Sheaves on Graphs and Their Homological Invariants

Joel Friedman*

April 13, 2011

Abstract

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. The maximum excess gives a simple interpretation of an important graph invariant, which will be used to study the Hanna Neumann Conjecture in a future paper.

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.

*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. jf@cs.ubc.ca or jf@math.ubc.ca. Research supported in part by an NSERC grant. Research done in part at the Centre Bernoulli, funded by the Swiss National Science Foundation.

1 Introduction

The main goal of this paper 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 a future paper we shall use the results of this paper to study this conjecture. Moreover, our methods will address 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 paper 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 paper 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 study of the SHNC. In more precise terms, 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 theorems in graph theory in our study of the SHNC.

This paper 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 paper. This discussion will be made more precise, with more background, in Section 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 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 paper. 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 paper 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 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 study 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 paper 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 the inequalities regarding the maximum excess that we need to study the SHNC.

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 future study 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 Betti number calculation (for sheaves on graphs), that includes more information, such as giving a simple interpretation of this number (the maximum excess) and a finite algorithm for computing it (pulling back and computing a twisted Betti number, which is a finite procedure).

We note that for the purpose of studying the SHNC, the main results needed from this paper 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.

The rest of this paper is organized as follows. In Section 2 we give precise definitions and statements of the theorems in this paper. In Section 3 we review part of what might be called “Galois theory of graphs” that we will use in this paper. In Section 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 the paper). In Section 5 we define the twisted cohomology and compute the twisted cohomology of the constant sheaf of a graph; we also interpret twisted cohomology in terms of Abelian covers. In Section 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 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 8 we prove Theorem 2.10, that says that the first twisted Betti number and the maximum excess agree after an appropriate pullback. In Section 9 we make some concluding remarks.

We wish to thank Laurant Bartholdi, for conversations and introducing us to the SHNC, and Avner Friedman, for comments on a draft of this paper. We thank Luc Illusie, for an inspiring discussion regarding our ideas involving sheaf theory, homology, and the SHNC. 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.

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 paper, and state or describe other results in this paper. In most of this paper we work with directed graphs (digraphs), which makes things notationally simpler; as we remark in Section 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.

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

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

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

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

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 2.1 *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.*

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 2.2 *We define the* zeroth *and* first homology groups of F *to be, respectively,*

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

We denote by $h_i(G, \mathcal{F})$ the dimension of $H_i(G, \mathcal{F})$ as an $\mathbb{F}$-vector space, and call it the i-th Betti number of $\mathcal{F}$. *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 $H_i(G)$, 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 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 $G' \to G''$; one example of such a morphism is a surjection whose kernel is what we call a ρ-kernel, which will be crucial to our study of the SHNC.

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

Definition 2.3 *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 2.4 *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

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

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 paper we introduce other pairs of invariants, essentially variations of h_0, h_1, that satisfy the same properties. Our study of the SHNC will be based on the fact that the "maximum excess" is part of such a pair. Let us make these notions precise.

Definition 2.5 *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 2.6 *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*

$$\alpha_0(\mathcal{F}) - \alpha_1(\mathcal{F}) = \chi(\mathcal{F}); \qquad (2)$$

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

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

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 (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 2.4.

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 2.7 *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*

$$\Gamma_{\rm ht}(U) = \bigoplus_{e \in E_G} \{w \in \mathcal{F}(e) \mid d_h(w), d_t(w) \in U\}; \tag{3}$$

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

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

Theorem 2.8 *The maximum excess is a first quasi-Betti number.*

Theorem 2.8 will be crucial to our study of the SHNC. Somewhat surprisingly, the statement of this theorem and all the necessary definitions do not involve any homology theory.

We shall show Theorem 2.8 by identifying the maximum excess with a certain "limit" Betti number.

2.3 Twisted Homology

One graph theoretic reformulation of the SHNC involves the reduced cyclicity defined in equation (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

$$\rho(G) = \lim_{n\to\infty} h_1(L_n)/n, \tag{4}$$

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 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 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

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

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

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

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

We easily see, akin to equation (4), that

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

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

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

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

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^{twist} and the maximum excess agree and equal $\rho(G)$. Notice also that it is immediate that h_1^{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^{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

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

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

$$\mathcal{U}(v) = \mathbb{F}^4, \quad \mathcal{U}(e_i) = \mathbb{F}^2 \quad \text{for } i = 1, 2, \tag{5}$$

and

$$d_h = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad d_t = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix}, \tag{6}$$

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

$$d_{\mathcal{U}^\psi} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & -\psi(e_2) & 0 \\ -\psi(e_1) & 0 & 0 & 1 \\ 0 & -\psi(e_1) & 0 & -\psi(e_2) \end{bmatrix}. \tag{7}$$

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

$$\nu_1 = \alpha - \psi(e_1)\gamma, \quad \nu_2 = \beta - \psi(e_1)\delta, \quad \nu_3 = \alpha - \psi(e_2)\beta, \quad \nu_4 = \gamma - \psi(e_2)\delta. \tag{8}$$

The fact that $h_1^{\mathrm{twist}}(\mathcal{U}) \neq 0$ 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, that

$$\mathrm{m.e.}(\mu^*\mathcal{F}) = \deg(\mu)\mathrm{m.e.}(\mathcal{F}). \tag{9}$$

The "unhappy 4-bundle" also shows that h_1^{twist} does not enjoy this "scaling by $\deg(\mu)$ under pullback" property. Indeed, $h_1^{\mathrm{twist}}(\mathcal{U}) = 1$; 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 $h_1^{\mathrm{twist}}(\phi^*\mathcal{U}) = 0$. 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

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

2.5 The Fundamental Lemma and Limit Homology

The following is the main and most difficult theorem in this paper; 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 2.10 *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

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

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

2.6 Limits and Limiting Betti Numbers

In this subsection we give a new interpretation to our main theorem, Theorem 2.10. 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 2.10 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 2.10 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 2.10 implies Theorem 2.8.

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.

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).

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 3.1 (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.

