consider the task of describing all Galois covering maps $\pi\colon K \to G$ with Galois group $\mathcal{G}$; consider also the task of giving a meaning to a "random" such Galois covering (i.e., describe a natural probability space whose atoms are such coverings). This can be done in a number of ways, via Galois coordinates or the monodromy map. Here we shall review these ideas and apply them. These ideas occur (in parts) in many places in the literature; see, for example, [**Fri08, Fri03, AL02, Fri93**].

Again, fix a graph, G, and a group, $\mathcal{G}$. By *Galois coordinates on G with values in $\mathcal{G}$* we mean a choice of $a_e \in \mathcal{G}$ for each $e \in E_G$. From the $\{a_e\}$ we build a covering map $\phi\colon K \to G$ by taking $V_K = V_G \times \mathcal{G}$ and taking $E_K = E_G \times \mathcal{G}$ with the head and tail, respectively, of an edge (e, a) being

$$h_K(e,a) = (h_G e, a_e a), \quad t_K(e,a) = (t_G e, a), \tag{1.10}$$

respectively. We define a $\mathcal{G}$ action on K via $g \in \mathcal{G}$ is the morphism such that for $P \in V_G \amalg E_G$ and $a \in \mathcal{G}$, g sends (P, a) to

$$(P,a)g = (P, ag); \tag{1.11}$$

in view of the fact that a_e multiplies to the left in equation (1.10), we see that the right multiplication of g on a in equation (1.11) actually defines a digraph morphism. Let ϕ be projection onto the first coordinate. Clearly ϕ is a Galois covering with Galois group $\mathcal{G}$.

Conversely, let $\phi\colon K \to G$ be any $\mathcal{G}$ Galois covering. We may identify V_K with $V_G \times \mathcal{G}$ by choosing for each $v \in V_G$ an element $v' \in V_K$ such that $\phi(v') = v$ and declaring v' to have coordinates $(v, 1)$ where 1 is the identity in $\mathcal{G}$; we say that v' is the *origin for v in K*; then for all $v'' \in V_K$ with $\phi(v'') = v$ there is a unique $g \in \mathcal{G}$ with $v'' = v'g$, and we declare v'' to have coordinates (v, g). For any $g' \in \mathcal{G}$ we have $v''g' = vgg'$ which has coordinates (v, gg'); hence g' acts on coordinates by right multiplication. Now choose an edge $e' \in E_K$, and let $e = \phi(e')$; there exist unique $a_{e'}, g \in \mathcal{G}$ for which the endpoints of e' have coordinates

$$te' = (te, g), \quad he' = (he, a_{e'}g).$$

But the $\mathcal{G}$ action on K then shows that for any g' we have

$$t(e'g') = (te, gg'), \quad h(e'g') = (he, a_{e'}gg').$$

It follows that $a_{e'}$ depends only on $e = \phi(e')$, i.e., $a_{e'} = a_{e'g'}$ for all $e' \in K$ and $g' \in \mathcal{G}$. In other words, there is a unique a_e for each $e \in E_G$ such that the ϕ fibres of e join $(t_G e, g)$ to $(h_G e, a_e g)$ for each $g \in \mathcal{G}$. In summary, for each choice of an element in the vertex fibres we get Galois coordinates (and conversely).

Notice that in setting the coordinates on V_K, if for $v \in V_G$ we choose a different origin, namely $v'g_v$ instead of v', then we have $v'g = (v'g_v)g_v^{-1}g$ for any $g \in \mathcal{G}$; it follows that the vertx $v'g$, which would have had coordinates (v, g) with v' as origin, will have coordinates $(v, g_v^{-1}g)$ with $v'g$ as origin. In particular, if for $e' \in V_K$ and $e = \phi(e)$ we have $te' = (te, g)$ and $he' = (he, a_e g)$ in one set of coordinates for some g, and the origins of te and he are respectively translated by g_{te} and g_{he}, then in the new coordinates

$$te' = (te, g_{te}^{-1}g), \quad he' = (he, g_{he}^{-1}a_e g).$$

Setting $g' = g_{te}^{-1}g$, it follows that in the new, translated coordinates we have $te' = (te, g')$ and $he' = (he, \widetilde{a}_e g')$, where

$$\widetilde{a}_e = g_{he}^{-1}a_e g_{te}.$$

So changing Galois coordinate origins as such amounts to a transformation of Galois coordinates

$$a_e \mapsto \widetilde{a}_e = g_{he}^{-1}a_e g_{te} \tag{1.12}$$

for a family $\{g_v\}_{v\in V_G}$ of $\mathcal{G}$ values indexed on V_G.

Galois coordinates give a nice model of a random Galois cover of a given graph with given Galois group—just choose the each a_e uniformly in $\mathcal{G}$, assuming $\mathcal{G}$ is finite, and independently over the $e \in E_G$. If one wants a model of a random cover, one that is not Galois, one often chooses V_K to have vertices $V_G \times \{1, \ldots, n\}$, where n is the degree of the cover, and chooses random matchings over each G edge (random permutations over self-loops); see, e.g., [**Fri08, Fri03, AL02**].

1.3.3. Walks and Monodromy.

Another type of coordinates for Galois coverings are the monodromy maps. For this we need to fix some notation regarding walks in a digraph.

Definition 1.12. *Let G be a digraph. By an* oriented edge *of G we mean a formal symbol e^+ or e^- where $e \in E_G$. We extend the head and tail map to oriented edges via $he^+ = te^- = he$ and $te^+ = he^- = te$. We say that the* inverse *of e^+ is e^- and vice versa. An* undirected walk *(or simply* walk*) in G is an alternating sequence of vertices and oriented edges $w = (v_0, f_1, v_1, f_2, v_2, \ldots, f_r, v_r)$*

with $hf_i = v_i$*,* $tf_i = v_{i-1}$ *for* $i = 1, \ldots, r$*; we call* r *its* length*; we say that* w *is* closed *if* $v_r = v_0$*; we say that* w *is* non-backtracking *or* reduced *if for each* $i = 1, \ldots, r-1$*,* f_i *and* f_{i+1} *are not inverses of each other.*

If G is a digraph and $v \in V_G$, then we define $\pi_1(G, v)$ to be the group of non-backtracking closed walks about v, where the group operation is concatenation of walks (which we reduce until they are non-backtracking). This, of course, is isomorphic to the usual fundamental group, $\pi_1(\widetilde{G}, v)$, where $\widetilde{G}$ is the geometric realization of G, where vertices of G correspond to points and edges of G correspond to unit intervals. If G is connected, then $\pi_1(G, v)$ is a free group on $h_1(G)$ generators. We may also describe $\pi_1(G, v)$ as the classes of closed walks about v, where two walks are equivalent if they reduce to the same non-backtracking word ("reduce" meaning repeatedly eliminating any two consecutive steps of the walk that traverse an edge and then its inverse).

Let $\phi\colon G' \to G$ be Galois with Galois group $\mathcal{G}$, with G connected, and let $\{a_e\}$ be Galois coordinates for ϕ. Extend the $\{a_e\}$ to be defined on oriented edges via $a_{e^+} = a_e$, $a_{e^-} = a_e^{-1}$. Fix a $v \in V_G$. Then for any closed walk, w, about v in G, we let e_i be the oriented edge traversed by w on the i-th step and set

$$\mathrm{Mndrmy}_{\phi,\{a_e\}}(w) = a_{e_k} \ldots a_{e_1},$$

where $\{a_e\}_{e \in E_G}$ are Galois coordinates on ϕ. We call $\mathrm{Mndrmy}_{\phi,\{a_e\}}$ the *monodromy* map with respect to $\{a_e\}$; it is a group morphism from $\pi_1(G, v)$ to $\mathcal{G}$. Conversely, given a group morphism

$$M\colon \pi_1(G, v) \to \mathcal{G}$$

with G connected, we can form a covering $\phi\colon G' \to G$ with Galois coordinates $\{a_e\}$ such that $\mathrm{Mndrmy}_{\phi,\{a_e\}} = M$; indeed, we let T be an undirected spanning tree for G, define $a_e = 1$ for $e \in E_T$ (where 1 denotes the identity in $\mathcal{G}$), and define a_e for $e \in E_G \setminus E_T$ by taking an element $\gamma \in \pi_1(G, v)$ composed entirely of E_T edges except for one edge e (traversed in the same orientation as e) and set $a_e = M(e)$; since $\pi_1(G, v)$ is a free group on $E_G \setminus E_T$, this implies that M is well-defined and equals $\mathrm{Mndrmy}_{\phi,\{a_e\}}$.

If we change Galois coordinates on ϕ, then according to equation (1.12) we get a conjugate element. Hence there is a natural map:

$$\mathrm{Mndrmy}_{\phi}\colon \pi_1(G, v) \to \mathrm{ConjClass}(\mathcal{G}).$$

If $v' \in V_G$ has a path, p, to v, then the map $\gamma \mapsto p\gamma p^{-1}$ gives a homomorphism $\pi_1(G, v) \to \pi_1(G, v')$, and the two monodromy maps, respectively, send γ and $p\gamma p^{-1}$ to the same conjugacy class; hence we get a map

$$\mathrm{Mndrmy}_{\phi}\colon \pi_1(G) \to \mathrm{ConjClass}(\mathcal{G})$$

independent of the base point (for G connected). Any notion defined on conjugacy classes of $\mathcal{G}$ becomes defined on $\pi_1(G)$ via monodromy. For example, if $\mathcal{G}$ is Abelian, then the conjugacy classes of $\mathcal{G}$ are the same as $\mathcal{G}$, and we get a homomorphism

$$\mathrm{Mndrmy}_{\phi}\colon \pi_1(G) \to \mathcal{A},$$

for any cover $\phi\colon G' \to G$ with Abelian Galois group $\mathcal{A}$ (compare this to the discussion of torsors in Section 5.2 of [**Fri93**]). We remark that if the monodromy map is onto $\mathcal{A}$, and G is connected then G' is connected; indeed, this means that any two vertices in the same fiber are connected, since any vertex in G' has a path to a vertex in any vertex fibre (lifted from the element of $\pi_1(G)$ that maps to the appropriate element of $\mathcal{A}$); hence we can connect any two vertices via a path.

1.3.4. Covering maps and ρ.

Here we describe a remarkable property of ρ under covering maps.

THEOREM 1.13. *For any covering map $\pi\colon K \to G$ of degree d, we have $\chi(K) = d\chi(G)$ and $\rho(K) = d\rho(G)$.*

PROOF. The claim on χ follows since $d = |V_K|/|V_G| = |E_K|/|E_G|$. To show the claim on ρ, it suffices to consider the case of G connected, the general case obtained by summing over connected components; but similarly it suffices to consider the case of K connected. In this case

$$\rho(G) = h_1(G) - 1 = -\chi(G) = -d\chi(K) = d\big(h_1(K) - 1\big) = d\rho(K).$$

$\square$

1.4. Sheaf Theory and Homology

In this section we define sheaves of vector spaces over a graph, G, and their homology groups, and give their basic properties. Then we explain the definitions and properties in terms of sheaf theory on Grothendieck topologies; in case G has no self-loops, we describe a topological space, $\mathrm{Top}(G)$, whose sheaves give an equivalent description of our notion of sheaf.

In the first subsection we describe everything in simple terms, giving some claims without proof; the reader can either prove them from scratch, or wait until the second subsection where we explain that all of these claims are special cases of well-known results.

1.4.1. Homology and Pullbacks.

The basic definitions of sheaves were given in Subsection 1.2.1. In this subsection we prove Theorem 1.4 and discuss pullbacks and related functors.

PROOF (OF THEOREM 1.4). By the "vertexwise and edgewise" nature of taking images and kernels, we see that we have a diagram

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & 0 & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathcal{F}_1(E) & \longrightarrow & \mathcal{F}_2(E) & \longrightarrow & \mathcal{F}_3(E) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & \mathcal{F}_1(V) & \longrightarrow & \mathcal{F}_2(V) & \longrightarrow & \mathcal{F}_3(V) & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

The theorem follows from the standard "delta" or "connecting" map in homological algebra, via the "snake lemma" (see [**Lan02, AM69, HS97**]). $\square$

Next we describe the functoriality of sheaves. For any sheaf, $\mathcal{F}$, on a graph, G, and any morphism $\phi\colon K \to G$ of directed graphs, recall from Subsection 1.2.4 that the "pullback" sheaf $\phi^*\mathcal{F}$ on K is defined via

$$(\phi^*\mathcal{F})(P) = \mathcal{F}(\phi(P)) \qquad \text{for all } P \in V_K \amalg E_K,$$

and for all $e \in E_K$,

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

If $\mathcal{F}$ is a sheaf on G and K is a subgraph of G, then there is a sheaf on G denoted $\mathcal{F}_K$ called "$\mathcal{F}$ restricted to K and extended by zero," defined by $(\mathcal{F}_K)(P)$ is 0 if $P \notin V_K \amalg E_K$, and otherwise $\mathcal{F}(P)$; the restriction maps are inherited from $\mathcal{F}$ (when 0 is not involved). Notice that in case $\mathcal{F} = \underline{\mathbb{F}}$, then we have

$$\underline{\mathbb{F}}_K(V_G) = \mathbb{F}^{V_K}, \quad \underline{\mathbb{F}}_K(E_G) = \mathbb{F}^{E_K}, \tag{1.13}$$

and $d = d_h - d_t$ is the standard incidence matrix of K; hence $H_i(\underline{\mathbb{F}}_K) \simeq H_i(K)$.

If $\phi\colon K \to G$ is an arbitrary map, and $\mathcal{F}$ a sheaf on K, there is a natural sheaf $\phi_!\mathcal{F}$ on G defined as follows:

$$(\phi_!\mathcal{F})(P) = \bigoplus_{Q \in \phi^{-1}(P)} \mathcal{F}(Q), \quad \forall P \in V_G \amalg E_G,$$

with the restriction maps induced from those of $\mathcal{F}$, i.e., $(\phi_!\mathcal{F})(h,e)$ is the sum of the maps taking, for $e' \in \phi^{-1}(e)$, the $\mathcal{F}(e')$ component of $(\phi_!\mathcal{F})(e)$ to the $\mathcal{F}(he')$ component of $(\phi_!\mathcal{F})(he)$ via the map $\mathcal{F}(h,e')$. The reader can now observe that

$$(\phi_!\mathcal{F})(V_G) \simeq \mathcal{F}(V_K), \quad (\phi_!\mathcal{F})(E_G) \simeq \mathcal{F}(E_K), \tag{1.14}$$

and $d_{\phi_!\mathcal{F}}$ is the same map as $d_{\mathcal{F}}$ modulo these isomorphisms; hence

$$H_i(\phi_!\mathcal{F}) \simeq H_i(\mathcal{F}) \tag{1.15}$$

for $i = 0, 1$. In Subsection 1.4.3 we prove that $\phi_!$ is the left adjoint of ϕ^*, and in particular the isomorphisms of homology groups above are immediate; in Subsection 1.4.4 we explain the role of $\phi_!$ in certain "vanishing theorems" (of sheaf invariants). We shall make special use of $\phi_!$ for ϕ étale in our approach to the SHNC (see Theorems 1.16 and 2.14).

If $\phi\colon K \to G$ is the inclusion of a subgraph, and $\mathcal{F}$ is a sheaf on G, then $\mathcal{F}_K$, defined before, equals $\phi_!\phi^*\mathcal{F}$. More generally we write $\mathcal{F}_K$ for $\phi_!\phi^*\mathcal{F}$ for arbitrary ϕ, provided that ϕ is understood in context. Since $\phi^*\underline{\mathbb{F}} = \underline{\mathbb{F}}$ for arbitrary ϕ, we always have $\underline{\mathbb{F}}_K = \phi_!\underline{\mathbb{F}}$. This observation, combined with equation (1.15), gives another proof that $H_i(\underline{\mathbb{F}}_K)$ is canonically isomorphic to $H_i(K)$ for $i = 0, 1$; this proof, based on adjoints, is less explicit than the proof based on equation (1.13) and the remarks just below it.

The tensor product of two sheaves on G is defined as the tensor product their values at each point and each vertex, as $\mathbb{F}$-vector spaces. Note that if $\phi\colon K \to G$ is an arbitrary morphism of digraphs, and $\mathcal{F}$ is a sheaf on G, we easily verify that

$$\mathcal{F}_K = \phi_!\phi^*\mathcal{F} = \mathcal{F} \otimes \underline{\mathbb{F}}_K,$$

and if $K' \to G$ is another morphism we have an isomophism of sheaves on G

$$\underline{\mathbb{F}}_K \otimes \underline{\mathbb{F}}_{K'} \simeq \underline{\mathbb{F}}_{K \times_G K'}. \tag{1.16}$$

Furthermore, if $L \to G$ is an arbitrary digraph morphism, we have an equality of sheaves on K,

$$\phi^*\underline{\mathbb{F}}_L = \underline{\mathbb{F}}_{K \times_G L}.$$

If $G' \subset G$, then there is a natural inclusion of sheaves on G, $\underline{\mathbb{F}}_{G'} \to \underline{\mathbb{F}}$ (but not generally any nonzero morphism from $\underline{\mathbb{F}} = \underline{\mathbb{F}}_G$ to $\underline{\mathbb{F}}_{G'}$).

If $\phi\colon K \to G$ is a morphism of graphs, and $\alpha\colon \mathcal{F}_1 \to \mathcal{F}_2$ is a morphism of sheaves on K, then we have natural a natural morphism

$$\phi_!\alpha\colon \phi_!\mathcal{F}_1 \to \phi_!\mathcal{F}_2,$$

that make $\phi_!$ a functor on the category of sheaves. Similarly for ϕ_*, and for the pullback, ϕ^* (which acts the other way, from sheaves and their morphisms on G to those on K).

1.4.2. Standard Sheaf Theories. In this subsection we explain the connections with classical sheaf theory on topological spaces. We then describe our definitions and particular choice of homology theory (and the role of $\phi_!$) in terms of the view of Grothendieck et al. ([**sga72a, sga72b, sga73, sga77**]).

First consider an arbitrary topological space on a finite set, X. Say that an open set, U, in X is *irreducible* if U is nonempty[2] and not the union of its proper subsets. It is known that the category of sheaves on X is equivalent to the category of presheaves on the irreducible open subsets; this can be proven directly—the essential idea is that if a set is not irreducible, then we can construct its value at a sheaf from those on its subsets; there is also a proof in Section 2.5 of [**Fri05**], where this fact follows easily from the Comparison Lemma of [**sga72a**], Exposé III, 4.1. As is pointed out in [**Fri05**], this theorem is valid for any finite *semitopological* Grothendieck topology, where semitopological means that the underlying category has only one morphism from any object to itself.

For example, if $X = \{A, B, C, D\}$ with irreducible open sets being $\{A\}$, $\{C\}$, $\{A, B, C\}$, and $\{A, D, C\}$. Then one can recover a sheaf on X (which has seven open sets) on the basis of its values on these four sets, and any presheaf on these four sets extends to a sheaf on X. We remark that X geometrically corresponds (see [**Fri05**]) to a circle, X, covered by two overlapping intervals, the intervals corresponding to $\{A, B, C\}$ and $\{A, D, C\}$. We have $h_i(X) = 1$ for $i = 0, 1$.

Let G be a digraph with no self-loops. In this case our sheaf theory agrees with a standard topological one. Namely, let $\mathrm{Top}(G)$ be the topological space on $V_G \amalg E_G$, whose open sets are subgraphs of G. There are two types of open irreducible sets: those of the form $\{v\}$ with $v \in V_G$, and those of the form $\{he, e, te\}$ with $e \in E_G$; for each e we have $\{he\}$ and $\{te\}$ are subsets of $\{he, e, te\}$, and hence a sheaf on $\mathrm{Top}(G)$ is determined by its values on the sets of type $\{v\}$ and $\{he, e, te\}$ and the restrictions from the values on $\{he, e, te\}$ to both $\{he\}$ and $\{te\}$. We therefore recover our definition of a sheaf on a graph (i.e., Definition 1.1).

Note that in the above $X = \{A, B, C, D\}$ definition, this is equivalent to $\mathrm{Top}(G)$ with $V_G = \{A, C\}$ and $E_G = \{B, D\}$ and any heads/tails correspondences making this a graph of two vertices joined by two edges.

Notice that the above construction also gives a space, $\mathrm{Top}(G)$, when G has self-loops. But this space has the wrong properties and homology groups. For example, if G has one vertex and one self-loop, then $h_i(G) = 1$ for $i = 0, 1$ as defined in the previous section; however, $\mathrm{Top}(G)$ amounts to one irreducible open lying in another

[2] If the empty set were considered irreducible, the subcategory of irreducible open sets would have an initial element, making the structure sheaf injective and giving the wrong homology groups. One can say that the empty set is the union of proper subsets, namely the empty union; as such the empty set is reducible "by definition."

(with only one inclusion, not the desired two), and we have $h_1(\mathrm{Top}(G)) = 0$. So we now give a Grothendieck topology for every digraph, G, that gives our sheaf and homology theory.

For each digraph, G, let $\mathrm{Cat}(G)$ be the category whose objects are $V_G \amalg E_G$ and where the $2|E_G|$ non-identity morphisms are given by $he \to e$ and $te \to e$ ranging over all $e \in E_G$ (with two distinct morphisms $he \to e$ and $te \to e$, even when $he = te$). Then a sheaf over $\mathrm{Cat}(G)$ with the *grossière topologie*, i.e., a presheaf over the category $\mathrm{Cat}(G)$, is just the notion of a sheaf given earlier. Again, if e is a self-loop, then this category has two morphisms between two distinct objects; it is easy to see that the category of sheaves over a graph with a self-loop cannot be equivalent to the category of sheaves over any topological space.

Notice that earlier definitions regarding sheaves on G and related matters often involve a P in $V_G \amalg E_G$, giving vertices and edges a somewhat equal treatment; this happens because V_G and E_G comprise the objects of $\mathrm{Cat}(G)$, and only the morphisms of $\mathrm{Cat}(G)$ distinguish them.

At this point we will use explain certain features of the homology theory we use here. The proofs are in or are easy consequences of material in [**Fri05**], and is mostly easily derivable from material in [**sga72a, sga72b, sga73, sga77**] (which contains a lot of other material...). We shall assume the reader is familiar with basic sheaf and cohomology theory found in any algebraic geometry text, such as [**Har77**], and we will just list a few points that are not standard, or where the finite graph situation is different. Let $\mathbf{Sh}(G)$ be the category of sheaves of vector spaces (over some fixed field, $\mathbb{F}$) on G.

(1) $\mathbf{Sh}(G)$ have enough projectives as well as injectives. (See [**Fri05**] for a simple characterization of all injectives or projectives.)
(2) If $u\colon K \to G$ is a morphism of graphs, the pullback, $u^*\colon \mathbf{Sh}(G) \to \mathbf{Sh}(K)$ is defined via
$$(u^*\mathcal{F})(P) = \mathcal{F}(u(P))$$
for $P \in V_K \amalg E_K$, with its natural restriction maps inherited from $\mathcal{F}$ (this is the same pullback defined in Subsections 1.2.4 and 1.4.1); u^* has a left adjoint, $u_!$ (defined in Subsection 1.4.1), and a right adjoint, u_* (see [**sga72a**], Exposé I, Proposition 5.1). In other words,
$$\mathrm{Hom}_G(\phi_!\mathcal{F}, \mathcal{L}) \simeq \mathrm{Hom}_K(\mathcal{F}, \phi^*\mathcal{L}) \quad \forall \mathcal{F} \in \mathbf{Sh}(G),\ \mathcal{L} \in \mathbf{Sh}(K), \tag{1.17}$$
and similarly for ϕ_*.
(3) As a consequence we have
$$\mathrm{Ext}^i_G(\phi_!\mathcal{F}, \mathcal{L}) \simeq \mathrm{Ext}^i_K(\mathcal{F}, \phi^*\mathcal{L}) \quad \forall \mathcal{F} \in \mathbf{Sh}(G),\ \mathcal{L} \in \mathbf{Sh}(K), \tag{1.18}$$
and similarly for ϕ_*.
(4) If $u\colon G' \to G$ is an inclusion of graphs, then $u_!\mathcal{F}$ is just $\mathcal{F}_{G'}$, i.e., the sheaf that is zero outside G' and $\mathcal{F}$ when restricted to G'.
(5) Any sheaf, $\mathcal{F}$, over G has an injective resolution
$$\left(\bigoplus_{v\in V_G} (k_v)_*\mathcal{F}(v) \oplus \bigoplus_{e\in E_G} (k_e)_*\mathcal{F}(e)\right) \longrightarrow \left(\bigoplus_{e\in E_G} (k_e)_*\big(\mathcal{F}(te) \oplus \mathcal{F}(he)\big)\right)$$
where for $P \in V_G \amalg E_G$, k_P denotes the morphism from the category, Δ_0, of one object and one (identity) morphism, to $\mathrm{Cat}(G)$ sending the object of Δ_0 to P. In our case, this means that for a vector space, W,

we have $(k_P)_*W$ has the value $W^{d(Q)}$ at Q, where $d(Q)$ is the number of morphisms from Q to P. For $\mathcal{F} = \underline{\mathbb{F}}$ this is homotopy equivalent to a simpler resolution, namely

$$\underline{\mathbb{F}} \to \bigoplus_{v\in V_G} (k_v)_*\mathbb{F} \to \bigoplus_{e\in E_G} (k_e)_*\mathbb{F} \tag{1.19}$$

(see the paragraph about greedy resolutions and "rank" order in Section 2.11 of [**Fri05**]).

(6) Similarly, any sheaf, $\mathcal{F}$, over G has a projective resolution

$$\left(\bigoplus_{e\in E_G} \big((k_{te})_!\mathcal{F}(e) \oplus (k_{he})_!\mathcal{F}(e)\big)\right) \longrightarrow \left(\bigoplus_{v\in V_G} (k_v)_!\mathcal{F}(v) \oplus \bigoplus_{e\in E_G} (k_e)_!\mathcal{F}(e)\right)$$

Again, $\underline{\mathbb{F}}$ (and numerous other sheaves encountered in practice) have a simpler ("rank" order) resolution:

$$\bigoplus_{v\in V_G} (k_v)_!\mathbb{F}^{d_v-1} \to \bigoplus_{e\in E_G} (k_e)_!\mathbb{F} \to \underline{\mathbb{F}}, \tag{1.20}$$

where d_v is the degree of v (the sum of the indegree and outdegree), and the $d_v - 1$ represents the fact that $\mathbb{F}^{d_v-1}$ is really the kernel of the map $\mathbb{F}^{d_v} \to \mathbb{F}$ which is addition of coordinates; similarly, in equation (1.19), the $\mathbb{F}$ in $(k_e)_!\mathbb{F}$ is really the cokernel of the diagonal inclusion $\mathbb{F} \to \mathbb{F}^2$, with the 2 in $\mathbb{F}^2$ coming from the fact that each edge is incident upon two vertices.

(7) This means that the derived functors, $\mathrm{Ext}^i(\mathcal{F}_1, \mathcal{F}_2)$, of $\mathrm{Hom}(\mathcal{F}_1, \mathcal{F}_2)$ can be computed as the cohomology groups of

$$\bigoplus_{v\in V_G} \mathrm{Hom}\big(\mathcal{F}_1(v), \mathcal{F}_2(v)\big) \oplus \bigoplus_{e\in E_G} \mathrm{Hom}\big(\mathcal{F}_1(e), \mathcal{F}_2(e)\big)$$
$$\longrightarrow \bigoplus_{e\in E_G} \mathrm{Hom}\big(\mathcal{F}_1(e), \mathcal{F}_2(te) \oplus \mathcal{F}_2(he)\big)$$

Now we can understand our choice of homology groups. From equations (1.19) and (1.20), we see that the constant sheaf, $\underline{\mathbb{F}}$, has a simple injective resolution but a more awkward projective resolution. So the homology theory that we've defined earlier amounts to

$$H_i(\mathcal{F}) = \big(\mathrm{Ext}^i(\mathcal{F}, \underline{\mathbb{F}})\big)^\vee,$$

where $^\vee$ denotes the dual space; we have

$$h_0(\mathcal{F}) - h_1(\mathcal{F}) = \chi(\mathcal{F}) = \dim\big(\mathcal{F}(V)\big) - \dim\big(\mathcal{F}(E)\big).$$

As an alternative, one could study the standard cohomology theory

$$H^i(\mathcal{F}) = \mathrm{Ext}^i(\underline{\mathbb{F}}, \mathcal{F}).$$

But we easily see that

$$\dim\big(H^0(\mathcal{F})\big) - \dim\big(H^1(\mathcal{F})\big) = \dim\big(\mathcal{F}(E)\big) - \sum_{v\in V_G} (d_v - 1)\dim\big(\mathcal{F}(v)\big).$$

This is another avenue to study, but does not seem to capture in a simple way the invariant $\rho = \rho(G)$ of a digraph, G.

We remark that we could reverse the role of open and closed sets in this discussion. Indeed, to any sheaf, $\mathcal{F}$, of finite dimensional $\mathbb{F}$-vector spaces on a finite category, $\mathcal{C}$, we can take the spaces dual to the $\mathcal{F}(P)$ for objects, P, of $\mathcal{C}$, thereby getting a sheaf, $\mathcal{F}^\vee$, defined on $\mathcal{C}^{\mathrm{opp}}$, the category opposite to $\mathcal{C}$ (i.e., the category obtained by reversing the arrows). Taking the opposite category has the effect of exchanging open and closed sets, exchanging projectives and injectives, etc.

Let us briefly explain the name "structure sheaf." Generally speaking, in sheaf theory each topological space or Grothendieck topology comes with a special sheaf called the "structure sheaf" that has several properties. One key property is that the "global sections" of a sheaf, $\mathcal{F}$, should reasonably be interpreted as the sheaf homomorphisms to $\mathcal{F}$ from the structure sheaf. This makes "global cosections," on which our homology theory is based, to be sheaf homomorphisms from $\mathcal{F}$ to the "structure sheaf." Hence we call $\underline{\mathbb{F}}$ the structure sheaf.

1.4.3. $\nu_!$, the left adjoint to ν^*. As mentioned in the previous subsection, if $\nu\colon G' \to G$ is an arbitrary graph morphism, then ν^* has a left adjoint, $\nu_!$. In this subsection we show that $\nu_!$ is the left adjoint to ν_*, based on the general construction given in [**sga72a**]. Although ν^* has a right adjoint, ν_*, for our homology theory it is $\nu_!$ that seems more important.

The general construction of $\nu_!$ is given in [**sga72a**], Exposé I, Proposition 5.1). Alternatively, the reader can simply take the $\nu_!$ that we describe and verify that it satisfies equation (1.17).

According to [**sga72a**], Exposé I, Proposition 5.1, given a sheaf, $\mathcal{F}$, on a graph G, i.e., a presheaf on $\mathrm{Cat}(G)$, the value $\nu_!\mathcal{F}(P)$ for $P \in V_G \amalg E_G$ is determined as follows: form the category I_ν^P whose objects are

$$\{(m,X) \mid X \in V_{G'} \amalg E_{G'},\ m\colon P \to \nu(X) \text{ is a morphism in } \mathrm{Cat}(G)\},$$

with a morphism from (m,X) to (m',X') being a morphism $\mu\colon X \to X'$ in $\mathrm{Cat}(G')$ such that $m' = \nu(\mu)m$; then the projection $(m,X) \mapsto X$ followed by $\mathcal{F}$ gives a contravariant functor from I_ν^P to $\mathbb{F}$-vector spaces, and we take the inductive limit in I_ν^P. It follows that if $e \in E_G$, then I_ν^e is category whose objects are (id_e, e') where e' lies over e, and id_e is the identity at e. It follows that

$$(\nu_!\mathcal{F})(e) = \bigoplus_{e' \in \nu^{-1}(e)} \mathcal{F}(e').$$

If $v \in V_G$, then I_ν^v contains the following:

(1) (id_v, v') for each v' over v;
(2) (μ, e') for every $e' \in E_{G'}$ over an $e \in E_G$ with $he = v$, with μ the morphism from v to e given by the head relation; and
(3) the same with "tail" replacing "heads."

We claim that each object (μ, e') has a unique morphism in I_ν^v to an element (id_v, v'), where $v' = he'$ in part (2) and $v' = te'$ in part (3). So the inductive limit for $(\nu_!\mathcal{F})(v)$ can be restricted to the subcategory of objects in part (1), and we again get a direct sum:

$$(\nu_!\mathcal{F})(v) = \bigoplus_{v' \in \nu^{-1}(v)} \mathcal{F}(v').$$

We leave it to the reader to verify that the restriction maps of $\nu_!\mathcal{F}$ are just the natural maps induced by $\mathcal{F}$.

Now we see that

$$(\nu_!\mathcal{F})(V_G) \simeq \mathcal{F}(V_{G'}), \quad (\nu_!\mathcal{F})(E_G) \simeq \mathcal{F}(E_{G'}),$$

with $d_{\nu_!\mathcal{F}}$ and $d_{\mathcal{F}}$ identified under the isomorphism. Hence they have the same homology groups, same adjacency matrix, etc. The main difference is that one is a sheaf on G, the other a sheaf on G'.

1.4.4. $\nu_!$ and Contagious Vanishing Theorems. In this section, we comment that vanishing of homology groups of a sheaf implies the vanishing certain homology groups of related sheaves. We call such results "contagious vanishing" theorems. This gives a nice use of the sheaves $\nu_!\underline{\mathbb{F}}$. Let us first explain our interest in such results, as motivated by the SHNC.

As mentioned before, we will show that the SHNC is implied by the vanishing maximum excess of a sheaf that we call a ρ-kernel. The ρ-kernel actually arises when considering a trivial and very special case of the SHNC; however it turns out that the vanishing of the maximum excess these ρ-kernels actually imply the entire SHNC. What happens is that the trivial case of the SHNC, when expressed as a short/long exact sequence, can be "tensored" with sheaves of the form $\nu_!\underline{\mathbb{F}}$; then a general "contagious vanishing theorem" implies that the maximum excess of the ρ-kernel tensored with $\nu_!\underline{\mathbb{F}}$ vanishes; this proves *all cases* of the SHNC. In other words, the vanishing of a homology group of a sheaf or of a related group can be more powerful than it first seems. Let us describe the underlying ideas, which are not specific to the SHNC.

Let $G' \subset G$ be digraphs, and let G be a sheaf on $\mathcal{F}$. Then we have an exact sequence
$$0 \to \mathcal{F}_{G'} \to \mathcal{F} \to \mathcal{F}/\mathcal{F}_{G'} \to 0.$$
Of course, when G has no self-loops, then this is a special case of the general short exact sequence
$$0 \to \mathcal{F}_U \to \mathcal{F} \to \mathcal{F}_Z \to 0,$$
where $\mathcal{F}$ is a sheaf on a topological space, U is an open subset, and Z is the closed complement (see [**Har77**], Chapter II, Exercise 1.19 or Chapter III, proof of Theorem 2.7). The long exact sequence implies that if $h_1(\mathcal{F}) = 0$, then $h_1(\mathcal{F}_{G'}) = 0$. Of course, the same is true of any first quasi-Betti number, and so we have the following simple but useful theorem.

