Sheaves on Graphs and a Proof of the Hanna Neumann Conjecture

Joel Friedman*

April 30, 2011

Abstract

The main goal of this paper is to prove the Hanna Neumann Conjecture; in fact, we prove a strengthened form of the conjecture. We study these conjectures using what we have called "sheaves on graphs" in [Fri]. We show that both conjectures are implied by the vanishing of a certain invariant, the "maximum excess," of certain sheaves that we call ρ-kernels.

Our approach involves "graph Galois theory," an analogue of classical Galois theory in the graph setting. We use it to construct the ρ-kernels. We use the symmetry in Galois theory to argue that if the Strengthened Hanna Neumann Conjecture is false, then the maximum excess of "most of" these ρ-kernels must be large. We then give an inductive argument to show that this is impossible.

1 Introduction

Howson, in [How54], showed that if $\mathcal{K}, \mathcal{L}$ are nontrivial, finitely generated subgroups of a free group, $\mathcal{F}$, then $\mathcal{K} \cap \mathcal{L}$ is finitely generated, and moreover that

$$\mathrm{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \le 2\,\mathrm{rank}(\mathcal{K})\,\mathrm{rank}(\mathcal{L}) - \mathrm{rank}(\mathcal{K}) - \mathrm{rank}(\mathcal{L}). \qquad (1)$$

*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.

Hanna Neumann, in [Neu56, Neu57] improved this bound to what is now called the Hanna Neumann Bound,

$$\mathrm{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \le 2 \Bigl(\mathrm{rank}(\mathcal{K}) - 1\Bigr) \Bigl(\mathrm{rank}(\mathcal{L}) - 1\Bigr); \tag{2}$$

furthermore, she conjectured that one can remove the factor of 2 in this bound, i.e., that

$$\mathrm{rank}(\mathcal{K} \cap \mathcal{L}) - 1 \le \Bigl(\mathrm{rank}(\mathcal{K}) - 1\Bigr) \Bigl(\mathrm{rank}(\mathcal{L}) - 1\Bigr); \tag{3}$$

this conjecture is now known as the Hanna Neumann Conjecture (or HNC). One goal of this paper is to prove the HNC. Moreover, we shall prove a strengthened form of the conjecture, first studied by Walter Neumann in [Neu90], known as the Strengthened Hanna Neumann Conjecture (or SHNC); we will state the strengthened conjecture in the next section.

Theorem 1.1 *The Hanna Neumann Conjecture and the Strengthened Hanna Neumann Conjecture hold.*

These conjectures have received considerable attention (see [Bur71, Imr77b, Imr77a, Ser83, Ger83, Sta83, Neu90, Tar92, Dic94, Tar96, Iva99, Arz00, DF01, Iva01, Kha02, MW02, JKM03, Neu07, Eve08, Min10]). However, our proof uses very different methods from the previous papers.

The main new idea in our approach to the SHNC is the notion of "sheaves on a graph," introduced in [Fri], and their homology groups and related invariants. The SHNC has a well-known reformulation in terms of graph theory, and we will prove this reformulation by proving a stronger theorem about certain sheaves that we call ρ-kernels. We anticipate that the ideas in this paper and [Fri] will have applications to other areas in graph theory, for various reasons; for one thing, sheaves vastly generalize algebraic graph theory (e.g., adjacency matrices, incidence matrices, etc.); for another, sheaves can express new ideas not found in classical graph theory; in particular, sheaf theory sometimes gives "additional morphisms between graphs" that don't exist in classical graph theory, and such morphisms are crucial to the construction of the ρ-kernels in this paper.

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. More precisely, we work with sheaves of finite dimensional vector spaces on a finite Grothendieck topology, $\mathrm{Top}(G)$, that we associate to a graph, G. In [Fri05] we studied such sheaves over a large class of finite Grothendieck topologies, motivated by complexity theory. While graphs represent a very special case of such finite Grothendieck topologies, the results in [Fri] give stronger results for the special case of graphs, and study new invariants. The HNC and SHNC were our motivation for most of the results in [Fri].

Our proof also makes use of an analogue in graph theory of the classical Galois theory for number fields. This "graph Galois theory" was studied in [Fri93, ST96], with parts of it known earlier in [Gro77]. This theory is crucial in simplifying the SHNC to a problem concerning certain the sheaves we call ρ-kernels.

We use a graph theoretic reformulation of the SHNC that involves a graph theoretic invariant, $\rho(G)$, of a graph, G, that we call its *reduced cyclicity*. In order to use sheaf theory to study the reduced cyclicity, ρ, one has to extend the definition of ρ to sheaves. A key idea in [Fri] is to realize that ρ is a certain limit of scaled first Betti numbers; this serves both to (1) define ρ on all sheaves, and (2) indicate that ρ should satisfy the same type of relations that all first Betti numbers do. What is remarkable is that this extension of ρ to sheaves, i.e., as a scaled limit, is equal to a what we call the "maximum excess" of a sheaf, that can be defined very simply and without any reference to homology theory. This simple definition makes the maximum excess look like quantities defined in matching theory or expanding graphs. Furthermore the maximum excess has other intriguing properties, such as it scales under pullbacks of covering maps, and it appears difficult to compute in general.