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

$$h_K(e,a) = (h_G e, a_e a), \quad t_K(e,a) = (t_G e, a), \tag{10}$$

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

$$(P,a)g = (P, ag); \tag{11}$$

in view of the fact that a_e multiplies to the left in equation (10), we see that the right multiplication of g on a in equation (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

$$a_e \mapsto \widetilde{a}_e = g_{he}^{-1}a_e g_{te} \tag{12}$$

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].

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 3.2 *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

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

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 (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.

3.4 Covering maps and ρ

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

Theorem 3.3 *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

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

$\square$

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.

4.1 Homology and Pullbacks

The basic definitions of sheaves were given in Subsection 2.1. In this subsection we prove Theorem 2.4 and discuss pullbacks and related functors.

Proof Of Theorem 2.4. By the "vertexwise and edgewise" nature of taking images and kernels, we see that we have a diagram

$$\begin{tikzcd}
 & 0 \arrow[d] & 0 \arrow[d] & 0 \arrow[d] & \\
0 \arrow[r] & \mathcal{F}_1(E) \arrow[r] \arrow[d] & \mathcal{F}_2(E) \arrow[r] \arrow[d] & \mathcal{F}_3(E) \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{F}_1(V) \arrow[r] \arrow[d] & \mathcal{F}_2(V) \arrow[r] \arrow[d] & \mathcal{F}_3(V) \arrow[r] \arrow[d] & 0 \\
 & 0 & 0 & 0 &
\end{tikzcd}$$

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

□

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 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

$$\underline{\mathbb{F}}_K(V_G) = \mathbb{F}^{V_K}, \quad \underline{\mathbb{F}}_K(E_G) = \mathbb{F}^{E_K}, \tag{13}$$

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 $\mathcal{F}(h,e')$. We shall make use of $\phi_!$ for ϕ étale in our approach to the SHNC. The reader can now observe that

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

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

$$H_i(\phi_!\mathcal{F}) \simeq H_i(\mathcal{F}) \tag{14}$$

for $i = 0, 1$. In Subsection 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 4.4 we explain the role of $\phi_!$ in certain "vanishing theorems" (of sheaf invariants).

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 (14), 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 (13) and the remarks just below it.

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

[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."

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 2.1).

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 (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 2.4 and 4.1); u^* has a left adjoint, $u_!$ (defined in Subsection 4.1), and a right adjoint, u_* (see [sga72a], Exposé I, Proposition 5.1). In other words,
$$\mathrm{Hom}_G(\phi_!\mathcal{F}, \mathcal{L}) \simeq \mathrm{Hom}_K(\mathcal{F}, \phi^*\mathcal{L}) \quad \forall \mathcal{F} \in \mathbf{Sh}(G),\ \mathcal{L} \in \mathbf{Sh}(K), \tag{15}$$
and similarly for ϕ_*.

3. As a consequence we have
$$\mathrm{Ext}^i_G(\phi_!\mathcal{F}, \mathcal{L}) \simeq \mathrm{Ext}^i_K(\mathcal{F}, \phi^*\mathcal{L}) \quad \forall \mathcal{F} \in \mathbf{Sh}(G),\ \mathcal{L} \in \mathbf{Sh}(K), \tag{16}$$
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

$$\underline{\mathbb{F}}\to\bigoplus_{v\in V_G}(k_v)_*\mathbb{F}\to\bigoplus_{e\in E_G}(k_e)_*\mathbb{F} \tag{17}$$

(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:

$$\bigoplus_{v\in V_G}(k_v)_!\mathbb{F}^{d_v-1}\to\bigoplus_{e\in E_G}(k_e)_!\mathbb{F}\to\underline{\mathbb{F}}, \tag{18}$$

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 (17), 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 (17) and (18), 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.

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 (15).

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'.

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 "induced 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 4.1 *If α_1 is any first quasi-Betti number for sheaves on a graph, G, and if $\alpha_1(\mathcal{F}) = 0$, 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 $\mathrm{m.e.}(\mathcal{F}_2) \le \mathrm{m.e.}(\mathcal{F}_3)$, then a simple homological explanation for this inequality would be that $\mathrm{m.e.}(\mathcal{F}_1) = 0$. But this would, in turn, imply that $\mathrm{m.e.}((\mathcal{F}_2)_{G'}) \le \mathrm{m.e.}((\mathcal{F}_3)_{G'})$ for all open subsets, G', of G, which could be a much stronger inequality (and is much stronger for the setting of the SHNC).

A slightly stronger "contagious vanishing" theorem says that that if $\mathrm{m.e.}(\mathcal{F}) = 0$ for a sheaf, $\mathcal{F}$, on a digraph, G, and $\nu\colon G' \to G$ is étale, then $\mathrm{m.e.}(\mathcal{F}_{G'}) = 0$ where $\mathcal{F}_{G'} = \nu_! \nu^* \mathcal{F}$. This follows easily once we prove equation (9) (as Theorem 6.5) and Theorem 2.10. Indeed, since ν is étale, it factors as an open inclusion $j\colon G' \to G''$ followed by a covering map $\mu\colon G'' \to G$. From Theorem 6.5 we will know that $\mathrm{m.e.}(\mathcal{F}) = 0$ implies $\mathrm{m.e.}(\mu^*\mathcal{F}) = 0$ which implies $\mathrm{m.e.}(\mathcal{F}') = 0$, where

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

using "contagious vanishing" for open inclusions. But equation (16), applied to twisted homology in the special case $\mathcal{L} = \underline{\mathbb{F}}$, and taking limits, shows that for any sheaf, $\mathcal{K}$, on G'' we have

$$\mathrm{m.e.}(\mu_!\mathcal{K}) = \mathrm{m.e.}(\mathcal{K}).$$

Hence $\mu_!\mathcal{F}'$ has maximum excess zero; but

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

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; we also describe a sheaf on a graph that we call the "unhappy 4-bundle," whose maximum excess is zero, but whose first twisted Betti number is one.

5.1 Remarks on the Definition

Twists and twisted homology were defined in Subsection 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).