THEOREM 1.14. *If α_1 is any first quasi-Betti number for sheaves of $\mathbb{F}$-vector spaces on a graph, G, and if $\alpha_1(\mathcal{F}) = 0$ for such a sheaf, $\mathcal{F}$, then for any subgraph, G', of G we have $\alpha_1(\mathcal{F}_{G'}) = 0$.*

The intuition is clear in case α_1 is h_1 or h_1^{twist} or the maximum excess, and $\mathcal{F} = \underline{\mathbb{F}}$: passing to a subgraph cannot increase the first Betti number or the reduced cyclicity of a graph.

One way in which a sheaf $\mathcal{F}_{G'}$ can naturally arise is when we take a short exact sequence of sheaves in G,
$$0 \to \mathcal{F}_1 \to \mathcal{F}_2 \to \mathcal{F}_3 \to 0,$$
and take the tensor product with $\underline{\mathbb{F}}_{G'}$; the tensor product preserves exactness (i.e., all higher Tor groups vanish in sheaves of vector spaces over graphs), so we get a new short exact sequence
$$0 \to \mathcal{F}_1 \otimes \underline{\mathbb{F}}_{G'} \to \mathcal{F}_2 \otimes \underline{\mathbb{F}}_{G'} \to \mathcal{F}_3 \otimes \underline{\mathbb{F}}_{G'} \to 0;$$
now note that for any sheaf, $\mathcal{F}$, on G we have
$$\mathcal{F} \otimes \underline{\mathbb{F}}_{G'} = \mathcal{F}_{G'}.$$

As a consequence, if one has an exact sequence of sheaves on G,

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

and one expects that $\text{m.e.}(\mathcal{F}_2) \le \text{m.e.}(\mathcal{F}_3)$, then a simple homological explanation for this inequality would be that $\text{m.e.}(\mathcal{F}_1) = 0$. But this would, in turn, imply that $\text{m.e.}((\mathcal{F}_2)_{G'}) \le \text{m.e.}((\mathcal{F}_3)_{G'})$ for all open subsets, G', of G, which could be a much stronger inequality (and is much stronger for the setting of the SHNC).

Let us state a slightly stronger "contagious vanishing" theorem that we shall apply to the maximum excess.

Definition 1.15. *By a* scaling *first quasi-Betti number, α_1, we mean a rule that, for some field, $\mathbb{F}$, and* any *digraph, G, assigns a non-negative real number to each sheaf of $\mathbb{F}$-vector spaces over G, such that*

1. *α_1 is a first quasi-Betti number when restricted to sheaves on G for any digraph, G;*
2. *for any covering map $\phi\colon K \to G$ of digraphs and any sheaf, $\mathcal{F}$, on G we have*
$$\alpha_1(\phi^*\mathcal{F}) = \alpha_1(\mathcal{F}) \deg(\phi);$$
and
3. *for an étale $\phi\colon K \to G$ and any sheaf, $\mathcal{F}$, on K we have*
$$\alpha_1(\phi_!\mathcal{F}) = \alpha_1(\mathcal{F}).$$

By the end of this chapter we will know that the maximum excess is a scaling first quasi-Betti number: condition (2) follows from Theorem 1.27; conditions (1) and (3) follow from Theorem 1.10 by taking limits; since condition (1) is almost immediate, we prove only condition (3).

For arbitrary $\phi\colon K \to G$, and arbitrary $\mu\colon G' \to G$, let $K' = G' \times_G K$, and let $\mu'\colon K' \to K$ and $\phi'\colon K' \to G'$ be the projections. We easily see (on each vertex and edge of G) a natural isomorphism

$$(\phi')_!(\mu')^*\mathcal{F} \simeq \mu^*\phi_!\mathcal{F}. \tag{1.21}$$

Using equation (1.15) we have

$$h_1\bigl((\mu')^*\mathcal{F}\bigr) = h_1\bigl((\phi')_!(\mu')^*\mathcal{F}\bigr) = h_1(\mu^*\phi_!\mathcal{F}).$$

Now we take $\mu\colon G' \to G$ to be a covering map; then $\mu'\colon K' \to K$ is a covering map of the same degree as μ; since ϕ is étale, so is $\phi'\colon K' \to G'$, and hence the girth of K' is at least that of G' (since any closed, non-backtracking walk on K' pushes down, via ϕ', to one on G' of equal length). Hence if the girth of G' is sufficiently large we have

$$\text{m.e.}(\mathcal{F}) = h_1\bigl((\mu')^*\mathcal{F}\bigr)/\deg(\mu') = h_1(\mu^*\phi_!\mathcal{F})/\deg(\mu) = \text{m.e.}(\phi_!\mathcal{F}).$$

Theorem 1.16. *Let α_1 be a scaling first quasi-Betti number. If $\alpha_1(\mathcal{F}) = 0$ for a sheaf, $\mathcal{F}$, on a digraph, G, and $\nu\colon G' \to G$ is étale, then $\alpha_1(\mathcal{F}_{G'}) = 0$ where $\mathcal{F}_{G'} = \nu_!\nu^*\mathcal{F} \simeq \mathcal{F} \otimes \underline{\mathbb{F}}_{G'}$.*

Proof. Since ν is étale, it factors as an open inclusion $j\colon G' \to G''$ followed by a covering map $\mu\colon G'' \to G$. Since α_1 scales, we have $\alpha_1(\mathcal{F}) = 0$ implies that $\alpha_1(\phi^*\mathcal{F})$. which implies $\alpha_1(\mathcal{F}') = 0$, where

$$\mathcal{F}' = (\mu^*\mathcal{F})_{G'} = j_!j^*\mu^*\mathcal{F}$$

by Theorem 1.14. Hence

$$\alpha_1(\mu_!\mathcal{F}') = \alpha_1(\mathcal{F}') = 0.$$

But

$$\mu_!\mathcal{F}' = \mu_! j_! j^* \mu^* \mathcal{F} \simeq \nu_! \nu^* \mathcal{F} = \mathcal{F}_{G'},$$

so

$$\alpha_1(\mathcal{F}_{G'}) = \alpha_1(\mu_!\mathcal{F}') = 0.$$

$\square$

Note that if $\phi\colon K \to G$ is not étale, then for a sheaf, $\mathcal{F}$, on K, the maximum excess of $\mathcal{F}$ and $\phi_!\mathcal{F}$ need not agree. For example, consider $\phi\colon B_2 \to B_1$, the unique morphism of digraphs, where B_i is the graph with one vertex and i self-loops, and $\mathcal{F} = \mathcal{U}$, the unhappy 4-bundle of Subsection 1.2.4. Then m.e.$(\mathcal{U}) = 0$ but m.e.$(\phi_!\mathcal{U}) = 1$, as the head/tail neighbourhood of the span of $\beta - \gamma$ is two dimensional. Moreover, notice how our proof that m.e.$(\mathcal{F}) =$ m.e.$(\phi_!\mathcal{F})$ for ϕ étale would fail to work for arbitrary ϕ: for arbitrary ϕ, not necessarily étale, we would not be able to assert that K' has large girth; in the example in this paragraph, ϕ takes two edges to one, and as a result K' always has girth at most two (and Abelian girth at most eight).

We finish with a remark that may be useful when generalizing sheaves to discrete structures beyond graphs. Equation (1.21) is known as a "base change" morphism. In more general contexts, there is usually a natural "base change" morphism

$$\mu^*\phi_!\mathcal{F} \to (\phi')_!(\mu')^*\mathcal{F}$$

that is not generally injective or surjective, not even for presheaves of vector spaces on finite categories; see [**Fri05**]. However if $\phi\colon K \to G$ is any digraph morphism, then ϕ determines a functor $\Phi\colon \mathrm{Cat}(K) \to \mathrm{Cat}(G)$ of the associated categories, and this functor is always "target liftable," i.e., for any morphism, α, of $\mathrm{Cat}(G)$ and object, P, in $\mathrm{Cat}(K)$ with $\Phi(P)$ being the target of α, there is a morphism, α', for $\mathrm{Cat}(K)$ whose target is P and with $\Phi(\alpha') = \alpha$. In [**Fri05**] we see that the "target liftable" property for ϕ (or, more precisely, Φ) guarantees the isomorphism in the base change morphism (actually [**Fri05**] speaks of the dual morphism $\mu^*\phi_*\mathcal{F} \to (\phi')_*(\mu')^*\mathcal{F}$ and "source liftable" in the discussion after Theorem 10.2 there). So if we, for example, spoke of graphs without requiring the edges to have both endpoints, then the simple example of [**Fri05**] shows that equation (1.21) would fail. Of course, if K is a subgraph obtained from G by deleting any number of edges, then the inclusion ϕ (or, more precisely, associated Φ) is not source liftable. Hence "target liftability" (and not "source liftability") should guide us in generalizing sheaves on graphs to more general discrete structures, if we wish to have similar theorems about analogues of the maximum excess being invariant under $\phi_!$ for étale ϕ.

1.5. Twisted Cohomology

In this section we describe a number of aspects of twisted homology, and give its relationship to the homology of pullbacks under Abelian covers. We show that the first twisted Betti number of the structure sheaf of a graph, G, agrees with $\rho(G)$. We then prove a number of related results, such as giving a condition under which the maximum excess agrees with the first Betti number.

1.5.1. Remarks on the Definition. Twists and twisted homology were defined in Subsection 1.2.3. In this subsection we make a few remarks on the definitions.

In our definition of twists, for symmetry we could have also specified a multiplier (like $\psi(e)$) for $\mathcal{F}^\psi(h,e)$, not just $\mathcal{F}^\psi(t,e)$; i.e., we could have defined a twists to be a map $E_G \times \{t,h\} \to \mathbb{F}'$. But there is no real need for a $\mathcal{F}^\psi(h,e)$ multiplier, since all twisted homology groups would be isomorphic.

Note that $h_i^{\mathrm{twist}}(\mathcal{F})$ could be alternatively described as the "generic dimension of $h_i(\mathcal{F}^\psi)$;" more precisely, there is a polynomial, f, in $\{\psi_e\}$ over $\mathbb{F}$ such that the dimension of $h_i(\mathcal{F}^\psi)$ for any fixed twist, ψ, with $\psi_e \in \mathbb{F}$, is $h_i^{\mathrm{twist}}(\mathcal{F})$ provided that $f(\psi) \neq 0$. Furthermore, for any particular $\psi \in \mathbb{F}^{E_G}$, the dimension of $h_i(\mathcal{F}^\psi)$ is at least the generic dimension. All these facts follow from the fact that the rank of a matrix is the size of the largest square submatrix whose determinant does not vanish. This discussion assumes either that $\mathbb{F}$ is infinite or that $\mathbb{F}$ is considered as embedded in an infinite or sufficiently large extension field of itself (it is not clear how to give an interesting meaning to "generic" when dealing with finite dimensional spaces over finite fields). Of course, the advantange of our original definition, which involves $\mathbb{F}(\psi)$ with the ψ being indeterminates, is that it gives a simple, usable definition for arbitrary $\mathbb{F}$, even when $\mathbb{F}$ is finite.

1.5.2. Twists and Abelian Coverings. We now wish to describe twisting as giving the homology of pullbacks under Abelian coverings. Given an Abelian group, $\mathcal{A}$, say that a field, $\mathbb{F}$, is a *Fourier field for* $\mathcal{A}$ if $\mathbb{F}$ contains $n = |\mathcal{A}|$ distinct n-th roots of 1 (which holds, for example, when the characteristic of $\mathbb{F}$ is relatively prime to n and $\mathbb{F}$ is algebraically closed). In this case, if $\mathcal{A}$, acts on a vector space, S, over a field, $\mathbb{F}$, then we have a canonical isomorphism

$$\bigoplus_\nu S^\nu \simeq S,$$

where $\nu\colon \mathcal{A} \to \mathbb{F}$ ranges over all characters on $\mathcal{A}$ and

$$S^\nu = \{s \in S \mid as = \nu(a)s \text{ for all } a \in \mathcal{A}\};$$

indeed, for each ν we have $S^\nu \subset S$, and these inclusions give a map from the direct sum of the S^ν to S; the inverse map, from S to the direct sum of the S^ν, is given as the sum of the maps from S to any particular S^ν via

$$s \mapsto (1/n) \sum_{\alpha \in \mathcal{A}} \nu^{-1}(\alpha)(\alpha s); \tag{1.22}$$

the values $1/n$ and $\nu^{-1}(\alpha)$ all lie in $\mathbb{F}$ for any $\mathbb{F}$ that is a Galois field for $\mathcal{A}$.

LEMMA 1.17. *Let $\phi\colon G' \to G$ be an Abelian covering map with Galois group $\mathcal{A}$. Let $\mathcal{F}$ be a sheaf of $\mathbb{F}$-vector spaces on G such that $\mathbb{F}$ is Fourier field for $\mathcal{A}$. Then*

$$H_i(\phi^*\mathcal{F}) \simeq \bigoplus_\psi \bigl(H_i(\phi^*\mathcal{F})\bigr)^\nu, \tag{1.23}$$

the sum is over all characters, ν, of $\mathcal{A}$. Let $\vec{a} = \{a_e\}_{e \in E_G}$ be any Galois coordinates for $\phi\colon G' \to G$, and for any character, ν, of $\mathcal{A}$, let $\nu(\vec{a})$ denote the $\mathbb{F}$-twist taking $e \in E_G$ to $\nu(a_e)$. Then for each ν we have

$$\bigl(H_i(\phi^*\mathcal{F})\bigr)^\nu \simeq H_i\bigl(\mathcal{F}^{\nu(\vec{a})}\bigr).$$

PROOF. We have an $\mathcal{A}$ action on $(\phi^*\mathcal{F})(E_{G'})$ via

$$(af)(e) = f(ea)$$

for all $a \in \mathcal{A}$, $f \in (\phi^*\mathcal{F})(E_{G'})$, and $e \in E_{G'}$. Similarly $(af)(v) = f(va)$ defines an $\mathcal{A}$ action on $(\phi^*\mathcal{F})(V_{G'})$. The map in equation (1.22) gives isomorphisms

$$(\phi^*\mathcal{F})(E) \to \bigoplus_\nu \big((\phi^*\mathcal{F})(E)\big)^\nu, \quad (\phi^*\mathcal{F})(V) \to \bigoplus_\nu \big((\phi^*\mathcal{F})(V)\big)^\nu,$$

and $d_{\phi^*\mathcal{F}}$ intertwines with these maps, which establishes equation (1.23). It remains to identify

$$\big(H_i(\phi^*\mathcal{F})\big)^\nu$$

with H_i of the appropriately twisted $\mathcal{F}$. So choose Galois coordinates, $\{a_e\}$, and therefore identify $V_{G'}$ with $V_G \times \mathcal{A}$ and $E_{G'}$ with $E_G \times \mathcal{A}$ so that

$$h(e,a) = (he, a_e a) \quad \text{and} \quad t(e,a) = (te, a)$$

(as in Subsection 1.3.2). Given an $f \in (\phi^*\mathcal{F})(E)$, define $\widetilde{f} \in \mathcal{F}(E)$ via

$$\widetilde{f}(e) = f(e, \mathrm{id}_\mathcal{A}),$$

where $\mathrm{id}_\mathcal{A}$ is the identity of $\mathcal{A}$ and we identify $E_{G'}$ with $E_G \times \mathcal{A}$ as above. Similarly define a linear map $f \mapsto \widetilde{f}$ from $(\phi^*\mathcal{F})(V)$ to $\mathcal{F}(V)$. Now consider

$$f \in \big(H_1(\phi^*\mathcal{F})\big)^\nu.$$

For all $v' \in V_{G'}$ we have

$$\sum_{e' \text{ s.t. } te'=v'} f(e') = \sum_{e' \text{ s.t. } he'=v'} f(e').$$

Taking $v' = (v, \mathrm{id}_\mathcal{A})$ yields

$$\sum_{te=v} f(e, \mathrm{id}_\mathcal{A}) = \sum_{he=v} f(e, a_e^{-1}) = \sum_{he=v} (a_e^{-1} f)(e, \mathrm{id}_\mathcal{A}),$$

which, since $f \in (H_1(\phi^*\mathcal{F}))^\nu$,

$$= \sum_{he=v} \nu(a_e^{-1}) f(e, \mathrm{id}_\mathcal{A}).$$

It follows that

$$\sum_{te=v} \widetilde{f}(e) = \sum_{he=v} \nu(a_e^{-1}) \widetilde{f}(e).$$

In other words, if we set $f'(e) = \nu(a_e^{-1})\widetilde{f}(e)$, then we have

$$\sum_{te=v} \nu(a_e) f'(e) = \sum_{he=v} f'(e).$$

Hence $f' \in H_1(\mathcal{F}^{\nu(\vec{a})})$. Clearly given f' we can reconstruct $\widetilde{f}$ and then f, namely

$$f(e,a) = \nu(a_e)\nu(a) f'(e).$$

Hence $f \mapsto f'$ is an isomorphism

$$\big((\phi^*\mathcal{F})(E_{G'})\big)^\nu \to \mathcal{F}^{\nu(\vec{a})}(E_G).$$

Furthermore we have an analogous map $f \mapsto \widetilde{f}$

$$\big((\phi^*\mathcal{F})(V_{G'})\big)^\nu \to \mathcal{F}^{\nu(\vec{a})}(V_G),$$

namely, $\widetilde{f}(v) = f(v, \mathrm{id}_{\mathcal{A}})$, which likewise is an isomorphism. Hence we get a commutative diagram:

$$\begin{tikzcd}
\big((\phi^*\mathcal{F})(E_{G'})\big)^\nu \arrow[r, "f \mapsto f'"] \arrow[d, "d_{\phi^*\mathcal{F}}"'] & \mathcal{F}^{\nu(\vec{a})}(E_G) \arrow[d, "d_{\mathcal{F}^{\nu(\vec{a})}}"] \\
\big((\phi^*\mathcal{F})(V_{G'})\big)^\nu \arrow[r, "f \mapsto \widetilde{f}"] & \mathcal{F}^{\nu(\vec{a})}(V_G)
\end{tikzcd}$$

Since the horizontal arrows are isomorphisms, this diagram sets up isomorphisms between the kernel and cokernel of the vertical arrows. Hence for $i = 0, 1$ we have

$$\big(H_i(\phi^*\mathcal{F})\big)^\nu \simeq H_i(\mathcal{F}^{\nu(\vec{a})})$$

□

Lemma 1.17 shows that if $\mathbb{F}$ is any infinite field, $\mathcal{F}$ is any sheaf on a digraph, G, and we take a random $\mathbb{Z}/p\mathbb{Z}$ cover $\mu\colon G' \to G$, then we have that $h_i(\mu^*\mathcal{F})/p$ tends to $h_i^{\mathrm{twist}}(\mathcal{F})$ in probability as $p \to \infty$.

Lemma 1.17 also shows that if $\mu\colon G' \to G$ is an Abelian cover with covering group $\mathcal{A}$, then $H_i(\mu^*\mathcal{F})$ is the sum of $|\mathcal{A}|$ groups, each isomorphic to an $H_i(\mathcal{F}^\psi)$ for a particular value of ψ, and hence of dimension at least $h_i^{\mathrm{twist}}(\mathcal{F})$. We conclude the following lemma.

Lemma 1.18. *If $\mu\colon G' \to G$ is any Abelian cover of G, and $\mathcal{F}$ is any sheaf on G, then*

$$h_i(\mu^*\mathcal{F}) \geq \deg(\mu) h_i^{\mathrm{twist}}(\mathcal{F}).$$

This can be viewed as an upper bound for $h_i^{\mathrm{twist}}(\mathcal{F})$. Now we note the trivial lower bound

$$h_1^{\mathrm{twist}}(\mathcal{F}) \geq -\chi(\mathcal{F}),$$

since $h_1^{\mathrm{twist}}(\mathcal{F})$ is the kernel of a matrix whose dimension of domain minus that of codomain is $-\chi(\mathcal{F})$.

If G is any connected digraph, then for any prime, p, we claim that G has an Abelian cover of degree p that is connected; indeed, just take the monodromy map to map any generator of $\pi_1(G)$ to $1 \in \mathbb{Z}/p\mathbb{Z}$ and use the remark at the end of Subsection 1.3.3. In this case we have $h_1(G') = 1-\chi(G') = 1-p\chi(G) = p\rho(G)+1$. But by Lemma 1.18 with $\mathcal{F} = \underline{\mathbb{F}}$ (so that $\mu^*\mathcal{F} = \underline{\mathbb{F}}$ on G') we have

$$h_1^{\mathrm{twist}}(\underline{\mathbb{F}}) \leq h_1(G', \underline{\mathbb{F}})/p = h_1(G')/p = \rho(G) + (1/p).$$

Letting $p \to \infty$ we conclude $h_1^{\mathrm{twist}}(\underline{\mathbb{F}}) \leq \rho(G)$. But the "trivial lower bound" gives

$$h_1^{\mathrm{twist}}(\underline{\mathbb{F}}) \geq -\chi(\underline{\mathbb{F}}) = \rho(G).$$

If G is not connected then we apply the above to each of its connected components and conclude the following theorem.

Theorem 1.19. *For any digraph, G, we have $\rho(G) = h_1^{\mathrm{twist}}(\underline{\mathbb{F}})$.*

1.5.3. The Maximum Excess Bound. Let $\mathcal{F}$ be a sheaf of $\mathbb{F}$-vector spaces on a digraph, G, and let $U \subset \mathcal{F}(V)$. Let $\psi = \{\psi(e)\}_{e \in E_G}$ be a twist of indeterminates. Then $d = d_{\mathcal{F}^\psi} \colon \mathcal{F}(E) \to \mathcal{F}(V)$ can be restricted as a morphism

$$\Gamma_{\mathrm{ht}}(U) \otimes_{\mathbb{F}} \mathbb{F}' \to U \otimes_{\mathbb{F}} \mathbb{F}'.$$

By the "trivial bound," the kernel of this morphism has dimension at least

$$\dim\bigl(\Gamma_{\mathrm{ht}}(U)\bigr) - \dim(U) = \mathrm{excess}(\mathcal{F}, U).$$

Hence the kernel of d has at least this dimension. This gives the following simple bound.

Lemma 1.20. *For any sheaf, $\mathcal{F}$, on a digraph, G, we have*

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

We wish to show that this holds with equality in certain cases; Theorem 1.10 says that equality will hold if $\mathcal{F}$ is pulled back appropriately.

Definition 1.21. *If $\mathcal{F}$ is a sheaf on a digraph, G, we say that $\mathcal{F}$ is* edge simple *if $\mathcal{F}(e)$ is of dimension 0 or 1 for each $e \in E_G$.*

Theorem 1.22. *Let $\mathbb{F}$ be an infinite field. Let $\mathcal{F}$ be an edge simple sheaf of $\mathbb{F}$-vector spaces on a digraph, G. Then*

$$h_1^{\mathrm{twist}}(\mathcal{F}) = \mathrm{m.e.}(\mathcal{F}).$$

Proof. Let $\{e_1, \ldots, e_r\} \subset E$ be the edges where $\mathcal{F}(e) \neq 0$. Let $\psi = \{\psi_i\}_{i=1,\ldots,r}$ be indeterminates, and let

$$\mathcal{F}(V)(\psi) = \bigl(\mathcal{F}(V)\bigr) \otimes_{\mathbb{F}} \mathbb{F}(\psi).$$

For each e_i choose a $w_i \in \mathcal{F}(e_i)$ with $w_i \neq 0$, and let

$$v_i = a_i + \psi_i b_i \in \mathcal{F}(V)(\psi), \quad \text{with} \quad a_i = \mathcal{F}(h, e_i)(w_i), \quad b_i = \mathcal{F}(t, e_i)(w_i).$$

Say that a v_j is *critical* for $v_1, \ldots, v_r$ if the span of $\{v_i\}_{i \neq j}$ is of dimension one less than $\{v_i\}_{i=1,\ldots,r}$. Let us first prove the lemma assuming that no vector is critical. Let r' be the dimension of the span of the v_i, so $h_1^{\mathrm{twist}}(\mathcal{F}) = r - r'$. In view of Lemma 1.20, suffices to show that

$$\mathrm{m.e.}(\mathcal{F}) \geq r - r'.$$

If $r - r' = 0$ there is nothing to prove. So we may assume $r' < r$.

We wish to show that there exists a $U \subset \mathcal{F}(V)$ such that

$$|\{i \mid a_i, b_i \in U\}| \geq \dim(U) + r - r'.$$

Let us first assume that for any I with $\{v_i\}_{i \in I}$ independent (over $\mathbb{F}(\psi)$) we also have that $\{a_i\}_{i \in I}$ are independent (over $\mathbb{F}$).

By reordering the v_i, we may assume that

$$v_1, v_2, \ldots v_{r'}$$

are linearly independent. Let A be the span of $a_1, \ldots, a_{r'}$. Consider that

$$(a_1 + \psi_1 b_1) \wedge \cdots \wedge (a_{r'+1} + \psi_{r'+1} b_{r'+1}) = 0. \tag{1.24}$$

Considering the constant coefficient (i.e., with no ψ_i's) of this wedge product, we have $a_1 \wedge \cdots \wedge a_{r'+1} = 0$, and therefore $a_{r'+1} \in A$; similarly considering the $\psi_{r'+1}$

coefficient shows that $b_{r'+1} \in A$. Replacing $v_{r'+1}$ with any v_s with $s > r'+1$ shows that

$$b_{r'+1}, \dots, b_r, a_1, \dots, a_r \in A.$$

In other words, we have shown that if U is the span of the $a_1, \dots, a_r$, we have that U is r' dimensional and contains any b_j such that j lies outside a set, I, such that $|I| = r'$ and $\{v_i\}_{i\in I}$ are independent. But no vector, v_i, is critical for $\{v_i\}$; hence for any j there is an I of size r' such that j lies outside I and $\{v_i\}_{i\in I}$ are independent. Hence $b_j \in U$ for any $j = 1, \dots, r$. Hence $\mathrm{excess}(\mathcal{F}, U) \geq r - r'$. This establishes the lemma when no vector, v_i, is critical, and when for all I, $\{v_i\}_{i\in I}$ are independent implies that $\{a_i\}_{i\in I}$ are as well.

Now let us establish the lemma assuming no vector, v_i, is critical but without assuming $\{v_i\}_{i\in I}$ independent implies $\{a_i\}_{i\in I}$ is independent. Note that since $\mathbb{F}$ is infinite, any generic set in $\mathbb{F}^n$ (i.e., complement of the set of zeros of a polynomial) is nonempty. For each I for which $\{v_i\}_{i\in I}$ is independent, we have

$$\bigwedge_{i\in I}(a_i + \psi_i b_i) \neq 0 \quad (\text{in } \Lambda^{|I|}(\mathcal{F}(V) \otimes_{\mathbb{F}} \mathbb{F}(\psi))\ ).$$

So for a generic set, G_I, of $\theta \in \mathbb{F}^r$ we have

$$\bigwedge_{i\in I}(a_i + \theta_i b_i) \neq 0.$$

So choose a $\theta \in \mathbb{F}^r$ in the intersection of all G_I for all I with $\{v_i\}_{i\in I}$ independent. Let $\widetilde{\psi} = \psi + \theta$ (where $\theta \in \mathbb{F}^r$ and ψ is a collection of r indeterminates), and let

$$\widetilde{v}_i = a_i + \widetilde{\psi}_i b_i = \widetilde{a}_i + \psi_i b_i,$$

where $\widetilde{a}_i = a_i + \theta b_i$. We have $\{v_i\}_{i\in I}$ is independent precisely when $\{\widetilde{v}_i\}_{i\in I}$ is, since they differ by a parameter translation, but whenever this holds we also have that the $\{\widetilde{a}_i\}_{i\in I}$ are independent. But we have already proven the lemma in this case, i.e., the case of $\widetilde{v}_i = \widetilde{a}_i + \psi_i b_i$, since each independent subset of $\{\widetilde{v}_i\}$ has the corresponding subset of $\{\widetilde{a}_i\}$ being independent. Hence we can apply the lemma to conclude that there is a subspace U of $\mathcal{F}(V)$ of dimension r', namely the span of the $\widetilde{a}_i$, such that

$$\widetilde{a}_1, \dots, \widetilde{a}_r, b_1, \dots, b_r \in U.$$

But a_i is an $\mathbb{F}$-linear combination of $\widetilde{a}_i$ and b_i, so $\widetilde{a}_i, b_i \in U$ also implies $a_i \in U$. Hence, again, $\mathrm{excess}(\mathcal{F}, U) \geq r - r'$.

Let us finish by proving the lemma in general, i.e., without the assumption that each v_i is critical. Again, let r' be the dimension of the span of $v_1, \dots, v_r$ as above. If some element of $v_1, \dots, v_r$ is critical, we may assume it is v_1; in this case, if some element of $v_2, \dots, v_r$ is critical for that set, we may assume it is v_2; continuing in this fashion, there is an s such that for all $i < s$, v_i is critical for $v_i, \dots, v_r$, and no element of $v_s, \dots, v_r$ is critical for that set. Consider the sheaf $\mathcal{F}'$ which agrees with $\mathcal{F}$ everywhere except that $\mathcal{F}'(e_i) = 0$ for $i < s$ (and so $\mathcal{F}$ and $\mathcal{F}'$ agree at all vertices and all e_i with $i \geq s$). Then $\{v_s, \dots, v_r\}$ is of size $r - s + 1$, but also the span of $\{v_s, \dots, v_r\}$ is of size $r' - s + 1$ (by the criticality of the v_i with $i < s$), and hence $h_1^{\mathrm{twist}}(\mathcal{F}') = r - r'$. But since no element of $v_s, \dots, v_r$ is critical for that set, the lemma holds for the case of $\mathcal{F}'$ (as shown by the end of the previous paragraph). We therefore construct a U such that $\mathrm{excess}(\mathcal{F}', U) \geq r - r'$. Since $\mathcal{F}'(V) \subset \mathcal{F}(V)$, we can view $U \subset \mathcal{F}(V)$ and it is clear that $\Gamma_{\mathrm{ht}}(U)$ in $\mathcal{F}'$ is a subset of $\Gamma_{\mathrm{ht}}(U)$ in $\mathcal{F}$. Hence

$$\mathrm{excess}(\mathcal{F}, U) \geq \mathrm{excess}(\mathcal{F}', U) = r - r'.$$

□

1.6. Maximum Excess and Supermodularity

In this section we prove that pulling back a sheaf via ϕ multiplies the maximum excess by $\deg(\phi)$. To prove this we will prove supermodularity of the excess function, which has a number of important consequences. Before discussing this, we develop some terminology and simple observations about what we call "compartmentalized subspaces;" this development will be used in this section and in Section 1.8. We finish this section with some additional remarks about the maximum excess.

1.6.1. Compartmentalized Subspaces. In this subsection we mention a few important definitions, and some simple theorems we will use regarding these definitions.

Definition 1.23. *Let W be a finite dimensional vector space over a field, $\mathbb{F}$. By a* decomposition *of W we mean an isomorphism a direct sum of vector spaces with W, i.e.,*

$$\pi \colon \bigoplus_{s \in S} W_s \to W.$$

For any $s \in S$ and any $v \in W_s$, let the extension of v of index s by zero, *denoted* extend(v, s), *to be the element of $\oplus_{s \in S} W_s$ that is v on W_s and zero on W_q with $q \neq s$. For $s \in S$ and a subspace $W' \subset W$, let* the portion of W' supported in s *be*

$$\mathrm{supportedIn}(s, W') = \Big\{ v \in W_s \mid \pi\big(\mathrm{extend}(v, s)\big) \in W' \Big\},$$

and let the compartmentalization of W' *be*

$$(W')_{\mathrm{comp}} = \pi \left(\bigoplus_{s \in S} \mathrm{supportedIn}(s, W') \right),$$

which is a subspace of W'. We say that a subspace W' is compartmentalized *if $(W')_{\mathrm{comp}} = W'$. We say that $w_1, \ldots, w_m \in W$ are* compartmentally distinct *if for any $s \in S$ there is at most one j between 1 and m for which the W_s component of w_j is non-zero.*

So $W' \subset W$ as above is compartmentalized iff W' is the image under π of a set of the form

$$\bigoplus_{s \in S} W'_s.$$

The intuitive point of the definition of compartmentalized subspaces is that certain constructions, such as maximum excess, are performed over the direct summands of a vector space; in some such constructions, the compartmentalized subspaces are the subspaces of key interest.

In this section we will use only these definitions. In Section 1.8, we use two simple observations about the situation of Definition 1.23. First, if $w_1, \ldots, w_m$ are compartmentally distinct, then $w_1, \ldots, w_m$ are linearly independent if (and only if) they are each non-zero. Second, $W' \subset W$ is compartmentalized only if (and if) there exist quotients, Q_s, of W_s for $s \in S$ such that π induces an isomorphism

$$\bigoplus_{s\in S} Q_s \to W/W'. \tag{1.25}$$

It will be helpful to formally combine these two observations into a theorem that follows immediately; we will use this theorem repeatedly in Section 1.8, in our proof of Theorem 1.10.

THEOREM 1.24. *Let W be a finite dimensional vector space with a decomposition. Let $w_1, \ldots, w_m$ be compartmentally distinct, and let $W' \subset W$ be a compartmentalized subspace of W. Then the images of $w_1, \ldots, w_m$ in W/W' are linearly independent (in W/W') iff they are nonzero (in W/W').*

Compartmentalization is a key to our definition of maximum excess. Indeed, for a sheaf, $\mathcal{F}$, on a digraph, G, both $\mathcal{F}(V)$ and $\mathcal{F}(E)$ are defined as direct sums, and hence come with natural decompositions. The head/tail neighbourhood is a compartmentalized space by its definition in equation (1.3); this is crucial to the resulting definition of excess and maximum excess, in Definition 1.7. Note that d_h, d_t (but not d in general) are "compartmentalized morphisms" in that they take vectors supported in one component of $\mathcal{F}(E)$ to those supported in one component of $\mathcal{F}(V)$. This means that with our definition of head/tail neighbourhood, for any $U \subset \mathcal{F}(V)$ and any twist, ψ, on G, the twisted differential, $d_{\mathcal{F}^\psi}$ takes $\Gamma_{\mathrm{ht}}(U) \otimes_{\mathbb{F}} \mathbb{F}(\psi)$ to $U \otimes_{\mathbb{F}} \mathbb{F}(\psi)$.