The graph theoretic reformulation of the SHNC amounts to an inequality between the reduced cyclicity of three graphs. We will formulate a stronger, simple conjecture about sheaves, roughly as follows. Relations between (co)homology groups of related spaces or sheaves are often described by ("long") exact sequences. Recall that any time a vector space in an exact sequence vanishes, the vector space to its right injects into its right neighbour; so a zero Betti number (i.e., dimension of the vector space) in an exact sequence, implies the next two dimensions on the right are nondecreasing. So our idea is to fit the graphs in the SHNC into an appropriate short exact sequence, and deduce the desired inequality by showing that the appropriate homology group vanishes. The maximum excess is only a "limiting" homology group dimension, but, as a limit, it satisfies the same inequalities that Betti numbers do. Using graph Galois theory, we are almost immediately lead to a short exact sequence appropriate for this approach to the SHNC;

the short exact sequence involves all the graphs involved in the SHNC, plus a sheaf we are forced to construct, which are the ρ-kernels. Hence the ρ-kernels, created by necessity in forming certain short exact sequences, are sheaves whose vanishing maximum excess implies the SHNC.

It turns out that some ρ-kernels have nonvanishing maximum excess. However, any graph, L, of interest to us in the SHNC, will have a family of associated ρ-kernels, and we will prove that the generic maximum excess in this family is zero, for each L. To do this we shall use Galois theory and symmetry to argue that if this generic maximum excess does not vanish then it is large, i.e., a multiple of the order of the associated Galois group. Then we will give an inductive argument, showing that if the generic maximum excess of ρ-kernels for L is positive, then the same is true when we remove some edge from L, provided that L has positive reduced cyclicity. The base case of the induction, when L has vanishing reduced cyclicity, is trivial to establish. Hence each L of interest in the SHNC has a ρ-kernel of vanishing maximum excess, and this establishes in SHNC.

We emphasize that our proof also uses the main theorem in [Fri], that shows that the "maximum excess," which is crucial to this paper, is actually a limiting "twisted" Betti number, and therefore satisfies certain inequalities arising from the long exact sequence(s) associated to any short exact sequence. However, in this paper we use only some of these inequalities, and it is quite possible that one can prove the subset of inequalities we need for this paper more simply and directly, without going through the homological approach of [Fri]; we discuss this a bit further in the conclusion of this paper.

So the proof we give of the SHNC has no direct reference to homology theory or Betti numbers; for this paper, we need only a few simple definitions—a sheaf on a graph and its "Euler characteristic" and "maximum excess"—and then we quote some results from [Fri] whose statement requires no homology theory. However, the proof of these results in [Fri] uses a lot of homology theory, and homology theory does shed more light on the results of this paper. In the concluding section of this paper we will give some optional remarks on how homology theory can help our understanding of the maximum excess and of our proof of the SHNC.

This paper and [Fri] show that the SHNC is not merely an attempt to improve an inequality by a factor of two; our study of the SHNC has lead to new ideas in sheaves on graphs that can be applied to graph theory, some of which appear in [Fri]. This came as a surprise to us at first, although it is perhaps less surprising in retrospect, for a number of reasons. First, the

HNC and SHNC seem to describe a fairly fundamental question in group theory (of how rank behaves under intersection). Second, the SHNC can be viewed as a graph theory question involving the reduced cyclicity, which is an interesting graph invariant (e.g. it scales under covering maps). Third, the SHNC, viewed in terms of the Galois theory of graphs (as in [Fri93]), has a simple homological explanation, as implied by the vanishing of a limit homology group. The vanishing of (co)homology groups has a vast literature and importance; the SHNC is an interesting and seemingly difficult result in the family of homology group vanishing theorems. Fourth, Lior Silberman has pointed out to us that the reduced cyclicity is the discrete analogue of L^2 Betti numbers; the L^2 Betti number was defined first by Atiyah ([Ati76]), and has been the subject of much surrounding the "Atiyah conjecture" (see [Lüc02]). Mineyev's article, [Min10], also makes a connection between the SHNC and L^2 Betti numbers.

At this point we can give more motivation for the use sheaf theory in this paper, i.e., why we do not just use graphs and their homology. Our reformulation of the SHNC begins by searching for a morphism involving the graphs of interest to the SHNC. In order for this morphism to exist, to be surjective, and to have a kernel, we must work with more general objects than graphs. In many topological situations, the topological spaces are sufficiently "robust" that one does not have to generalize the objects. However, in non-Hausdorff spaces, such as graphs or those in algebraic geometry, many geometric notions, such as "connect two points with a path," "form a cone," etc., don't make sense or are very awkward to implement. So for graphs we use sheaf theory, which is a simple (co)homology theory that is adapted to our spaces, but general and expressive enough for appropriate surjections and kernels to exist.