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

$$s \mapsto (1/n) \sum_{\alpha \in \mathcal{A}} \nu^{-1}(\alpha)(\alpha s); \tag{19}$$

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 5.1 *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*

$$H_i(\phi^*\mathcal{F}) \simeq \bigoplus_{\psi} \Bigl(H_i(\phi^*\mathcal{F})\Bigr)^{\nu}, \tag{20}$$

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 (19) gives isomorphisms

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

and $d_{\phi^*\mathcal{F}}$ intertwines with these maps, which establishes equation (20). It remains to identify

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

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 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

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

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

Lemma 5.1 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 5.1 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 5.2 *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 3.3. In this case we have $h_1(G') = 1-\chi(G') = 1 - p\chi(G) = p\rho(G) + 1$. But by Lemma 5.2 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 5.3 *For any digraph, G, we have $\rho(G) = h_1^{\mathrm{twist}}(\underline{\mathbb{F}})$.*

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 5.4 *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 2.10 says that equality will hold if $\mathcal{F}$ is pulled back appropriately.

Definition 5.5 *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 5.6 *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 5.4, 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

$$(a_1 + \psi_1 b_1) \wedge \cdots \wedge (a_{r'+1} + \psi_{r'+1} b_{r'+1}) = 0. \tag{21}$$

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

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

In other words, we have shown that if U is the span of the $a_1, \ldots, a_r$, 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 $j = 1, \ldots, r$. 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

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

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 $v_1, \ldots, v_r$ as above. If some element of $v_1, \ldots, v_r$ is critical, we may assume it is v_1; in this case, if some element of $v_2, \ldots, v_r$ 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 $v_i, \ldots, v_r$, and no element of $v_s, \ldots, v_r$ 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 $\{v_s, \ldots, v_r\}$ is of size $r - s + 1$, but also the span of $\{v_s, \ldots, v_r\}$ 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 $v_s, \ldots, v_r$ 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

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

□

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 8.

6.1 Compartmentalized Subspaces

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

Definition 6.1 *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* $\operatorname{extend}(v,s)$, *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*

$$\operatorname{supportedIn}(s, W') = \Bigl\{ v \in W_s \mid \pi\bigl(\operatorname{extend}(v,s)\bigr) \in W' \Bigr\},$$

and let the compartmentalization of W' *be*

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

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 8, we use two simple observations about the situation of Definition 6.1. 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

$$\bigoplus_{s\in S} Q_s \to W/W'. \tag{22}$$

It will be helpful to formally combine these two observations into a theorem that follows immediately; we will use this theorem repeatedly in Section 8, in our proof of Theorem 2.10.

Theorem 6.2 *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 (3); this is crucial to the resulting definition of excess and maximum excess, in Definition 2.7. 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 $\Gamma_{\rm ht}(U) \otimes_{\mathbb{F}} \mathbb{F}(\psi)$ to $U \otimes_{\mathbb{F}} \mathbb{F}(\psi)$.

6.2 Supermodularity and Its Consequences

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

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

and hence

$$\text{m.e.}(\mathcal{F}) \geq \max\Big(0, -\chi(\mathcal{F})\Big).$$

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

Theorem 6.3 *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 (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.

□

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

Theorem 6.4 *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.,*

$$\mathrm{excess}(U_1) + \mathrm{excess}(U_2) \le \mathrm{excess}(U_1 \cap U_2) + \mathrm{excess}(U_1 + U_2) \tag{23}$$

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

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

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

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

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

$$\dim(U_1) + \dim(U_2) = \dim(U_1 \cap U_2) + \dim(U_1 + U_2). \tag{24}$$

On the other hand

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

and

$$\Gamma_{ht}(U_1 + U_2) \supset \Gamma_{ht}(U_1) + \Gamma_{ht}(U_2); \tag{25}$$

hence

$$\dim\Bigl(\Gamma_{ht}(U_1)\Bigr)+\dim\Bigl(\Gamma_{ht}(U_2)\Bigr) \le \dim\Bigl(\Gamma_{ht}(U_1\cap U_2)\Bigr)+\dim\Bigl(\Gamma_{ht}(U_1+U_2)\Bigr). \tag{26}$$

Combining equations (24) and (26) yields equation (23). 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 (26) and hence (25) must hold with equality.

□

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

Theorem 6.5 *Let $\phi\colon G' \to G$ be a covering map of graphs, and let $\mathcal{F}$ be a sheaf on G. Then*

$$\text{m.e.}(\phi^*\mathcal{F}) = \deg(\phi)\ \text{m.e.}(\mathcal{F}). \tag{27}$$

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 6.4 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

$$\text{excess}\Bigl(\mathcal{F}', \phi^{-1}(T)\Bigr) = \deg(\phi)\text{excess}(\mathcal{F}, T). \tag{28}$$

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

$$\text{m.e.}(\mathcal{F}') \ge \deg(\phi)\text{m.e.}(\mathcal{F}). \tag{29}$$

It remains to prove the reverse inequality in order to establish equation (27); note that if we do so, then the second statement of the theorem follows from equation (28).

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

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

In summary,

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

From equation (29), 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 3.1), 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

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

and since ν is Galois we have

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

It follows that

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

$\square$

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 2.10. We shall illustrate these theorems on the unhappy 4-bundle, which gives great insight into our proof of Theorem 2.10 that we give in Section 8.

Let $\mathbb{Z}$ be the set of integers, and let $\mathbb{Z}_{\geq 0}$ 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 7.1 *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 2.10 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

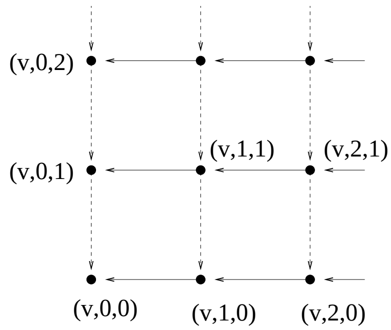

Figure 1: $B_2[\mathbb{Z}_{\geq 0}]$.

just the usual d-dimensional non-negative integer lattice, depicted in Figures 1 and 2. If $G' \to G$ is a covering 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 2.4 (just beneath equation (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 $e \in E_K$. 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 $i = 0, 1$.