1.6.2. Supermodularity and Its Consequences. First we make some simple remarks on the maximum excess. For any sheaf, $\mathcal{F}$, we have

$$\mathrm{excess}(\mathcal{F}, 0) = 0, \quad \mathrm{excess}(\mathcal{F}, \mathcal{F}(V)) = -\chi(\mathcal{F}),$$

and hence

$$\mathrm{m.e.}(\mathcal{F}) \geq \max\bigl(0, -\chi(\mathcal{F})\bigr).$$

We now show that if U achieves the maximum excess of $\mathcal{F}$, then U must be compartmentalized.

THEOREM 1.25. *Let the maximum excess of a sheaf, $\mathcal{F}$, on a digraph, G, be achieved on a space $U \subset \mathcal{F}(V)$. Then U is compartmentalized with respect to the identification π given by*

$$\pi\colon \bigoplus_{v\in V_G} \mathcal{F}(v) \to \mathcal{F}(V).$$

PROOF. For $e \in E_G$ and $w \in \mathcal{F}(e)$, if we have $d_t w \in U$, then

$$d_t w = \pi\Bigl(\mathrm{extend}\bigl(\mathcal{F}(t,e)w, te\bigr)\Bigr) \in U_{\mathrm{comp}};$$

similarly if $d_h w \in U$, then $d_h w \in U_{\mathrm{comp}}$. Hence, in view of equation (1.3), we have

$$\Gamma_{\mathrm{ht}}(U_{\mathrm{comp}}) = \Gamma_{\mathrm{ht}}(U).$$

Hence, if U_{comp} is a proper subspace of U, then

$$\mathrm{excess}(\mathcal{F}, U_{\mathrm{comp}}) < \mathrm{excess}(\mathcal{F}, U).$$

So if U maximizes the excess, then $U_{\mathrm{comp}} = U$; i.e., U is compartmentalized. $\square$

The main results in this section stem from the following easy theorem.

THEOREM 1.26. *Let $\mathcal{F}$ be a sheaf on a graph, G. Then the excess, as a function of $U \subset \mathcal{F}(V)$, is supermodular, i.e.,*

$$\operatorname{excess}(U_1) + \operatorname{excess}(U_2) \leq \operatorname{excess}(U_1 \cap U_2) + \operatorname{excess}(U_1 + U_2) \tag{1.26}$$

for all $U_1, U_2 \subset \mathcal{F}(V)$. It follows that the maximizers of the excess function of $\mathcal{F}$,

$$\operatorname{maximizers}(\mathcal{F}) = \{U \subset \mathcal{F}(V) \mid \operatorname{excess}(U) = \text{m.e.}(\mathcal{F})\},$$

is a sublattice of the set of subsets of $\mathcal{F}(V)$, i.e., is closed under intersection and sum (and therefore has a unique maximal element and a unique minimal element). Finally, if U_1, U_2 are maximizers of the excess function of $\mathcal{F}$, then

$$\Gamma_{\rm ht}(U_1 + U_2) = \Gamma_{\rm ht}(U_1) + \Gamma_{\rm ht}(U_2).$$

PROOF. We use the fact that if A_1, A_2 are any subspaces of an $\mathbb{F}$-vector space, then

$$\dim(A_1) + \dim(A_2) = \dim(A_1 \cap A_2) + \dim(A_1 + A_2).$$

In particular, for $U_1, U_2 \subset \mathcal{F}(V)$ we have

$$\dim(U_1) + \dim(U_2) = \dim(U_1 \cap U_2) + \dim(U_1 + U_2). \tag{1.27}$$

On the other hand

$$\Gamma_{ht}(U_1 \cap U_2) = \Gamma_{ht}(U_1) \cap \Gamma_{ht}(U_2)$$

and

$$\Gamma_{ht}(U_1 + U_2) \supset \Gamma_{ht}(U_1) + \Gamma_{ht}(U_2); \tag{1.28}$$

hence

$$\dim\bigl(\Gamma_{ht}(U_1)\bigr) + \dim\bigl(\Gamma_{ht}(U_2)\bigr) \leq \dim\bigl(\Gamma_{ht}(U_1 \cap U_2)\bigr) + \dim\bigl(\Gamma_{ht}(U_1 + U_2)\bigr). \tag{1.29}$$

Combining equations (1.27) and (1.29) yields equation (1.26). It follows that if U_1 and U_2 are maximizers of the excess function of $\mathcal{F}$, then so are $U_1 \cap U_2$ and $U_1 + U_2$, and equations (1.29) and hence (1.28) must hold with equality. $\square$

The supermodularity has a number of important consequences. We list two such theorem below.

THEOREM 1.27. *Let $\phi \colon G' \to G$ be a covering map of graphs, and let $\mathcal{F}$ be a sheaf on G. Then*

$$\text{m.e.}(\phi^*\mathcal{F}) = \deg(\phi)\ \text{m.e.}(\mathcal{F}). \tag{1.30}$$

Furthermore, if the maximum excess of $\mathcal{F}$ is achieved at $U \subset \mathcal{F}(V_G)$, then the maximum excess of $\phi^*\mathcal{F}$ is achieved at $\phi^{-1}(U)$.

PROOF. Our proof uses Theorem 1.26 and Galois theory. Let $\mathcal{F}' = \phi^*\mathcal{F}$. If $T \subset \mathcal{F}(V)$ is compartmentalized, $T = \oplus_{v \in V_G} T_v$, let

$$\phi^{-1}(T) = \bigoplus_{v' \in V_{G'}} T_{\phi(v')} \subset \mathcal{F}'(V_{G'}).$$

Since ϕ is a covering map, the number of preimages of any element of $V_G \amalg E_G$ is $\deg(\phi)$, and hence

$$\operatorname{excess}\bigl(\mathcal{F}', \phi^{-1}(T)\bigr) = \deg(\phi) \operatorname{excess}(\mathcal{F}, T). \tag{1.31}$$

Taking T to maximize the excess of $\mathcal{F}$ we get

$$\text{m.e.}(\mathcal{F}') \geq \deg(\phi)\, \text{m.e.}(\mathcal{F}). \tag{1.32}$$

It remains to prove the reverse inequality in order to establish equation (1.30); note that if we do so, then the second statement of the theorem follows from equation (1.31).

First let us assume that ϕ is Galois, with Galois group $\mathrm{Gal}(\phi)$. Each $g \in \mathrm{Gal}(\phi)$ is a morphism $g\colon K \to K$. Let $\mathcal{F}' = \phi^*\mathcal{F}$. There is a natural map $\iota_g\colon g^*\mathcal{F}' \to \mathcal{F}'$, since for every $P \in V_{G'} \amalg E_{G'}$ we have $\mathcal{F}'(P) = \mathcal{F}'(Pg)$ (note that this really is equality of vector spaces; they both equal $\mathcal{F}(\phi(P))$, by definition). So ι_g gives automorphism on $\mathcal{F}'(E_{G'})$ and $\mathcal{F}'(V_{G'})$. For any $U \subset (\phi^*\mathcal{F})(V)$, any element of $\mathrm{Gal}(\phi)$ preserves $\dim(U)$ and $\dim(\Gamma_{ht}(U))$, and hence the excess. It follows that for all $g \in \mathrm{Gal}(\phi)$, ι_g takes maximizers$(\phi^*\mathcal{F})$ to itself. Hence if W is the unique maximal element of the maximizers, then W is invariant under ι_g for all $g \in \mathrm{Gal}(\phi)$; this means that if $W = \oplus_{v' \in V(G')} W_{v'}$ and

$$\widetilde{W} = \bigoplus_{v \in V_G} \left(\sum_{v' \in \phi^{-1}(v)} W_{v'} \right),$$

then ($W_{v'} = W_{v''}$ if $\phi(v') = \phi(v'')$ and) $W = \phi^{-1}(\widetilde{W})$. Hence

$$\begin{aligned} \text{m.e.}(\mathcal{F}') &= \text{excess}(W) \\ &= \deg(\phi)\ \text{excess}_{\mathcal{F}}(\widetilde{W}) \le \deg(\phi)\ \text{m.e.}(\mathcal{F}). \end{aligned}$$

In summary,

$$\text{m.e.}(\mathcal{F}') \le \deg(\phi)\ \text{m.e.}(\mathcal{F}).$$

From equation (1.32), it follows that the above inequality holds with equality.

It remains to prove the equality when $\phi\colon G' \to G$ is not Galois. By the Normal Extension Theorem of Galois graph theory (i.e., Theorem 1.11), there exists a $\nu\colon L \to G'$ be such that $\phi\nu$ (and hence ν) is Galois. Since $\phi\nu$ is Galois, we have

$$\text{m.e.}(\nu^*\phi^*\mathcal{F}) = \deg(\phi\nu)\,\text{m.e.}(\mathcal{F}),$$

and since ν is Galois we have

$$\text{m.e.}(\nu^*(\phi^*\mathcal{F})) = \deg(\nu)\,\text{m.e.}(\phi^*\mathcal{F}).$$

It follows that

$$\text{m.e.}(\phi^*\mathcal{F}) = \deg(\phi)\,\text{m.e.}(\mathcal{F}).$$

$\square$

1.6.3. Additional Remarks on the Maximum Excess. Here we make some additional remarks on the maximum excess, either for later use or to provide some more intuition about it.

We mention that $\text{m.e.}(\mathcal{F}) + \chi(\mathcal{F})$ can be viewed as a generalization of the "number of acyclic components" of a graph; for example, for the sheaf $\underline{\mathbb{F}}$ on G we have

$$\text{m.e.}(\underline{\mathbb{F}}) + \chi(\underline{\mathbb{F}}) = \rho(G) + |V_G| - |E_G| = h_0^{\text{acyclic}}(G)$$

equals the number of "acyclic components" of G, i.e. the number of connected components of G that have no cycles, i.e., that are isolated vertices or trees. A similar remark holds for $\underline{\mathbb{F}}$ replaced by $\underline{\mathbb{F}}_K$ and G replaced by K, for any map $K \to G$.

We shall make use of the following alternate interpretation of the maximum excess.

THEOREM 1.28. *For any sheaf, $\mathcal{F}$, on a digraph, G, the maximum excess of $\mathcal{F}$ is the same as*

$$\max_{\mathcal{F}' \subset \mathcal{F}} -\chi(\mathcal{F}'),$$

i.e., the maximum value of minus the Euler characteristic over all subsheaves, $\mathcal{F}'$, of $\mathcal{F}$.

PROOF. Each compartmentalized $U \subset \mathcal{F}(V)$ along with $\Gamma_{\mathrm{ht}}(U)$ determines a subsheaf $\mathcal{F}'$ whose Euler characteristic is minus the excess of U. Conversely, for any subsheaf $\mathcal{F}' \subset \mathcal{F}$ we have $U = \mathcal{F}'(V)$ satisfies

$$\dim(\mathcal{F}') = \dim(U), \qquad \mathcal{F}'(E) \subset \Gamma_{\mathrm{ht}}(H).$$

Hence the excess of U is at least minus the Euler characteristic of $\mathcal{F}'$. $\square$

The above theorem has a simple graph theoretic analogue, namely that

$$\rho(G) = \max_{H \subset G} -\chi(H).$$

One can easily prove this directly (with $\rho(G) = -\chi(H)$ when H consists of all cyclic connected components of G) or use Theorem 1.28.

We remark that it is easy to give a direct proof that the maximum excess satisfies some of the properties of a first quasi-Betti number. For example, it is immediate that for sheaves $\mathcal{F}_1, \mathcal{F}_2$ on a graph, G, we have

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

As another example, if $\mathcal{F}_1 \to \mathcal{F}_2$ is an injection, then Theorem 1.28 shows that

$$\mathrm{m.e.}(\mathcal{F}_1) \le \mathrm{m.e.}(\mathcal{F}_2).$$

It is quite conceivable that all of the "first quasi-Betti number" properties of the maximum excess have simple, direct proofs that avoid using Theorem 1.10. However, we find that Theorem 1.10, that implies that the maximum excess is a limiting twisted Betti number, is extremely useful in providing intuition about the maximum excess.

1.7. h_1^{twist} and the Universal Abelian Covering

For a digraph, G, we will study its maximum Abelian covering, $\pi\colon G[\mathbb{Z}] \to G$, which is an infinite graph, and show that for a sheaf $\mathcal{F}$, on G, we have $H_1^{\mathrm{twist}}(\mathcal{F})$ is non-zero iff there is a non-zero element of $H_1(\pi^*\mathcal{F})$ that is of finite support. This is crucial to our proof of Theorem 1.10. We shall illustrate these theorems on the unhappy 4-bundle, which gives great insight into our proof of Theorem 1.10 that we give in Section 1.8.

Let $\mathbb{Z}$ be the set of integers, and let $\mathbb{Z}_{\ge 0}$ be the set of non-negative integers. For a set, S, we use $\mathbb{Z}^S$ to denote the set of functions from S to $\mathbb{Z}$. We define the *rank* of an $n \in \mathbb{Z}^S$ to be

$$\mathrm{rank}(n) = \sum_{s \in S} n(s)$$

(in this paper S will always be finite, so the summation makes sense).

Given a digraph, G, let $G[\mathbb{Z}]$ be the infinite digraph with

$$V_{G[\mathbb{Z}]} = V_G \times \mathbb{Z}^{E_G}, \quad E_{G[\mathbb{Z}]} = E_G \times \mathbb{Z}^{E_G},$$

with heads and tails maps given for each $e \in E_G$ and $n \in \mathbb{Z}^{E_G}$ by

$$h_{G[\mathbb{Z}]}(e,n) = (h_G e, n), \quad t_{G[\mathbb{Z}]}(e,n) = (t_G e, n+\delta_e),$$

where $\delta_e \in \mathbb{Z}^{E_G}$ is 1 at e and 0 elsewhere. Projection onto the first component gives an infinite degree covering map $\pi\colon G[\mathbb{Z}] \to G$. For a vertex, (v,n), or an edge, (e,n), of $G[\mathbb{Z}]$, we define its *rank* to be the rank of n.

DEFINITION 1.29. *For a digraph, G, we define the* universal Abelian covering *of G to be $\pi\colon G[\mathbb{Z}] \to G$ described in the previous paragraph.*

It is not important to us, but easy to verify, that π factors uniquely through any connected Abelian covering of G. Abelian coverings have been studied in numerous works, including [**FT05, FMT06**].

We similarly define $G[\mathbb{Z}_{\geq 0}]$, with $\mathbb{Z}_{\geq 0}$ replacing $\mathbb{Z}$ everywhere; $G[\mathbb{Z}_{\geq 0}]$ can be viewed as a subgraph of $G[\mathbb{Z}]$.

Our approach to Theorem 1.10 involves the properties of the graphs $G[\mathbb{Z}_{\geq 0}]$, so let us consider some examples. If B_d denotes the bouquet of d self-loops, i.e., the digraph with one vertex and d edges, then $B_d[\mathbb{Z}_{\geq 0}]$ is just the usual d-dimensional non-negative integer lattice, depicted in Figures 1 and 2. If $G' \to G$ is a covering

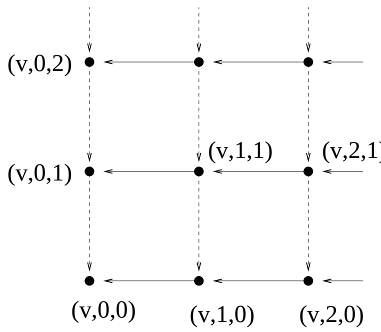

FIGURE 1. $B_2[\mathbb{Z}_{\geq 0}]$.

map of degree d, then $G'[\mathbb{Z}] \to G[\mathbb{Z}]$ and $G'[\mathbb{Z}_{\geq 0}] \to G[\mathbb{Z}_{\geq 0}]$ are both covering maps. However, for $d > 1$ and $|E_G| \geq 1$, we have $|E_{G'}| > |E_G|$, and the covering will be of infinite degree.

Now consider $G'[\mathbb{Z}_{\geq 0}]$, where $\phi\colon G' \to B_2$ is the degree two cover of B_2 discussed with the unhappy 4-bundle in Subsection 1.2.4 (just beneath equation (1.9)). As we see, and illustrated in Figure 3, $G'[\mathbb{Z}_{\geq 0}]$ has no cycle of length four. As we shall see, the fact that $h_1^{\mathrm{twist}}(\mathcal{U}) = 1$ is a result, in a sense, of the cycles of length four in $B_2[\mathbb{Z}_{\geq 0}]$; the fact that these cycles "open up" to non-closed walks in $G'[\mathbb{Z}_{\geq 0}]$ is partly why $h_1^{\mathrm{twist}}(\phi^*\mathcal{U}) = 0$.

Now we define homology groups on graphs of the form $G[\mathbb{Z}]$ and $G[\mathbb{Z}_{\geq 0}]$, and, more generally, any infinite graph. If K is a infinite graph that is locally finite (i.e., each vertex is incident upon a finite number of edges), we can still define a sheaf (of finite dimensional vector spaces over a field, $\mathbb{F}$) just as before. Hence a sheaf, $\mathcal{F}$, on K as a collection of a finite dimensional $\mathbb{F}$-vector space, $\mathcal{F}(P)$ for each $P \in V_K \amalg E_K$, along with restriction maps $\mathcal{F}(h,e)$ and $\mathcal{F}(t,e)$ for each $e \in E_K$.

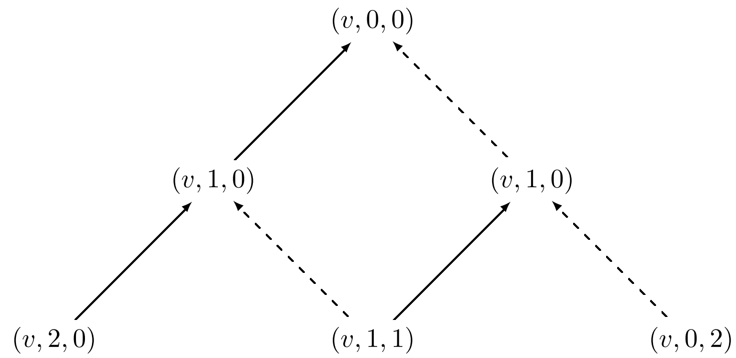

FIGURE 2. First part of $B_2[\mathbb{Z}_{\geq 0}]$. Notice the cycle of length four.

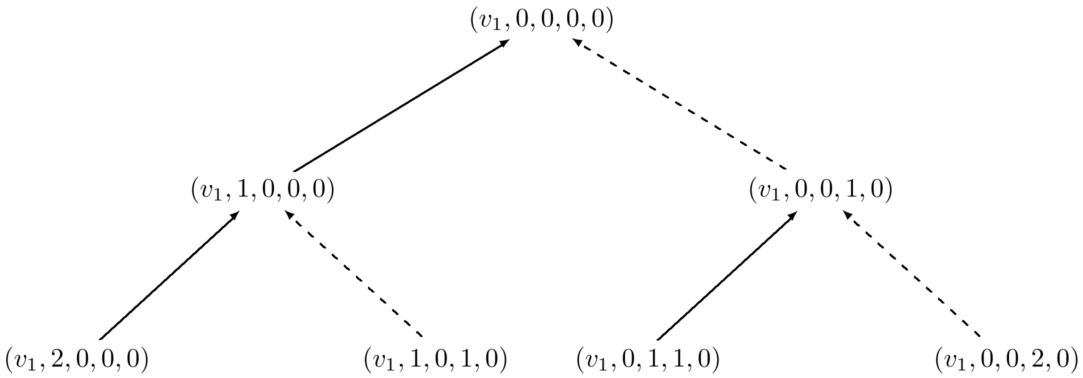

FIGURE 3. First part of $G'[\mathbb{Z}_{\geq 0}]$ near $(v,\vec{0})$. No cycles of length four. The four $(\mathbb{Z}_{\geq 0})^{E_{G'}}$ coordinates are, in order, $e_1^1, e_1^2, e_2^1, e_2^2$ where e_i^j lies over $e_i \in E_{B_2}$ and are described in the last equations of Subsection 1.2.4 that give the ν_i^j.

We shall define

$$\mathcal{F}^{\oplus}(V) = \bigoplus_{v \in V_G} \mathcal{F}(v), \quad \text{and} \quad \mathcal{F}^{\Pi}(V) = \prod_{v \in V_G} \mathcal{F}(v),$$

which generally differ, $\mathcal{F}^{\oplus}(V)$ being the subset of $\mathcal{F}^{\Pi}(V)$ of elements $\{f_v\}_{v \in V_G}$ that are supported (i.e., nonzero) on only finitely many v. Similarly we define $\mathcal{F}^{\oplus}(E)$ and $\mathcal{F}^{\Pi}(E)$. Then $d = d_h - d_t$ can be viewed as a map $\mathcal{F}^{\Pi}(E) \to \mathcal{F}^{\Pi}(V)$ or, respectively, $\mathcal{F}^{\oplus}(E) \to \mathcal{F}^{\oplus}(V)$, and their cokernels and kernels are respectively denoted $H_i^{\Pi}(\mathcal{F})$ and $H_i^{\oplus}(\mathcal{F})$ for $i = 0, 1$.

If $\mathcal{F}$ is a sheaf on G, and $\pi\colon G[\mathbb{Z}] \to G$ the universal Abelian covering, then $\pi^*\mathcal{F}$ is a sheaf on $G[\mathbb{Z}]$.

The following simple but important observation explains our interest in the universal Abelian covering.

LEMMA 1.30. *Let $\mathcal{F}$ be a sheaf on G, and $\pi\colon G[\mathbb{Z}] \to G$ the universal Abelian covering. Then $H_1^{\mathrm{twist}}(\mathcal{F})$ is non-trivial iff $H_1^{\oplus}(\pi^*\mathcal{F})$ is non-trivial. If so, there is a non-zero $w \in H_1^{\oplus}(\pi^*\mathcal{F})$ that is supported on $G[\mathbb{Z}_{\geq 0}]$.*

PROOF. For each $e \in E_G$, let $\mathcal{F}(e)$ be of dimension d_e and have basis $f_{e,1}, \ldots, f_{e,d_e}$. Let

$$a_{e,i} = \mathcal{F}(h,e)f_{e,i} \in \mathcal{F}(he), \quad b_{e,i} = \mathcal{F}(t,e)f_{e,i} \in \mathcal{F}(te).$$

We have $h_1^{\mathrm{twist}}(\mu^*\mathcal{F}) \geq 1$ iff the vectors

$$a_{e,i} + \psi(e)b_{e,i}$$

are linear dependent over $\mathbb{F}(\psi)$, where ψ is a collection of indeterminates indexed on E_G. This holds iff there are rational functions $c_{e,i} \in \mathbb{F}(\psi)$ for each $e \in E_G$ and $i = 1, \ldots, d_e$ such that

$$\sum_{e \in E_G} \sum_{i=1}^{d_e} c_{e,i}(\psi)(a_{e,i} + \psi(e)b_{e,i}) = 0, \tag{1.33}$$

where not all $c_{e,i}$ are zero. We may multiply the denominators of the $c_{e,i}(\psi)$ to assume that they are polynomials, not all zero. We may write

$$c_{e,i}(\psi) = \sum_{n \in (\mathbb{Z}_{\geq 0})^{E_G}} c_{e,i,n}\psi^n,$$

where $c_{e,i,n} \in \mathbb{F}$ and

$$\psi^n = \prod_{e \in E_G} \psi^{n(e)}(e).$$

In summary, we see that $h_1^{\mathrm{twist}}(\mathcal{F}) \neq 0$ iff there exist $c_{e,i,n} \in \mathbb{F}$, with $c_{e,i,n} = 0$ for all but finitely many n, such that

$$\sum_{n \in (\mathbb{Z}_{\geq 0})^{E_G}} \sum_{e,i} \psi^n c_{e,i,n}(a_{e,i} + \psi_e b_{e,i}) = 0 \tag{1.34}$$

and not all the $c_{e,i,n} = 0$. But equation (1.34) is equivalent to saying that

$$w_{(e,n)} = \sum_{i=1}^{d_e} c_{e,i,n} f_{e,i}$$

is a non-zero element of $H_1^{\oplus}(\pi^*\mathcal{F})$. Hence $h_1^{\mathrm{twist}}(\mathcal{F}) \neq 0$ iff $H_1^{\oplus}(\pi^*\mathcal{F}) \neq 0$.

□

The following is a simple graph theoretic definition that is crucial to our proof of Lemma 1.32.

DEFINITION 1.31. *The* Abelian girth *of a digraph graph, G, is the girth of $G[\mathbb{Z}]$.*

Since $G[\mathbb{Z}] \to G$ is a covering map, the girth of $G[\mathbb{Z}]$, which is the Abelian girth of G, is at least the girth of G. Note also that B_1, the digraph with one vertex and one edge (a self-loop), has girth one but infinite Abelian girth, i.e., $G[\mathbb{Z}]$ is a two-sided infinite path and has no cycles. Similarly B_2, the digraph with one vertex and two edges, has girth one but Abelian girth four.

1.8. Proof of Theorem 1.10

We begin with the following lemma that is one of the (if not the) technical core of this chapter.

LEMMA 1.32. *Let $\mathcal{F}$ be a sheaf on a digraph, G. Let $\mu\colon G' \to G$ be a covering map such that G' is of Abelian girth greater than*

$$2\Bigl(\dim\bigl(\mathcal{F}(V)\bigr) + \dim\bigl(\mathcal{F}(E)\bigr)\Bigr).$$

Then $h_1^{\mathrm{twist}}(\mu^*\mathcal{F}) > 0$ implies that m.e.$(\mathcal{F}) > 0$*.*

In Subsection 1.8.7, the last subsection of this section, we use this lemma to prove Theorem 1.10. The rest of the subsections of this section will be devoted to proving the lemma; our proof, whose basic idea is fairly simple, requires a lot of new notation and definitions.

1.8.1. Outline of the Proof of Lemma 1.32. Consider the hypotheses of Lemma 1.32. Let $\pi\colon G'[\mathbb{Z}] \to G'$ be the universal Abelian cover of G', and let $\mathcal{F}' = \mu^*\mathcal{F}$. We assume $h_1^{\mathrm{twist}}(\mathcal{F}') \geq 1$, and we wish to prove that m.e.$(\mathcal{F}) \geq 1$. According to Lemma 1.30, there exists a nonzero $w \in H_1^{\oplus}(\pi^*\mathcal{F}')$ supported in $G'[\mathbb{Z}_{\geq 0}]$; fix such a w.

Let us introduce some notation to explain the idea behind the proof. For $e \in E_G$, we may identify $\mathcal{F}(e)$ with the subspace of $\mathcal{F}(E)$ supported in e, i.e., consisting of vectors whose $\mathcal{F}(e')$ component vanishes for $e' \neq e$ (this subspace is the image of $\mathcal{F}(e)$ under $u \mapsto \mathrm{extend}(u,e)$). If $f \in E_{G'[\mathbb{Z}]}$, then we let w_f be the f-component of w (as done in the proof of Lemma 1.30), so $w_f \in (\pi^*\mathcal{F}')(f)$; but $(\pi^*\mathcal{F}')(f)$ equals $\mathcal{F}(\mu\pi f)$, and can therefore be identified with the subset of $\mathcal{F}(E)$ supported in $\mu\pi f$; let $\overline{w_f}$ be the element of $\mathcal{F}(E)$ corresponding to w_f. For $F \subset E_{G'[\mathbb{Z}]}$, set

$$C(F) = \mathrm{span}\{\overline{w_f} \mid f \in F\} \subset \mathcal{F}(E),$$
$$A(F) = \mathrm{span}\{d_{\mathcal{F},h}\overline{w_f} \mid f \in F\} = d_{\mathcal{F},h}C(F) \subset \mathcal{F}(V),$$

and

$$B(F) = \mathrm{span}\{d_{\mathcal{F},t}\overline{w_f} \mid f \in F\} = d_{\mathcal{F},t}C(F) \subset \mathcal{F}(V).$$

Our idea is to construct an increasing sequence of subgraphs, $U_1 \subset \cdots \subset U_r = U$, of $G[\mathbb{Z}_{\geq 0}]$, and set $F_i = E_{U_i}$, so that $F = F_r$ satisfies

$$\dim\bigl(A(F) + B(F)\bigr) \leq \dim\bigl(C(F)\bigr) - 1. \tag{1.35}$$

At this point we have

$$\mathrm{excess}\bigl(\mathcal{F}, A(F) + B(F)\bigr) \geq 1$$

and the lemma is established.

The subgraphs $U_1, \ldots, U_r$ will be selected in "phases." In the first phase we choose $U_1, \ldots, U_{k_1}$ for some integer $k_1 \geq 1$. We will show that

$$\dim\bigl(A(F_{k_1})\bigr) \leq \dim\bigl(C(F_{k_1})\bigr) - k_1. \tag{1.36}$$

This inequality is worse than equation (1.35) because it doesn't involve $B(F_{k_1})$; however, it is possibly better, in that the right-hand-side has a $-k_1$ and we may have $k_1 > 1$.

The i-th phase will select $U_{k_{i-1}+1}, U_{k_{i-1}+2}, \ldots, U_{k_i}$ for some integer $k_i \geq k_{i-1}$. (Hence we set $k_0 = 0$ for consistency and convenience.) The third, fifth, and all odd numbered phases will be called C-phases, for a reason that will become clear (see equations (1.43) and (1.58) and nearby discussion); the C-phases select their U_i in a similar way. The second phase will be called a B-phase; in this phase we choose $U_{k_1+1}, \ldots, U_{k_2}$ to derive an equality akin to equation (1.36) that involves $B(F_{k_1})$ (namely equation (1.56)); unfortunately, the inequality no longer involves $A(F_{k_1})$ and $C(F_{k_2})$, rather it involves $A(F_{k_2})$ and $C(F_{k_2})$. The fourth, sixth, and all even numbered phases will be called B-phases, because of the way in which their U_i are selected (see equation (1.59)).

After the first two phases, i.e., the first C-phase and first B-phase, each subsequent phase, alternating between C-phases and B-phases, allows us to write an inequality akin to equation (1.35) or (1.36). The inequality after the i-th phase will involve the values of A, B, C at $F_{k_i}, F_{k_{i-1}}, F_{k_{i-2}}$; roughly speaking, as i gets larger, the values of A, B, or C on $F_{k_i}, F_{k_{i-1}}, F_{k_{i-2}}$ must "converge," since these are subspaces of finite dimensional spaces $\mathcal{F}(V)$ and $\mathcal{F}(E)$. At the point of "convergence" (more precisely, when either equation (1.60) or (1.61) hold) our phases end after completing the i-th phase, whereupon taking $r = k_i$ we will have that $F = F_r$ satisfies equation (1.35) and we are done.

Now we give the details. The construction of the U_i and the inequalities we prove involve definitions of what we call "stars" and "star union data," given in Subsection 1.8.2. We shall describe the first and second phase, respectively, in detail in Subsections 1.8.3 and 1.8.5, respectively. In Subsection 1.8.4 we state and prove a number of facts used in Subsections 1.8.3 and 1.8.5 in greater generality; we hope that this greater generality will clarify the proofs. In Subsection 1.8.6 we finish the proof of Lemma 1.32. As mentioned before, in Subsection 1.8.7, we use Lemma 1.32 to prove Theorem 1.10.

1.8.2. Star Union Data. We now fix some graph theoretic notions to describe the U_i, F_i, and related concepts. For a vertex, u, of $G'[\mathbb{Z}_{\geq 0}]$, let the *star at* v, denoted $\mathrm{Star}(u)$, be the subgraph of $G'[\mathbb{Z}_{\geq 0}]$ consisting of those edges of $G'[\mathbb{Z}_{\geq 0}]$ whose head is u and of those vertices that are the endpoints of these edges (the star at u is easily seen to be a tree, since $G'[\mathbb{Z}_{\geq 0}]$ has no self-loops or multiple edges).

DEFINITION 1.33. *For any sequence $v = (v_1, \ldots, v_j)$ of vertices of $G'[\mathbb{Z}_{\geq 0}]$, we define the* star union *of v to be the union of the stars at $v_1, \ldots, v_j$. Furthermore, to any such sequence $v = (v_1, \ldots, v_j)$ we associate the following data, $(U_i, F_i, I_i, X_i)_{i=1,\ldots,j}$, that we call* star union data*: for positive integer $i \leq j$ we associate*

1. *the i-th* star union*, U_i, which is the star union of $(v_1, \ldots, v_i)$;*
2. *the i-th* edge set*, $F_i = E_{U_i}$;*
3. *the i-th* interior edge set*, $I_i \subset F_i$, the set of edges in U_i whose tail is one of $v_1, \ldots, v_i$;*
4. *the i-th* interior vertex set*, $\{v_1, \ldots, v_i\}$; and*
5. *the i-th* exterior vertex set*, $X_i = V_{U_i} \setminus \{v_1, \ldots, v_i\}$.*

N.B.: Throughout the rest of this section, the variables U_i, F_i, I_i, X_i and terminology of Definition 1.33, will refer to star union data with respect to the variable $v = (v_1, \ldots, v_j)$, where j will change during the section. Our goal is to construct $v = (v_1, \ldots, v_r)$ such that $F = F_r$ satisfies equation (1.35), but to do so will construct v in phases, and during any part of any phase the variables U_i, F_i, I_i, X_i refer to the portion of v constructed so far (which limits i to be at most j for the current value of j)

1.8.3. The First C-Phase. We remind the reader that, as explained at the end of Subsection 1.8.2, U_i, F_i, I_i, X_i are assumed to refer to star union data derived from a sequence $v = (v_1, v_2, \ldots)$, at any stage of its construction.