The rest of this paper is organized as follows. In Section 2 we describe the SHNC and previous work on the HNC and SHNC, including some resolved special cases of the SHNC. In Section 3 we give a common graph theoretic reformulation of the SHNC. In Section 4 we give the definition of a sheaf and some of its associated invariants, and state some theorems from [Fri] that we will use in our proof of the SHNC. In Section 5 we describe applications of "graph Galois theory" to simplifying the SHNC in a way that leads to the construction of ρ-kernels. In Section 6 we construct ρ-kernels and prove that if their maximum excesses vanish then the SHNC holds; we also describe what we call "k-th power kernels," which generalize ρ-kernels and which will be necessary to prove our main theorems about the generic maximum excess of ρ-kernels. In Section 7 we use symmetry to prove that the generic maximum excess of a certain type of k-th power kernel is divisible by the order of a group associated to the class of kernels. In Section 8 we prove some comparison theorems about how the maximum excesses of different classes of k-th power kernels compare; the main theorem that we prove shows that if ρ-kernels associated to a graph, L, have positive generic maximum excess, then the same is true of a graph L' that consist of L with one edge discarded, provided that $\rho(L') = \rho(L) - 1$. In Section 9 we briefly combine a number of theorems of previous sections to argue that if the SHNC does not hold, then for some L we have the class of ρ-kernels associated to L have positive generic maximum excess, which by the results of Section 8 means that the same is true for some L with $\rho(L) = 0$, which we easily show is impossible. This establishes the SHNC. In Section 10 we make some concluding remarks. In Appendix A we show that vanishing maximum excess of enough ρ-kernels implies the SHNC, but we do so just using elementary graph theory; this shows that a lot of the sheaf theory can be translated into direct graph theoretic terms; it also shows that a simple sheaf theoretic calculation may translate into a much longer graph theoretic calculation.

We wish to thank the following people: Laurant Bartholdi, for conversations and introducing us to this conjecture; 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 The Strengthened Hanna Neumann Conjecture

In this section we state the SHNC and comment on previous work, including some established special cases of the SHNC.

Walter Neumann, in [Neu90], showed that the Hanna Neumann Bound, i.e., equation (2), could be strengthened to

$$\sigma(\mathcal{K}, \mathcal{L}) \leq 2 \ \mathrm{rk}_{-1}(\mathcal{K}) \ \mathrm{rk}_{-1}(\mathcal{L}),$$

where $\mathrm{rk}_n(\mathcal{G})$ denotes $\max(\mathrm{rank}(\mathcal{G})+n,0)$, and where

$$\sigma(\mathcal{K},\mathcal{L}) = \sum_{\mathcal{K}x\mathcal{F}\in\mathcal{K}\backslash\mathcal{F}/\mathcal{L}} \mathrm{rk}_{-1}(\mathcal{K}\cap x^{-1}\mathcal{L}x),$$

the summation being over the double coset, $\mathcal{K}\backslash\mathcal{F}/\mathcal{L}$, representatives, x; taking x be the identity in the summation shows that

$$\mathrm{rk}_{-1}(\mathcal{K}\cap\mathcal{L}) \le \sigma(\mathcal{K},\mathcal{L}),$$

so that Walter Neumann's above bound strengthens the Hanna Neumann Bound. Walter Neumann further formulated the conjecture that

$$\sigma(\mathcal{K},\mathcal{L}) \le \mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}), \tag{4}$$

now known as the Strengthened Hanna Neumann Conjecture (or SHNC).

The HNC and SHNC have been studied in many papers, including [Bur71, Imr77b, Imr77a, Ser83, Ger83, Sta83, Neu90, Tar92, Dic94, Tar96, Iva99, Arz00, DF01, Iva01, Kha02, MW02, JKM03, Neu07, Eve08, Min10]. For the rest of this section we review previous work on these conjectures.

One collection of results on the problem involves general bounds on $\sigma(\mathcal{K},\mathcal{L})$ or $\mathrm{rk}_1(\mathcal{K}\cap\mathcal{L})$. It turns out that all general bounds we know for the HNC, i.e., on $\mathrm{rk}_1(\mathcal{K}\cap\mathcal{L})$, also are known to hold for $\sigma(\mathcal{K},\mathcal{L})$. Also, all bounds we know are of the form

$$\sigma(\mathcal{K},\mathcal{L}) \le 2\ \mathrm{rank}(\mathcal{K})\ \mathrm{rank}(\mathcal{L}) + c_1\ \mathrm{rank}(\mathcal{K}) + c_2\ \mathrm{rank}(\mathcal{L}) + c_3$$

for ranks $\mathcal{K},\mathcal{L}$ sufficiently large, where c_1,c_2,c_3 are constants depending on the bound; thus all improvements of Howson's original bound are in the lower order terms, i.e., in the c_i's. The improved bounds on $\sigma(\mathcal{K},\mathcal{L})$ after [How54, Neu56, Neu57] include the bound

$$2\ \mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) - \min(\mathrm{rk}_{-1}(\mathcal{K}),\mathrm{rk}_{-1}(\mathcal{L}))$$

of Burns in [Bur71][1], the bound

$$\mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) + \max(\mathrm{rk}_{-2}(\mathcal{K})\ \mathrm{rk}_{-2}(\mathcal{L}) - 1, 0)$$

[1]Bounds appearing before [Neu90] are stated as bounds on $\mathrm{rk}_{-1}(\mathcal{K}\cap\mathcal{L})$, but actually give bounds on $\sigma(\mathcal{K},\mathcal{L})$ as well.

of Tardos [Tar92, Tar96], and, what is the best bound prior to ours,

$$\mathrm{rk}_{-1}(\mathcal{K})\ \mathrm{rk}_{-1}(\mathcal{L}) + \mathrm{rk}_{-3}(\mathcal{K})\ \mathrm{rk}_{-3}(\mathcal{L}) \tag{5}$$

of Dicks and Formanek in [DF01].