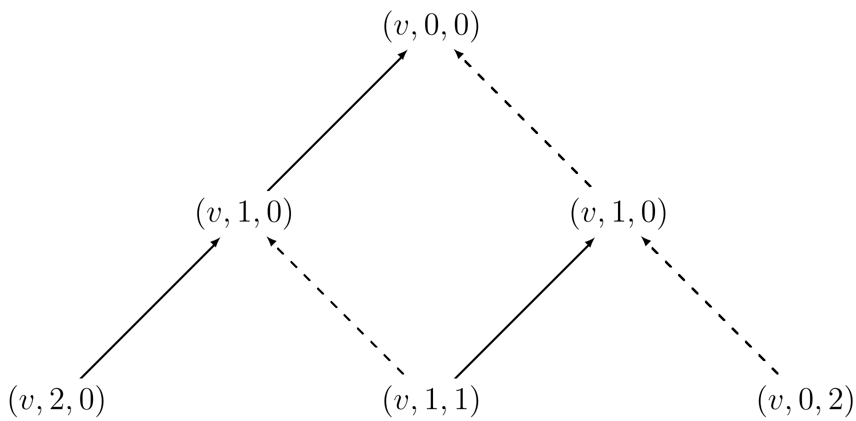

Figure 2: First part of $B_2[\mathbb{Z}_{\geq 0}]$. Notice the cycle of length four.

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 7.2 *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

$$\sum_{e \in E_G} \sum_{i=1}^{d_e} c_{e,i}(\psi)(a_{e,i} + \psi(e)b_{e,i}) = 0, \tag{30}$$

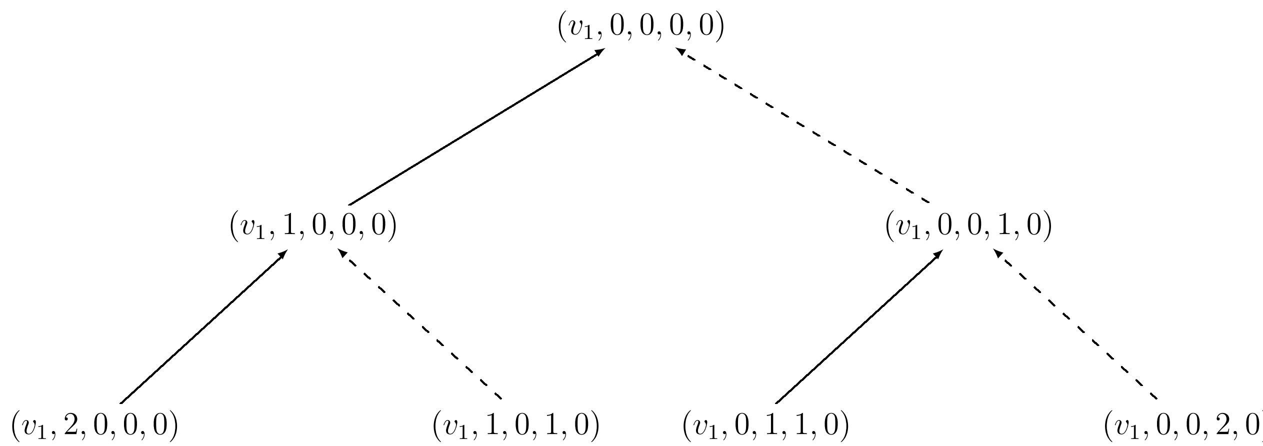

Figure 3: First part of $G'[\mathbb{Z}_{\geq 0}]$ near $(v, \vec{0})$. No cycles of length four. The four $(\mathbb{Z}_{\geq 0})^{E_{G'}}$ coordinates are, in order, $e_1^1, e_1^2, e_2^1, e_2^2$ where e_i^j lies over $e_i \in E_{B_2}$ and are described in the last equations of Subsection 2.4 that give the ν_i^j.

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

$$\sum_{n \in (\mathbb{Z}_{\geq 0})^{E_G}} \sum_{e,i} \psi^n c_{e,i,n}(a_{e,i} + \psi_e b_{e,i}) = 0 \tag{31}$$

and not all the $c_{e,i,n} = 0$. But equation (31) 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 $H_1^{\oplus}(\pi^* \mathcal{F}) \neq 0$.

$\square$

The following is a simple graph theoretic definition that is crucial to our proof of Lemma 8.1.

Definition 7.3 *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.

8 Proof of Theorem 2.10

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

Lemma 8.1 *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*

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

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

In Subsection 8.7, the last subsection of this section, we use this lemma to prove Theorem 2.10. 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.

8.1 Outline of the Proof of Lemma 8.1

Consider the hypotheses of Lemma 8.1. Let $\pi\colon G'[\mathbb{Z}] \to G'$ be the universal Abelian cover of G', and let $\mathcal{F}' = \mu^*\mathcal{F}$. We assume $h_1^{\rm twist}(\mathcal{F}') \geq 1$, and we wish to prove that $\mathrm{m.e.}(\mathcal{F}) \geq 1$. According to Lemma 7.2, 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 7.2), 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

$$\dim\Bigl(A(F) + B(F)\Bigr) \leq \dim\Bigl(C(F)\Bigr) - 1. \tag{32}$$

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

$$\dim\Bigl(A(F_{k_1})\Bigr) \leq \dim\Bigl(C(F_{k_1})\Bigr) - k_1. \tag{33}$$

This inequality is worse than equation (32) 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 (40) and (55) 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 (33) that involves $B(F_{k_1})$ (namely equation (53)); 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 (56)).

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 (32) or (33). 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 (57) or (58) hold) our phases end after completing the i-th phase, whereupon taking $r = k_i$ we will have that $F = F_r$ satisfies equation (32) 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 8.2. We shall describe the first and second phase, respectively, in detail in Subsections 8.3 and 8.5, respectively. In Subsection 8.4 we state and prove a number of facts used in Subsections 8.3 and 8.5 in greater generality; we hope that this greater generality will clarify the proofs. In Subsection 8.6 we finish the proof of Lemma 8.1. As mentioned before, in Subsection 8.7, we use Lemma 8.1 to prove Theorem 2.10.