Choose any edge, e_1, of minimal rank with $\overline{w_{e_1}} \neq 0$ and let $v_1 = he_1$ and let $\rho = \mathrm{rank}(v_1)$. We claim

$$\dim(A(F_1)) + 1 \leq \dim(C(F_1));$$

indeed, if v_1 is the tail of an edge, f, then $\overline{w_f} = 0$, by the minimal rank of e_1. Hence

$$\sum_{e \text{ s.t. } he=v_1} d_h \overline{w_e} \;=\; \sum_{e \text{ s.t. } te=v_1} d_t \overline{w_e} \;=0. \tag{1.37}$$

Consider the set

$$E^1 = \{e \mid he = v_1 \text{ and } \overline{w_e} \neq 0\} \subset E_{G'[\mathbb{Z}]_{\geq 0}}.$$

We claim that

$$\dim\bigl(C(F_1)\bigr) = |E^1|; \tag{1.38}$$

indeed

$$\mathcal{F}(E) = \bigoplus_{e \in E_G} \mathcal{F}(e),$$

and since $\mu\pi\colon G'[\mathbb{Z}_{\geq 0}] \to G$ is a covering map, for each $f \in E_G$ there is at most one $e \in E_{G'[\mathbb{Z}_{\geq 0}]}$ such that $\mu e = f$ and $he = v_1$. Hence each nonzero w_e with $e \in F_1$ is taken to its own component of $\mathcal{F}(E)$. So in the terminology of Subsection 1.6.1, the nonzero w_e are compartmentally distinct, and hence independent, by Theorem 1.24. Hence equation (1.38) holds. By contrast, equation (1.37) shows that the $d_h\overline{w_e}$ with $e \in E^1$ sum to zero and are therefore dependent; hence

$$\dim\bigl(A(F_1)\bigr) \leq |E^1| - 1,$$

and so

$$\dim\bigl(A(F_1)\bigr) \leq \dim\bigl(C(F_1)\bigr) - 1. \tag{1.39}$$

Assume that there is an $e_2 \in E_{G'[\mathbb{Z}_{\geq 0}]}$ for which $\mathrm{rank}(e_2) = \rho$ and $\overline{w_{e_2}} \notin C(F_1)$. In this case the first phase continues; we fix any such e_2, set $v_2 = he_2$. We claim that

$$\dim\bigl(A(F_2)/A(F_1)\bigr) \leq \dim\bigl(C(F_2)/C(F_1)\bigr) - 1. \tag{1.40}$$

Indeed, let E^2 be the number set of e such that $he = v_2$ and $\overline{w_e} \notin C(F_1)$ (i.e., $\overline{w_e}$ is non-zero modulo $C(F_1)$). Note that $C(F_1)$ is compartmentalized. Also, the $\overline{w_e}$ with $e \in E^2$ are compartmentally distinct (by the same argument as used for E^1, which is true when e ranges over the edges of any star). Hence, by Theorem 1.24, the $\overline{w_e}$ with $e \in E^2$ are linearly independent in $\mathcal{F}(E)/C(F_1)$. Hence

$$\dim\bigl(C(F_2)/C(F_1)\bigr) = |E^2|.$$

However, as with E^1 we have

$$\sum_{e \in E^2} d_h \overline{w_e} = 0,$$

since v_2 has rank ρ (so $\overline{w_e} = 0$ for all e with $te = v_2$). But if $he = v_2$ and $e \notin E^2$, then $\overline{w_e} \in C(F_1)$ and so $A(\{e\}) \in A(F_1)$. Hence

$$\sum_{e \in E^2} d_h \overline{w_e} \in A(F_1),$$

It follows that

$$\dim\bigl(A(F_2)/A(F_1)\bigr) \le |E^2| - 1.$$

This establishes equation (1.40), and adding that equation to equation (1.39) gives

$$\dim\bigl(A(F_2)\bigr) \le \dim\bigl(C(F_2)\bigr) - 2.$$

If there is an e_3 such that $\mathrm{rank}(e_3) = \rho$ and $\overline{w_{e_3}} \notin C(F_2)$, then the first phase continues, with $v_3 = he_3$, and we have

$$\dim\bigl(A(F_3)\bigr) \le \dim\bigl(C(F_3)\bigr) - 3.$$

We similarly find e_i and set $v_i = he_i$ for each positive integer i for which there is an e_i of rank ρ with $\overline{w_{e_i}} \notin C(F_{i-1})$; for any such i we have

$$\dim\bigl(A(F_i)\bigr) \le \dim\bigl(C(F_i)\bigr) - i. \tag{1.41}$$

But for any such i we have

$$\dim\bigl(C(F_i)\bigr) \ge i; \tag{1.42}$$

hence for any such i we have $i \le \dim(\mathcal{F}(E))$, and so for some $k_1 \le \dim(\mathcal{F}(E))$ this process stops at $i = k_1$, i.e., we construct $e_1, \ldots, e_{k_1}$ of rank ρ with $\overline{w_{e_i}} \notin C(F_{i-1})$ for $i = 2, \ldots, k_1$, but $C(F_{k_1})$ contains all $\overline{w_e}$ for $\mathrm{rank}(e) = \rho$. This is the end of the first phase.

A concise way to describe the first phase is that we choose any minimal $v_1, \ldots, v_{k_1}$ of rank ρ such that

$$\forall e \in E_{G[\mathbb{Z}_{\ge 0}]} \text{ of rank } \rho, \quad \overline{w_e} \in C(F_{k_1}), \tag{1.43}$$

where minimal means that if we discard any v_i from $v_1, \ldots, v_{k_1}$ then equation (1.43) does not hold. We call this a C-phase because the equation (1.43) involves a "C," as will all odd numbered phases. Notice that equation (1.41) is somewhat similar to our desired equation (1.35); one big difference is that equation (1.41) makes no mention of B, but only of A and C.

1.8.4. Moseying Sequences. Before describing the second phase, i.e., the first B-phase, we wish to organize the inequalities we will need into a number of lemmas. Furthermore, we will usually state these lemmas in a slightly more general context; this will help illustrate exactly what assumptions are being used.

We consider the setup and notation of the first two paragraphs of Subsection 1.8.1, which fixes $\mathcal{F}$, $\mu\colon G' \to G$, $\pi\colon G'[\mathbb{Z}_{\ge 0}] \to G'$, $w \in H_1^{\oplus}(\pi^*\mu^*\mathcal{F})$, and defines $\overline{w_f}$ for any $f \in E_{G'[\mathbb{Z}_{\ge 0}]}$, and defines $A(F), B(F), C(F)$ for any $F \subset E_{G'[\mathbb{Z}_{\ge 0}]}$.

We will work with a sequence of vertices, $v = (v_1, \ldots, v_s)$, of $G'[\mathbb{Z}_{\ge 0}]$, but we will not assume the v_i are constructed by our phases. Instead, we will be careful to write down our assumptions on the v_i in a way that will make clear which of their properties is used when and how. Our central definition in this general context will be that of a "moseying sequence."

Definition 1.34. *By a* moseying sequence of length s for G' *we mean a sequence* $v = (v_1, \ldots, v_s)$ *of distinct vertices of* $G'[\mathbb{Z}]$ *for which* $\mathrm{rank}(v_{i+1}) - \mathrm{rank}(v_i)$ *is* 0 *or* 1 *for each* i*; if this difference is* 1 *we say that* v jumps at i. *We define star union data,* U_i, F_i, I_i, X_i *as in Subsection 1.8.2. For ease of notation we define* U_0, F_0, I_0, X_0 *to be empty (i.e.,* U_0 *is the empty graph,* F_0, I_0, X_0 *the empty set).*

Moseying sequences are our basic object of study.

Definition 1.35. *A moseying sequence,* v*, of length* s *is* of increasing dimension *if the integers*

$$n_i = \dim\bigl(C(F_i)\bigr) + \dim\bigl(B(I_i)\bigr)$$

satisfy

$$0 = n_0 < n_1 < n_2 < \cdots < n_s.$$

Lemma 1.36. *Let* v *be a moseying sequence of length* s *of increasing dimension for a digraph,* G'*. Then*

$$s \leq \dim\bigl(\mathcal{F}(E)\bigr) + \dim\bigl(\mathcal{F}(V)\bigr).$$

Furthermore, for any $i \leq s$*,* U_i *has no cycles provided that the girth of* $G'[\mathbb{Z}]$ *is at least* $2i + 1$*.*

Proof. The first statement is clear. For the second statement, assume, to the contrary, that U_i has a cycle. U_i is the union of stars, which are trees of diameter two. If c is a cycle in U_i of minimal length, then it traverses each vertex at most once. But every vertex of c not appearing in v must be a leaf (i.e., tail of an edge) of a star, and hence followed by (and preceded by) a vertex in v. Hence the length of c is at most twice i. Hence $G'[\mathbb{Z}]$ has a cycle of length at most $2i$, contradicting the hypotheses of the lemma. □

The inequality in equation (1.41), derived after the first C-phase, will be built up along further phases to eventually give equation (1.35). However, to express these later phase inequalities, we shall need some graph theoretic notions, such as the "overdegree" and "capacity" that we now define.

Definition 1.37. *Let* v *be a moseying sequence of length* s *for* G'*. For any* $u \in V_{G'[\mathbb{Z}]}$ *we define the* stable outdegree of u, *denoted* $\mathrm{sod}(u)$*, to be the outdegree of* u *in* U_s*. (If* v *is not a vertex of* U_s*, we define its outdegree in* U_s *to be zero.)*

Note that the outdegree of u in U_{j-1}, viewed as a function of j, does not change as soon as $\mathrm{rank}(v_j) \geq \mathrm{rank}(u)$; indeed, the edges that affect the outdegree of u are the edges of rank equal to $\mathrm{rank}(u) - 1$, and such edges come from stars about vertices of $\mathrm{rank}(u) - 1$. Hence, for any j with $1 \leq j \leq s$, we have

$$\mathrm{rank}(v_j) \geq \mathrm{rank}(u) \implies \mathrm{sod}(u) = \mathrm{outdeg}(U_{j-1}, u), \tag{1.44}$$

where $\mathrm{outdeg}(G, w)$ denotes the outdegree of w in G. In particular,

$$\mathrm{sod}(v_j) = \mathrm{outdeg}(U_{j-1}, v_j)$$

for all $j = 1, \ldots, s$.

Definition 1.38. *Let* v *be a moseying sequence of length* s *for* G'*. By the overdegree of* U_i*, for an integer,* i *with* $1 \leq i \leq s$*, we mean*

$$\mathrm{Over}(U_i) = \sum_{v \in X_i} \bigl(\mathrm{outdeg}(U_i, v) - 1\bigr),$$

Notice that for any i, the overdegree of U_i is non-negative, since each exterior vertex of U_i is the tail of some edge in U_i, and hence has outdegree at least one.

DEFINITION 1.39. *Let v be a moseying sequence of length s for G'. For non-negative integer, $i \le s$, we define the* capacity of U_i *to be*

$$\mathrm{Cap}(U_i) = h_0(U_i) + \mathrm{Over}(U_i).$$

Note that for $i \ge 1$, $h_0(U_i) \ge 1$, since U_i is nonempty, and $\mathrm{Over}(U_i) \ge 0$; hence for $i \ge 1$ we have $\mathrm{Cap}(U_i) \ge 1$. Our fundamental inequalities will use the capacity.

LEMMA 1.40. *Let v be a moseying sequence of length s for G'. Assume that U_j has no cycles for some $j \le s$. Then for any non-negative integers $i \le j$ we have*

$$\mathrm{Cap}(U_j) = \mathrm{Cap}(U_i) - \sum_{m=i+1}^{j} \bigl(\mathrm{sod}(v_m) - 1\bigr)$$

PROOF. It suffices to prove the lemma for $j = i+1$, for then the general lemma follows by induction on $j - i$.

So assume $j = i + 1$, and set $\rho = \mathrm{rank}(v_{i+1})$. Let p_0 and p_1, respectively, be the number of vertices of rank ρ and $\rho + 1$, respectively, in which the star of v_{i+1} intersects U_i; so p_0 is 1 or 0 according to whether or not $v_{i+1} \in V_{U_i}$, and p_1 is the number of tails of edges in $\mathrm{Star}(v_{i+1})$ that lie in U_i; let $p = p_0 + p_1$. First, note that since $U_{i+1} = U_i \cup \mathrm{Star}(v_{i+1})$, we have

$$\chi(U_{i+1}) = \chi(U_i) + \chi\bigl(\mathrm{Star}(v_{i+1})\bigr) - \chi\bigl(U_i \cap \mathrm{Star}(v_{i+1})\bigr);$$

since U_i, U_{i+1}, and any star have $h_1 = 0$, in the above equation we may replace each χ with h_0, and conclude that

$$h_0(U_{i+1}) = h_0(U_i) + h_0\bigl(\mathrm{Star}(v_{i+1})\bigr) - h_0\bigl(U_i \cap \mathrm{Star}(v_{i+1})\bigr);$$

since $U_i \cap \mathrm{Star}(v_{i+1})$ contains no edges, it has p connected components (p isolated vertices), and hence

$$h_0(U_{i+1}) = h_0(U_i) + 1 - p. \tag{1.45}$$

Second, note that each of the p_1 tails of edges of the star adds one to its degree in U_{i+1} over that of U_i; the remaining tails of star edges have degree one in U_{i+1}. This means that U_{i+1} gains p_1 over U_i in the overdegree contribution from vertices of rank $\rho + 1$. Third, note that $p_0 = 1$ iff $v_{i+1} \in V_{U_i}$ iff v_{i+1} contributes

$$\mathrm{outdeg}(U_i, v_{i+1}) - 1 = \mathrm{sod}(v_{i+1}) - 1$$

to the overdegree of U_i; if so, this contribution is lost in U_{i+1}, since v_{i+1} becomes an interior vertex. Hence if $p_0 = 0$ we have

$$\mathrm{Over}(U_{i+1}) = \mathrm{Over}(U_i) + p_1$$

and if $p_0 = 1$ we have

$$\mathrm{Over}(U_{i+1}) = \mathrm{Over}(U_i) + p_1 - (\mathrm{sod}(v_{i+1}) - 1);$$

in both cases we may write

$$\mathrm{Over}(U_{i+1}) = \mathrm{Over}(U_i) + p - \mathrm{sod}(v_{i+1}).$$

Combining this with equation (1.45) yields

$$\mathrm{Cap}(U_{i+1}) = \mathrm{Cap}(U_i) + 1 - \mathrm{sod}(v_{i+1}),$$

which proves the lemma for $j = i+1$ and therefore, as explained earlier, for all $j > i$. $\square$

LEMMA 1.41. *Let v be a moseying sequence of length s for G'. Assume that v jumps at an integer $i < s$, but not at $i+1, i+2, \ldots, k$ for some integer $k \le s$. (We adopt the convention that v jumps at i if $i = 0$.) Assume that for each edge, e, of $G'[\mathbb{Z}]$ of rank at most $\operatorname{rank}(v_i)$ we have $\overline{w_e} \in C(F_i)$. Then for any j with $i+1 \le j \le k$ we have*

$$\dim\Big(A(F_k)/\big(A(F_j)+B(F_i)\big)\Big) \le \dim\big(C(F_k)/C(F_j)\big) - (k-j). \tag{1.46}$$

We remark that the assumptions of this lemma are highly restrictive; to apply this to our phases, $i+1$ (or v_{i+1}) will have to be the beginning of a B-phase, and k (or v_k) will lie either in that B-phase or the C-phase immediately thereafter. Also, if v jumps somewhere between $i+1$ and k, then we cannot expect equation (1.46) to hold unless $B(F_i)$ is replaced with $B(F_{i'})$ for an $i' > i$.

PROOF. For $j = k$ the lemma is immediate. Let us first establish the case $k = j+1$; the general case will then easily follow by induction on $k - j$. Let $\rho = \operatorname{rank}(v_i)$.

Consider that

$$\sum_{te=v_{j+1}} d_t \overline{w_e} = \sum_{he=v_{j+1}} d_h \overline{w_e}.$$

We have $d_t\overline{w_e} \in B(F_i)$ for all e with $te = v_{j+1}$, and, more generally, for any e of rank ρ, since $\overline{w_e} \in C(F_i)$. Hence

$$\sum_{he=v_{j+1}} d_h \overline{w_e} \in B(F_i). \tag{1.47}$$

Now, as before, let E' be those e with $he = v_{j+1}$ and $\overline{w_e} \notin C(F_j)$, and let E'' be the same but with $\overline{w_e} \in C(F_j)$. We have

$$\dim\big(C(F_{j+1})/C(F_j)\big) = |E'|,$$

since $C(F_j)$ is a compartmentalized subspace of $\mathcal{F}(E)$; yet for $e \in E''$ we have $d_h\overline{w_e} \in A(F_j)$ and hence

$$\sum_{e\in E''} d_h \overline{w_e} \in A(F_j),$$

which implies, along with equation (1.47) that

$$\sum_{e\in E'} d_h\overline{w_e} = \sum_{he=v_{j+1}} d_h\overline{w_e} - \sum_{e\in E''} d_h\overline{w_e} \in B(F_i) + A(F_j).$$

Hence the $d_h\overline{w_e}$ ranging over $e \in E'$ are linearly depedent modulo $A(F_j) + B(F_i)$, and so

$$\dim\Big(A(F_{j+1}) \;/\; \big(A(F_j)+B(F_i)\big)\Big) \le |E'| - 1.$$

Hence

$$\dim\Big(A(F_{j+1}) \;/\; \big(A(F_j)+B(F_i)\big)\Big) \le \dim\big(C(F_{j+1})/C(F_j)\big) - 1. \tag{1.48}$$

This establishes the case $k = j+1$ of the lemma.

The general case of the lemma now follows from the fact that F_j and hence $C(F_j)$ are increasing in j, and hence

$$\dim\bigl(C(F_k)/C(F_j)\bigr) = \sum_{m=j}^{k-1} \dim\bigl(C(F_{m+1})/C(F_m)\bigr);$$

similarly the spaces $A(F_j)$ modulo $B(F_i)$, i.e., viewed as subspaces of $\mathcal{F}(V)/B(F_i)$, are increasing in j, and hence

$$\dim\Bigl(A(F_k) \;/\; \bigl(A(F_j) + B(F_i)\bigr)\Bigr) = \sum_{m=j}^{k-1} \dim\Bigl(A(F_{m+1}) \;/\; \bigl(A(F_m) + B(F_i)\bigr)\Bigr).$$

Hence applying equation (1.48) with m replacing j and m over the range $j, j+1, \ldots, k-1$ yields the lemma. $\square$

LEMMA 1.42. *Let v be a moseying sequence of length s for G'. Then for non-negative integers $i \le j \le s$ we have*

$$\dim\bigl(B(I_j)/B(I_i)\bigr) \le \sum_{m=i+1}^{j} \mathrm{sod}(v_m).$$

PROOF. Clearly $B(I_j)/B(I_i)$ is at most the size of $I_j \setminus I_i$. But an edge, e, of $G'[\mathbb{Z}]$, lies in $I_j \setminus I_i$ (viewing $I_i \subset I_j$ as subsets of $E_{G'[\mathbb{Z}]}$) precisely when $te = v_m$ for some m between $i+1$ and j; furthermore, for each such m, the number of e with $te = v_m$ in U_j is $\mathrm{outdeg}(U_j, v_m)$. Hence

$$\dim\bigl(B(I_j)/B(I_i)\bigr) \le \sum_{m=i+1}^{j} \mathrm{outdeg}(U_j, v_m).$$

But $\mathrm{outdeg}(U_j, v_m) = \mathrm{sod}(v_m)$, either by definition, if $j = s$ or, if $j < s$, in view of equation (1.44) and the fact that $\mathrm{rank}(v_{j+1}) \ge \mathrm{rank}(v_m)$. Hence the lemma follows. $\square$

1.8.5. The First B-Phase. At this point we have finished the first C-phase, having constructed $v_1, \ldots, v_{k_1}$. If

$$B(F_{k_1}) \subset A(F_{k_1}), \tag{1.49}$$

then we are done, for then $F = F_{k_1}$ satisfies equation (1.35), in view of equation (1.41) with $i = k_1$. In this case we end our phases, and Lemma 1.32 is finished in this case. Otherwise $B(F_{k_1})$ is not entirely contained in $A(F_{k_1})$. At this point we enter the second phase; the rough idea is to generate an inequality similar to equation (1.41), but which involves $B(F_{k_1})$; this will come at the expense of making the A and C terms involve F_{k_2} as opposed to F_{k_1}.

We will choose $v_{k_1+1}, \ldots, v_{k_2}$ minimal with

$$B(F_{k_1}) \subset A(F_{k_1}) + B(I_{k_2}), \tag{1.50}$$

which we do as follows: choose any $e \in F_{k_1}$ with $d_t\overline{w_e} \notin A(F_{k_1})$, and set $v_{k_1+1} = te$; then $d_t\overline{w_e} \in B(I_{k_1+1})$; then choose any $e' \in F_{k_1}$ with $d_t\overline{w_{e'}} \notin A(I_{k_1}) + B(I_{k_1+1})$ and take $v_{k_1+2} = te'$ if such an e' exists; continuing on in this fashion we generate a new vertices v_i until we reach a vertex v_{k_2} such that

$$\forall e \in F_{k_1}, \quad d_t\overline{w_e} \in A(F_{k_1}) + B(I_{k_2});$$

such a point is reached, since we have proper containments

$$A(F_{k_1}) \subset A(F_{k_1}) + B(I_{k_1+1}) \subset A(F_{k_1}) + B(I_{k_1+2}) \subset \cdots \tag{1.51}$$

which are subsets of the finite dimensional space $\mathcal{F}(V)$. Hence this point is reached with

$$k_2 - k_1 \le \dim\bigl(\mathcal{F}(V)\bigr),$$

and since $k_1 \le \dim(\mathcal{F}(V))$ (see equation (1.42) and the discussion below it), we have

$$k_2 \le \dim\bigl(\mathcal{F}(V)\bigr) + \dim\bigl(\mathcal{F}(E)\bigr). \tag{1.52}$$

The choice of $v_{k_1+1}, \ldots, v_{k_2}$ comprises the second phase; we call this a (the first) B-phase because of the prominence of the letter "B" in equation (1.50). Now we combine a number of inequalities from Subsection 1.8.4 to prove a sequel to equation (1.41).

First, Lemma 1.40 with $j = v_{2k}$ and $i = 0$ (for which the lemma is still valid) shows that

$$\mathrm{Cap}(U_{k_2}) = k_2 - \sum_{m=1}^{k_2} \mathrm{sod}(v_m) \tag{1.53}$$

(note that U_{k_2} has no cycles, using Lemma 1.36). Second, Lemma 1.41 with $k = k_2$ and $i = j = k_1$ yields

$$\dim\bigl(A(F_{k_2})/(A(F_{k_1}) + B(F_{k_1}))\bigr) \le \dim\bigl(C(F_{k_2})/C(F_{k_1})\bigr) - (k_2 - k_1). \tag{1.54}$$

Third, we have $I_{k_1} = \emptyset$ since $v_1, \ldots, v_{k_1}$ are all of rank ρ. Hence Lemma 1.42 with $j = k_2$ and $i = k_1$ gives

$$\dim\bigl(B(I_{k_2})\bigr) = \dim\bigl(B(I_{k_2})/B(I_{k_1})\bigr) \le \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i). \tag{1.55}$$

We have now established three inequalities in equations (1.53), (1.54), and (1.55). We now establish a simple inequality to describe the end of the first B-phase.

Equations (1.55) and (1.41) with $i = k_1$ imply that

$$\dim\bigl(A(F_{k_1}) + B(I_{k_2})\bigr) \le \dim\bigl(C(F_{k_1})\bigr) - k_1 + \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i),$$

and in view of equation (1.50) this implies that

$$\dim\bigl(A(F_{k_1}) + B(F_{k_1})\bigr) \le \dim\bigl(C(F_{k_1})\bigr) - k_1 + \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i),$$

Equation (1.54) added to this gives

$$\begin{aligned} &\dim\bigl(A(F_{k_2}) + B(F_{k_1})\bigr) \\ &\le \dim\bigl(C(F_{k_2})\bigr) - k_2 + \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i) \\ &= \dim\bigl(C(F_{k_2})\bigr) - k_2 + \sum_{i=1}^{k_2} \mathrm{sod}(v_i) \end{aligned}$$

(since $\mathrm{sod}(v_i) = 0$ for $i = 1, \ldots, k_1$)

$$= \dim\bigl(C(F_{k_2})\bigr) - \sum_{i=1}^{k_2}\bigl(\mathrm{sod}(v_i) - 1\bigr).$$

Then using equation (1.53) we get

$$\dim\bigl(A(F_{k_2}) + B(F_{k_1})\bigr) \le \dim\bigl(C(F_{k_2})\bigr) - \mathrm{Cap}(U_{k_2}). \tag{1.56}$$

This equation is all we need to know about the B-phase we have just finished.

If

$$B(F_{k_2}) \subset A(F_{k_2}) + B(F_{k_1}), \tag{1.57}$$

then our phases are over and we easily establish Lemma 1.32: indeed, we have

$$\dim\bigl(A(F_{k_2}) + B(F_{k_2})\bigr) = \dim\bigl(A(F_{k_2}) + B(F_{k_1})\bigr) \le \dim\bigl(C(F_{k_2})\bigr) - 1$$

since $\mathrm{Cap}(U_{k_2}) \ge 1$ (indeed, $h_0(U_{k_2}) \ge 1$ and the overdegree is non-negative). Hence we have established equation (1.35) with $F = F_{k_2}$ and we are done.

Otherwise we undergo a second C-phase, possibly a second B-phase, possibly a third C-phase, etc. So for $i = 2, 3, \ldots$, the $(2i-1)$-th phase, or i-th C-phase, adds vertices $v_{k_{2i-2}+1}, \ldots, v_{k_{2i-1}}$ of rank $\rho + i - 1$ so that

$$\forall e \in E_{G[\mathbb{Z}_{\ge 0}]} \text{ of rank } \rho + i - 1, \quad \overline{w_e} \in C(F_{k_{2i-1}}) \tag{1.58}$$

(for $j \ge k_{2i-1} + 1$ we successively add a vertex v_j which is the head of an edge, e, of rank $\rho + i - 1$ for which $\overline{w_e} \notin C(F_j)$, augmenting j until no such edges exist); the $(2i)$-th phase, or the i-th B-phase, adds $v_{k_{2i-1}+1,\ldots,k_{2i}}$ so that

$$B(F_{k_{2i-1}}) \subset A(F_{k_{2i-1}}) + B(I_{k_{2i}}); \tag{1.59}$$

as in the first B-phase, the i-th B-phase selects its vertices by choosing an $e \in F_{k_{2i-1}}$ for which

$$d_t\overline{w_e} \notin A(F_{k_{2i-1}}) + B(I_{k_{2i-1}}),$$

setting $v_{k_{2i-1}+1} = te$; then choosing an $e' \in F_{k_{2i-1}}$ for which

$$d_t\overline{w_{e'}} \notin A(F_{k_{2i-1}}) + B(I_{k_{2i-1}+1}),$$

setting $v_{k_{2i-1}+2} = te'$; then repeating this procedure until reaching $v_{k_{2i}}$ such that for all $e \in F_{k_{2i-1}}$ we have

$$d_t\overline{w_e} \in A(F_{k_{2i-1}}) + B(I_{k_{2i}}),$$

whereupon equation (1.59) holds (minimally, i.e., it would fail to hold if we omitted any vertex, v_m, added during this phase).

The phases end either at the end of a C-phase or B-phase as follows: the phases end at the j-th C-phase for $j \ge 1$ when

$$B(F_{k_{2j-1}}) \subset A(F_{k_{2j-1}}) + B(F_{k_{2j-3}}) \tag{1.60}$$

(with $k_{-1} = 0$ and so $F_{k_{-1}} = \emptyset$ for the case $j = 1$), which restricts to equation (1.49) for $j = 1$; the phases end at the j-th B-phase for $j \ge 1$ when

$$B(F_{k_{2j}}) \subset A(F_{k_{2j}}) + B(F_{k_{2j-1}}), \tag{1.61}$$

which restricts to equation (1.57) for $j = 1$. In the next subsection show that one of these two conditions eventually holds for some finite j, and that $F = F_r$ with $r = k_{2j}$ satisfies equation (1.35). We already have all the main inequalities needed to prove this, and just need to apply them to the phases beyond the second phase.

1.8.6. End of the Proof of Lemma 1.32.

PROOF OF LEMMA 1.32. Now we claim that, for all $i \geq 1$, at the end of the i-th C-phase we have

(1.62)
$$\dim\bigl(A(F_{k_{2i-1}}) + B(F_{k_{2i-3}})\bigr) \leq \dim\bigl(C(F_{k_{2i-1}})\bigr) - \mathrm{Cap}(U_{k_{2i-2}}) - (k_{2i-1} - k_{2i-2})$$

(for $i = 1$ we understand that $k_{-1} = k_0 = 0$ and $F_0 = \emptyset$), and that, for all $i \geq 1$, at the end of the i-th B-phase we have

$$\dim\bigl(A(F_{k_{2i}}) + B(F_{k_{2i-1}})\bigr) \leq \dim\bigl(C(F_{k_{2i}})\bigr) - \mathrm{Cap}(U_{k_{2i}}). \tag{1.63}$$

We shall prove these by induction. To do so, first note that after i phases we produce a sequence $v = (v_1, \ldots, v_{k_i})$ that is of increasing dimension, since each v_m of a C-phase increases $\dim(C(F_m))$ by at least one, and each v_m of a B-phase increases $\dim(B(F_m))$ by at least one. Hence, according to Lemma 1.36,

$$k_i \leq \dim\bigl(\mathcal{F}(V)\bigr) + \dim\bigl(\mathcal{F}(E)\bigr), \tag{1.64}$$

and U_{k_i} contains no cycles, using the hypotheses of Lemma 1.32.

Let us also note that the phases eventually end. Indeed, if $k_{2j} = k_{2j-1}$, then according to equation (1.61) we finish. Hence, we are not done by the j-th B-phase we have

$$k_{2j} > k_{2j-1} \geq k_{2j-2} > k_{2j-3} \geq \cdots \geq k_2 > k_1 \geq 1,$$

so $k_{2j} \geq j + 1$; in view of equation (1.64), the total number of phases is less than

$$2\Bigl(\dim\bigl(\mathcal{F}(V)\bigr) + \dim\bigl(\mathcal{F}(E)\bigr)\Bigr).$$

Equation (1.63) has been established for $i = 1$ in equation (1.56). So let us first show that equation (1.63) implies equation (1.62) with i replaced by $i + 1$.

So assume equation (1.63) for some $i \geq 1$. By Lemma 1.41, since v jumps at k_{2i-1} but does not jump thereafter until k_{2i+1}, we have

$$\dim\bigl(A(F_{k_{2i+1}})/(A(F_{k_{2i}}) + B(F_{k_{2i-1}}))\bigr)$$
$$\leq \dim\bigl(C(F_{k_{2i+1}})/C(F_{k_{2i}})\bigr) - (k_{2i+1} - k_{2i}).$$

Adding this to equation (1.63) yields

$$\dim\bigl(A(F_{k_{2i+1}}) + B(F_{k_{2i-1}})\bigr) \leq \dim\bigl(C(F_{k_{2i+1}})\bigr) - \mathrm{Cap}(U_{k_{2i}}) - (k_{2i+1} - k_{2i}).$$

This is equation (1.62), with i replaced by $i + 1$.

Finally assume equation (1.62) for some value of $i \geq 1$; we shall conclude that equation (1.63) holds for the same value of i. By Lemma 1.42 we have

$$\dim\bigl(B(I_{k_{2i}})/B(I_{k_{2i-2}})\bigr) \leq \sum_{m=k_{2i-2}+1}^{k_{2i}} \mathrm{sod}(v_m).$$

This implies that

(1.65)
$$\dim\Bigl(\bigl(A(F_{k_{2i-1}}) + B(I_{k_{2i}})\bigr) \;/\; \bigl(A(F_{k_{2i-1}}) + B(I_{k_{2i-2}})\bigr)\Bigr) \leq \sum_{m=k_{2i-2}+1}^{k_{2i}} \mathrm{sod}(v_m).$$

In view of equation (1.59), and since $I_{k_{2i}} \subset F_{k_{2i-1}}$, we have

$$A(F_{k_{2i-1}}) + B(F_{k_{2i-1}}) = A(F_{k_{2i-1}}) + B(I_{k_{2i}}); \tag{1.66}$$

similarly we have

$$A(F_{k_{2i-3}}) + B(F_{k_{2i-3}}) = A(F_{k_{2i-3}}) + B(I_{k_{2i-2}})$$

and therefore

$$A(F_{k_{2i-1}}) + B(F_{k_{2i-3}}) = A(F_{k_{2i-1}}) + B(I_{k_{2i-2}}) \tag{1.67}$$

Given equations (1.66) and (1.67), equation (1.65) can be rewritten as

$$\dim\Big(\big(A(F_{k_{2i-1}}) + B(F_{k_{2i-1}})\big) \;/\; \big(A(F_{k_{2i-1}}) + B(F_{k_{2i-3}})\big)\Big) \leq \sum_{m=k_{2i-2}+1}^{k_{2i}} \operatorname{sod}(v_m). \tag{1.68}$$

Adding this to equation (1.62) gives

$$\dim\big(A(F_{k_{2i-1}}) + B(F_{k_{2i-1}})\big)$$

$$\leq \dim\big(C(F_{k_{2i-1}})\big) - \operatorname{Cap}(U_{k_{2i-2}}) - (k_{2i-1} - k_{2i-2}) + \sum_{m=k_{2i-2}+1}^{k_{2i}} \operatorname{sod}(v_m)$$

$$= \dim\big(C(F_{k_{2i-1}})\big) - \operatorname{Cap}(U_{k_{2i}}) + (k_{2i} - k_{2i-1})$$

in view of Lemma 1.40 with i, j respectively set to k_{2i-2}, k_{2i}. Adding this to Lemma 1.41 with i, j, k respectively replaced with $k_{2i-2}, k_{2i-1}, k_{2i}$ yields

$$\dim\big(A(F_{k_{2i}}) + B(F_{k_{2i-1}})\big) \leq \dim\big(C(F_{k_{2i}})\big) - \operatorname{Cap}(U_{k_{2i}}).$$

This proves equation (1.63).

At this point we have established equations (1.62) and (1.63), and the fact that the phases eventually end. Now we claim that Lemma 1.32 easily follows. Indeed, if our phases end at the j-th B-phase, then

$$B(F_{k_{2j}}) \subset A(F_{k_{2j}}) + B(F_{k_{2j-1}}),$$

and so equation (1.63) gives

$$\dim\big(A(F_{k_{2j}}) + B(F_{k_{2j}})\big) \leq \dim\big(C(F_{k_{2j}})\big) - \operatorname{Cap}(U_{k_{2j}}).$$

Since $U_{k_{2j}}$ is non-empty, its capacity is at least one, and hence $F = F_r$ with $r = k_{2j}$ satisfies equation (1.35). Similarly, if our phases end at the j-th C-phase, then

$$B(F_{k_{2j-1}}) \subset A(F_{k_{2j-1}}) + B(F_{k_{2j-3}}),$$

and so equation (1.62) gives

$$\dim\big(A(F_{k_{2j-1}}) + B(F_{k_{2j-1}})\big) \leq \dim\big(C(F_{k_{2j-1}})\big) - 1,$$

since

$$\operatorname{Cap}(U_{k_{2j-2}}) + (k_{2j-1} - k_{2j-2}) \geq 1$$

(for $j = 1$ this follows since $k_1 > 0$, and for $j \geq 2$ this follows since $U_{k_{2j-2}}$ is nonempty). Hence, similarly, $F = F_r$ with $r = k_{2j-1}$ satisfies equation (1.35). $\square$

1.8.7. Proof of Theorem 1.10.

PROOF (OF THEOREM 1.10). First we will verify Theorem 1.10 in some special cases.

Lemma 1.32 establishes Theorem 1.10 in the case where $\mathrm{m.e.}(\mathcal{F}) = 0$.