Another collection of results concerns special cases of the HNC and SHNC that are resolved. To be precise, say that the "HNC holds for $(\mathcal{K},\mathcal{L})$" if equation (3) holds, and say that $\mathcal{K}$ is *universal for the HNC* if for any $\mathcal{L}$, the HNC holds for $(\mathcal{K},\mathcal{L})$. Similarly for the SHNC and equation (4). Similar to before, all results we know that resolve special cases of the HNC also resolve those cases of the SHNC. Note that any finitely generated free group, $\mathcal{F}$, is a subgroup of $\mathcal{F}_2$, the free group on two generators, so we are free to assume that $\mathcal{F}=\mathcal{F}_2$ in the HNC. Here are some results on special cases of the SHNC that are easy to describe in group theoretic terms:

1. $\mathcal{K}$ is universal for the SHNC if it is of rank at most three ([DF01]), in view of equation (5), with rank two settled earlier by Tardos ([Tar92]);

2. $\mathcal{K}$ is universal for the SHNC if it is positively generated (see [Kha02, MW02, Neu07]);

3. $\mathcal{K}$ is universal for the SHNC for "most" $\mathcal{K}$ (see [Neu90, JKM03]);

4. the SHNC holds either for $(\mathcal{K},\mathcal{L})$ or for $(\mathcal{K},\mathcal{L}')$ for any $\mathcal{K},\mathcal{L}$ that are subgroups of $\mathcal{F}_2$, where $\mathcal{L}'$ is obtained from $\mathcal{L}$ by the map taking each generator of $\mathcal{F}_2$ to its inverse (see [JKM03]).

The result of item (3) on "most" groups, of Walter Neumann ([Neu90]), and some additional results on the SHNC, such as Corollary 3.2 of [MW02], are easier to describe using a graph theoretic formulation of the SHNC that we give in the next section. It is also known that the SHNC is related to the coherence problem in one-relator groups ([Wis05]).

3 Graph Theoretic Formulation of the SHNC

The goal of this section is to describe an equivalent formulation of the HNC and SHNC in graph theoretic terms involving fibre products; this formulation is implicit in [How54], but more explicit in [Imr77b, Imr77a, Ger83, Sta83, Neu90] and other references in [Dic94]. There is another equivalent reformulation of the SHNC by Dicks in [Dic94], known as the "amalgamated graph conjecture," which we do not discuss here.

We must 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 graph is assumed to be finite, i.e., the vertex and edge sets are finite.

Recall that a morphism of digraphs, $\nu\colon K \to G$, is a pair $\nu = (\nu_V, \nu_E)$ of maps $\nu_V\colon V_K \to V_G$ and $\nu_E\colon E_K \to E_G$ such that $t_G\nu_E = \nu_V t_K$ and $h_G\nu_E = \nu_V h_K$. We can usually drop the subscripts from ν_V and ν_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 $\nu_1\colon G_1 \to G$ and $\nu_2\colon G_2 \to G$ has

$$V_K = \{(v_1, v_2) \mid v_i \in V_{G_i},\ \nu_1 v_1 = \nu_2 v_2\},$$

$$E_K = \{(e_1, e_2) \mid e_i \in E_{G_i},\ \nu_1 e_1 = \nu_2 e_2\},$$
$$t_K = (t_{G_1}, t_{G_2}), \quad \text{and} \quad h_K = (h_{G_1}, h_{G_2}).$$

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 set theoretic projections on V_K and E_K. We say that $\nu\colon K \to G$ is a covering map (respectively, étale[2]) 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) to those of $\nu(v)$, and a bijection (respectively, injection) of outgoing edges of v to those of $\nu(v)$. We say that $\nu\colon K \to G$ is of degree d or is d-to-1 if the number of preimages of any vertex or edge of G equals d; if so, we write $[K\colon G]$ for the degree, d; if G is connected and ν is a covering map, then ν is d-to-1 for some d.

By a *bicoloured digraph*, or simply a *bigraph*, we mean a directed graph, G, such that each edge is coloured (or labelled) either "1" or "2." It is also equivalent to giving a directed graph homomorphism $\nu\colon G \to B_2$, where B_2 is the graph with one vertex and two self-loops, one coloured "1" and the other "2." If, moreover, ν is étale, we call ν or (somewhat abusively) G an

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

étale bigraph, which means that G is a bigraph such that no vertex has two incident edges, both incoming or both outgoing, of the same colour.