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 8.2 *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 8.2, 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 (32), 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)

8.3 The First C-Phase

We remind the reader that, as explained at the end of Subsection 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

$$\sum_{e \text{ s.t. } he=v_1} d_h\overline{w_e} = \sum_{e \text{ s.t. } te=v_1} d_t\overline{w_e} = 0. \tag{34}$$

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

$$\dim\Big(C(F_1)\Big) = |E^1|; \tag{35}$$

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 6.1, the nonzero w_e are compartmentally distinct, and hence independent, by Theorem 6.2. Hence equation (35) holds. By contrast, equation (34) shows that the $d_h\overline{w_e}$ with $e \in E^1$ sum to zero and are therefore dependent; hence

$$\dim\Big(A(F_1)\Big) \le |E^1| - 1,$$

and so

$$\dim\Big(A(F_1)\Big) \le \dim\Big(C(F_1)\Big) - 1. \tag{36}$$

Assume that there is an $e_2 \in E_{G'[\mathbb{Z}_{\ge 0}]}$ 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

$$\dim\Big(A(F_2)/A(F_1)\Big) \le \dim\Big(C(F_2)/C(F_1)\Big) - 1. \tag{37}$$

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 6.2, the $\overline{w_e}$ with $e \in E^2$ are linearly independent in $\mathcal{F}(E)/C(F_1)$. Hence

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

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

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

This establishes equation (37), and adding that equation to equation (36) gives

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

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

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

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

$$\dim\Bigl(A(F_i)\Bigr) \leq \dim\Bigl(C(F_i)\Bigr) - i. \tag{38}$$

But for any such i we have

$$\dim\Bigl(C(F_i)\Bigr) \geq i; \tag{39}$$

hence for any such i we have $i \leq \dim(\mathcal{F}(E))$, and so for some $k_1 \leq \dim(\mathcal{F}(E))$ 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

$$\forall e \in E_{G[\mathbb{Z}_{\geq 0}]} \text{ of rank } \rho, \quad \overline{w_e} \in C(F_{k_1}), \tag{40}$$

where minimal means that if we discard any v_i from $v_1, \ldots, v_{k_1}$ then equation (40) does not hold. We call this a C-phase because the equation (40) involves a "C," as will all odd numbered phases. Notice that equation (38) is somewhat similar to our desired equation (32); one big difference is that equation (38) makes no mention of B, but only of A and C.

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 8.1, which fixes $\mathcal{F}$, $\mu\colon G' \to G$, $\pi\colon G'[\mathbb{Z}_{\geq 0}] \to G'$, $w \in H_1^{\oplus}(\pi^*\mu^*\mathcal{F})$, and defines $\overline{w_f}$ for any $f \in E_{G'[\mathbb{Z}_{\geq 0}]}$, and defines $A(F), B(F), C(F)$ for any $F \subset E_{G'[\mathbb{Z}_{\geq 0}]}$.

We will work with a sequence of vertices, $v = (v_1, \ldots, v_s)$, of $G'[\mathbb{Z}_{\geq 0}]$, 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 8.3 *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 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 8.4 *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 8.5 *Let* v *be a moseying sequence of length* s *of increasing dimension for a digraph,* G'*. Then*

$$s \le \dim\Bigl(\mathcal{F}(E)\Bigr) + \dim\Bigl(\mathcal{F}(V)\Bigr).$$

Furthermore, for any $i \le s$*,* 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 (38), derived after the first C-phase, will be built up along further phases to eventually give equation (32). 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 8.6 *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

$$\mathrm{rank}(v_j) \geq \mathrm{rank}(u) \implies \mathrm{sod}(u) = \mathrm{outdeg}(U_{j-1}, u), \tag{41}$$

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 8.7 *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*

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

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

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

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

Lemma 8.9 *Let v be a moseying sequence of length s for G'. Assume that U_j has no cycles for some $j \leq s$. Then for any non-negative integers $i \leq j$ we have*

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

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

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

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

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

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

$$h_0(U_{i+1}) = h_0(U_i) + 1 - p. \tag{42}$$

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 (42) 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 8.10 *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 $\mathrm{rank}(v_i)$ we have $\overline{w_e} \in C(F_i)$. Then for any j with $i+1 \le j \le k$ we have*

$$\dim\Big(A(F_k)/\big(A(F_j)+B(F_i)\big)\Big) \le \dim\Big(C(F_k)/C(F_j)\Big) - (k-j). \tag{43}$$

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 (43) 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 $\rho = \mathrm{rank}(v_i)$.

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

$$\sum_{he=v_{j+1}} d_h\overline{w_e} \in B(F_i). \tag{44}$$

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 (44) 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

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

Hence

$$\dim\Big(A(F_{j+1}) \;\big/\; \big(A(F_j) + B(F_i)\big)\Big) \le \dim\Big(C(F_{j+1})/C(F_j)\Big) - 1. \tag{45}$$

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

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

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

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

Hence applying equation (45) with m replacing j and m over the range $j, j+1, \ldots, k-1$ yields the lemma.

$\square$

Lemma 8.11 *Let v be a moseying sequence of length s for G'. Then for non-negative integers $i \le j \le s$ we have*

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

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

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

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 (41) and the fact that $\mathrm{rank}(v_{j+1}) \geq \mathrm{rank}(v_m)$. Hence the lemma follows.

$\square$

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

$$B(F_{k_1}) \subset A(F_{k_1}), \tag{46}$$

then we are done, for then $F = F_{k_1}$ satisfies equation (32), in view of equation (38) with $i = k_1$. In this case we end our phases, and Lemma 8.1 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 (38), 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

$$B(F_{k_1}) \subset A(F_{k_1}) + B(I_{k_2}), \tag{47}$$

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

$$A(F_{k_1}) \subset A(F_{k_1}) + B(I_{k_1+1}) \subset A(F_{k_1}) + B(I_{k_1+2}) \subset \cdots \tag{48}$$

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

$$k_2 - k_1 \leq \dim\Bigl(\mathcal{F}(V)\Bigr),$$

and since $k_1 \leq \dim(\mathcal{F}(V))$ (see equation (39) and the discussion below it), we have

$$k_2 \leq \dim\Bigl(\mathcal{F}(V)\Bigr) + \dim\Bigl(\mathcal{F}(E)\Bigr). \tag{49}$$

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 (47). Now we combine a number of inequalities from Subsection 8.4 to prove a sequel to equation (38).