DEFINITION 1.43. *A sheaf, $\mathcal{E}$, on a digraph, G, is* edge supported *if $\mathcal{E}(V) = 0$.*

For an edge supported sheaf, $\mathcal{E}$, it is immediate that for any covering map $\phi\colon G' \to G$ we have

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{E}) = \mathrm{m.e.}(\phi^*\mathcal{E}) = \deg(\phi)\dim(\mathcal{E}(E)).$$

This establishes Theorem 1.10 in the case where $\mathcal{F}$ is edge supported and ϕ is any covering map.

Next we introduce a type of sheaf which will be an important tool.

DEFINITION 1.44. *A sheaf, $\mathcal{F}$, on a graph G, is said to be* tight *if the maximum excess of $\mathcal{F}$ occurs at and only at $\mathcal{F}(V)$.*

LEMMA 1.45. *For any sheaf, $\mathcal{F}$, on a digraph, G, there is a tight sheaf, $\mathcal{F}'$, that is a subsheaf of $\mathcal{F}$, such that* $\mathrm{m.e.}(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F})$. *Furthermore, let $\mathcal{F}' \subset \mathcal{F}$ be sheaves on a graph, G, with $-\chi(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F})$ (which includes the situation in the previous sentence); then we have* $\mathrm{m.e.}(\mathcal{F}/\mathcal{F}') = 0$.

PROOF. Let $\mathcal{F}$ be a sheaf on G, and let $U \subset \mathcal{F}(V)$ be the minimum subspace of $\mathcal{F}(V)$ on which the maximum excess occurs. Let $\mathcal{F}'$ be the subsheaf of $\mathcal{F}$ such that $\mathcal{F}'(V) = U$ and $\mathcal{F}'(E) = \Gamma_{\mathrm{ht}}(U)$. We have that $\mathrm{m.e.}(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F})$ and the maximum excess of $\mathcal{F}'$ occurs at and only at $\mathcal{F}'(V)$ (by the minimality of U). This establishes the first sentence in the lemma. In particular

$$\mathrm{m.e.}(\mathcal{F}) = \mathrm{m.e.}(\mathcal{F}') = -\chi(\mathcal{F}').$$

For the second sentence of the lemma, we claim that $\mathcal{F}/\mathcal{F}'$ has maximum excess zero, for if not then we have compartmentalized

$$U \subset \mathcal{F}(V)/\mathcal{F}'(V), \quad W \subset \mathcal{F}(E)/\mathcal{F}'(E)$$

with $d_h W, d_t W \subset U$ and $\dim(U) < \dim(W)$. So let U' be the inverse image of U in $\mathcal{F}(V)$ (under the map $\mathcal{F}(V) \to \mathcal{F}(V)/\mathcal{F}'(V)$), and W' that of W in $\mathcal{F}(E)$. We have that U' and W' are compartmentalized. If $w' \in W'$, we claim that $d_{h,\mathcal{F}}w'$ must lie in U'; indeed, $[w']$, the class of w' in $\mathcal{F}(V)/\mathcal{F}'(V)$, is taken to U via $d_{h,\mathcal{F}/\mathcal{F}'}$, and we have a commutative diagram

$$\begin{array}{ccc} \mathcal{F}(E) & \longrightarrow & \mathcal{F}(E)/\mathcal{F}'(E) \\ \big\downarrow & & \big\downarrow \\ \mathcal{F}(V) & \longrightarrow & \mathcal{F}(V)/\mathcal{F}'(V) \end{array}$$

and particular elements

$$\begin{array}{ccc} w' \in W' & \longmapsto & [w'] \in W \\ \big\downarrow & & \big\downarrow \\ d_{h,\mathcal{F}}w' & \longmapsto & [d_{h,\mathcal{F}}w'] = d_{h,\mathcal{F}/\mathcal{F}'}[w'] \in U \end{array}$$

Hence $[d_{h,\mathcal{F}}w']$, the class of $d_{h,\mathcal{F}}w'$ in $\mathcal{F}(V)/\mathcal{F}'(V)$, lies in U and hence $d_{h,\mathcal{F}}w'$ lies in U'. Similarly $d_{t,\mathcal{F}}w'$ lies in U', and hence $W' \subset \Gamma_{\rm ht}(U')$. Since U', W' are compartmentalized, it follows that

$$\begin{aligned}\mathrm{excess}(\mathcal{F}, U') &\geq \dim(W') - \dim(U') \\ &= \dim(W) + \dim(\mathcal{F}'(E)) - \dim(U) - \dim(\mathcal{F}'(V)).\end{aligned}$$

Since $\dim(\mathcal{F}'(E)) - \dim(\mathcal{F}'(V)) = -\chi(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F}') = \mathrm{m.e.}(\mathcal{F})$, the above displayed equation implies that

$$\mathrm{excess}(\mathcal{F}, U') \geq \dim(W) - \dim(U) + \mathrm{m.e.}(\mathcal{F}) \geq 1 + \mathrm{m.e.}(\mathcal{F})$$

which is a contradiction. □

Returning to the proof of Theorem 1.10, we claim that it suffices to establish it for tight sheaves; indeed, consider an arbitrary sheaf, $\mathcal{F}$, and apply Lemma 1.45 to obtain a sheaf tight sheaf, $\mathcal{F}'$, as described in the lemma. For any map $\phi\colon G' \to G$, we have an exact sequence

$$0 \to \phi^*\mathcal{F}' \to \phi^*\mathcal{F} \to \phi^*(\mathcal{F}/\mathcal{F}') \to 0.$$

We have that $\mathcal{F}/\mathcal{F}'$ has maximum excess zero, and hence so does $\phi^*(\mathcal{F}/\mathcal{F}')$; by Lemma 1.32,

$$h_1^{\rm twist}\big(\phi^*(\mathcal{F}/\mathcal{F}')\big) = 0$$

provided that ϕ is a covering map with the Abelian girth of G' at least

$$\begin{aligned}&2\Big(\dim\big((\mathcal{F}/\mathcal{F}')(V)\big) + \dim\big((\mathcal{F}/\mathcal{F}')(E)\big)\Big) + 1 \\ \leq\ &2\Big(\dim\big(\mathcal{F}(V)\big) + \dim\big(\mathcal{F}(E)\big)\Big) + 1.\end{aligned}$$

In this case we get in the long exact sequence beginning

$$0 \to H_1^{\rm twist}(\phi^*\mathcal{F}') \to H_1^{\rm twist}(\phi^*\mathcal{F}) \to H_1^{\rm twist}(\phi^*(\mathcal{F}/\mathcal{F}')) \to \cdots$$

amounts to

$$0 \to H_1^{\rm twist}(\phi^*\mathcal{F}') \to H_1^{\rm twist}(\phi^*\mathcal{F}) \to 0,$$

or

$$H_1^{\rm twist}(\phi^*\mathcal{F}') \simeq H_1^{\rm twist}(\phi^*\mathcal{F}).$$

Hence to prove Theorem 1.10 for all $\mathcal{F}$ of a given maximum excess, it suffices to prove it for those of the $\mathcal{F}$ that are tight.

We finish the proof by induction on $\mathrm{m.e.}(\mathcal{F})$ via a second exact sequence.

Lemma 1.46. *Let $\mathcal{F}$ be a tight sheaf on a graph, G, of maximum excess at least one. Then there exists a subsheaf, $\mathcal{F}''$, of $\mathcal{F}$, such that*

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

and such that $\mathcal{F}/\mathcal{F}''$ is edge supported and $\dim((\mathcal{F}/\mathcal{F}'')(E)) = 1$.

Proof. Let $\mathcal{F}''$ be any subsheaf such that $\mathcal{F}''(V) = \mathcal{F}(V)$ and $\mathcal{F}''(E)$ is a codimension one subspace of $\mathcal{F}(E)$. Then $\mathcal{F}/\mathcal{F}''$ is edge supported with the dimension of $(\mathcal{F}/\mathcal{F}'')(E)$ equal one. We claim that, furthermore, the maximum excess of $\mathcal{F}''$ is $\mathrm{m.e.}(\mathcal{F}) - 1$; indeed this excess is achieved by $\mathcal{F}''(V) = \mathcal{F}(V)$; furthermore, for any U properly contained in $\mathcal{F}''(V) = \mathcal{F}(V)$ we have

$$\mathrm{excess}(\mathcal{F}'', U) \leq \mathrm{excess}(\mathcal{F}, U) \leq \mathrm{m.e.}(\mathcal{F}) - 1.$$

□

We now prove Theorem 1.10 by induction upon m.e.$(\mathcal{F})$. The base case, m.e.$(\mathcal{F}) = 0$, was established in Lemma 1.32. Assume that we have established that Theorem 1.10 holds whenever m.e.$(\mathcal{F}) \le k$ for some integer $k \ge 0$. We wish to prove Theorem 1.10 for all $\mathcal{F}$ of maximum excess $k+1$, and we know it suffices to do so when $\mathcal{F}$ is tight. So let $\mathcal{F}$ be a tight sheaf of maximum excess of $k+1$, and let $\mathcal{F}''$ be any subsheaf as in Lemma 1.46. Then Theorem 1.10 holds for $\mathcal{F}''$, since $\mathcal{F}''$ has maximum excess k; so for $\phi\colon G' \to G$ of girth greater than

$$\begin{aligned} & 2\Big(\dim\big(\mathcal{F}''(V)\big) + \dim\big(\mathcal{F}''(E)\big)\Big) \\ \le\ & 2\Big(\dim\big(\mathcal{F}(V)\big) + \dim\big(\mathcal{F}(E)\big)\Big) \end{aligned}$$

we have

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') = \mathrm{m.e.}(\phi^*\mathcal{F}'') = \deg(\phi)k. \tag{1.69}$$

Since, by the construction of $\mathcal{F}''$ in Lemma 1.46, we have

$$\chi(\mathcal{F}'') = \chi(\mathcal{F}) + 1;$$

by tightness of $\mathcal{F}$ we have $\chi(\mathcal{F}) = -k-1$ and hence

$$-\chi(\mathcal{F}'') = k = \mathrm{m.e.}(\mathcal{F}'');$$

hence

$$h_0^{\mathrm{twist}}(\phi^*\mathcal{F}'') = \chi(\phi^*\mathcal{F}'') + h_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') = \deg(\phi)(-k) + \mathrm{m.e.}(\phi^*\mathcal{F}'')$$

$$= \deg(\phi)(-k) + \deg(\phi)(k) = 0.$$

We have a short exact sequence

$$0 \to \phi^*\mathcal{F}'' \to \phi^*\mathcal{F} \to \phi^*(\mathcal{F}/\mathcal{F}'') \to 0,$$

which yields the long exact sequence

$$0 \to H_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') \to H_1^{\mathrm{twist}}(\phi^*\mathcal{F}) \to H_1^{\mathrm{twist}}(\phi^*(\mathcal{F}/\mathcal{F}'')) \to 0,$$

since $h_0^{\mathrm{twist}}(\phi^*\mathcal{F}'') = 0$. Hence

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}) = h_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') + h_1^{\mathrm{twist}}(\phi^*(\mathcal{F}/\mathcal{F}'')). \tag{1.70}$$

But according to Lemma 1.46, $\mathcal{F}/\mathcal{F}''$ is edge supported, and we therefore know that Theorem 1.10 holds for $\mathcal{F}/\mathcal{F}''$ for any covering map, ϕ, and hence

$$h_1^{\mathrm{twist}}(\phi^*(\mathcal{F}/\mathcal{F}'')) = \deg(\phi)\mathrm{m.e.}(\mathcal{F}/\mathcal{F}'') = \deg(\phi).$$

Therefore equations (1.69) and (1.70) shows that

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}) = \deg(\phi)(k+1) = \mathrm{m.e.}(\phi^*\mathcal{F})$$

This establishes Theorem 1.10 for all tight $\mathcal{F}$ with m.e.$(\mathcal{F}) = k+1$.

Hence, by induction on the maximum excess of $\mathcal{F}$, Theorem 1.10 holds for all sheaves, $\mathcal{F}$, on G.

$\square$

1.9. Concluding Remarks

In this section we conclude with a few remarks about the results in this chapter and ideas for further research.

We would like to know how much we can prove about the maximum excess without appealing to homology theory. Our main application of homology theory to the maximum excess was Theorem 1.10, which implies that the maximum excess is a first quasi-Betti number. But part of the proof of Theorem 1.10, namely Subsection 1.8.7, involved a lot of direct reasoning about the maximum excess and short exact sequences. While we believe that the interaction between twisted homology and maximum excess is interesting, we also think that a treatment of maximum excess without homology might give some new insights into the maximum excess.

The maximum excess gives an interpretation of the limit of

$$h_i^{\rm twist}(\phi^*\mathcal{F})/\deg(\phi)$$

over covering maps $\phi\colon G' \to G$ for a sheaf, $\mathcal{F}$, of $\mathbb{F}$-vector spaces on a digraph G. It would be interesting to have an interpretation of

$$\lim_{\phi} \frac{\dim(\mathrm{Ext}^i(\phi^*\mathcal{F}, \phi^*\mathcal{G}))}{\deg(\phi)}$$

for any sheaves $\mathcal{F}, \mathcal{G}$; the maximum excess gives the interpretation in the special case where $\mathcal{G}$ is the structure sheaf, $\underline{\mathbb{F}}$, in which case the Ext groups reduce to (duals of) homology groups. We would also be interesting in generalizations of this to a wider class of settings, such as an arbitrary finite category, or an interesting subclass such as semitopological categories (defined as categories where any morphism of an object to itself must be the identity morphism; see [**Fri05**]).

We would also be interested in knowing if there is a good algorithm for computing the maximum excess of a sheaf exactly, or even just giving interesting upper and lower bounds on it. This would also be interesting for certain types of sheaves. For example, it would be interesting to know classes of sheaves for which the first twisted Betti number equals the maximum excess, in addition to edge simple sheaves of Theorem 1.22.

Notice that if G is an undirected graph, all the discussion in this chapter goes through. Either one can orient each edge and use the notation in this chapter, or just rewrite the notation in this chapter without reference to heads or tails. We see that the distinction between heads and tails is never essential. For example, rather than having twists at the tails of edges, we can have them at the heads and tails of edges. Rather than define a canonical $d = d_{\mathcal{F}}$ to define homology, we simply define homology as

$$\mathrm{Ext}^i(\mathcal{F}, \underline{\mathbb{F}})^\vee,$$

which, by the injective resolution of $\underline{\mathbb{F}}$, becomes the homology groups of

$$\cdots \to 0 \to \oplus_e \mathcal{F}(e) \to \oplus_v \mathcal{F}(v) \to 0,$$

where each $\mathcal{F}(e)$ is really

$$(\mathcal{F}(e))^2/\Delta_e$$

where Δ_e is the diagonal in $(\mathcal{F}(e))^2$ (see the discussion regarding equation (1.19) that appears just below equation (1.20)). Choosing an identification of $(\mathcal{F}(e))^2/\Delta_e$ with $\mathcal{F}(e)$ via $(a,b) \mapsto a-b$ or $(a,b) \mapsto b-a$ amounts to choosing an orientation for

e. The price of giving a "canonical" treatment of the undirected case, i.e., avoiding edge orientations, is that one has to work with $(\mathcal{F}(e))^2/\Delta_e$ instead of $\mathcal{F}(e)$.

CHAPTER 2

The Hanna Neumann Conjecture

2.1. Introduction

Howson, in [**How54**], showed that if $\mathcal{K}, \mathcal{L}$ are nontrivial, finitely generated subgroups of a free group, $\mathcal{F}$, then $\mathcal{K} \cap \mathcal{L}$ is finitely generated, and moreover that

$$\operatorname{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \leq 2 \operatorname{rank}(\mathcal{K}) \operatorname{rank}(\mathcal{L}) - \operatorname{rank}(\mathcal{K}) - \operatorname{rank}(\mathcal{L}). \tag{2.1}$$

Hanna Neumann, in [**Neu56, Neu57**] improved this bound to what is now called the Hanna Neumann Bound,

$$\operatorname{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \leq 2 \big(\operatorname{rank}(\mathcal{K}) - 1\big) \big(\operatorname{rank}(\mathcal{L}) - 1\big); \tag{2.2}$$

furthermore, she conjectured that one can remove the factor of 2 in this bound, i.e., that

$$\operatorname{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \leq \big(\operatorname{rank}(\mathcal{K}) - 1\big) \big(\operatorname{rank}(\mathcal{L}) - 1\big); \tag{2.3}$$

this conjecture is now known as the Hanna Neumann Conjecture (or HNC). One goal of this chapter is to prove the HNC. Moreover, we shall prove a strengthened form of the conjecture, first studied by Walter Neumann in [**Neu90**], known as the Strengthened Hanna Neumann Conjecture (or SHNC); we will state the strengthened conjecture in the next section.

Theorem 2.1. *The Hanna Neumann Conjecture and the Strengthened Hanna Neumann Conjecture hold.*

These conjectures have received considerable attention (see [**Bur71, Imr77b, Imr77a, Ser83, Ger83, Sta83, Neu90, Tar92, Dic94, Tar96, Iva99, Arz00, DF01, Iva01, Kha02, MW02, JKM03, Neu07, Eve08, Min10**]). However, our proof uses very different methods from the previous papers.

The main new idea in our approach to the SHNC is to reduce it to the vanishing maximum excess of a type of sheaf we call a *ρ-kernel*. Although this was described in the introduction to this paper, we can be a bit more precise here in view of the developments in Chapter 1. The SHNC has a well-known reformulation in terms of an inequality involving the reduced cyclicity of graphs; we shall reformulate this in terms of two graphs, the inequality now saying that the reduced cyclicity of one graph is less than that of another. This would follow if we can (1) realize both graphs as sheaves over some base graph, G, (2) find a surjection of the first onto the second, and (3) show that the kernel (a sheaf) has vanishing maximum excess, in view of the fact that the maximum excess is a first quasi-Betti number than reduces to the reduced cyclicity on a sheaves associated to graphs. We shall use Galois graph theory to carry this out; then the base graph, G, will be a Cayley graph; the resulting kernels (for the SHNC) will be called ρ-kernels. It is interesting to note the surjections we use for the SHNC don't generally exist as surjections of graphs, rather only as surjections of sheaves; hence in representing graphs as sheaves, the "additional morphisms" we get are crucial to the construction of ρ-kernels and hence to our proof of the SHNC.

It turns out that some ρ-kernels have nonvanishing maximum excess (at least if one defines ρ-kernels in a broad sense). However, any graph, L, of interest to us in the SHNC, will have a family of associated ρ-kernels, and we will prove that the generic maximum excess in this family is zero, for each L. To do this we shall use Galois theory and symmetry to argue that if this generic maximum excess does not vanish then it is large, i.e., a multiple of the order of the associated Galois group. Then we will give an inductive argument, showing that if the generic maximum excess of ρ-kernels for L is positive, then the same is true when we remove some edge from L, provided that L has positive reduced cyclicity. The base case of the induction, when L has vanishing reduced cyclicity, is trivial to establish. Hence each L of interest in the SHNC has a ρ-kernel of vanishing maximum excess, and this establishes in SHNC.

We emphasize that our proof of the SHNC also uses Theorem 1.10, that implies that the "maximum excess" is first quasi-Betti number. However, it is quite possible that one can prove the inequalities we need for the SHNC for without Theorem 1.10. Furthermore, in the conclusion to this chapter we will give a slight variant of our proof of the SHNC that will avoid any reference to Theorem 1.10 or any homology theories (although this would require the lengthy combinatorial argument of Appendix A). However, regardless of how we present the proof, we shall explain in the conclusion that the homology theory can provide valuable insight into the maximum excess.

This paper shows that the SHNC is not merely an attempt to improve an inequality by a factor of two; our study of the SHNC has lead to new ideas in sheaves on graphs that can be applied to graph theory. This came as a surprise to us at first, although it is perhaps less surprising in retrospect, for a number of reasons. First, the HNC and SHNC seem to describe a fairly fundamental question in group theory (of how rank behaves under intersection). Second, the SHNC can be viewed as a graph theory question involving the reduced cyclicity, which is an interesting graph invariant (e.g. it scales under covering maps). Third, the SHNC, viewed in terms of the Galois graph theory, has a simple homological explanation, namely the vanishing of a limit homology group. The vanishing of (co)homology groups has a vast literature and importance; the SHNC is an interesting and seemingly difficult result in the family of homology group vanishing theorems. Fourth, Lior Silberman has pointed out to us that the reduced cyclicity is the discrete analogue of L^2 Betti numbers; the L^2 Betti number was defined first by Atiyah ([**Ati76**]), and has been the subject of much surrounding the "Atiyah conjecture" (see [**Lüc02**]). Mineyev's article, [**Min10**], also makes a connection between the SHNC and L^2 Betti numbers.

At this point we can give more motivation for the use sheaf theory in this chapter, i.e., why we do not just use graphs and their homology. Our reformulation of the SHNC begins by searching for a morphism involving the graphs of interest to the SHNC. In order for this morphism to exist, to be surjective, and to have a kernel, we must work with more general objects than graphs. In many topological situations, the topological spaces are sufficiently "robust" that one does not have to generalize the objects. However, in non-Hausdorff spaces, such as graphs or those in algebraic geometry, many geometric notions, such as "connect two points with a path," "form a cone," etc., don't make sense or are very awkward to implement. So for graphs we use sheaf theory, which is a simple (co)homology theory that is adapted to our spaces, but general and expressive enough for appropriate surjections and kernels to exist. Of course, it is possible that there are other reasonable frameworks that one could use instead of sheaves.

The rest of this chapter is organized as follows. In Section 2.2 we describe the SHNC and previous work on the HNC and SHNC, including some resolved special cases of the SHNC. In Section 2.3 we give a common graph theoretic reformulation of the SHNC. In Section 2.4 we describe applications of "graph Galois theory" to simplifying the SHNC in a way that leads to the construction of ρ-kernels; this builds on some of the Galois graph theory described ealier, in Section 1.3. In Section 2.5 we construct ρ-kernels and prove that if their maximum excesses vanish then the SHNC holds; we also describe what we call "k-th power kernels," which generalize ρ-kernels and which will be necessary to prove our main theorems about the generic maximum excess of ρ-kernels. In Section 2.6 we use symmetry to prove that the generic maximum excess of a certain type of k-th power kernel is divisible by the order of a group associated to the class of kernels. In Section 2.7 we prove some comparison theorems about how the maximum excesses of different classes of k-th power kernels compare; the main theorem that we prove shows that if ρ-kernels associated to a graph, L, have positive generic maximum excess, then the same is true of a graph L' that consist of L with one edge discarded, provided that $\rho(L') = \rho(L) - 1$. In Section 2.8 we briefly combine a number of theorems of previous sections to argue that if the SHNC does not hold, then for some L we have the class of ρ-kernels associated to L have positive generic maximum excess, which by the results of Section 2.7 means that the same is true for some L with $\rho(L) = 0$, which we easily show is impossible. This establishes the SHNC. In Section 2.9 we make some concluding remarks, including a variant of our proof that avoids Theorem 1.10 and any reference to homology theory. Such a proof will require Appendix A, where we show that vanishing maximum excess of enough ρ-kernels implies the SHNC, but we do so just using elementary graph theory; this shows that a lot of the sheaf theory can be translated into direct graph theoretic terms; it also shows that a simple sheaf theoretic calculation may translate into a much longer graph theoretic calculation.

2.2. The Strengthened Hanna Neumann Conjecture

In this section we state the SHNC and comment on previous work on the HNC and SHNC, including some established special cases of these conjectures.

Walter Neumann, in [**Neu90**], showed that the Hanna Neumann Bound, i.e., equation (2.2), could be strengthened to

$$\sigma(\mathcal{K},\mathcal{L}) \le 2\ \mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}),$$

where $\mathrm{rk}_n(\mathcal{G})$ denotes $\max(\mathrm{rank}(\mathcal{G})+n,0)$, and where

$$\sigma(\mathcal{K},\mathcal{L}) = \sum_{\mathcal{K}x\mathcal{F}\in\mathcal{K}\backslash\mathcal{F}/\mathcal{L}} \mathrm{rk}_{-1}(\mathcal{K}\cap x^{-1}\mathcal{L}x),$$

the summation being over the double coset, $\mathcal{K}\backslash\mathcal{F}/\mathcal{L}$, representatives, x; taking x be the identity in the summation shows that

$$\mathrm{rk}_{-1}(\mathcal{K}\cap\mathcal{L}) \le \sigma(\mathcal{K},\mathcal{L}),$$

so that Walter Neumann's above bound strengthens the Hanna Neumann Bound. Walter Neumann further formulated the conjecture that

$$\sigma(\mathcal{K},\mathcal{L}) \le \mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}), \tag{2.4}$$

now known as the Strengthened Hanna Neumann Conjecture (or SHNC). For the rest of this section we review previous work on the HNC and SHNC.

One collection of results on the problem involves general bounds on $\sigma(\mathcal{K},\mathcal{L})$ or $\mathrm{rk}_1(\mathcal{K}\cap\mathcal{L})$. It turns out that all general bounds we know for the HNC, i.e., on $\mathrm{rk}_1(\mathcal{K}\cap\mathcal{L})$, also are known to hold for $\sigma(\mathcal{K},\mathcal{L})$. Also, all bounds we know are of the form

$$\sigma(\mathcal{K},\mathcal{L}) \le 2\ \mathrm{rank}(\mathcal{K})\ \mathrm{rank}(\mathcal{L}) + c_1\ \mathrm{rank}(\mathcal{K}) + c_2\ \mathrm{rank}(\mathcal{L}) + c_3$$

for ranks $\mathcal{K},\mathcal{L}$ sufficiently large, where c_1, c_2, c_3 are constants depending on the bound; thus all improvements of Howson's original bound are in the lower order terms, i.e., in the c_i's. The improved bounds on $\sigma(\mathcal{K},\mathcal{L})$ after [**How54, Neu56, Neu57**] include the bound

$$2\ \mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) - \min(\mathrm{rk}_{-1}(\mathcal{K}), \mathrm{rk}_{-1}(\mathcal{L}))$$

of Burns in [**Bur71**][1], the bound

$$\mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) + \max(\mathrm{rk}_{-2}(\mathcal{K})\ \mathrm{rk}_{-2}(\mathcal{L}) - 1, 0)$$

of Tardos [**Tar92, Tar96**], and, what is the best bound prior to ours,

$$\mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) + \mathrm{rk}_{-3}(\mathcal{K})\ \mathrm{rk}_{-3}(\mathcal{L}) \tag{2.5}$$

of Dicks and Formanek in [**DF01**].

Another collection of results concerns special cases of the HNC and SHNC that are resolved. To be precise, say that the "HNC holds for $(\mathcal{K},\mathcal{L})$" if equation (2.3) holds, and say that $\mathcal{K}$ is *universal for the HNC* if for any $\mathcal{L}$, the HNC holds for $(\mathcal{K},\mathcal{L})$. Similarly for the SHNC and equation (2.4). Similar to before, all results we know that resolve special cases of the HNC also resolve those cases of the SHNC. Note that any finitely generated free group, $\mathcal{F}$, is a subgroup of $\mathcal{F}_2$, the free group on two generators, so we are free to assume that $\mathcal{F}=\mathcal{F}_2$ in the HNC and SHNC. Here are some results on special cases of the SHNC that are easy to describe in group theoretic terms:

(1) $\mathcal{K}$ is universal for the SHNC if it is of rank at most three ([**DF01**]), in view of equation (2.5), with rank two settled earlier by Tardos ([**Tar92**]);
(2) $\mathcal{K}$ is universal for the SHNC if it is positively generated (see [**Kha02, MW02, Neu07**]);
(3) $\mathcal{K}$ is universal for the SHNC for "most" $\mathcal{K}$ (see [**Neu90, JKM03**]);
(4) the SHNC holds either for $(\mathcal{K},\mathcal{L})$ or for $(\mathcal{K},\mathcal{L}')$ for any $\mathcal{K},\mathcal{L}$ that are subgroups of $\mathcal{F}_2$, where $\mathcal{L}'$ is obtained from $\mathcal{L}$ by the map taking each generator of $\mathcal{F}_2$ to its inverse (see [**JKM03**]).

[1] Bounds appearing before [**Neu90**] are stated as bounds on $\mathrm{rk}_{-1}(\mathcal{K}\cap\mathcal{L})$, but actually give bounds on $\sigma(\mathcal{K},\mathcal{L})$ as well.

The result of item (3) on "most" groups, of Walter Neumann ([**Neu90**]), and some additional results on the SHNC, such as Corollary 3.2 of [**MW02**], are easier to describe using a graph theoretic formulation of the SHNC that we give in the next section. It is also known that the SHNC is related to the coherence problem in one-relator groups ([**Wis05**]).

2.3. Graph Theoretic Formulation of the SHNC

The goal of this section is to describe an equivalent formulation of the HNC and SHNC in graph theoretic terms involving fibre products; this formulation is implicit in [**How54**], but more explicit in [**Imr77b, Imr77a, Ger83, Sta83, Neu90**] and other references in [**Dic94**]. There is another equivalent reformulation of the SHNC by Dicks in [**Dic94**], known as the "amalgamated graph conjecture," which we do not discuss here.

By a *bicoloured digraph*, or simply a *bigraph*, we mean a directed graph, G, such that each edge is coloured (or labelled) either "1" or "2." It is also equivalent to giving a directed graph homomorphism $\nu\colon G \to B_2$, where B_2 is the graph with one vertex and two self-loops, one coloured "1" and the other "2." If, moreover, ν is étale, we call ν or (somewhat abusively) G an étale bigraph, which means that G is a bigraph such that no vertex has two incident edges, both incoming or both outgoing, of the same colour.

Given a digraph, G, recall the definition of the reduced cyclicity from equation (1.1), where $\mathrm{conn}(G)$ denotes the connected components of G; set

$$\rho'(G) = \max_{X \in \mathrm{conn}(G)} \Big(\max\big(0, h_1(X) - 1\big) \Big).$$

The HNC is equivalent to

$$\rho'(K \times_{B_2} L) \le \rho(K)\rho(L) \tag{2.6}$$

for all étale bigraphs K and L; the SHNC is equivalent to

$$\rho(K \times_{B_2} L) \le \rho(K)\rho(L) \tag{2.7}$$

for all étale bigraphs K and L (see [**How54, Imr77b, Imr77a, Ger83, Sta83, Neu90, Dic94**]). We shall work with this form of the SHNC. Again, we say the HNC or SHNC, respectively, holds for a pair of étale bigraphs, (K, L), if equation (2.6) or (2.7), respectively, holds; and again, we say that K is *universal* for the HNC or SHNC, respectively, if for any L the same conjecture holds for (K, L).

Let us briefly explain the connection between the group theoretic formulations of the HNC and SHNC and the graph theoretic formulations. Given generators, g_1, g_2, for the free group, $\mathcal{F}_2$, for each subgroup $\mathcal{K} \subset \mathcal{F}_2$, there is a canonically associated étale bigraph, K; K is given by constructing the Schreier coset graph, $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, and letting K be the "core" of $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, i.e., its smallest subgraph containing all reduced loops based at the vertex $\mathcal{K}$ (see [**MW02**] or the references in the previous paragraph); $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$ with directed edges labelled either g_1 or g_2 is a (typically infinite degree) covering of B_2, and K, a finite subgraph of $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, is therefore an étale bigraph. If $\mathcal{K}, \mathcal{L}$ are subgroups of $\mathcal{F}_2$, and K, L the corresponding étale bigraphs, then each component of $K \times_{B_2} L$ corresponds to the graph associated to $\mathcal{K} \cap x^{-1}\mathcal{L}x$ ranging over double coset representatives, x.

Theorem 2.1 will be proven by the following equivalent theorem.

THEOREM 2.2. *The Strengthened Hanna Neumann Conjecture holds. That is, if $K \to B_2$ and $L \to B_2$ are two étale bigraphs over B_2, then*

$$\rho(K \times_{B_2} L) \le \rho(K)\rho(L). \tag{2.8}$$

Equation (2.8) is tight in that if either K or L is a covering of B_2 (i.e., has all vertices of degree four), then the inequality is satisfied with equality.

2.4. Galois and Covering Theory in the SHNC

In this section develop more aspects of Galois theory, in addition to those given in Section 1.3. This will later lead us to sheaves we call ρ-kernels. Let us now give definitions and state the main theorems to be developed in this section.

DEFINITION 2.3. *By the* Cayley bigraph on a group, $\mathcal{G}$, with generators g_1 and g_2, *denoted* $G = \mathrm{Cayley}(\mathcal{G}; g_1, g_2)$, *we mean the étale bigraph, G, where $V_G = \mathcal{G}$ and $E_G = \mathcal{G} \times \{1, 2\}$ (as sets), such that for each $g \in \mathcal{G}$ and $i = 1, 2$, the edge (g, i) has colour i, tail g, and head $g_i g$.*

We reduce the SHNC to the special case of subgraphs of a Cayley graph, as follows.

THEOREM 2.4. *To prove Theorem 2.2, the SHNC, it suffices verify the SHNC on all pairs, (L, L'), such that L, L' are subgraphs of the same Cayley bigraph. In particular, to prove the SHNC it suffices to show that any subgraph of a Cayley bigraph is universal for the SHNC.*

The following simplifications of the SHNC on subgraphs of Cayley graphs will help solidify the connection between the SHNC and ρ-kernels of the next section.

THEOREM 2.5. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Then*

1. *L is universal for the SHNC if for any étale $L' \to G$ we have (L, L') satisfies the SHNC (with L' inheriting the edge colouring from G, i.e., from the composition $L' \to G$ followed by $G \to B_2$);*
2. *for any étale $L' \to G$ we have*
$$L \times_{B_2} L' \simeq (L\mathcal{G}) \times_G L',$$
where
$$L\mathcal{G} = \coprod_{g \in \mathcal{G}} Lg.$$

Before giving Galois theory we quickly describe the remarkable reason for the strong connection between the SHNC and covering and Galois theory. Since its proof is so short, we give it here as well.

THEOREM 2.6. *For any covering map $\pi\colon K \to G$ of degree d, we have $\chi(K) = d\chi(G)$ and $\rho(K) = d\rho(G)$.*

PROOF. The claim on χ follows since $d = |V_K|/|V_G| = |E_K|/|E_G|$. To show the claim on ρ, it suffices to consider the case of G connected, the general case obtained by summing over connected components; but similarly it suffices to consider the case of K connected. In this case

$$\rho(G) = h_1(G) - 1 = -\chi(G) = -d\chi(K) = d\big(h_1(K) - 1\big) = d\rho(K).$$

□

From this theorem it follows that if $\widetilde{K} \to K$ and $\widetilde{L} \to L$ are covering maps of étale bigraphs, then

$$\rho(\widetilde{K} \times_{B_2} \widetilde{L}) - \rho(\widetilde{K})\rho(\widetilde{L}) = [\widetilde{K}\colon K]\ [\widetilde{L}\colon L]\bigl(\rho(K \times_{B_2} L) - \rho(K)\rho(L)\bigr);$$

hence (K, L) satisfy the SHNC iff $(\widetilde{K}, \widetilde{L})$ do. This means that to study the SHNC, one can always pass to covers of the bigraphs of interest.