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

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

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

$$\rho(G) = \sum_{X \in \mathrm{conn}(G)} \max\Bigl(0, h_1(X) - 1\Bigr), \tag{6}$$

which we call the *reduced cyclicity of* G, and

$$\rho'(G) = \max_{X \in \mathrm{conn}(G)} \biggl(\max\Bigl(0, h_1(X) - 1\Bigr)\biggr).$$

The HNC is equivalent to

$$\rho'(K \times_{B_2} L) \le \rho(K)\rho(L) \tag{7}$$

for all étale bigraphs K and L; the SHNC is equivalent to

$$\rho(K \times_{B_2} L) \le \rho(K)\rho(L) \tag{8}$$

for all étale bigraphs K and L (see [How54, Imr77b, Imr77a, Ger83, Sta83, Neu90, Dic94]). We shall work with this form of the SHNC. Again, we say the HNC or SHNC, respectively, holds for a pair of étale bigraphs, (K, L), if equation (7) or (8), respectively, holds; and again, we say that K is *universal* for the HNC or SHNC, respectively, if for any L the same conjecture holds for (K, L).

Let us briefly explain the connection between the group theoretic formulations of the HNC and SHNC and the graph theoretic formulations. Given generators, g_1, g_2, for the free group, $\mathcal{F}_2$, for each subgroup $\mathcal{K} \subset \mathcal{F}_2$, there is a canonically associated étale bigraph, K; K is given by constructing the Schreier coset graph, $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, and letting K be the "core" of $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, i.e., its smallest subgraph containing all reduced loops based at the vertex $\mathcal{K}$ (see [MW02] or the references in the previous paragraph); $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$ with directed edges labelled either g_1 or g_2 is a (typically infinite degree) covering of B_2, and K, a finite subgraph of $\mathrm{Sch}(\mathcal{F}_2, \mathcal{K}, \{g_1, g_2\})$, is therefore an étale bigraph. If $\mathcal{K}, \mathcal{L}$ are subgroups of $\mathcal{F}_2$, and K, L the corresponding étale bigraphs, then each component of $K \times_{B_2} L$ corresponds to the graph associated to $\mathcal{K} \cap x^{-1}\mathcal{L}x$ ranging over double coset representatives, x.

Theorem 1.1 will be proven by the following equivalent theorem.

Theorem 3.1 *The Strengthened Hanna Neumann Conjecture holds. That is, if $K \to B_2$ and $L \to B_2$ are two étale bigraphs over B_2, then*

$$\rho(K \times_{B_2} L) \le \rho(K)\rho(L). \tag{9}$$

Equation (9) is tight in that if either K or L is a covering of B_2 (i.e., has all vertices of degree four), then the inequality is satisfied with equality.

4 Sheaves and Maximum Excess

In this section we will define sheaves, some invariants of sheaves, and then the inequalities from [Fri] that we need in this paper.

4.1 Sheaves

Definition 4.1 *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 whose values are* U *at each vertex and edge, and whose restriction maps are all 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 explained in [Fri].*

The field, $\mathbb{F}$, is arbitrary here; at some later times we may place restrictions on $\mathbb{F}$, for example, that it be algebraically closed.

To a sheaf, $\mathcal{F}$, on a graph, 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.

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

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

For any sheaf, $\mathcal{F}$, on a graph, G, and any morphism $\phi\colon K \to G$ of directed graphs, we define a "pullback" sheaf $\phi^*\mathcal{F}$ on K 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 $\phi\colon K \to G$ is an arbitrary map, and $\mathcal{F}$ a sheaf on K, there is a natural sheaf $\phi_!\mathcal{F}$ on G defined as follows:

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

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

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

that $d_{\phi_!\mathcal{F},t}$ is the same map as $d_{\mathcal{F},t}$ modulo these isomorphisms, and that the same holds with h replacing t. It follows that $\mathcal{F}$ and $\phi_!\mathcal{F}$ agree on any invariant defined solely in terms of d_t and d_h.

It turns out (see [Fri]) that $\phi_!$ is the left adjoint of ϕ^*.

If $\phi\colon K \to G$ is a morphism of digraphs, we will often write $\underline{\mathbb{F}}_K$ for $\phi_!\underline{\mathbb{F}}$, when confusion is unlikely to arise (which requires ϕ to be understood in context). If $K' \to G$ is another morphism, then we have a natural isomorphism of sheaves over G,

$$\underline{\mathbb{F}}_K \otimes \underline{\mathbb{F}}_{K'} \simeq \underline{\mathbb{F}}_{K\times_G K'}, \tag{10}$$

where $K \times_G K' \to G$ is the fibre product of $K \to G$ and $K' \to G$. If $L \to G$ is an arbitrary digraph morphism, we have an equality of sheaves on K,

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

Finally, if $G' \subset G$ with $j\colon G' \to G$ being the inclusion, then $j_!\underline{\mathbb{F}} = \underline{\mathbb{F}}_{G'}$ is the sheaf whose values are $\mathbb{F}$ on each edge and vertex of G', and 0 elsewhere, and whose restriction maps with nonzero source and nonzero target, hence from $\mathbb{F}$ to $\mathbb{F}$, are the identity maps.