First, Lemma 8.9 with $j = v_{2k}$ and $i = 0$ (for which the lemma is still valid) shows that

$$\mathrm{Cap}(U_{k_2}) = k_2 - \sum_{m=1}^{k_2} \mathrm{sod}(v_m) \tag{50}$$

(note that U_{k_2} has no cycles, using Lemma 8.5). Second, Lemma 8.10 with $k = k_2$ and $i = j = k_1$ yields

$$\dim\Big(A(F_{k_2})/(A(F_{k_1}) + B(F_{k_1}))\Big) \leq \dim\Big(C(F_{k_2})/C(F_{k_1})\Big) - (k_2 - k_1). \tag{51}$$

Third, we have $I_{k_1} = \emptyset$ since $v_1, \ldots, v_{k_1}$ are all of rank ρ. Hence Lemma 8.11 with $j = k_2$ and $i = k_1$ gives

$$\dim\Big(B(I_{k_2})\Big) = \dim\Big(B(I_{k_2})/B(I_{k_1})\Big) \leq \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i). \tag{52}$$

We have now established three inequalities in equations (50), (51), and (52). We now establish a simple inequality to describe the end of the first B-phase.

Equations (52) and (38) with $i = k_1$ imply that

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

and in view of equation (47) this implies that

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

Equation (51) added to this gives

$$\dim\Big(A(F_{k_2}) + B(F_{k_1})\Big)$$

$$\leq \dim\Big(C(F_{k_2})\Big) - k_2 + \sum_{i=k_1+1}^{k_2} \mathrm{sod}(v_i)$$

$$= \dim\Big(C(F_{k_2})\Big) - k_2 + \sum_{i=1}^{k_2} \mathrm{sod}(v_i)$$

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

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

Then using equation (50) we get

$$\dim\Big(A(F_{k_2}) + B(F_{k_1})\Big) \leq \dim\Big(C(F_{k_2})\Big) - \mathrm{Cap}(U_{k_2}). \tag{53}$$

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

If

$$B(F_{k_2}) \subset A(F_{k_2}) + B(F_{k_1}), \tag{54}$$

then our phases are over and we easily establish Lemma 8.1: indeed, we have

$$\dim\Big(A(F_{k_2}) + B(F_{k_2})\Big) = \dim\Big(A(F_{k_2}) + B(F_{k_1})\Big) \leq \dim\Big(C(F_{k_2})\Big) - 1$$

since $\mathrm{Cap}(U_{k_2}) \geq 1$ (indeed, $h_0(U_{k_2}) \geq 1$ and the overdegree is non-negative). Hence we have established equation (32) 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

$$\forall e \in E_{G[\mathbb{Z}_{\geq 0}]} \text{ of rank } \rho + i - 1, \quad \overline{w_e} \in C(F_{k_{2i-1}}) \tag{55}$$

(for $j \geq k_{2i-1} + 1$ 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

$$B(F_{k_{2i-1}}) \subset A(F_{k_{2i-1}}) + B(I_{k_{2i}}); \tag{56}$$

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 (56) 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 $j \geq 1$ when

$$B(F_{k_{2j-1}}) \subset A(F_{k_{2j-1}}) + B(F_{k_{2j-3}}) \tag{57}$$

(with $k_{-1} = 0$ and so $F_{k_{-1}} = \emptyset$ for the case $j = 1$), which restricts to equation (46) for $j = 1$; the phases end at the j-th B-phase for $j \geq 1$ when

$$B(F_{k_{2j}}) \subset A(F_{k_{2j}}) + B(F_{k_{2j-1}}), \tag{58}$$

which restricts to equation (54) 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 (32). We already have all the main inequalities needed to prove this, and just need to apply them to the phases beyond the second phase.

8.6 End of the Proof of Lemma 8.1

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

$$\dim\Bigl(A(F_{k_{2i-1}}) + B(F_{k_{2i-3}})\Bigr) \leq \dim\Bigl(C(F_{k_{2i-1}})\Bigr) - \mathrm{Cap}(U_{k_{2i-2}}) - (k_{2i-1} - k_{2i-2}) \tag{59}$$

(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

$$\dim\Bigl(A(F_{k_{2i}}) + B(F_{k_{2i-1}})\Bigr) \leq \dim\Bigl(C(F_{k_{2i}})\Bigr) - \mathrm{Cap}(U_{k_{2i}}). \tag{60}$$

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 8.5,

$$k_i \leq \dim\Bigl(\mathcal{F}(V)\Bigr) + \dim\Bigl(\mathcal{F}(E)\Bigr), \tag{61}$$

and U_{k_i} contains no cycles, using the hypotheses of Lemma 8.1.

Let us also note that the phases eventually end. Indeed, if $k_{2j} = k_{2j-1}$, then according to equation (58) 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 (61), the total number of phases is less than

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

Equation (60) has been established for $i = 1$ in equation (53). So let us first show that equation (60) implies equation (59) with i replaced by $i+1$.

So assume equation (60) for some $i \geq 1$. By Lemma 8.10, 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 (60) 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 (59), with i replaced by $i+1$.