For the rest of this section we describe a number of aspects of what we call Galois graph theory and use it for prove Theorems 2.4 and 2.5.

2.4.1. Galois Theory of Graphs. Here we further develop the Galois theory of graphs discussed in Section 1.3. We remind the reader that in this article Galois group actions, when written multiplicatively (i.e., not viewed as functions or morphisms) will be written on the right, since our Cayley graphs are written with its generators acting on the left.

The following fact is an analogue of a standard and surprisingly useful fact in descent theory (as in [**Del77**]); it is also surprisingly useful for the SHNC, despite the fact that it is trivial.

THEOREM 2.7. *Let $\pi\colon K \to G$ be Galois. Then*

$$K \times_G K = \coprod_{\sigma \in \mathrm{Aut}(K/G)} K_\sigma$$

where K_σ is the subgraph of $K \times_G K$ given via

$$V_{K_\sigma} = \{(v, vg) \mid v \in V_K,\ g \in \mathrm{Aut}(K/G)\},$$
$$E_{K_\sigma} = \{(e, eg) \mid e \in E_K,\ g \in \mathrm{Aut}(K/G)\}.$$

Each K_σ is isomorphic to K.

(See [**Fri93**], and compare with the identical formula for fields in [**Del77**], Section I.5.1).

COROLLARY 2.8. *In Theorem 2.7, let us further assume that we have morphisms $K_1 \to K$ and $K_2 \to K$. Then*

$$K_1 \times_G K_2 \simeq \coprod_{\sigma \in \mathrm{Aut}(K/G)} K_1 \times_K (K_2\sigma).$$

PROOF.

$$\begin{aligned} K_1 \times_G K_2 &\simeq (K_1 \times_K K) \times_G (K \times_K K_2) \simeq K_1 \times_K (K \times_G K) \times_K K_2 \\ &\simeq \coprod_\sigma \bigl(K_1 \times_K K_\sigma \times_K K_2\bigr) \simeq \coprod_\sigma \bigl(K_1 \times_K (K_2\sigma)\bigr). \end{aligned}$$

□

There are many extensions to this basic theory. We mention one interesting example.

Assume, for simplicity, that G is connected. If $K \to G$ is Galois and factors as $K \to K' \to G$, then $K \to K'$ is Galois, with Galois group being the subgroup of $\mathrm{Aut}(K/G)$ fixing any vertex or edge fiber of $K \to K'$; hence $\mathrm{Aut}(K/K')$ is a subgroup of $\mathrm{Aut}(K/G)$. Conversely, a subgroup of $\mathrm{Aut}(K/G)$ divides the vertices and edges of K into orbits, giving a graph K' (whose vertices and edges are these orbits) and a factorization $K \to K' \to G$. Furthermore, for an intermediate cover

$K \to K' \to G$, $K \to K'$ is always Galois (since $\mathrm{Aut}(K/K')$ has the right cardinality), and $K' \to G$ is Galois iff the subgroup of $\mathrm{Aut}(K/G)$ fixing $K \to K'$ fibers is a normal subgroup of $\mathrm{Aut}(K/G)$. See [**ST96**] for details.

If $K \to G$ is Galois and factors as $K \to K' \to G$,

$$K \times_G K' = \coprod_{g \in \mathrm{Aut}(K/G)/\mathrm{Aut}(K/K')} K_g, \tag{2.9}$$

where

$$V_{K_g} = \{(v, [v]g) \mid v \in V_K\}, \quad E_{K_g} = \{(e, [e]g) \mid e \in E_K\},$$

where $[v], [e]$ respectively denote the images of v, e, respectively, in K'; each K_g is isomorphic to K. Special cases of this statement include the trivial case $K' = G$ and the case $K' = K$ stated earlier.

2.4.2. Base Change. There are a number of easy "stability under base change" results; these say that in a digram arising from arbitrary digraph morphisms $L \to B$ and $M \to B$,

$$\begin{tikzcd} L \times_B M \arrow[r] \arrow[d] & M \arrow[d] \\ L \arrow[r] & B \end{tikzcd}$$

if $L \to B$ has a certain property, then so does $L \times_B M \to M$. Just from the construction of the fibre product, we easily see that the following classes of morphisms are stable under base change: étale morphisms, covering morphism, and Galois morphisms (and many others that we won't need, such as open inclusions, morphisms that are d–to–1 for some fixed d, etc.).

2.4.3. Etale Factorization. In this subsection we shall prove that any étale map factorizes as an open inclusion followed by a covering map. This will easily establish Theorem 2.4.

We define an *open inclusion* to be any inclusion $H \to G$ of a subgraph, H, in a graph, G. We say the inclusion is *dense* if $V_H = V_G$; this agrees with the topological notion, i.e., the closure of G in H is H, under the topological view of G in [**Fri11b**].

LEMMA 2.9. *Let $\pi\colon G \to B$ be an étale map. Then π factors as an open inclusion, $\iota\colon G \to G'$, followed by a covering map, $\pi'\colon G' \to B$. If the vertex fibres of π (i.e., $\pi^{-1}(v)$ over all $v \in V_B$) are all of the same size, i.e., $\pi_V\colon V_G \to V_B$ is d-to-1 for some d, then we may assume ι is dense; if in addition G is connected, then we may assume G' is connected.*

A variant of the first sentence of this theorem is called Marshall Hall's theorem in [**Sta83**].

PROOF. By adding isolated vertices to G we may assume π_V is d-to-1 for some d. Extend G to G' and π to $\pi'\colon G' \to B$ by completing each $\pi^{-1}(e)$ to a perfect bipartite matching of the vertices over the tail of e and those over the head of e (for a self-loop we view these two sets of vertices as disjoint). Clearly $\pi'\colon G' \to B$ is a covering map. If π_V was originally d-to-1 for some d, then G' is obtained by adding only edges, so G is dense in G'; if furthermore G is connected, then the G' obtained by adding only edges is, of course, still connected. $\square$

Here is an easy, but vital, observation.

LEMMA 2.10. *If $S \to B_2$ is a Galois map with Galois group, $\mathcal{G}$, then S is isomorphic to a Cayley bigraph on the group $\mathcal{G} = \mathrm{Gal}(S/B_2)$.*

PROOF. Choose any $v_0 \in V_S$ to be the "origin." The association $g \mapsto v_0 g$ sets up an identification of $\mathcal{G}$ with V, by definition of a Galois covering map, since there is a unique vertex of B_2 and hence a singe vertex fibre in S. Since $S \to B_2$ is a covering map, the vertex v_0 is the tail of a unique colour 1 edge, e, whose head is $v_0 g_1$ for a unique $g_1 \in \mathcal{G}$. For any g we have eg has tail $v_0 g$ and head $v_0 g_1 g$. It follows that identifying V with $\mathcal{G}$ means that there is an edge $(g, g_1 g)$ (i.e., whose tail is g and head is $g_1 g$) of colour 1 for each $g \in \mathcal{G}$. Similarly for edges of colour 2, and this sets up an isomorphism between S and $\mathrm{Cayley}(\mathcal{G}; g_1, g_2)$. $\square$

PROOF (OF THEOREM 2.4). Let $G \to B_2$ and $K \to B_2$ be étale maps. Let these étale maps factor as open inclusions followed by covering maps as $G \to \widetilde{G} \to B$ and $K \to \widetilde{K} \to B$. Let S be a Galois cover of $\widetilde{G} \times_B \widetilde{K}$. Consider $G' = G \times_{\widetilde{G}} S$, which admits a natural map to G (namely projection onto the first component), and similarly $K' = K \times_{\widetilde{K}} S$. We claim that $G' \to G$ is a covering map; indeed, by stability under base change (see Subsection 2.4.2), since $\widetilde{K} \to B_2$ is a covering map, so is $\widetilde{G} \times_B \widetilde{K} \to \widetilde{G}$; since $S \to \widetilde{G} \times_B \widetilde{K}$ is a covering map, so is $S \to \widetilde{G}$; hence, by base change so is $G' \to G$. Similarly $K' \to K$ is a covering map. According to Theorem 2.6 and the discussion below, the SHNC is satisfied at (G, K) iff it is satisfied at (G', K'). But G', K' are subgraphs of S, and S is a Galois cover of B_2, and therefore a Cayley bigraph. $\square$

Although we shall not need it, we mention that the idea in this last proof can be extended from $G \to B_2$ and $K \to B_2$ to an arbitrary number of étale maps, $L_i \to B_2$, and gives the following interesting fact.

THEOREM 2.11. *For any étale bigraphs $L_1, \ldots, L_k$, there are covering maps $L'_i \to L_i$ and a Cayley bigraph, S, such that each L'_i is a subgraph of S that is dense (i.e., L'_i has the same vertex set as S).*

2.4.4. The Proof of Theorem 2.5. We finish this section with the proof of Theorem 2.5.

PROOF OF THEOREM 2.5. Claim (1) of the theorem is a simple base change argument: if L is a subgraph of a Cayley bigraph, G, and $L' \to B_2$ is any étale bigraph, let $L'' = L' \times_{B_2} G$. Then, by base change (see Subsection 2.4.2), $L'' \to G$ is étale and $L'' \to L'$ is a covering map. Then (L, L') satisfies the SHNC iff (L, L'') does. Hence L is universal for the SHNC iff (L, L'') satisfies the SHNC for all étale bigraphs, L'', whose colouring map factor as $L'' \to G \to B_2$.

Claim (2) is an immediate consequence of Corollary 2.8, with G, K, K_1, K_2 respectively replaced by B_2, G, L, L', noting that $\mathrm{Aut}(G/B_2) = \mathcal{G}$.

$\square$

2.5. ρ-kernels

In this section we introduce a collection of sheaves that are central to our proof of the SHNC. They are called ρ-kernels. Before defining them, we motivate their construction by showing how their study is connected to the SHNC. First we need to set some notation on Cayley graphs.

2.5.1. Sheaves on Cayley graphs. Let $G = \mathrm{Cayley}(\mathcal{G}; g_1, g_2)$ be a Cayley bigraph on a group, $\mathcal{G}$. Recall that since our generators act on the left, e.g., the colour 1 edges are of the form $(g, g_1 g)$, the Galois group of G is $\mathcal{G}$ acting on the right. Now we define a right action of $\mathcal{G}$ on sheaves on G. We shall state this in slightly more general terms. This is completely straightforward and mildly tedious, but convenient in this section and vital to Section 2.6.

DEFINITION 2.12. *We say that a group, $\mathcal{G}$,* acts *on a digraph, G,* on the right*, if associated to each $g \in \mathcal{G}$ is an isomorphism π_g, of G such that $\pi_{g_1 g_2} = \pi_{g_2}\pi_{g_1}$ for all $g_1, g_2 \in G$. We will identify g with π_g if no confusion can arise. If L is a subgraph of G, we write Lg for the image of L under g (i.e., under π_g); similarly if $P \in V_G \amalg E_G$, Pg denotes the image of P under g.*

Of course, if G is a Cayley bigraph on a group, $\mathcal{G}$, then $\mathcal{G}$ acts on G on the right.

THEOREM 2.13. *Let a group, $\mathcal{G}$, act on a digraph, G, on the right. Then each element of $\mathcal{G}$ acts naturally as a functor on sheaves, via the association $g \mapsto \pi^*_{g^{-1}}$, such that*

1. *$\mathcal{G}$ acts on the right, i.e., if we write $\mathcal{F}g$ for $\pi^*_{g^{-1}}\mathcal{F}$ for any sheaf, $\mathcal{F}$, on G, then for any $g_1, g_2 \in \mathcal{G}$ we have $\mathcal{F}g_1g_2 = (\mathcal{F}g_1)g_2$, and similarly with the sheaf $\mathcal{F}$ replaced by a morphism of sheaves;*
2. *for each sheaf, $\mathcal{F}$, on G, any $g \in \mathcal{G}$, and any $P \in V_G \amalg E_G$ we have*
$$(\mathcal{F}g)(P) = \mathcal{F}(Pg^{-1});$$
and
3. *for each subgraph $L \subset G$ and field, $\mathbb{F}$, we have*
$$\underline{\mathbb{F}}_L\, g = \underline{\mathbb{F}}_{Lg}.$$

In Section 2.6 it will be important to use the fact that for each $g \in \mathcal{G}$, $\pi^*_{g^{-1}}$ is a *functor*, i.e., it acts (compatibly) on morphisms of sheaves as well as on sheaves.

PROOF. For item (1), we recall that for any $u\colon G' \to G$, u^* is a functor on sheaves, and for composable morphisms of digraphs, u_1, u_2, we have $(u_1u_2)^* = u_2^*u_1^*$; hence, since $\mathcal{G}$ acts on the right, for any $g_1, g_2 \in \mathcal{G}$ we have
$$\pi^*_{g_1^{-1}}\pi^*_{g_2^{-1}} = (\pi_{g_2^{-1}}\pi_{g_1^{-1}})^* = (\pi_{g_1^{-1}g_2^{-1}})^* = \pi^*_{(g_2g_1)^{-1}}$$
and so $g \mapsto \pi^*_{g^{-1}}$ is defines an action on sheaves and morphisms of sheaves that acts on the right.

Item (2) follows immediately from the definition of the pullback. Item (3) follows since for all $P \in V_G \amalg E_G$ and $L \subset G$ and $g \in \mathcal{G}$ we have
$$(\underline{\mathbb{F}}_L\, g)(P) = \underline{\mathbb{F}}_L(Pg^{-1});$$
but $Pg^{-1} \in L$ iff $P \in Lg$, so
$$\underline{\mathbb{F}}_{Lg}(P) = \underline{\mathbb{F}}_L(Pg^{-1}) = (\underline{\mathbb{F}}_L\, g)(P).$$
Hence $\underline{\mathbb{F}}_{Lg} = \underline{\mathbb{F}}_L\, g$. □

Given a sheaf, $\mathcal{F}$, on G we define
$$\mathcal{F}\mathcal{G} = \bigoplus_{g \in \mathcal{G}} \mathcal{F}g.$$

In particular, for $L \subset G$, if we set

$$L\mathcal{G} = \coprod_{g \in \mathcal{G}} Lg$$

(akin to the notation in Theorem 2.5) then

$$\underline{\mathbb{F}}_L \mathcal{G} \simeq \underline{\mathbb{F}}_{L\mathcal{G}}.$$

2.5.2. Kernels and the SHNC. The following theorem summarizes our approach to the SHNC.

THEOREM 2.14. *Let L be a subgraph of a Cayley bigraph, G. Assume there is an exact sequence*

$$0 \to \mathcal{K} \to \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^{\rho(L)} \to 0 \tag{2.10}$$

such that m.e.$(\mathcal{K}) = 0$. *Then L is universal for the SHNC.*

PROOF. According to Theorem 2.5, it suffices to show that for each étale $u\colon L'' \to G$ we have that (L, L'') satisfies the SHNC. Tensoring equation (2.10) with $\mathbb{F}_{L''} = u_! \underline{\mathbb{F}}$ gives

$$0 \to \mathcal{K} \otimes \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_{L''}^{\rho(L)} \to 0. \tag{2.11}$$

Since m.e.$(\mathcal{K}) = 0$ and $u\colon L'' \to G$ is étale, Theorem 1.16 and the discussion before it implies that

$$\text{m.e.}(\mathcal{K} \otimes \underline{\mathbb{F}}_{L''}) = 0.$$

Since the maximum excess is a first quasi-Betti number, this and equation (2.11) implies that

$$\text{m.e.}(\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''}) \leq \text{m.e.}(\underline{\mathbb{F}}_{L''}^{\rho(L)}) = \rho(L)\rho(L'').$$

But

$$\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''} \simeq \underline{\mathbb{F}}_{(L\mathcal{G}) \times_G L''} \simeq \underline{\mathbb{F}}_{L \times_{B_2} L''}$$

(using equation (1.16) and Corollary 2.8), and so we have

$$\rho(L \times_{B_2} L'') = \text{m.e.}(\underline{\mathbb{F}}_{L \times_{B_2} L''}) = \text{m.e.}(\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''}) \leq \rho(L)\rho(L'').$$

□

2.5.3. Definition and Existence of ρ-Kernels and k-th Power Kernels. We begin with some notation to describe the kernels we introduce here and study throughout the rest of this paper.

Let G be a Cayley bigraph on a group, $\mathcal{G}$. For any integer $k \geq 0$, let $\mathbb{F}^{k \times \mathcal{G}}$ be the set of $k \times |\mathcal{G}|$ matrices with entries $m_{ig} \in \mathbb{F}$ indexed over $i = 1, \ldots, k$ and $g \in \mathcal{G}$. If $M \in \mathbb{F}^{k \times \mathcal{G}}$, then we can view M as a map from $\mathbb{F}^{\mathcal{G}}$ to $\mathbb{F}^k$. Then M gives rise to a morphism of constant sheaves

$$\underline{M}\colon \underline{\mathbb{F}}^{\mathcal{G}} \to \underline{\mathbb{F}}^k.$$

For any $L \subset G$ and $g \in \mathcal{G}$, we have an inclusion $\underline{\mathbb{F}}_{Lg} \to \underline{\mathbb{F}}$, which gives us an inclusion

$$\underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}} \mathcal{G} \simeq \underline{\mathbb{F}}^{\mathcal{G}}.$$

Thus we get a monomorphism

$$\iota_{L\mathcal{G}}\colon \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^{\mathcal{G}},$$

and, for any $M \in \mathbb{F}^{k\times\mathcal{G}}$, a composite morphism

$$\underline{M}\iota_{L\mathcal{G}} \colon \underline{\mathbb{F}}_L\mathcal{G} \to \underline{\mathbb{F}}^k.$$

We shall often write ι instead of $\iota_{L\mathcal{G}}$, since the subscript $L\mathcal{G}$ can be inferred from the source (even if two different ι's are involved).

Definition 2.15. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be a field. For any integer $k \geq 0$, we say that $M \in \mathbb{F}^{k\times\mathcal{G}}$ is* L-surjective *if the map, $\underline{M}\iota_{L\mathcal{G}}$ is surjective. If so, we that its kernel, $\mathcal{K} = \mathcal{K}_M(L, G, \mathcal{G})$, is a k-th power kernel for $(L, G, \mathcal{G})$; if, in addition, $k = \rho(L)$, we also say that $\mathcal{K}$ is a* ρ-kernel for *$(L, G, \mathcal{G})$.*

Note that when kernels are defined in category theory, i.e., for a category with a zero morphism, then a kernel is defined only up to (unique) isomorphism. However, for sheaves on a graph, we can define the kernel of a morphism $\mathcal{F}_1 \to \mathcal{F}_2$ uniquely, as the subsheaf of $\mathcal{F}_1$ that is the kernel. Hence we will speak of *the* kernel of a morphism, or *its* kernel, for convenience; when we say "a kernel" we shall mean the category theory notion, i.e., any morphism $\mathcal{K} \to \mathcal{F}_1$ that is the equalizer of $\mathcal{F}_1 \to \mathcal{F}_2$ and the zero morphism.

Note that we could also define k-th power kernels when $\underline{M}\iota_{L\mathcal{G}}$ is not surjective, as the element of the derived category (see [**GM03**]) as a single shift of the mapping cone of $\underline{M}\iota_{L\mathcal{G}}$; we shall not pursue this here.

The important point to notice is that if $k \leq \rho(L)$, "most" matrices $M \in \mathbb{F}^{k\times\mathcal{G}}$ are L-surjective. We now demonstrate this, in a rather explicit fashion.

Definition 2.16. *We say that $M \in \mathbb{F}^{k\times\mathcal{G}}$ is* totally linearly independent *(or just* totally independent*) if every subset $\mathcal{G}'$ of $\mathcal{G}$ of size k we have $\{m^g\}_{g\in\mathcal{G}'}$ is linearly independent, where m^g denotes the column of M corresponding to $g \in \mathcal{G}$.*

Lemma 2.17. *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Then the number of vertices of L and the number of edges of either colour in L are all at least $\rho(L)$.*

Proof. Adding vertices and edges to a graph does not decrease its reduced cyclicity (i.e., its ρ). So if P is an edge of colour 2, let L' be L union all vertices of G and all edges of colour 1. Then $\rho(L') \geq \rho(L)$ and L' has the same number of edges of colour 2 as L. But if we discard the edges of colour 2 from L' we are left with a union of cycles, for which $\rho = 0$, and discarding one edge decreases ρ by at most one (given equation (1.1)). Hence the number of edges of colour 2 in L' is at least $\rho(L)$, and so the same is true of the number of colour 2 edges in L.

Similarly L must have at least $\rho(L)$ edges of colour 1. Finally, since each vertex of L is the head of at most one edge of colour 1, the number of vertices is also at least $\rho(L)$. $\square$

Now we wish to describe ρ-kernels, both as a kernel of a sheaf morphism and, alternatively, by explicitly giving their values and restrictions.

Definition 2.18. *Fix a subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$. Fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ for an integer $k \geq 0$. For a subset $T \subset \mathcal{G}$, the* T-free subspace of *$\ker(M)$ we mean the set*

$$\mathrm{Free}_T = \mathrm{Free}_T(M) = \{\vec{a} \in \ker(M) \mid \forall g \notin T,\ a_g = 0\}.$$

A free subspace of $\ker(M)$ *is a subspace that is* T*-free for some* $T \subset \mathcal{G}$*. For* $P \in V_G \amalg E_G$*, we set*

$$\mathcal{G}_L(P) = \{g \in \mathcal{G} \mid P \in Lg\}.$$

In the above definition, if $M \in \mathbb{F}^{k\times\mathcal{G}}$ is totally independent, then for all $T \subset \mathcal{G}$ we have

$$\dim(\mathrm{Free}_T) = \max(0, |T| - k). \tag{2.12}$$

LEMMA 2.19. *Let* L *be a subgraph of a Cayley bigraph,* G*, on a graph* $\mathcal{G}$*. Let* $M \in \mathbb{F}^{k\times\mathcal{G}}$ *be totally independent, for some* $k \le \rho(L)$*. Then*

$$\underline{M}\iota\colon \underline{\mathbb{F}}_L\mathcal{G} \to (\underline{\mathbb{F}}_G)^k$$

is surjective. Furthermore, if $\mathcal{K}_M$ *denotes its kernel, then for each* $P \in V_G \amalg E_G$ *we have*

$$\mathcal{K}_M(P) = \mathrm{Free}_{\mathcal{G}_L(P)}(M),$$

in the notation of Definition 2.18, and the restriction maps for $\mathcal{K}_M$ *are the inclusions. In particular,*

$$\dim\bigl(\mathcal{K}_M(P)\bigr) = n_P - \rho(L),$$

where $n_P = |\mathcal{G}_L(P)|$*. (We shall sometimes write* $\mathcal{K}_M$ *as* $\mathcal{K}_M(L)$ *or* $\mathcal{K}_M(L, G, \mathcal{G})$ *to emphasize* $\mathcal{K}_M$*'s dependence upon* L*,* G*, and* $\mathcal{G}$*.)*

PROOF. For each $g \in \mathcal{G}$ we have $\underline{\mathbb{F}}_L g = \underline{\mathbb{F}}_{Lg}$. Hence for each $P \in V_G \amalg E_G$, we have

$$(\underline{\mathbb{F}}_L g)(P) = (\underline{\mathbb{F}}_{Lg})(P) = \begin{cases} \mathbb{F} & \text{if } P \in Lg \\ 0 & \text{if } P \notin Lg \end{cases}$$

Hence

$$(\underline{\mathbb{F}}_L\mathcal{G})(P) \simeq \bigoplus_{g\in\mathcal{G}_L(P)} \mathbb{F},$$

and the image of $\underline{M}\iota$ in $(\mathbb{F}_G)^k$ at P is the span of the subcollection of the n_P columns of M corresponding to the elements of $\mathcal{G}_L(P) \subset \mathcal{G}$. Since G is a Cayley bigraph, n_P is either the number of vertices, edges of colour 1, or edges of colour 2 in L. By Lemma 2.17 we have $n_P \ge \rho(L)$, and hence this subcollection of n_P vectors in M spans $\mathbb{F}^k$. Hence $\underline{M}\iota$ is surjective at P, and its kernel, $\mathrm{Free}_{\mathcal{G}_L(P)}$, is of dimension $n_P - k$. The restriction maps on $\underline{\mathbb{F}}_L\mathcal{G}$ are, component by component, those of the individual $\underline{\mathbb{F}}_L g$ over all $g \in \mathcal{G}$, and those are just inclusions; since $\mathcal{K}$ is a subsheaf of $\underline{\mathbb{F}}_L\mathcal{G}$, we have that $\mathcal{K}$ inherits those restriction maps. $\square$

Note that it is easy to see, even with $\mathcal{G} = \mathbb{Z}/3\mathbb{Z}$ and L consisting of five edges, that there need not be any graph theoretic surjections $L\mathcal{G} \to G^{\rho(L)}$, where $G^{\rho(L)}$ is $\rho(L)$ disjoint copies of G; so in passing from the graphs $L\mathcal{G}$ and $G^{\rho(L)}$ to the sheaves $\underline{\mathbb{F}}_L\mathcal{G}$ and $\mathbb{F}^{\rho(L)}$, there exists a surjection of sheaves that does not arise from any surjection of graphs. So an added benefit of working with sheaves (aside from using them to form kernels useful in studying the SHNC) is that sheaves give "additional surjections" that don't exist in graph theory.

2.6. Symmetry and Algebra of the Excess

In this section we make some general observations about the maximum excess of k-th power kernels. The main observation is that given $(L, G, \mathcal{G})$ as usual, the maximum excess of $\mathcal{K}_M(L)$ for generic $M \in \mathbb{F}^{k\times\mathcal{G}}$ is divisible by $|\mathcal{G}|$, where by "generic" we mean for M in some subset of $\mathbb{F}^{k\times\mathcal{G}}$ that contains a nonempty, Zariski open subset of $\mathbb{F}^{k\times\mathcal{G}}$. Let us outline this argument.

First, in Subsection 2.6.1, we will show that for any $M \in \mathbb{F}^{k\times\mathcal{G}}$ and $g \in \mathcal{G}$ we have $\mathcal{K}_M(L)g \simeq \mathcal{K}_{Mg}(L)$ where Mg is obtained by an appropriate action of g on the columns of M. This means that if $\mathcal{F}$ is the maximal (or minimal) excess maximizer for $\mathcal{K}_M(L)$, then $\mathcal{F}g$ is isomorphic to the maximal (or, respectively, minimal) excess maximizer for $\mathcal{K}_{Mg}(L)$. It may be helpful, albeit somewhat fanciful, to understand this symmetry via two "observers" looking at the exact sequence

$$0 \to \mathcal{K}_M \to \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{M\iota} \underline{\mathbb{F}}^k \to 0,$$

one who examines this at $P \in V_G \amalg E_G$, and the other at Pg, for g fixed and P varying; for example, $\underline{\mathbb{F}}_L\mathcal{G}$ "looks" the same to both observers, except that its summands appear permuted from one observer to the other.

In Subsection 2.6.2 we discuss the generic maximum excess of $\mathcal{K}_M = \mathcal{K}_M(L)$ with L fixed and $M \in \mathbb{F}^{k\times\mathcal{G}}$ variable (and $G, \mathcal{G}, k$ fixed). The key to this discussion is considering what we call "dimension profiles," which we now define.

Definition 2.20. *By a* dimension profile *on a bigraph, G, we mean a function*

$$n\colon V_G \amalg E_G \to \mathbb{Z}_{\geq 0}.$$

For any such n, we set

$$\chi(n) = \sum_{v\in V_G} n(v) - \sum_{e\in E_G} n(e); \qquad |n| = \sum_{P\in V_G\amalg E_G} n(P).$$

Any sheaf, $\mathcal{F}$, on G determines a dimension profile, $\dim(\mathcal{F})$*, as the function $P \mapsto$* $\dim(\mathcal{F}(P))$*. For any dimension profile, n, of a Cayley bigraph, G, on a group, $\mathcal{G}$, any subgraph, $L \subset G$, any field, $\mathbb{F}$, and any $k \geq 0$, let*

$$\mathcal{M}(n) = \mathcal{M}(n, L, G, \mathcal{G}, \mathbb{F}, k)$$

be the set of $M \in \mathbb{F}^{k\times\mathcal{G}}$ for which $\mathcal{K}_M = \mathcal{K}_M(L, G, \mathcal{G})$ exists (i.e., M is L-surjective) and has a subsheaf, $\mathcal{F}$, with $\dim(\mathcal{F}) = n$.

We easily see that for all n, $\mathcal{M}(n)$ is a constructible subset of $\mathbb{F}^{k\times\mathcal{G}}$. Let $\mathcal{N}$ be the set of n for which $\mathcal{M}(n)$ is generic, i.e., contains a Zariski open subset of $\mathbb{F}^{k\times\mathcal{G}}$; since $\mathcal{M}(n)$ is constructible, it is equivalent to say that its Zariski closure is $\mathbb{F}^{k\times\mathcal{G}}$. The generic (in M) value of the maximum excess of $\mathcal{K}_M$ is the largest value of $-\chi(n)$ among those $n \in \mathcal{N}$; let $\mathcal{N}'$ be the subset of $n \in \mathcal{N}$ which attain this largest $-\chi(n)$ value. Since $\mathcal{G}$ is finite, for any $n \in \mathcal{N}$ there is a generic set of M such that $Mg \in \mathcal{M}(n)$ for all $g \in \mathcal{G}$. Hence, if $n \in \mathcal{N}'$ is chosen with $|n|$ at a maximum value (or minimum value), then by the uniqueness of the maximum (or minimum) maximizer of the excess, the symmetry $\mathcal{K}_M g \simeq \mathcal{K}_{Mg}$ implies that $n(P) = n(Pg^{-1})$ for all $P \in V_G \amalg E_G$ and $g \in \mathcal{G}$. Hence the generic maximum excess of $\mathcal{K}_M$, which equals $-\chi(n)$ for any $n \in \mathcal{N}'$, is divisible by $\mathcal{G}$.

We wish to remark that the generic maximum excess is not generally attained by all $M \in \mathbb{F}^{k\times\mathcal{G}}$. For example, our approach to the SHNC is based on the fact that the generic maximum excess of a ρ-kernel is zero, i.e., for $k = \rho(L)$ (see

Theorem 2.33). However, if M is zero in one column, but totally independent in the others, then M will still be L-surjective provided that L has at least $\rho(L)+1$ edges of each colour. (A simple example of such an L can be obtained by deleting one edge of each colour from $\mathrm{Cayley}(\mathbb{Z}/m\mathbb{Z}; 1, 1)$ with $m \geq 2$.) In such a situation, the fact that M has a column of zeros implies that $\mathcal{K}_M$ has $\mathbb{F}_L g$ as a subsheaf (more precisely, a direct summand) for some $g \in \mathcal{G}$, and hence the maximum excess of $\mathcal{K}_M$ will be at least $\rho(L)$. Hence any L that has at least $\rho(L)+1$ edges of each colour, and for which $\rho(L) > 0$, has a ρ-kernel of positive maximum excess. Hence it is essential to study the maximum excess of $\mathcal{K}_M$ with some restrictions on M, i.e., requiring some special properties of M; in our case, these properties restrict M to some generic subset of $\mathbb{F}^{\rho(L)\times\mathcal{G}}$.

2.6.1. Symmetry of k-th Power Kernels. The point of this subsection is to establish the following symmetry of k-th power kernels.

THEOREM 2.21. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be an arbitrary field. Let k be an arbitrary non-negative integer and $M \in \mathbb{F}^{k\times\mathcal{G}}$. Let Mg be the matrix (described earlier) whose g' column, for $g' \in \mathcal{G}$, is the $g'g^{-1}$ column of M. Then M is L-surjective iff Mg is L-surjective, and if so then $\mathcal{K}_M(L)g \simeq \mathcal{K}_{Mg}(L)$.*

PROOF. We begin our discussion of symmetry with a somewhat pedantic, but important, point. If $\mathcal{A}$ is a category in which finite direct sums exists, such as an additive category, and $\{A_s\}_{s\in S}$ is a family of objects in the category indexed upon a finite set, S, then their direct sum comes with projections

$$f_r\colon \bigoplus_{s\in S} A_s \to A_r$$

for each $r \in S$. If $\pi\colon S \to S$ is a permutation, then we have a "component permuting map," $P = P(\pi)$, given by

$$P(\pi)\left(\bigoplus_{s\in S} A_s\right) = \bigoplus_{s\in S} A_{\pi(s)},$$

The two direct sums in this last equation are isomorphic, but not equal (e.g., the direct sum on the right-hand-side has the projection f'_r whose target is $A_{\pi(r)}$, not to A_r, for each $r \in \mathcal{G}$). We shall need to keep the seemingly unimportant operator $P = P(\pi)$ in mind in order to make things precise for this subsection. If $A_\bullet$ is any direct sum indexed on $\mathcal{G}$, then we easily see $P(\pi_2)(P(\pi_1)A_\bullet) = P(\pi_2\circ\pi_1)A_\bullet$.

Again, let $\mathbb{F}$ be a field, $k \geq 0$ an integer, L a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and $M \in \mathbb{F}^{k\times\mathcal{G}}$ that is L-surjective. We have an exact sequence.

$$0 \longrightarrow \mathcal{K} \longrightarrow \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g' \xrightarrow{M\iota} \underline{\mathbb{F}}^k \longrightarrow 0. \tag{2.13}$$

For a $g \in \mathcal{G}$, applying $\pi^*_{g^{-1}}$, of Theorem 2.13, to this sequence gives an exact sequence:

$$0 \longrightarrow \mathcal{K}g \longrightarrow \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\pi^*_{g^{-1}}(M\iota)} \underline{\mathbb{F}}^k g \longrightarrow 0. \tag{2.14}$$

We have $\underline{\mathbb{F}}g = \underline{\mathbb{F}}$ since $\underline{\mathbb{F}}$ is a constant sheaf (note the we mean that the two are equal, not merely isomorphic). Note that $\pi^*_{g^{-1}}$ acts on sheaves by renaming the vertices of G, so it acts on $\underline{M}\iota$ only by permuting sheaf inclusions $\underline{\mathbb{F}}_{Lg''} \to \underline{\mathbb{F}}$ for various values of g''; in other words,

$$\pi^*_{g^{-1}}(\underline{M}\iota_{\underline{\mathbb{F}}_L\mathcal{G}}) = \underline{M}\iota',$$

where ι' is ι with the source $\underline{\mathbb{F}}_L\mathcal{G}g$. Hence we may write equation (2.14) as

$$0 \longrightarrow \mathcal{K}g \xrightarrow{j} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\underline{M}\iota'} \underline{\mathbb{F}}^k \longrightarrow 0, \tag{2.15}$$

where j is an inclusion.