We now define morphisms of sheaves, which are fundamental to many aspects of sheaf theory, including studying sheaf invariants via exact sequences.

Definition 4.2 *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]).

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

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

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

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

4.2 Maximum Excess of a Sheaf

Let $\mathcal{F}$ be a sheaf on a graph, 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

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

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

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

In [Fri] it is shown that if $U \subset \mathcal{F}(V)$ attains the maximum excess, then U is *compartmentalized*, meaning that it is a direct sum over each $v \in V_G$ of a subspace of $\mathcal{F}(v)$ (see [Fri]). In [Fri] it is also shown that the excess is a supermodular function, i.e., for any $U_1, U_2 \subset \mathcal{F}(V)$ we have

$$\operatorname{excess}(\mathcal{F}, U_1) + \operatorname{excess}(\mathcal{F}, U_2) \le \operatorname{excess}(\mathcal{F}, U_1 \cap U_2) + \operatorname{excess}(\mathcal{F}, U_1 + U_2)$$

(this is easy to show). It follows that the *maximizers* of $\mathcal{F}$, i.e., the subspaces $U \subset \mathcal{F}(V)$ on which the excess is maximized, is closed under intersection and addition. In particular, each sheaf $\mathcal{F}$ has a unique maximum (or maximal) maximizer (namely the sum of all the maximizers) and unique minimum (or minimal) maximizer (namely the intersection of all the maximizers).

It is not hard to see that

1. for the structure sheaf, $\underline{\mathbb{F}}$, we have
$$\text{m.e.}(\underline{\mathbb{F}}) = \rho(G);$$
2. for sheaves $\mathcal{F}_1, \mathcal{F}_2$ we have
$$\text{m.e.}(\mathcal{F}_1 \oplus \mathcal{F}_2) = \text{m.e.}(\mathcal{F}_1) + \text{m.e.}(\mathcal{F}_2);$$
and

3. if $\mathcal{E}$ is a sheaf that is *edge supported*, meaning that $\mathcal{E}(V) = 0$, then

$$\mathrm{m.e.}(\mathcal{E}) = \dim\Bigl(\mathcal{E}(E)\Bigr) = -\chi(\mathcal{E})$$

(and similarly the maximum excess a vertex supported sheaf, defined similarly, is zero);

the reader can either consult [Fri] or prove these from scratch.

Definition 4.3 *A sequence of real numbers, $x_0, \ldots, x_r$ is a* triangular sequence *if for any $i = 1, \ldots, r-1$ we have*

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

The following theorem is a consequence of the main theorem of [Fri].

Theorem 4.4 *Let G be a digraph. Consider a "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). Then the sequence

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

is a sequence of non-negative integers that is triangular.

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

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

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

In [Fri], a number of interesting questions about the properties of the maximum excess, alluded to in the introduction, are discussed, such as trying to find an efficient algorithm to find the maximum excess of a sheaf (the algorithm given there is exponential in the "dimension" of a sheaf).

The last result of [Fri] we shall use (in the subsection there on "contagious vanishing theorems") is the following observation:

Lemma 4.5 *Let $\mathcal{F}$ be a sheaf on a digraph, G, of maximum excess zero, and let $u\colon K \to G$ be an étale morphism. Then $\mathcal{F} \otimes \underline{\mathbb{F}}_K$ has maximum excess zero, where $\underline{\mathbb{F}}_K$ denotes $u_!\underline{\mathbb{F}}$.*

Let us mention an alternate interpretation of the maximum excess.

Theorem 4.6 *For any sheaf, $\mathcal{F}$, on a digraph, G, the maximum excess of $\mathcal{F}$ is the same as*

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

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

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

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

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

$\square$

This theorem has a simple graph theoretic analogue, namely that

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

One can easily prove this directly (with $\rho(G) = -\chi(H)$ when H consists of all cyclic connected components of G) or use Theorem 4.6.

5 Galois and Covering Theory in the SHNC

In this section we establish a number of important definitions and facts concerning graph 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). Our main use of Galois theory in the SHNC is to reduce the its study to subgraphs of a Cayley graph. This will later lead us to sheaves we call ρ-kernels. Let us now give definitions and state the main theorems to be developed in this section.

Definition 5.1 *By the* Cayley bigraph on a group, $\mathcal{G}$, with generators g_1 and g_2, *denoted* $G = \mathrm{Cayley}(\mathcal{G}; g_1, g_2)$, *we mean the étale bigraph,* G, *where* $V_G = \mathcal{G}$ *and* $E_G = \mathcal{G} \times \{1,2\}$ *(as sets), such that for each* $g \in \mathcal{G}$ *and* $i = 1, 2$, *the edge* (g, i) *has colour* i, *tail* g, *and head* $g_i g$.

We reduce the SHNC to the special case of subgraphs of a Cayley graph, as follows.

Theorem 5.2 *To prove Theorem 3.1, the SHNC, it suffices verify the SHNC on all pairs,* (L, L'), *such that* L, L' *are subgraphs of the same Cayley bigraph. In particular, to prove the SHNC it suffices to show that any subgraph of a Cayley bigraph is universal for the SHNC.*

The following simplifications of the SHNC on subgraphs of Cayley graphs will help solidify the connection between the SHNC and ρ-kernels of the next section.