Finally assume equation (59) for some value of $i \geq 1$; we shall conclude that equation (60) holds for the same value of i. By Lemma 8.11 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

$$\dim\Bigl(\Bigl(A(F_{k_{2i-1}}) + B(I_{k_{2i}})\Bigr) \Big/ \Bigl(A(F_{k_{2i-1}}) + B(I_{k_{2i-2}})\Bigr)\Bigr) \leq \sum_{m=k_{2i-2}+1}^{k_{2i}} \mathrm{sod}(v_m). \tag{62}$$

In view of equation (56), and since $I_{k_{2i}} \subset F_{k_{2i-1}}$, we have

$$A(F_{k_{2i-1}}) + B(F_{k_{2i-1}}) = A(F_{k_{2i-1}}) + B(I_{k_{2i}}); \tag{63}$$

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

$$A(F_{k_{2i-1}}) + B(F_{k_{2i-3}}) = A(F_{k_{2i-1}}) + B(I_{k_{2i-2}}) \tag{64}$$

Given equations (63) and (64), equation (62) can be rewritten as

$$\dim\Big(\big(A(F_{k_{2i-1}})+B(F_{k_{2i-1}})\big) \;/\; \big(A(F_{k_{2i-1}})+B(F_{k_{2i-3}})\big)\Big) \leq \sum_{m=k_{2i-2}+1}^{k_{2i}} \operatorname{sod}(v_m). \tag{65}$$

Adding this to equation (59) gives

$$\begin{aligned} &\dim\Big(A(F_{k_{2i-1}}) + B(F_{k_{2i-1}})\Big) \\ &\leq \dim\Big(C(F_{k_{2i-1}})\Big) - \operatorname{Cap}(U_{k_{2i-2}}) - (k_{2i-1} - k_{2i-2}) + \sum_{m=k_{2i-2}+1}^{k_{2i}} \operatorname{sod}(v_m) \\ &= \dim\Big(C(F_{k_{2i-1}})\Big) - \operatorname{Cap}(U_{k_{2i}}) + (k_{2i} - k_{2i-1}) \end{aligned}$$

in view of Lemma 8.9 with i, j respectively set to k_{2i-2}, k_{2i}. Adding this to Lemma 8.10 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 (60).

At this point we have established equations (59) and (60), and the fact that the phases eventually end. Now we claim that Lemma 8.1 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 (60) 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 (32). 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 (59) gives

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

since

$$\mathrm{Cap}(U_{k_{2j-2}})+(k_{2j-1}-k_{2j-2})\ge 1$$

(for $j=1$ this follows since $k_1>0$, and for $j\ge 2$ this follows since $U_{k_{2j-2}}$ is nonempty). Hence, similarly, $F=F_r$ with $r=k_{2j-1}$ satisfies equation (32).

□

8.7 Proof of Theorem 2.10

Proof of Theorem 2.10, first proof. First we will verify Theorem 2.10 in some special cases.

Lemma 8.1 establishes Theorem 2.10 in the case where $\mathrm{m.e.}(\mathcal{F})=0$.

Definition 8.12 *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 2.10 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 8.13 *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 8.14 *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

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

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

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

and particular elements

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

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 $W' \subset \Gamma_{\text{ht}}(U')$. Since U', W' are compartmentalized, it follows that

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

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

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

which is a contradiction.

$\square$

Returning to the proof of Theorem 2.10, we claim that it suffices to establish it for tight sheaves; indeed, consider an arbitrary sheaf, $\mathcal{F}$, and apply Lemma 8.14 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 8.1,

$$h_1^{\mathrm{twist}}\Bigl(\phi^*(\mathcal{F}/\mathcal{F}')\Bigr) = 0$$

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

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

In this case we get in the long exact sequence beginning

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

amounts to

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

or

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

Hence to prove Theorem 2.10 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 8.15 *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

$$\mathrm{excess}(\mathcal{F}'', U) \leq \mathrm{excess}(\mathcal{F}, U) \leq \mathrm{m.e.}(\mathcal{F}) - 1.$$

$\square$

We now prove Theorem 2.10 by induction upon $\mathrm{m.e.}(\mathcal{F})$. The base case, $\mathrm{m.e.}(\mathcal{F}) = 0$, was established in Lemma 8.1. Assume that we have established that Theorem 2.10 holds whenever $\mathrm{m.e.}(\mathcal{F}) \leq k$ for some integer $k \geq 0$. We wish to prove Theorem 2.10 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 8.15. Then Theorem 2.10 holds for $\mathcal{F}''$, since $\mathcal{F}''$ has maximum excess k; so for $\phi\colon G' \to G$ of girth greater than

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

we have

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') = \mathrm{m.e.}(\phi^*\mathcal{F}'') = \deg(\phi)k. \tag{66}$$

Since, by the construction of $\mathcal{F}''$ in Lemma 8.15, 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

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}) = h_1^{\mathrm{twist}}(\phi^*\mathcal{F}'') + h_1^{\mathrm{twist}}(\phi^*(\mathcal{F}/\mathcal{F}'')). \tag{67}$$

But according to Lemma 8.15, $\mathcal{F}/\mathcal{F}''$ is edge supported, and we therefore know that Theorem 2.10 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 (66) and (67) shows that

$$h_1^{\mathrm{twist}}(\phi^*\mathcal{F}) = \deg(\phi)(k+1) = \mathrm{m.e.}(\phi^*\mathcal{F})$$

This establishes Theorem 2.10 for all tight $\mathcal{F}$ with $\mathrm{m.e.}(\mathcal{F}) = k+1$.

Hence, by induction on the maximum excess of $\mathcal{F}$, Theorem 2.10 holds for all sheaves, $\mathcal{F}$, on G.

□

9 Concluding Remarks

In this section we conclude with a few remarks about the results in this paper 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 2.10, which implies that the maximum excess is a first quasi-Betti number. But part of the proof of Theorem 2.10, namely Subsection 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

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

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 5.6.

Notice that if G is an undirected graph, all the discussion in the paper goes through. Either one can orient each edge and use the notation in this paper, or just rewrite the notation in this paper 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 (17) that appears just below equation (18)). 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 $\mathcal{F}(e)$.

References

[AL02] Alon Amit and Nathan Linial. Random graph coverings I: General theory and graph connectivity. *Combinatorica*, 22(1):1–18, 2002.

[AM69] M. F. Atiyah and I. G. Macdonald. *Introduction to commutative algebra*. Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969.