Also, we have

$$P(\pi_{g^{-1}})\left(\bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g\right) = \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, \pi_{g^{-1}}(g')g = \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g' = \underline{\mathbb{F}}_L\mathcal{G}.$$

Hence from equation (2.15) we get a sequence

$$0 \longrightarrow \mathcal{K}g \xrightarrow{j} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{P(\pi_{g^{-1}})} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{P(\pi_g)} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\underline{M}\iota'} \underline{\mathbb{F}}^k \longrightarrow 0,$$

and hence an exact sequence (since $P(\pi_g)$ and $P(\pi_{g^{-1}})$ are isomorphisms)

$$0 \longrightarrow \mathcal{K}g \xrightarrow{P(\pi_{g^{-1}})\circ j} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\underline{M}\iota' P(\pi_g)} \underline{\mathbb{F}}^k \longrightarrow 0, \tag{2.16}$$

with j being the inclusion in equation (2.15). But clearly

$$\underline{M}\iota' P(\pi_g) = \underline{M}P(\pi_g)\,\iota_{\underline{\mathbb{F}}_L\mathcal{G}} = \underline{M\pi_g}\,\iota_{\underline{\mathbb{F}}_L\mathcal{G}},$$

where π_g is viewed as operating vectors in $\mathbb{F}^{\mathcal{G}}$ sending $\alpha \in \mathbb{F}^{\mathcal{G}}$, viewed as a function $\alpha\colon \mathcal{G} \to \mathbb{F}$ to $\pi_g\alpha$ given by

$$g' \mapsto \alpha(\pi_g(g')) = \alpha(g'g).$$

Hence setting $Mg = M\pi_g$, we get a short exact sequence

$$0 \longrightarrow \mathcal{K}g \longrightarrow \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\underline{Mg}\,\iota} \underline{\mathbb{F}}^k \xrightarrow{0}. \tag{2.17}$$

Hence $\mathcal{K}_{Mg}(L)$ is, up to isomorphism, just $\mathcal{K}g$.

To complete the proof of the theorem, it remains to find that permutation that brings the columns of M to those of Mg. If $M = \{m_{i,g'}\}$ give M's entries, then for any $w \in \mathbb{F}^{\mathcal{G}}$, the i-th component of $(Mg)w = M(\pi_g w)$ is

$$(M(\pi_g w))_i = \sum_{g'\in\mathcal{G}} m_{i,g'}(\pi_g w)_{g'} = \sum_{g'\in\mathcal{G}} m_{i,g'}w_{g'g} = \sum_{g''\in\mathcal{G}} m_{i,g''g^{-1}}w_{g''}.$$

Hence the i, g'' entry of Mg is $m_{i,g''g^{-1}}$, so the g'' column of Mg is the $g''g^{-1}$ of M. $\square$

We wish to make a few comments on equation (2.17) and how we derived it. First, kernels, in category theory, are defined only up to isomorphism; this is why we can "forget about" $P(\pi_{g^{-1}})j$ in equation (2.16); it is only important to know that this arrow gives an exact sequence there and in equation (2.17).

Note that the two actions of $g \in \mathcal{G}$ in equation (2.17) are right $\mathcal{G}$ actions on the exact sequence. To see this, first note that

$$(Mg)g' = (M\pi_g)\pi_{g'} = M\pi_{gg'} = Mgg'.$$

Then note that if we take the procedure for going from equation (2.13) to equation (2.17) and then do the same procedure with g replaced by g', then we easily see (paying close attention to the order of the P's, the π's, and the ι's) that we get the same equation as equation (2.13) with $\mathcal{K}$ replaced by $\mathcal{K}gg'$ and M replaced by Mgg'.

We wish to comment on something that seems a bit contradictory. The map $g \mapsto P(\pi_{g^{-1}})$ is a left action, and so it may seem strange that our forgotten monomorphism $P(\pi_{g^{-1}})j$ in equation (2.16) involves a left action. But note that if we apply g' to equation (2.17) we get

$$0 \longrightarrow \mathcal{K}gg' \xrightarrow{\pi^*_{g'^{-1}}(P(\pi_{g^{-1}})\circ j)} \underline{\mathbb{F}}_L\mathcal{G}g' \xrightarrow{\pi^*_{g'^{-1}}(\underline{M}\iota' P(\pi_g))} \underline{\mathbb{F}}^k \longrightarrow 0,$$

and hence an exact sequence

$$0 \longrightarrow \mathcal{K}gg' \xrightarrow{\alpha} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\beta} \underline{\mathbb{F}}^k \longrightarrow 0,$$

where

$$\begin{aligned} \alpha &= P(\pi_{g'^{-1}}) \circ \pi^*_{g'^{-1}}(P(\pi_{g^{-1}}) \circ j) = P(\pi_{g'^{-1}})P(\pi_{g^{-1}})j', \\ \beta &= \pi^*_{g'^{-1}}(\underline{M}\iota' P(\pi_g))P(\pi_{g'}), \end{aligned}$$

for an inclusion j'. Examining α we see that $P(\pi_{g'^{-1}})$ is applied to the left of $P(\pi_{g^{-1}})$, whose product equals $P(\pi_{(gg')^{-1}})$, so that g' appears to the right of g.

A similar remark applies for the column permuting rule taking M to Mg: $g \mapsto \pi_{g^{-1}}$ is a left action, not a right action. However, if $f\colon \mathcal{G} \to T$ is any function from $\mathcal{G}$ to a set, T, then defining a function fg via $(fg)(g') = f(g'g^{-1})$ defines a *right* action of $\mathcal{G}$ on functions from $\mathcal{G}$ to T; indeed, for $f\colon \mathcal{G} \to T$ and $g, g_1, g_2 \in \mathcal{G}$ we have

$$\bigl((fg_1)g_2)\bigr)(g) = (fg_1)(gg_2^{-1}) = f(gg_2^{-1}g_1^{-1}) = f\bigl(g(g_1g_2)^{-1}\bigr) = \bigl(f(g_1g_2)\bigr)(g).$$

So the left action $g \mapsto \pi_{g^{-1}}$ turns into a right action when it acts on the argument of a function.

We finish this subsection with a corollary of Theorem 2.21 that is our sole application of the theorem.

Corollary 2.22. *Let n be a dimension profile for Cayley bigraph, G, on a group, $\mathcal{G}$. Let L be a subgraph of G, let $\mathbb{F}$ be a field, and let $k \geq 0$ be an integer. Then for any $g \in \mathcal{G}$, we have*

$$\mathcal{M}(ng, L, G, \mathcal{G}, k) = \mathcal{M}(n, L, G, \mathcal{G}, k)g,$$

where ng is given by

$$(ng)(P) = n(Pg^{-1})$$

for all $P \in V_G \amalg E_G$.

(We easily check that the action $g \to ng$ in this corollary is a right action, similar to the above discussion of the action on functions from $\mathcal{G}$ to a set, T.)

Proof. Let $g \in \mathcal{G}$ and $M \in \mathcal{M}(n)$. Then there exists an $\mathcal{F} \subset \mathcal{K}_M$ such that $\dim(\mathcal{F}) = n$. Then we have $\mathcal{F}g \subset \mathcal{K}_M g$ and we have $\dim(\mathcal{F}g) = \dim(\mathcal{F})g$, since

$$\dim\bigl((\mathcal{F}g)(P)\bigr) = \dim\bigl(\mathcal{F}(Pg^{-1})\bigr)$$

for all $P \in V_G \amalg E_G$. But we have an isomorphism $\iota_g \colon \mathcal{K}_M g \to \mathcal{K}_{Mg}$ of sheaves on G; so on the one hand we have $\iota_g \mathcal{F} \subset \mathcal{K}_{Mg}$, and on the other hand, since isomorphisms preserve the dimension profile, we have $Mg \in \mathcal{M}(n')$ where

$$n' = \dim(\iota_g \mathcal{F} g) = \dim(\mathcal{F} g) = ng.$$

Hence $M \in \mathcal{M}(n)$ implies that $Mg \in \mathcal{M}(ng)$. Applying this observation to M replaced with Mg and g replaced with g^{-1} (or simply reversing the argument in this proof) shows the converse. Hence $\mathcal{M}(n)g = \mathcal{M}(ng)$. $\square$

2.6.2. Generic Maximum Excess. If $\mathbb{F}$ is a field and $r \geq 1$ an integer, then by a *generic* subset of $\mathbb{F}^r$ we mean a subset that contains

$$\{(x_1, \ldots, x_r) \in \mathbb{F}^r \mid p(x_1, \ldots, x_r) \neq 0\}$$

for some nonzero polynomial, p. Algebraic geometry and generic subsets are most commonly discussed (at least on the most basic level) under the assumption that $\mathbb{F}$ is algebraically closed. Under this situation, all generic sets are nonempty; this remains true if $\mathbb{F}$ is infinite, or if the polynomial, p, above is fixed and $\mathbb{F}$ is finite but sufficiently large.

In order to have a sensible definition of generic and to conform to the algebraic geometric literature, we will freely assume that $\mathbb{F}$ is algebraically closed. However, the theorems we obtain in this section and the next will be valid for any infinite field or "sufficiently large" finite field, $\mathbb{F}$, by applying these theorems to the algebraic closure of $\mathbb{F}$, finding the associated polynomials, p, to the generic sets of interest, and determining how large $\mathbb{F}$ needs to be so that the generic sets are nonempty. The reader may find it amusing to note that in all our discussion of generic sets and generic conditions, all that we ultimately care about is that certain of these generic sets are nonemtpy (e.g., that there is at least one ρ-kernel for $(L, G, \mathcal{G})$ with vanishing maximum excess).

Let us review some notation in algebraic geometry; see [**Har77**], Chapter 1, Section 1. Let us assume that $\mathbb{F}$ is algebraically closed. Let $\mathbb{A}^N = \mathbb{A}^N(\mathbb{F})$, where N is an integer or a set or a product thereof, denote affine N space over $\mathbb{F}$, i.e., the set $\mathbb{F}^N$, with its usual Zariski topology. (When we speak of topological notions on $\mathbb{F}^N$ we mean those of $\mathbb{A}^N(\mathbb{F})$; in the literature $\mathbb{A}^N(\mathbb{F})$ connotes $\mathbb{F}^N$ viewed as a topological space, or scheme, etc.) Recall that a *locally closed* set is the intersection of an open and closed set (i.e., a subset of $\mathbb{A}^N$ determined as the zeros of some polynomials and complement of the zeros of some other polynomials), and a *constructible* set on $\mathbb{A}^N$ amounts to a finite disjoint union of locally closed sets (see [**Har77**], Exercise II.3.18).

LEMMA 2.23. *Let $\mathbb{F}$ be an algebraically closed field, $k \geq 0$ an integer, and L a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. For each $n \colon V_G \amalg E_G \to \mathbb{Z}_{\geq 0}$, $\mathcal{M}(n) = \mathcal{M}(n, L, G, \mathcal{G}, k)$ is a constructible set.*

PROOF. We introduce $|n|\,|\mathcal{G}|$ indeterminates as follows: for each $P \in V_P \amalg E_P$, and $i = 1, \ldots, n(P)$, let $x_{P,i}$ be a vector of indeterminates indexed on $\mathcal{G}$ (there are $|n|$ vector variables $x_{P,i}$, for a total of $|n|\,|\mathcal{G}|$ indeterminates). We note that $M \in \mathcal{M}(n)$ precisely when one can find a solution for M and $x_{P,i}$ to the conditions

(1) M is L-surjective; i.e., for each $P \in V_G \amalg E_G$, $\mathbb{F}^k$ is spanned by the columns of M corresponding to the elements of $\mathcal{G}_P(L)$;
(2) for all P and i we have that $x_{P,i}$ has zero components outside of $\mathcal{G}_L(P)$;

(3) for all P and i, $Mx_{P,i} = 0$;
(4) for all P, $x_{P,1}, \ldots, x_{P,n_P}$ are linearly independent;
(5) for all $e \in E_G$ and all i we have that $x_{e,i}, x_{te,1}, x_{te,2}, \ldots, x_{te,n_{te}}$ are linearly dependent, and similarly with he replacing te.

The dependence or independence or spanning of vectors reduces to the vanishing or nonvanishing of determinants of the vectors' coordinates. Hence all the above equations give us a collection of polynomials $f_i \in \mathbb{F}[M,x]$ (polynomials in the entries of M and the $x_{P,i}$'s) and $\widetilde{f}_j \in \mathbb{F}[M,x]$ such that $M \in \mathcal{M}(n)$ iff for some x we have $(M,x) \in C$, where C is the set of (M,x) for which $f_i(M,x) = 0$ for all relevant i and $\widetilde{f}_j(M,x) \neq 0$ for all relevant j; hence C is constructible. But $M \in \mathcal{M}(n)$ iff $(M,x) \in C$ for some x; hence $\mathcal{M}(n)$ is the image of C under the projection

$$\mathbb{A}^{k\times\mathcal{G}} \times \mathbb{A}^{|n|\times\mathcal{G}} \to \mathbb{A}^{k\times\mathcal{G}}.$$

But any projection from an affine space to another by omitting some of the coordinates has the property that it takes constructible sets to constructible sets (see Exercise II.3.19 of [**Har77**] or Theorem 3.16 of [**Har92**], noting that such a projection is both regular and of finite type). Hence $\mathcal{M}(n)$, the image of C, is constructible. □

We recall that a generic subset, S, of some affine space, $\mathbb{F}^T$, is a subset that contains a nonemtpy Zariski open subset of the space; if S is constructible, then S is generic iff its Zariski closure is the affine space.

Next we claim that $\mathcal{M}(n)$ is generic in $\mathbb{F}^{k\times\mathcal{G}}$ for at least one n, provided that $k \le \rho(L)$, and that $\mathcal{M}(n) = \emptyset$ for all but finitely many n. Indeed, for any totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ we have that M is L-surjective (for $(L,G,\mathcal{G},\mathbb{F},k)$), and the zero sheaf, $\mathcal{Z}$, has $\dim(\mathcal{Z}) = 0$. Hence the Zariski closure of $\mathcal{M}(0)$ is $\mathbb{F}^{k\times\mathcal{G}}$. Furthermore, $\mathcal{K}_M(P)$, for any $P \in V_G \amalg E_G$, is of dimension at most $|\mathcal{G}| - k$; hence $\mathcal{M}(n) = \emptyset$ unless $|n(P)| \le |\mathcal{G}| - k$ for all $P \in V_G \amalg E_G$, and there are only finitely many such n.

Definition 2.24. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be an algebraically closed field. Let $k \le \rho(L)$ be a non-negative integer. We say that $n\colon V_G \amalg E_G \to \mathbb{Z}_{\ge 0}$ is* generic *for $(L,G,\mathcal{G},\mathbb{F},k)$ if the Zariski closure of $\mathcal{M}(n)$ is $\mathbb{A}^{k\times\mathcal{G}}(\mathbb{F})$. We define the* generic maximum excess *of $(L,G,\mathcal{G},\mathbb{F},k)$ to be the largest value of $-\chi(n)$ for which n is generic. We define n to be a* maximal profile *(respectively,* minimal profile*) of $(L,G,\mathcal{G},\mathbb{F},k)$ if n is generic, $-\chi(n)$ equals the generic maximum excess, and there is no $n' \neq n$ which is generic with $-\chi(n') = -\chi(n)$ and $n'(P) \ge n(P)$ (respectively $n'(P) \le n(P)$) for all $P \in V_G \amalg E_G$.*

Theorem 2.25. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and $k \le \rho(L)$ an integer. There is a unique maximal profile and a unique minimal profile for $(L,G,\mathcal{G},\mathbb{F},k)$. Furthermore, if n is either the maximal or minimal profile, and $P \in V_G \amalg E_G$, then $ng = n$ for all $g \in \mathcal{G}$ (in the notion of Corollary 2.22). In particular $-\chi(n)$ is divisible by $|\mathcal{G}|$.*

Actually, the proof below shows that the theorem is still true when $k > \rho(L)$, provided that L has at least k edges of each colour, so that a totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ is L-surjective.

Proof. Let n_1, n_2 be two maximal profiles for $(L,G,\mathcal{G},\mathbb{F},k)$. Let us show that $n_1 = n_2$. Consider the subset, S, of $\mathbb{F}^{k\times\mathcal{G}}$, M, such that $M \in \mathcal{M}(n_i)$ for $i = 1,2$

and m.e.$(\mathcal{K}_M) = -\chi(n_1) = -\chi(n_2)$. Clearly

$$S = \mathcal{M}(n_1) \cap \mathcal{M}(n_2) \cap \bigcap_{n \text{ s.t. } -\chi(n) > -\chi(n_1)} \overline{\mathcal{M}(n)},$$

where $\overline{\mathcal{M}(n)}$ denotes the complement of $\mathcal{M}(n)$. But if $-\chi(n) > -\chi(n_1)$ then, by assumption, n is not generic, and hence S is the intersection of a finite number of generic subsets of $\mathbb{F}^{k\times\mathcal{G}}$; hence S is a generic subset of $\mathbb{F}^{k\times\mathcal{G}}$, as well. But any element, $M \in S$, has subsheaves $\mathcal{F}_1, \mathcal{F}_2$, of $\mathcal{K}_M$ which obtain the maximum excess of $\mathcal{K}_M$ and with $\dim(\mathcal{F}_i) = n_i$ for $i = 1, 2$. But then $\mathcal{F} = \mathcal{F}_1 + \mathcal{F}_2$ also achieves the maximum excess and has $\deg(\mathcal{F}) \geq n_i$ for $i = 1, 2$. Hence

$$S \subset \bigcup_{n \text{ s.t. } -\chi(n) = -\chi(n_1),\ n \geq N} \mathcal{M}(n),$$

where $N = \max(n_1, n_2)$. Since the union on the right-hand-side is a finite union of constructible sets, the closure of one of these sets is $\mathbb{F}^{k\times\mathcal{G}}$. Hence there is an n with $-\chi(n) = -\chi(n_1)$ and $n \geq N = \max(n_1, n_2)$ such that n is generic; but if $n_1 \neq n_2$, then n does not equal either of them and is at least as big as either, which contradicts the maximality of the n_i, $i = 1, 2$. Hence $n_1 = n_2$, and the maximal profile is unique.

We argue similarly for the minimal profile, replacing $\mathcal{F}_1 + \mathcal{F}_2$ with $\mathcal{F}_1 \cap \mathcal{F}_2$.

Let n be the maximal profile for $(L, G, \mathcal{G}, \mathbb{F}, k)$ (now known to be unique). Since $\mathcal{M}(n)$ is a generic subset of $\mathbb{F}^{k\times\mathcal{G}}$, so is $\mathcal{M}(ng) = \mathcal{M}(n)g$ for any $g \in \mathcal{G}$. But then ng is also a maximal profile, since clearly $\chi(n) = \chi(ng)$ and $|n| = |ng|$. Hence $n = ng$ for all $g \in \mathcal{G}$. The same is true of the minimal profile.

It follows that the maximal (or minimal) profile, n, is invariant under $\mathcal{G}$, and hence has the same value on all the vertices, on all the edges of colour 1, and on all the edges of colour 2. Hence $-\chi(n)$ is divisible by $|\mathcal{G}|$ for the maximal (or minimal) profile, and hence the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ is divisible by $|\mathcal{G}|$. ☐

2.7. Variability of k-th Power Kernels

The main goal of this section is to prove the following theorem.

THEOREM 2.26. *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and let $k \leq \rho(L)$ be a positive integer. Then the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ is at most that of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$, and we have equality iff both excesses are zero.*

As a consequence we get the following theorem.

THEOREM 2.27. *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and let $k \leq \rho(L)$ be a positive integer. Let L' be obtained from L by removing a single edge. Then the generic maximum excess of $(L', G, \mathcal{G}, \mathbb{F}, k-1)$ is at least that of $(L, G, \mathcal{G}, \mathbb{F}, k)$.*

(As before, this theorem is also true if $k > \rho(L)$, provided that L has at least k edges of each colour, so that a totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ is L-surjective.)

A second goal of this section is to establish some general relations between kernels $\mathcal{K} = \mathcal{K}_M(L)$ as M and L vary. We shall derive two interesting, short exact sequences. First we establish a short exact sequence

$$0 \to \mathcal{K}_M(L) \to \mathcal{K}_{M'}(L) \to \underline{\mathbb{F}} \to 0, \tag{2.18}$$

for any $M \in \mathbb{F}^{k\times\mathcal{G}}$ and M' obtained from M by deleting the last row. Second we establish a short exact sequence

$$0 \to \mathcal{K}_{M'}(L') \to \mathcal{K}_{M'}(L) \to \mathcal{E} \to 0, \tag{2.19}$$

with L, L' as in Theorem 2.27, $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ such that $\mathcal{K}_{M'}(L')$ exists (i.e., M' is L'-surjective), and $\mathcal{E}$ a sheaf with $\mathcal{E}(V) = 0$ and $\mathcal{E}(E)$ of dimension $|\mathcal{G}|$.

Equation (2.19) will be used along with Theorem 2.25 to show that Theorem 2.26 implies Theorem 2.27.

Theorem 2.26 will not be proven with short exact sequences, but rather with a careful analysis of the unique minimal maximizer of the excess of $\mathcal{K}_M(L)$ and of that of $\mathcal{K}_{M'}(L)$. The sequence in equation (2.18) gives a relationship between $\mathcal{K}_M(L)$ and $\mathcal{K}_{M'}(L)$, but we don't know how to directly use this to conclude anything interesting about the two sheaves, such as the result of Theorem 2.26.

At this point we will divide our discussion into subsections. In Subsection 2.7.1, we will discuss the exact sequences related to our proof. In Subsection 2.7.2 we give the main observation of how the maximum excess changes in passing to subsheaves, and give an intuitive reason why the generic maximum excess of $\mathcal{K}_{M'}(L)$, as above, should be strictly greater than that of $\mathcal{K}_M(L)$ provided that these numbers don't both vanish. In Subsection 2.7.3 we mimic the notation of Section 2.6 to include a discussion of $\mathcal{K}_{M'}(L)$ as above and make our arguments precise, finishing the proof of Theorem 2.26; this will easily yield Theorem 2.27.

2.7.1. Variability as Exact Sequences.

Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. For any non-negative integer, $k \le \rho(L)$, we have that a generic $M \in \mathbb{F}^{k\times\mathcal{G}}$ gives rise to a short exact sequence

$$0 \to \mathcal{K}_M(L) \to \underline{\mathbb{F}}_L\mathcal{G} \to \underline{\mathbb{F}}^k \to 0. \tag{2.20}$$

First we considering the variance of this equation in M; in other words, fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ such that

$$\underline{M\iota}\colon \underline{\mathbb{F}}_L\mathcal{G} \to \underline{\mathbb{F}}^k$$

is surjective. Then we have an exact sequence given in equation (2.20); if $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ is M with its last row deleted, we have a similar exact sequence

$$0 \to \mathcal{K}_{M'}(L) \to \underline{\mathbb{F}}_L\mathcal{G} \to \underline{\mathbb{F}}^{k-1} \to 0. \tag{2.21}$$

Notice that this discussion and everything below will remain essentially the same if, more generally, M' is taken to be M followed by any surjective map $\mathbb{F}^k \to \mathbb{F}^{k-1}$. In any event, we get a digram:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] \arrow[d, dashed] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "\underline{M\iota}"] \arrow[d, "\simeq"] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{K}_{M'}(L) \arrow[r] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "\underline{M'\iota}"] & \underline{\mathbb{F}}^{k-1} \arrow[r] & 0
\end{tikzcd}$$

where the dotted arrow from $\mathcal{K}_{M'}(L)$ to $\mathcal{K}_M(L)$ is inferred from the solid arrows; furthermore, given that the solid horizontal arrows consist of an isomorphism and epimorphism, we infer that the dotted arrow is a monomorphism. We then complete the diagram to obtain a diagram

$$\begin{tikzcd}
 & & & 0 \arrow[d] & \\
 & 0 \arrow[d] & 0 \arrow[d] & \underline{\mathbb{F}} \arrow[d] & \\
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] \arrow[d] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "M\iota"] \arrow[d] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{K}_{M'}(L) \arrow[r] \arrow[d] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "M'\iota"] \arrow[d] & \underline{\mathbb{F}}^{k-1} \arrow[r] \arrow[d] & 0 \\
 & \mathcal{K}_{M'}(L)/\mathcal{K}_M(L) \arrow[d] & 0 & 0 & \\
 & 0 & & &
\end{tikzcd}$$

A simple diagram chase shows that the nonzero upper right sheaf, $\underline{\mathbb{F}}$, and the nonzero lower left sheaf, $\mathcal{K}_{M'}(L)/\mathcal{K}_M(L)$, are isomorphic. Hence we obtain the short exact sequence in equation (2.18).

An analogous exact sequence can be obtained by varying L in equation (2.20). Let $L' \subset L$ be a subgraph of L. Fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ that induces a surjection $\underline{\mathbb{F}}_{L'}\mathcal{G} \to \underline{\mathbb{F}}^k$. Then we get a diagram:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{K}_M(L') \arrow[r] \arrow[d, dashed] & \underline{\mathbb{F}}_{L'}\mathcal{G} \arrow[r] \arrow[d] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d, "\simeq"] & 0 \\
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r] & \underline{\mathbb{F}}^k \arrow[r] & 0
\end{tikzcd}$$

Since $\underline{\mathbb{F}}_{L'}\mathcal{G} \to \underline{\mathbb{F}}_L\mathcal{G}$ is an injection, and the last vertical arrow is an isomorphism, the inferred dotted arrow is an injection. We therefore add a bottom row to the diagram and infer from the 3×3 Lemma that

$$\mathcal{K}_M(L)/\mathcal{K}_M(L') \simeq (\underline{\mathbb{F}}_L/\underline{\mathbb{F}}_{L'})\mathcal{G}.$$

In particular, if L' is obtained from L by removing m edges, then we infer equation (2.19) (with M here replaced by M', and k implicit here replaced by $k-1$), where $\mathcal{E}$ is a sheaf with $\mathcal{E}(V) = 0$ and $\mathcal{E}(E)$ being of dimension $m|\mathcal{G}|$.

2.7.2. Maximum Excess and Subsheaves. The goal of this section is to explain the main idea we will use to prove Theorem 2.26; the formal proof will be given in the subsection after this one.

The following theorem gives a number of ways of demonstrating whether or not a sheaf and one of its subsheaves have the same maximum excess.

THEOREM 2.28. *Let $\mathcal{F}' \subset \mathcal{F}$ be sheaves on a digraph, G. Let $U \subset \mathcal{F}(V)$ be the minimal maximizer of the excess of $\mathcal{F}$, and let $U' \subset \mathcal{F}'(V)$ be the minimal maximizer of the excess of $\mathcal{F}'$. Then*

$$\text{m.e.}(\mathcal{F}') \le \text{m.e.}(\mathcal{F}), \tag{2.22}$$

with equality iff $U = U'$ *and*

$$\Gamma_{\rm ht}(\mathcal{F}, U') = \Gamma_{\rm ht}(\mathcal{F}', U').$$

We already know equation (2.22) is true, since the maximum excess is a quasi-Betti number; the novelty of this theorem is that we have a simple condition to characterize when equality holds.

PROOF. Since $U' \subset \mathcal{F}(V)$ and

$$\Gamma_{\rm ht}(\mathcal{F}', U') \subset \Gamma_{\rm ht}(\mathcal{F}, U'),$$

we have that

$$\text{m.e.}(\mathcal{F}') = \text{excess}(\mathcal{F}', U') \le \text{excess}(\mathcal{F}, U') \le \text{m.e.}(\mathcal{F});$$

hence equality holds in equation (2.22) iff

$$\text{excess}(\mathcal{F}', U') = \text{excess}(\mathcal{F}, U') = \text{m.e.}(\mathcal{F}).$$

The first equality holds iff

$$\Gamma_{\rm ht}(\mathcal{F}, U') = \Gamma_{\rm ht}(\mathcal{F}', U').$$

The second equality holds iff U' is also a maximizer for $\mathcal{F}$. But since U is the minimal maximizer for $\mathcal{F}$, this implies that $U \subset U'$; but this means that $U \subset U' \subset \mathcal{F}'(V)$, so U is a maximizer for $\mathcal{F}'$, and hence $U' \subset U$ (since U' is the minimal maximizer for $\mathcal{F}'$). Hence $U = U'$. $\square$

Theorem 2.28 gives us a number of ways to conclude that equation (2.22) holds with strict inequality in certain situations. For example, consider a subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$, and consider a value, k for which

$$\text{m.e.}\bigl(\mathcal{K}_M(L)\bigr) > 0$$

for a generic $M \in \mathbb{F}^{k\times\mathcal{G}}$. Let M' be obtained from M by removing its bottom row, and consider the minimal maximizer, $U = U(M') \subset \mathcal{K}_{M'}(L)(V)$ of the excess of $\mathcal{K}_{M'}(L)$. We have that $\mathcal{K}_M \subset \mathcal{K}_{M'}$, and hence

$$\text{m.e.}\bigl(\mathcal{K}_M(L)\bigr) = \text{m.e.}\bigl(\mathcal{K}_{M'}(L)\bigr)$$

implies that $U(M')$, which is generically nonzero, lies entirely in $\mathcal{K}_M(L)(V)$. But if $w \in \mathcal{K}_{M'}(L)(V)$ is any nonzero vector supported on $v \in V_G$, then we may identify w with the corresponding element of $\mathcal{K}_{M'}(L)(v)$, and so

$$w \in (\underline{\mathbb{F}}_L\mathcal{G})(v) \simeq \bigoplus_{g\in\mathcal{G}_L(v)} \mathbb{F}_g,$$

where $\mathbb{F}_g$ denotes a copy of $\mathbb{F}$. In other words, w is given by its $\mathcal{G}$ components, which are (zero outside of $\mathcal{G}_L(v)$ and are) elements of $\mathbb{F}$. Hence, if we add a generic extra row to M' on the bottom, to form M, the row, $\vec{m} = (m_g)_{g\in\mathcal{G}}$ will (generically in $\vec{m}$) satisfy

$$\sum_{g\in\mathcal{G}} w_g m_g \neq 0. \tag{2.23}$$

Hence $w \notin \mathcal{K}_M(L)(V)$ generically, and therefore the minimal maximizers for $\mathcal{K}_M(L)$ and $\mathcal{K}_{M'}(L)$ will generically be different. Hence, by Theorem 2.28, we have

$$\text{m.e.}\bigl(\mathcal{K}_{M'}(L)\bigr) \ge 1 + \text{m.e.}\bigl(\mathcal{K}_M(L)\bigr)$$

for generic M (and M' obtained from M by deleting its bottom row). This argument will establish Theorem 2.26; all we need to do is to make this rigourous.

2.7.3. Proof of Theorems 2.26 and 2.27. In this subsection we precisely state the idea in the last subsection as Theorem 2.30 and use it to prove Theorem 2.26. We then easily conclude Theorem 2.27.

Let $\mathbb{F}$ be a field, $\mathcal{G}$ a group, and $k \geq 1$ an integer. If $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ and $\vec{m} \in \mathbb{F}^{\mathcal{G}}$, we define $\mathrm{merge}(M', \vec{m})$ to be the element of $\mathbb{F}^{k\times\mathcal{G}}$ whose first $k-1$ rows consist of M' and whose k-th row consists of $\vec{m}$.

DEFINITION 2.29. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Let $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ be a matrix, for some integer $k \geq 1$, that is L-surjective and whose kernel, $\mathcal{K}_{M'}(L)$, has nonzero maximum excess. Define the* redundancy *of M', denoted* $\mathrm{redund}(M')$, *to be the subset of $\mathbb{F}^{\mathcal{G}}$ consisting of $\vec{m}$ such that $M = \mathrm{merge}(M', \vec{m})$ is L-surjective, and such that*

$$\mathrm{m.e.}\big(\mathcal{K}_{M'}(L)\big) = \mathrm{m.e.}\big(\mathcal{K}_M(L)\big).$$

THEOREM 2.30. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field and k a positive integer. Let $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ be a matrix that is L-surjective, and whose kernel, $\mathcal{K}_{M'}(L)$, has nonzero maximum excess. Then the redundancy of M' lies in a proper subspace of $\mathbb{F}^{\mathcal{G}}$.*

PROOF. This follows the argument of the last subsection. If U is the minimal maximizer of $\mathcal{K}_{M'}(L)$, then U is nonzero because the maximum excess is nonzero. Hence there exists a $w \in U$ supported at $v \in V_G$ with $w \neq 0$. So if $\mathrm{merge}(M', \vec{m})$ is L-surjective, we have $w \notin \mathcal{K}_M(L)$ if equation (2.23) holds. Since $w \neq 0$, equation (2.23) holds for all $\vec{m}$ outside of a proper subspace of $\mathbb{F}^{\mathcal{G}}$. □

PROOF OF THEOREM 2.26. If the generic maximum excesses were equal, then for a nonempty Zariski open subset, U, of $\mathbb{F}^{k\times\mathcal{G}}$, we would have for all $M \in U$ the maximum excess of $\mathcal{K}_M(L)$ is the same as that of $\mathcal{K}_{M'}(L)$, where M' is obtained from M by discarding its bottom row. Since U is nonempty and Zariski open, we have a polynomial, $p = p(M', \vec{m})$ such that

$$p(M', \vec{m}) \neq 0$$

implies that $(M', \vec{m}) \in U$. Write

$$p(M', \vec{m}) = \sum_{n \in \mathbb{Z}_{\geq 0}^{\mathcal{G}}} q_n(M')\vec{m}^n,$$

and fix any n such that $q_n \neq 0$. Then $q_n(M') \neq 0$ for all $M' \in U'$ for a nonempty Zariski open subset, U', of $\mathbb{F}^{(k-1)\times\mathcal{G}}$. For any fixed $M' \in U'$ we have $p(M', \vec{m})$ is a nonzero polynomial in $\vec{m}$; hence for fixed $M' \in U'$ we have that $(M', \vec{m}) \in U$ for a Zariski open subset of $\vec{m}$ in $\mathbb{F}^{\mathcal{G}}$.