Theorem 5.3 *Let* L *be a subgraph of a Cayley bigraph,* G, *on a group,* $\mathcal{G}$. *Then*

1. L *is universal for the SHNC if for any étale* $L' \to G$ *we have* (L, L') *satisfies the SHNC (with* L' *inheriting the edge colouring from* G, *i.e., from the composition* $L' \to G$ *followed by* $G \to B_2$*);*

2. *for any étale* $L' \to G$ *we have*

$$L \times_{B_2} L' \simeq (L\mathcal{G}) \times_G L',$$

where

$$L\mathcal{G} = \coprod_{g \in \mathcal{G}} Lg.$$

Before giving Galois theory we quickly describe the remarkable reason for the strong connection between the SHNC and covering and Galois theory. Since its proof is so short, we give it here as well.

Theorem 5.4 *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

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

$\square$

From this theorem it follows that if $\widetilde{K} \to K$ and $\widetilde{L} \to L$ are covering maps of étale bigraphs, then

$$\rho(\widetilde{K} \times_{B_2} \widetilde{L}) - \rho(\widetilde{K})\rho(\widetilde{L}) = [\widetilde{K}\colon K]\ [\widetilde{L}\colon L]\Big(\rho(K \times_{B_2} L) - \rho(K)\rho(L)\Big);$$

hence (K, L) satisfy the SHNC iff $(\widetilde{K}, \widetilde{L})$ do. This means that to study the SHNC, one can always pass to covers of the bigraphs of interest.

For the rest of this section we describe a number of aspects of what we call Galois graph theory and use it for prove Theorems 5.2 and 5.3.

5.1 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 a covering map 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 $\mathrm{Aut}(K/G)$ that acts simply (i.e., a non-identity $g \in \mathcal{G}$ acts without fixed edges or vertices) and transitively on each of the vertex and edge fibres of π; we declare $\mathcal{G}$ to be the Galois group. Allowing for $\pi\colon K \to G$ to be Galois in this generalized sense, i.e. for K not necessarily connected, does not change any of the theorems below, although some care is required when K is not connected; for example, 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, K_1, K_2 are connected and the $[K_i\colon G_i]$ are relatively prime).

Theorem 5.5 (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 $v_i \ne v_j$ for all i, j with $i \ne j$. 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.

Our next fact is an analogue of a standard and surprisingly useful fact in descent theory (as in [Del77]); it is also surprisingly useful for the SHNC, despite the fact that it is trivial.

Theorem 5.6 *Let $\pi\colon K \to G$ be Galois. Then*

$$K \times_G K = \coprod_{\sigma \in \mathrm{Aut}(K/G)} K_\sigma$$

where K_σ is the subgraph of $K \times_G K$ given via

$$V_{K_\sigma} = \{(v, vg) \mid v \in V_K,\ g \in \mathrm{Aut}(K/G)\},$$

$$E_{K_\sigma} = \{(e, eg) \mid e \in E_K,\ g \in \mathrm{Aut}(K/G)\}.$$

Each K_σ is isomorphic to K.

(See [Fri93], and compare with the identical formula for fields in [Del77], Section I.5.1).

Corollary 5.7 *In Theorem 5.6, let us further assume that we have morphisms $K_1 \to K$ and $K_2 \to K$. Then*

$$K_1 \times_G K_2 \simeq \coprod_{\sigma \in \mathrm{Aut}(K/G)} K_1 \times_K (K_2\sigma).$$

Proof

$$K_1 \times_G K_2 \simeq (K_1 \times_K K) \times_G (K \times_K K_2) \simeq K_1 \times_K (K \times_G K) \times_K K_2$$

$$\simeq \coprod_\sigma \Big(K_1 \times_K K_\sigma \times_K K_2\Big) \simeq \coprod_\sigma \Big(K_1 \times_K (K_2\sigma)\Big).$$

□

There are many extensions to this basic theory. We mention one interesting example.

Assume, for simplicity, that G is connected. If $K \to G$ is Galois and factors as $K \to K' \to G$, then $K \to K'$ is Galois, with Galois group being the subgroup of $\mathrm{Aut}(K/G)$ fixing any vertex or edge fiber of $K \to K'$; hence

$\mathrm{Aut}(K/K')$ is a subgroup of $\mathrm{Aut}(K/G)$. Conversely, a subgroup of $\mathrm{Aut}(K/G)$ divides the vertices and edges of K into orbits, giving a graph K' (whose vertices and edges are these orbits) and a factorization $K \to K' \to G$. Furthermore, for an intermediate cover $K \to K' \to G$, $K \to K'$ is always Galois (since $\mathrm{Aut}(K/K')$ has the right cardinality), and $K' \to G$ is Galois iff the subgroup of $\mathrm{Aut}(K/G)$ fixing $K \to K'$ fibers is a normal subgroup of $\mathrm{Aut}(K/G)$. See [ST96] for details.