[Arz00] G. N. Arzhantseva. A property of subgroups of infinite index in a free group. *Proc. Amer. Math. Soc.*, 128(11):3205–3210, 2000.

[Ati76] M. F. Atiyah. Elliptic operators, discrete groups and von Neumann algebras. In *Colloque "Analyse et Topologie" en l'Honneur de Henri Cartan (Orsay, 1974)*, pages 43–72. Astérisque, No. 32–33. Soc. Math. France, Paris, 1976.

[Bur71] Robert G. Burns. On the intersection of finitely generated subgroups of a free group. *Math. Z.*, 119:121–130, 1971.

[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.

[DF01] Warren Dicks and Edward Formanek. The rank three case of the Hanna Neumann conjecture. *J. Group Theory*, 4(2):113–151, 2001.

[Dic94] Warren Dicks. Equivalence of the strengthened Hanna Neumann conjecture and the amalgamated graph conjecture. *Invent. Math.*, 117(3):373–389, 1994.

[Eve08] Brent Everitt. Graphs, free groups and the Hanna Neumann conjecture. *J. Group Theory*, 11(6):885–899, 2008.

[FMT06] Joel Friedman, Ram Murty, and Jean-Pierre Tillich. Spectral estimates for abelian cayley graphs. *J. Comb. Theory Ser. B*, 96(1):111–121, 2006.

[Fri93] Joel Friedman. Some geometric aspects of graphs and their eigenfunctions. *Duke Math. J.*, 69(3):487–525, 1993.

[Fri03] Joel Friedman. Relative expanders or weakly relatively Ramanujan graphs. *Duke Math. J.*, 118(1):19–35, 2003.

[Fri05] Joel Friedman. Cohomology of grothendieck topologies and lower bounds in boolean complexity. 2005. `http://www.math.ubc.ca/~jf`, also `http://arxiv.org/abs/cs/0512008`, to appear.

[Fri06] Joel Friedman. Cohomology of grothendieck topologies and lower bounds in boolean complexity ii. 2006. `http://www.math.ubc.ca/~jf`, also `http://arxiv.org/abs/cs/0604024`, to appear.

[Fri07] Joel Friedman. Linear transformations in boolean complexity theory. In *CiE '07: Proceedings of the 3rd conference on Computability in Europe*, pages 307–315, Berlin, Heidelberg, 2007. Springer-Verlag.

[Fri08] Joel Friedman. A proof of Alon's second eigenvalue conjecture and related problems. *Mem. Amer. Math. Soc.*, 195(910):viii+100, 2008.

[FT05] Joel Friedman and Jean-Pierre Tillich. Generalized Alon–Boppana theorems and error-correcting codes. *SIAM J. Discret. Math.*, 19:700–718, July 2005.

[Ger83] S. M. Gersten. Intersections of finitely generated subgroups of free groups and resolutions of graphs. *Invent. Math.*, 71(3):567–591, 1983.

[Gro77] Jonathan L. Gross. Every connected regular graph of even degree is a Schreier coset graph. *J. Combinatorial Theory Ser. B*, 22(3):227–232, 1977.

[Har77] Robin Hartshorne. *Algebraic geometry*. Springer-Verlag, New York, 1977. Graduate Texts in Mathematics, No. 52.

[HS97] P. J. Hilton and U. Stammbach. *A course in homological algebra*, volume 4 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1997.

[Imr77a] Wilfried Imrich. On finitely generated subgroups of free groups. *Arch. Math. (Basel)*, 28(1):21–24, 1977.

[Imr77b] Wilfried Imrich. Subgroup theorems and graphs. In *Combinatorial mathematics, V (Proc. Fifth Austral. Conf., Roy. Melbourne Inst. Tech., Melbourne, 1976)*, pages 1–27. Lecture Notes in Math., Vol. 622. Springer, Berlin, 1977.

[Iva99] S. V. Ivanov. On the intersection of finitely generated subgroups in free products of groups. *Internat. J. Algebra Comput.*, 9(5):521–528, 1999.

[Iva01] S. V. Ivanov. Intersecting free subgroups in free products of groups. *Internat. J. Algebra Comput.*, 11(3):281–290, 2001.

[JKM03] Toshiaki Jitsukawa, Bilal Khan, and Alexei G. Myasnikov. On the Hanna Neumann conjecture, 2003.

[Kha02] Bilal Khan. Positively generated subgroups of free groups and the Hanna Neumann conjecture. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 155–170. Amer. Math. Soc., Providence, RI, 2002.

[Lan02] Serge Lang. *Algebra*, volume 211 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, third edition, 2002.

[Lüc02] Wolfgang Lück. *L^2-invariants: theory and applications to geometry and K-theory*, volume 44 of *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]. Springer-Verlag, Berlin, 2002.

[Min10] Igor Mineyev. The topology and analysis of the Hanna Neumann Conjecture. March 2010. Preprint. Available at `http://www.math.uiuc.edu/~mineyev/math/art/shnc.pdf`.

[MW02] J. Meakin and P. Weil. Subgroups of free groups: a contribution to the Hanna Neumann conjecture. In *Proceedings of the Conference on Geometric and Combinatorial Group Theory, Part I (Haifa, 2000)*, volume 94, pages 33–43, 2002.

[Neu90] Walter D. Neumann. On intersections of finitely generated subgroups of free groups. In *Groups—Canberra 1989*, volume 1456 of *Lecture Notes in Math.*, pages 161–170. Springer, Berlin, 1990.

[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(1):133–137, 1983.

[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.

[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.

[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.

[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.

[ST96] H. M. Stark and A. A. Terras. Zeta functions of finite graphs and coverings. *Adv. Math.*, 121(1):124–165, 1996.

[Sta83] John R. Stallings. Topology of finite graphs. *Invent. Math.*, 71(3):551–565, 1983.

[Tar92] Gábor Tardos. On the intersection of subgroups of a free group. *Invent. Math.*, 108(1):29–36, 1992.

[Tar96] Gábor Tardos. Towards the Hanna Neumann conjecture using Dicks' method. *Invent. Math.*, 123(1):95–104, 1996.