On the other hand, assuming that the maximum excesses in Theorem 2.26 are not both zero, the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$ is positive. Hence $\mathcal{K}_{M'}(L)$ has positive maximum excess for all M' in some nonempty, Zariski open subset, U'', of $\mathbb{F}^{(k-1)\times\mathcal{G}}$. But by Theorem 2.30, for any $M' \in U''$ we have that $(M', \vec{m}) \notin U$ for $\vec{m}$ outside of a proper subspace of $\mathbb{F}^{\mathcal{G}}$. But U' and U'' must intersect (being two nonempty, Zariski open subsets of an irreducible variety), which gives a contradiction. □

Proof of Theorem 2.27. Let the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ be m_k, that of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$ be m_{k-1}, and that of $(L', G, \mathcal{G}, \mathbb{F}, k-1)$ be m'_{k-1}. Since $k \le \rho(L)$ and hence $k-1 \le \rho(L')$ (we can see $\rho(L') \ge \rho(L) - 1$ from equation (1.1)), we have that m_k, m_{k-1}, m'_{k-1} are all multiples of $|\mathcal{G}|$. The theorem is immediate if $m_k = 0$, so we may assume $m_k > 0$. In this case Theorem 2.26 implies that $m_{k-1} > m_k$, and since these numbers are both multiples of $|\mathcal{G}|$, we have

$$m_{k-1} \ge m_k + |\mathcal{G}|. \tag{2.24}$$

But the exact sequence in equation (2.19) shows that for any $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ we have

$$\text{m.e.}\big(\mathcal{K}_{M'}(L')\big) \ge \text{m.e.}\big(\mathcal{K}_{M'}(L)\big) - |\mathcal{G}|. \tag{2.25}$$

Let U, U', respectively, are the subsets of $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ at which $\mathcal{K}_{M'}(L), \mathcal{K}_{M'}(L')$, respectively, attain their generic value; hence U, U' are generic, and therefore so is $U \cap U'$. Then applying equation (2.25) to any $M' \in U \cap U'$ implies that

$$m'_{k-1} \ge m_{k-1} - |\mathcal{G}|.$$

Combining this with equation (2.24) gives $m'_{k-1} \ge m_k$, which proves the theorem. □

2.8. Proof of the SHNC

In this section we prove the SHNC. At this point we have all the tools we need, except for one small detail.

Lemma 2.31. *Let L be an arbitrary étale bigraph with $\rho(L) > 0$. Then there exists an edge, $e \in E_L$, such that the graph, L', obtained by removing e from L has $\rho(L') = \rho(L) - 1$.*

Proof. For each $F \in E_L$ let L_F denote the subgraph of L obtained by removing the edges in F from L. It is easy to see that for each $e \in E_L$ we have that $\rho(L_{\{e\}})$ is either $\rho(L)$ or $\rho(L) - 1$; this can be seen from equation (1.1), since removing e from its connected component of L leaves h_1 the same or reduces it by one; alternatively, we can see this from the exact sequence

$$0 \to \underline{\mathbb{F}}_{L_{\{e\}}} \to \underline{\mathbb{F}}_L \to \underline{\mathbb{F}}_{L/L_{\{e\}}} \to 0,$$

using the fact that $L/L_{\{e\}}$ is (edge supported and) of maximum excess one.

For any two subgraphs, L', L'', of L we have an exact sequence

$$0 \to \underline{\mathbb{F}}_{L'\cap L''} \to \underline{\mathbb{F}}_{L'} \oplus \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_{L'\cup L''} \to 0.$$

Hence

$$\rho(L' \cap L'') \ge \rho(L') + \rho(L'') - \rho(L' \cup L'').$$

Taking F', F'' to be disjoint subsets of E_L, we see that if $\rho(L_{F'}) = \rho(L_{F''}) = \rho(L)$, then setting $L' = L_{F'}$, $L'' = L_{F''}$ yields

$$\rho(L_{F'\cup F''}) \ge \rho(L),$$

and so $\rho(L_{F'\cup F''}) = \rho(L)$. Hence, if $\rho(L_F) = \rho(L)$ for all $F \subset E_L$ of size one, then by induction we can show this holds for $F \subset E_L$ of any size, which is impossible (since removing all the edges of a graph leaves it with $\rho = 0$). Hence there is at least one $e \in E_L$ for which $\rho(L_{\{e\}}) = \rho(L) - 1$. □

Of course, one can give a purely graph theoretic proof of Lemma 2.31; we now sketch such a proof. From equation (1.1), it suffices to show that if L is connected with $h_1(L) \geq 2$ then we can remove an edge from L and reduce h_1 by one. We claim that it suffices to take any edge that remains after we repeatedly prune the leaves of L.

DEFINITION 2.32. *Let L be a subgraph of a Cayley bigraph, G on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Then by the* generic maximum excess of the ρ-kernel of type $(L, G, \mathcal{G}, \mathbb{F})$ *we mean the generic maximum excess of* $(L, G, \mathcal{G}, \mathbb{F}, \rho(L))$.

THEOREM 2.33. *Let L be a subgraph of a Cayley bigraph, G on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Then the generic maximum excess of the ρ-kernel of type $(L, G, \mathcal{G})$ is zero.*

PROOF. Fix G and $\mathcal{G}$ and let us prove the theorem for all L by induction on $\rho(L)$. The base case $\rho(L) = 0$ follows by definition, since the exact sequence

$$0 \to \mathcal{K} \to \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^0 \to 0$$

implies

$$\text{m.e.}(\mathcal{K}) \leq \text{m.e.}(\underline{\mathbb{F}}_L \mathcal{G}) = \sum_{g \in \mathcal{G}} \rho(Lg) = 0$$

(since $\rho(Lg) = \rho(L) = 0$ for all $g \in \mathcal{G}$). The inductive step of our induction on $\rho(L)$ is immediate from Theorem 2.27 applied to any L' obtained from L by removing a single edge so that $\rho(L') = \rho(L) - 1$; the existence of such an L' is given by Lemma 2.31. $\square$

PROOF OF THEOREM 2.1, THE SHNC. By the graph theoretic reformulation of the SHNC, it suffices to show Theorem 2.2. By Theorem 2.4 it suffices to show that any subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$, is universal for the SHNC. But by Theorem 2.33, there exists a ρ-kernel, $\mathcal{K} = \mathcal{K}_M(L)$ for $(L, G, \mathcal{G}, \mathbb{F})$ with vanishing maximum excess, for any algebraically closed $\mathbb{F}$. Hence we apply Theorem 2.14 to conclude that L is universal for the SHNC. $\square$

2.9. Concluding Remarks

We finish this paper with a few concluding remarks.

In this chapter we have made no explicit reference to homology theories. In [**Fri11b**] we have used the twisted homology to prove that the maximum excess is a first quasi-Betti number; hence the theorems in this paper ostensibly rely on homology theories. However, we think it quite possible that one may able to prove that the maximum excess is a first quasi-Betti number directly, or give a direct proof of the inequalities we made use of in this paper. For example, if $\mathcal{F}' \to \mathcal{F}$ is a monomorphism, then since the maximum excess is a first quasi-Betti number we know that

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

But this inequality is clear from the subsheaf formulation of maximum excess in Theorem 1.28.

We remark that we first proved the SHNC using twisted homology theory, and then rewrote our proofs to use only maximum excess. In fact, twisted homology theory offers some additional intuition regarding the maximum excess. Twisted homology theory shows that (after pulling back appropriately, see [**Fri11b**]), the maximum excess can be interpreted as the dimension of a certain vector space of "twisted harmonic one-forms" of the sheaf. When this dimension is $d > 0$, one can impose d' linear conditions on the twisted harmonic one-forms and still have a $d-d'$ dimensional space of one-forms. This is how we view the variance in L of $\mathcal{K}_{M'}(L)$, as in the exact sequence of equation (2.19): to take a space of one-forms on $\mathcal{K}_{M'}(L)$ and obtain a one-form in $\mathcal{K}_{M'}(L')$, one has to impose $|\mathcal{G}|\,|E_L \setminus E_{L'}|$ conditions on the one-forms, namely the conditions that they vanish on the edges in $L\mathcal{G}$ that do not lie in $L'\mathcal{G}$. Of course, one has to pullback by an appropriate covering map to make this rigourous (see [**Fri11b**]), but all the edge counts and dimension counts scale appropriately under any covering.

The k-th power kernels in this paper are subsheaves of the constant sheaf $\underline{\mathbb{F}\mathcal{G}} \simeq \underline{\mathbb{F}}^{\mathcal{G}}$. We believe that subsheaves of constant sheaves satisfy some stronger properties than general sheaves, regarding their homological invariants and maximum excess. It would nice to study this further.

Finally, we give a variant of our proof of the SHNC that involves no homology theory and, in particular, avoids any use of Theorem 1.10. As before, let L be any subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be a field. First note that using Appendix A, we can show that if there is a ρ-kernel for $(L, G, \mathcal{G})$ with vanishing maximum excess, then the SHNC holds for all pairs (L, L'), with L' any subgraph of G; Appendix A makes no use of homology or Theorem 1.10. (Appendix A is a bit tedious and long, however avoids use of applying Theorem 1.16 with α_1 being the maximum excess, and the proof that the maximum excess is a scaling first quasi-Betti number used Theorem 1.10.) Furthermore, if the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, \rho(L))$ were greater than zero, then it would be at least $|\mathcal{G}|$. Then, by induction, for $n = 1, \ldots, \rho(L)$ we have that the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, \rho(L)-n)$ would be at least $|\mathcal{G}|(1+n)$, in view of Theorems 2.25 and 2.26 (which makes no use of homology or Theorem 1.10). But this is impossible for $n = \rho(L)$, since the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, 0)$ is $|\mathcal{G}|\rho(L)$, because a 0-th power kernel is plainly just $\underline{\mathbb{F}}_L\,\mathcal{G}$, which has maximum excess $\rho(L)|\mathcal{G}|$. Hence the SHNC holds for all pairs of subgraphs of Cayley bigraphs, and hence holds for all pairs of étale bigraphs, by Theorem 2.4.

APPENDIX A

A Direct View of ρ-Kernels

In this appendix we give a direct combinatorial proof that the SHNC follows the vanishing maximum excess of some ρ-kernel for each triple $(L, G, \mathcal{G})$.

In this section we give a direct proof that the vanishing generic maximum excess of ρ-kernels for all subgraphs, L, of any Cayley graph, G, implies the SHNC. We shall not use exact sequences. We shall require a few definitions, and some calculations to follow. While this gives some extra intuition about ρ-kernels, this section is not essential to the proof of the SHNC; we shall omit some of the easy but tedious graph theoretic details.

Definition A.1. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. As usual, for $P \in V_G \amalg E_G$, let $\mathcal{G}_L(P)$ be the set of $g \in \mathcal{G}$ such that Lg contains P. By a* vertex family *on $(L, G, \mathcal{G})$ we mean a function, $\mathcal{U}$, from V_G to $\mathcal{P}(\mathcal{G})$, the power set (i.e., set of subsets) of $\mathcal{G}$, such that for all $v \in V_G$ we have $\mathcal{U}(v) \subset \mathcal{G}_L(v)$. Similarly, an* edge family *on $(L, G, \mathcal{G})$ is a function $\mathcal{W}\colon E_G \to \mathcal{P}(\mathcal{G})$ such that $\mathcal{W}(e) \subset \mathcal{G}_L(e)$ for all $e \in E_G$. A vertex family, $\mathcal{U}$, and edge family, $\mathcal{W}$, are* compatible *if for all $e \in E_G$ we have $\mathcal{W}(e) \subset \mathcal{U}(te) \cap \mathcal{U}(he)$. Given a vertex family, $\mathcal{U}$, the* induced edge family*, $\mathcal{U}_E$, of $\mathcal{U}$ is the edge family $\mathcal{U}_E$ given by*

$$\mathcal{U}_E(e) = \mathcal{U}(te) \cap \mathcal{U}(he) \cap \mathcal{G}_L(e).$$

The following lemmas motivate the above definitions; we omit their proofs, which are almost immediate.

Lemma A.2. *To each vertex family, $\mathcal{U}$ on $(L, G, \mathcal{G})$, and compatible edge family, $\mathcal{W}$, on $(L, G, \mathcal{G})$, there is a subgraph $H \subset L \times_{B_2} G$ determined via*

$$V_H = \{(vg^{-1}, v) \mid g \in \mathcal{U}(v)\}, \qquad E_H = \{(eg^{-1}, e) \mid g \in \mathcal{W}(e)\};$$

conversely, any subgraph $H \subset L \times_{B_2} G$ arises from a unique vertex family, $\mathcal{U}$, and compatible edge family, $\mathcal{W}$.

Lemma A.3. *For any vertex family, $\mathcal{U}$, on $(L, G, \mathcal{G})$ and compatible edge family, $\mathcal{W}$, we have*

$$\mathcal{W}(e) \subset \mathcal{U}_E(e).$$

In other words, $\mathcal{U}_E$ is the "largest" edge family compatible with $\mathcal{U}$.

Definition A.4. *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, let $M \in \mathbb{F}^{\rho(L)\times\mathcal{G}}$ be totally linearly independent, and let $\mathcal{K} = \mathcal{K}_M$ be the resulting ρ-kernel. By a* straight *subspace of $\mathcal{K}(V)$ we mean a subspace*

$$U = \sum_{v \in V_G} U(v) \in \mathcal{K}(V),$$

such that for each $v \in V_G$, we have

$$U(v) = \mathrm{Free}_{\mathcal{U}(v)} \tag{A.1}$$

for some $\mathcal{U}(v) \subset \mathcal{G}$, *with notation as in Definition 2.18.*

Our goal for the rest of this section is to prove the following theorem.

THEOREM A.5. *Let* L *be a subgraph of a Cayley graph,* G, *on a group,* $\mathcal{G}$. *The following conditions are equivalent:*

(1) *for all* $L' \subset G$, *the SHNC holds for* (L, L');
(2) *for every vertex family,* $\mathcal{U}$, *on* $(L, G, \mathcal{G})$ *we have*
$$\sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} \le \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)};$$
and
(3) *for some or any field,* $\mathbb{F}$, *and some or any totally independent* $M \in \mathbb{F}^{\rho(L) \times \mathcal{G}}$, *every straight subspace of* $\mathcal{K}(V)$ *with* $\mathcal{K} = \mathcal{K}_M(L)$ *has excess zero.*

We know by Theorem 2.4 that the SHNC holds iff it holds for all pairs (L, L') that are subgraphs of a Cayley graph, G. Hence Theorem 2.33, that implies condition (3) of this theorem, implies the SHNC.

PROOF. Conditions (2) and (3) are easily seen to be equivalent via equation (A.1) and equation (2.12).

If $\mathcal{U}$ is any vertex family on $(L, G, \mathcal{G})$, let the *positive set* of $\mathcal{U}$ consist of those $v \in V_G$ for which $|\mathcal{U}(v)| > \rho(L)$ and of those $e \in E_G$ for which $|\mathcal{U}_E(e)| > \rho(L)$. We easily see that the positive set forms a subgraph, L', of G, and that the pairs (Pg^{-1}, P) such that P is in the positive set and g lies in $\mathcal{U}(P)$ or $\mathcal{U}_E(P)$ (as is appropriate), forms a subgraph, H, of $L \times_{B_2} L'$. We see that
$$-\chi(H) = \sum_{e \in E_{L'}} |\mathcal{U}_E(e)| - \sum_{v \in V_{L'}} |\mathcal{U}(v)|$$
$$= \rho(L)\chi(L') + \sum_{e \in E_{L'}} \bigl(|\mathcal{U}_E(e)| - \rho(L)\bigr) - \sum_{v \in V_{L'}} \bigl(|\mathcal{U}(v)| - \rho(L)\bigr)$$
$$= \rho(L)\chi(L') + \sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)}.$$
Hence we may write
$$-\chi(H) - \rho(L)\chi(L') = \sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)}. \tag{A.2}$$
This equation is the main ingredient in the equivalence of conditions (1) and (3). Let us now state some graph theoretic lemmas that will firmly establish this equivalence.

LEMMA A.6. *For any digraphs* $H \subset G$, *we have*
$$-\chi(H) \le \rho(G);$$
and equality holds if H *consists of all connected components,* X, *of* G *with* $h_1(X) > 0$ *and any of those with* $h_1(X) = 0$.

PROOF. The statement about equality is clear from the definition of ρ in equation (1.1). The inequality can be established graph theoretically by induction on the number of vertices and edges in G that are not in H. Alternatively, see the end of Section 1.6. □

LEMMA A.7. *Let $\mathcal{U}$ be a vertex family for $(L, G, \mathcal{G})$, where L is a subgraph of a Cayley digraph, G, on a group, $\mathcal{G}$. Assume that*

$$\sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)} > 0. \tag{A.3}$$

Then there is a vertex family, $\mathcal{U}'$, which satisfies this inequality with $\mathcal{U}$ replaced by $\mathcal{U}'$, for which the positive set of $\mathcal{U}'$, L', satisfies $-\chi(L') = \rho(L')$.

PROOF. For any subgraph, $Y \subset G$ and vertex family $\mathcal{W}$ on $(L, G, \mathcal{G})$, set

$$f(\mathcal{W}, Y) = \sum_{e \in E_Y} |\mathcal{W}_E(e)|_{\rho(L)} - \sum_{v \in V_Y} |\mathcal{W}(v)|_{\rho(L)}.$$

Then clearly we have

$$f(\mathcal{U}, L') = \sum_{X \in \mathrm{conn}(L')} f(\mathcal{U}, X),$$

where $\mathrm{conn}(L')$ is the set of connected components of L'. But equation (A.3) says that $f(\mathcal{U}, G) > 0$, and clearly $f(\mathcal{U}, L') = f(\mathcal{U}, G)$. Hence we have $f(\mathcal{U}, X) > 0$ for some connected component, X, of L'; fix any such X.

We claim that $\rho(X) = -\chi(X)$. Since X is connected, this is true unless $\chi(X) = 1$; so it suffices to show that $\chi(X) = 1$ is impossible. If $\chi(X) = 1$, then by repeatedly pruning the leaves of X, i.e., deleting a vertex of degree one and its incident edge from X, we arrive at an isolated vertex. But if Y is any subgraph of G with a vertex, $v \in V_Y$, of degree one, and incident edge $e \in E_Y$, and if Y' is Y with v and e discarded, we claim that $f(\mathcal{U}, Y') \geq f(\mathcal{U}, Y)$; indeed, $\mathcal{U}_E(e) \subset \mathcal{U}(v)$, so

$$f(\mathcal{U}, Y') = f(\mathcal{U}, Y) + |\mathcal{U}(v)|_{\rho(L)} - |\mathcal{U}(e)|_{\rho(L)} \geq f(\mathcal{U}, Y).$$

Hence, by repeatedly pruning X we are left with X'' that is a single vertex with no edges, so $f(\mathcal{U}, X'') \geq f(\mathcal{U}, X) > 0$. But clearly $f(\mathcal{U}, X'') \leq 0$ for X'' consisting of a single vertex. Hence $\chi(X) = 1$ is impossible, and so $\chi(X) \leq 0$ and hence $\rho(X) = -\chi(X)$.

For any vertex family, $\mathcal{V}$ of $(L, G, \mathcal{G})$ and any subgraph $Y \in G$ define a vertex family $\mathcal{V}|_Y$ via

$$\mathcal{V}|_Y(v) = \begin{cases} \mathcal{V}(v) & \text{if } v \in V_Y, \\ \emptyset & \text{otherwise.} \end{cases}$$

for all $v \in V_G$. We easily see that

$$\mathcal{V}(e) \subset (\mathcal{V}|_L)_E(e) \tag{A.4}$$

for all $e \in E_L$. Hence

$$f(\mathcal{V}|_Y, G) = f(\mathcal{V}|_Y, Y) \geq f(\mathcal{V}, Y),$$

using equation (A.4). In particular, for $\mathcal{V} = \mathcal{U}$ and $Y = X$ we have

$$f(\mathcal{U}|_X, G) \geq f(\mathcal{U}, X) > 0.$$

So we take $\mathcal{U}' = \mathcal{U}|_X$ and let L' be its positive set. Then $f(\mathcal{U}', L') = f(\mathcal{U}', G) > 0$, and L' consists of X plus possibly some addition edges, so L' is connected and $\chi(L') \leq \chi(X) \leq 0$, so $\rho(L') = -\chi(L')$. □

At this point condition (1) of Theorem A.5 easily implies condition (2). For if condition (2) does not hold, then for some $\mathcal{U}$, and with L' given as its positive set, we may assume $\rho(L') = -\chi(L')$ we have

$$\sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)} > 0$$

and hence

$$\rho(L \times_{B_2} L') \geq -\chi(H) > -\rho(L)\chi(L') = \rho(L)\rho(L').$$

Hence the SHNC is false on a pair of subgraphs of G.

It remains to show that condition (2) of Theorem A.5 implies condition (1). Again, we need some graph theoretic considerations.

Lemma A.8. *Assume the SHNC is false on a pair of subgraphs, (L, L'), of a Cayley bigraph G on a group $\mathcal{G}$. Then there is a subgraph, $L'' \subset L'$, such that*

1. *the SHNC is false on (L, L''),*
2. *L'' is connected,*
3. *$-\chi(L'') = \rho(L'')$, and*
4. *there is a subgraph, $H \subset L \times_{B_2} L''$ such that $-\chi(H) = \rho(L \times_{B_2} L'')$, and if $\mathcal{U}$ is the vertex family associated to H, then we have*

$$f(\mathcal{U}, G) > 0.$$

If condition (1) of Theorem A.5 is false, then the hypothesis of the above lemma holds; but item (4) of the lemma means that condition (2) of Theorem A.5 is false. Hence we conclude this section, and the proof of Theorem A.5 with the proof of the above lemma.

Proof. So assume the SHNC is false on a pair of subgraphs, (L, L'). Similar to before, the SHNC must therefore be false on (L, L''), where L'' is some connected component of L. Fix such a connected component, L''.

We cannot have $\rho(L'') = 0$, for otherwise $\rho(L \times_{B_2} L'') = 0$ and the SHNC is not false on (L, L''). Hence we have $L'' \subset L'$ is connected and $\rho(L'') > 0$, whereupon we have $-\chi(L'') = \rho(L'')$.

Now let $L'' \subset L'$ be a minimal subgraph of L' (with respect to inclusion of subgraphs) with the properties that L'' is connected, $\rho(L'') > 0$, and the SHNC is false on (L, L''). We shall show that the lemma holds with this subgraph, L''; we have already established the first three items in the conclusion.

Then take any $H \subset L \times_{B_2} L''$ such that $-\chi(H) = \rho(L \times_{B_2} L'')$. We now make a number of claims regarding L'' that follow from the minimality of L''.

First, we claim that L'' has no leaves, i.e., no vertices of degree one. Otherwise, if v is a vertex of degree one, and e is its incident edge, then there are at least as many vertices in H over v as there are over v. So letting L''' be L'' with v and e discarded, we see that $\rho(L''') = \rho(L'')$; but if H' consists of the vertices and edges of H that do not lie over e or v, then H' is a subgraph of H and $\rho(H') = \rho(H)$ (since we obtain H' from H by discarding isolated vertices over v or vertices over v along with their single incident edges, lying over e). Hence the SHNC would fail also on (L, L'''), contradicting the minimality of L''.

Second, we claim that over each $e \in E_{L''}$ we there are at least $\rho(L) + 1$ edges in H; if not, we delete e from L'', obtaining $L''' \subset L''$, and delete the at most $\rho(L)$

edges over e from H, obtaining H' that lies over L'''; this yields a strict subgraph, L''' of L'' such that

$$\begin{aligned}\rho(L \times_{B_2} L''') \geq -\chi(H') \geq -\chi(H) - \rho(L) &= \rho(L \times_{B_2} L'') - \rho(L) \\ &> \rho(L)\bigl(\rho(L'') - 1\bigr).\end{aligned}$$

But since L'' is pruned, we have $\rho(L''') = \rho(L'') - 1$. So, once again, we have the SHNC fails on (L, L''') for some a proper subgraph, L''', of L''; this contradicts the minimality of L''.

Third, we claim that over each $v \in V_{L''}$ there are at least $\rho(L) + 1$ vertices in H. Indeed, if v is incident upon some edge, e, in L'', then e has at least $\rho(L) + 1$ vertices in H above it, so v does as well. If v is isolated in L'', i.e., incident upon no edge, then L'' consists of only v, since L'' is connected; but this contradicts the fact that $\rho(L'') > 0$.

To H is associated a vertex family, $\mathcal{U}$, and an edge family, $\mathcal{W}$. According to the three claims established in the previous three paragraphs, we have

$$v \in V_{L''} \implies |\mathcal{U}(v)| \geq \rho(L) + 1, \qquad e \in E_{L''} \implies |\mathcal{W}(e)| \geq \rho(L) + 1.$$

Clearly also

$$v \notin V_{L''} \implies \mathcal{U}(v) = \emptyset, \qquad e \notin E_{L''} \implies \mathcal{W}(e) = \emptyset.$$

It follows that, as before

$$\begin{aligned} f(\mathcal{U}, G) &\geq \sum_{e \in E_{L''}} |\mathcal{W}(e)|_{\rho(L)} - \sum_{v \in V_{L''}} |\mathcal{U}(e)|_{\rho(L)} \\ &= -\chi(H) - \rho(L)(|E_{L''}| - |V_{L''}|) = \rho(L \times_{B_2} L'') - \rho(L)\rho(L''),\end{aligned}$$

using the fact that L'' is connected and $\rho(L'') > 0$ (so that $\rho(L'') = |E_{L''}| - |V_{L''}|$). Hence $f(\mathcal{U}, G) > 0$, which shows item (4) in the conclusion of the lemma. $\square$

$\square$

Bibliography

[AL02] Alon Amit and Nathan Linial, *Random graph coverings I: General theory and graph connectivity*, Combinatorica **22** (2002), no. 1, 1–18.

[AM69] M. F. Atiyah and I. G. Macdonald, *Introduction to commutative algebra*, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969. MR 0242802 (39 #4129)

[Arz00] G. N. Arzhantseva, *A property of subgroups of infinite index in a free group*, Proc. Amer. Math. Soc. **128** (2000), no. 11, 3205–3210. MR MR1694447 (2001b:20040)

[Ati76] M. F. Atiyah, *Elliptic operators, discrete groups and von Neumann algebras*, Colloque "Analyse et Topologie" en l'Honneur de Henri Cartan (Orsay, 1974), Soc. Math. France, Paris, 1976, pp. 43–72. Astérisque, No. 32–33. MR MR0420729 (54 #8741)

[Bur71] Robert G. Burns, *On the intersection of finitely generated subgroups of a free group.*, Math. Z. **119** (1971), 121–130. MR MR0279166 (43 #4892)

[Del77] P. Deligne, *Cohomologie étale*, Lecture Notes in Mathematics, Vol. 569, Springer-Verlag, Berlin, 1977, Séminaire de Géométrie Algébrique du Bois-Marie SGA $4\frac{1}{2}$, Avec la collaboration de J. F. Boutot, A. Grothendieck, L. Illusie et J. L. Verdier. MR MR0463174 (57 #3132)

[DF01] Warren Dicks and Edward Formanek, *The rank three case of the Hanna Neumann conjecture*, J. Group Theory **4** (2001), no. 2, 113–151. MR MR1812321 (2002e:20051)

[Dic94] Warren Dicks, *Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture*, Invent. Math. **117** (1994), no. 3, 373–389. MR MR1283723 (95c:20034)

[Eve08] Brent Everitt, *Graphs, free groups and the Hanna Neumann conjecture*, J. Group Theory **11** (2008), no. 6, 885–899. MR MR2466915

[FMT06] Joel Friedman, Ram Murty, and Jean-Pierre Tillich, *Spectral estimates for abelian cayley graphs*, J. Comb. Theory Ser. B **96** (2006), no. 1, 111–121.

[Fri93] Joel Friedman, *Some geometric aspects of graphs and their eigenfunctions*, Duke Math. J. **69** (1993), no. 3, 487–525. MR 94b:05134

[Fri03] ———, *Relative expanders or weakly relatively Ramanujan graphs*, Duke Math. J. **118** (2003), no. 1, 19–35. MR MR1978881 (2004m:05165)

[Fri05] ———, *Cohomology of grothendieck topologies and lower bounds in boolean complexity*, `http://www.math.ubc.ca/~jf`, also `http://arxiv.org/abs/cs/0512008`, to appear.

[Fri06] ———, *Cohomology of grothendieck topologies and lower bounds in boolean complexity ii*, `http://www.math.ubc.ca/~jf`, also `http://arxiv.org/abs/cs/0604024`, to appear.

[Fri07] ———, *Linear transformations in boolean complexity theory*, CiE '07: Proceedings of the 3rd conference on Computability in Europe (Berlin, Heidelberg), Springer-Verlag, 2007, pp. 307–315.

[Fri08] ———, *A proof of Alon's second eigenvalue conjecture and related problems*, Mem. Amer. Math. Soc. **195** (2008), no. 910, viii+100. MR MR2437174

[Fri11a] Joel Friedman, *Sheaves on graphs and a proof of the hanna neumann conjecture*, available at `http://arxiv.org/pdf/1105.0129v1` and at `http://www.math.ubc.ca/~jf`.

[Fri11b] ———, *Sheaves on graphs and their homological invariants*, available at `http://arxiv.org/pdf/1104.2665v1` and at `http://www.math.ubc.ca/~jf`.

[FT05] Joel Friedman and Jean-Pierre Tillich, *Generalized Alon–Boppana theorems and error-correcting codes*, SIAM J. Discret. Math. **19** (2005), 700–718.

[Ger83] S. M. Gersten, *Intersections of finitely generated subgroups of free groups and resolutions of graphs*, Invent. Math. **71** (1983), no. 3, 567–591. MR MR695907 (85m:05037b)

[GM03] Sergei I. Gelfand and Yuri I. Manin, *Methods of homological algebra*, second ed., Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003. MR MR1950475 (2003m:18001)

[Gro77] Jonathan L. Gross, *Every connected regular graph of even degree is a Schreier coset graph*, J. Combinatorial Theory Ser. B **22** (1977), no. 3, 227–232. MR MR0450121 (56 #8419)

[Har77] Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York, 1977, Graduate Texts in Mathematics, No. 52. MR 57 #3116

[Har92] Joe Harris, *Algebraic geometry*, Graduate Texts in Mathematics, vol. 133, Springer-Verlag, New York, 1992, A first course. MR 1182558 (93j:14001)

[How54] A. G. Howson, *On the intersection of finitely generated free groups*, J. London Math. Soc. **29** (1954), 428–434. MR MR0065557 (16,444c)

[HS97] P. J. Hilton and U. Stammbach, *A course in homological algebra*, second ed., Graduate Texts in Mathematics, vol. 4, Springer-Verlag, New York, 1997. MR 1438546 (97k:18001)

[Imr77a] Wilfried Imrich, *On finitely generated subgroups of free groups*, Arch. Math. (Basel) **28** (1977), no. 1, 21–24. MR MR0439941 (55 #12822)

[Imr77b] ———, *Subgroup theorems and graphs*, Combinatorial mathematics, V (Proc. Fifth Austral. Conf., Roy. Melbourne Inst. Tech., Melbourne, 1976), Springer, Berlin, 1977, pp. 1–27. Lecture Notes in Math., Vol. 622. MR MR0463016 (57 #2980)

[Iva99] S. V. Ivanov, *On the intersection of finitely generated subgroups in free products of groups*, Internat. J. Algebra Comput. **9** (1999), no. 5, 521–528. MR MR1719719 (2000k:20023)

[Iva01] ———, *Intersecting free subgroups in free products of groups*, Internat. J. Algebra Comput. **11** (2001), no. 3, 281–290. MR MR1847180 (2002e:20052)

[JKM03] Toshiaki Jitsukawa, Bilal Khan, and Alexei G. Myasnikov, *On the Hanna Neumann conjecture*, 2003, Available as `http://arxiv.org/abs/math/0302009`.

[Kha02] Bilal Khan, *Positively generated subgroups of free groups and the Hanna Neumann conjecture*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 155–170. MR MR1921710 (2003e:20027)

[Lan02] Serge Lang, *Algebra*, third ed., Graduate Texts in Mathematics, vol. 211, Springer-Verlag, New York, 2002. MR 1878556 (2003e:00003)

[Lüc02] Wolfgang Lück, *L^2-invariants: theory and applications to geometry and K-theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics], vol. 44, Springer-Verlag, Berlin, 2002. MR MR1926649 (2003m:58033)

[Min10] Igor Mineyev, *The topology and analysis of the Hanna Neumann Conjecture*, Preprint. Available at `http://www.math.uiuc.edu/~mineyev/math/art/shnc.pdf`.

[Min11a] ———, *Groups, graphs, and the Hanna Neumann Conjecture*, Preprint. Available at `http://www.math.uiuc.edu/~mineyev/math/art/gr-gr-shnc.pdf`.

[Min11b] ———, *Submultiplicativity and the Hanna Neumann Conjecture*, Preprint. Available at `http://www.math.uiuc.edu/~mineyev/math/art/submult-shnc.pdf`.

[MW02] J. Meakin and P. Weil, *Subgroups of free groups: a contribution to the Hanna Neumann conjecture*, Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000), vol. 94, 2002, pp. 33–43. MR MR1950872 (2003k:20028)

[Neu56] Hanna Neumann, *On the intersection of finitely generated free groups*, Publ. Math. Debrecen **4** (1956), 186–189. MR MR0078992 (18,11f)

[Neu57] ———, *On the intersection of finitely generated free groups. Addendum*, Publ. Math. Debrecen **5** (1957), 128. MR MR0093537 (20 #61)

[Neu90] Walter D. Neumann, *On intersections of finitely generated subgroups of free groups*, Groups—Canberra 1989, Lecture Notes in Math., vol. 1456, Springer, Berlin, 1990, pp. 161–170. MR MR1092229 (92b:20026)

[Neu07] Walter D. Neumann, *A short proof that positive generation implies the Hanna Neumann Conjecture*, 2007, Available as `http://arxiv.org/abs/math/0702395`, to appear.

[Ser83] Brigitte Servatius, *A short proof of a theorem of Burns*, Math. Z. **184** (1983), no. 1, 133–137. MR MR711734 (85c:20019) [sga72a] *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. MR 50 #7130

[sga72b] *Théorie des topos et cohomologie étale des schémas. Tome 2*, 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. 270. MR 50 #7131

[sga73] *Théorie des topos et cohomologie étale des schémas. Tome 3*, Springer-Verlag, Berlin, 1973, 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 P. Deligne et B. Saint-Donat, Lecture Notes in Mathematics, Vol. 305. MR 50 #7132

[sga77] *Cohomologie l-adique et fonctions L*, Lecture Notes in Mathematics, Vol. 589, Springer-Verlag, Berlin, 1977, Séminaire de Géometrie Algébrique du Bois-Marie 1965–1966 (SGA 5), Edité par Luc Illusie. MR MR0491704 (58 #10907)

[ST96] H. M. Stark and A. A. Terras, *Zeta functions of finite graphs and coverings*, Adv. Math. **121** (1996), no. 1, 124–165. MR MR1399606 (98b:11094)

[Sta83] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR MR695906 (85m:05037a)

[Tar92] Gábor Tardos, *On the intersection of subgroups of a free group*, Invent. Math. **108** (1992), no. 1, 29–36. MR MR1156384 (93c:20048)

[Tar96] ———, *Towards the Hanna Neumann conjecture using Dicks' method*, Invent. Math. **123** (1996), no. 1, 95–104. MR MR1376247 (97b:20029)

[Wis05] Daniel T. Wise, *The coherence of one-relator groups with torsion and the Hanna Neumann conjecture*, Bull. London Math. Soc. **37** (2005), no. 5, 697–705. MR MR2164831 (2006f:20037)