If $K \to G$ is Galois and factors as $K \to K' \to G$,

$$K \times_G K' = \coprod_{g \in \mathrm{Aut}(K/G)/\mathrm{Aut}(K/K')} K_g, \tag{11}$$

where

$$V_{K_g} = \{(v, [v]g) \mid v \in V_K\}, \quad E_{K_g} = \{(e, [e]g) \mid e \in E_K\},$$

where $[v], [e]$ respectively denote the images of v, e, respectively, in K'; each K_g is isomorphic to K. Special cases of this statement include the trivial case $K' = G$ and the case $K' = K$ stated earlier.

5.2 Base Change

There are a number of easy "stability under base change" results; these say that in a digram arising from arbitrary digraph morphisms $L \to B$ and $M \to B$,

$$\begin{tikzcd} L \times_B M \arrow[r] \arrow[d] & M \arrow[d] \\ L \arrow[r] & B \end{tikzcd}$$

if $L \to B$ has a certain property, then so does $L \times_B M \to M$. Just from the construction of the fibre product, we easily see that the following classes of morphisms are stable under base change: étale morphisms, covering morphism, and Galois morphisms (and many others that we won't need, such as open inclusions, morphisms that are d–to–1 for some fixed d, etc.).

5.3 Etale Factorization

In this subsection we shall prove that any étale map factorizes as an open inclusion followed by a covering map. This will easily establish Theorem 5.2.

We define an *open inclusion* to be any inclusion $H \to G$ of a subgraph, H, in a graph, G. We say the inclusion is *dense* if $V_H = V_G$; this agrees with the topological notion, i.e., the closure of G in H is H, under the topological view of G in [Fri].

Lemma 5.8 *Let $\pi\colon G \to B$ be an étale map. Then π factors as an open inclusion, $\iota\colon G \to G'$, followed by a covering map, $\pi'\colon G' \to B$. If the vertex fibres of π (i.e., $\pi^{-1}(v)$ over all $v \in V_B$) are all of the same size, i.e., $\pi_V\colon V_G \to V_B$ is d-to-1 for some d, then we may assume ι is dense; if in addition G is connected, then we may assume G' is connected.*

A variant of the first sentence of this theorem is called Marshall Hall's theorem in [Sta83].
Proof By adding isolated vertices to G we may assume π_V is d-to-1 for some d. Extend G to G' and π to $\pi'\colon G' \to B$ by completing each $\pi^{-1}(e)$ to a perfect bipartite matching of the vertices over the tail of e and those over the head of e (for a self-loop we view these two sets of vertices as disjoint). Clearly $\pi'\colon G' \to B$ is a covering map. If π_V was originally d-to-1 for some d, then G' is obtained by adding only edges, so G is dense in G'; if furthermore G is connected, then the G' obtained by adding only edges is, of course, still connected.

$\square$

Here is an easy, but vital, observation.

Lemma 5.9 *If $S \to B_2$ is a Galois map with Galois group, $\mathcal{G}$, then S is isomorphic to a Cayley bigraph on the group $\mathcal{G} = \mathrm{Gal}(S/B_2)$.*

Proof Choose any $v_0 \in V_S$ to be the "origin." The association $g \mapsto v_0 g$ sets up an identification of $\mathcal{G}$ with V, by definition of a Galois covering map, since there is a unique vertex of B_2 and hence a singe vertex fibre in S. Since $S \to B_2$ is a covering map, the vertex v_0 is the tail of a unique colour 1 edge, e, whose head is $v_0 g_1$ for a unique $g_1 \in \mathcal{G}$. For any g we have eg has tail $v_0 g$ and head $v_0 g_1 g$. It follows that identifying V with $\mathcal{G}$ means that there is an edge $(g, g_1 g)$ (i.e., whose tail is g and head is $g_1 g$) of colour 1 for each $g \in \mathcal{G}$. Similarly for edges of colour 2, and this sets up an isomorphism between S and $\mathrm{Cayley}(\mathcal{G}; g_1, g_2)$.

$\square$

Proof of Theorem 5.2 Let $G \to B_2$ and $K \to B_2$ be étale maps. Let these étale maps factor as open inclusions followed by covering maps as $G \to \widetilde{G} \to B$ and $K \to \widetilde{K} \to B$. Let S be a Galois cover of $\widetilde{G} \times_B \widetilde{K}$. Consider $G' = G \times_{\widetilde{G}} S$, which admits a natural map to G (namely projection onto the first component), and similarly $K' = K \times_{\widetilde{K}} S$. We claim that $G' \to G$ is a covering map; indeed, by stability under base change (see Subsection 5.2), since $\widetilde{K} \to B_2$ is a covering map, so is $\widetilde{G} \times_B \widetilde{K} \to \widetilde{G}$; since $S \to \widetilde{G} \times_B \widetilde{K}$ is a covering map, so is $S \to \widetilde{G}$; hence, by base change so is $G' \to G$. Similarly $K' \to K$ is a covering map. According to Theorem 5.4 and the discussion below, the SHNC is satisfied at (G, K) iff it is satisfied at (G', K'). But G', K' are subgraphs of S, and S is a Galois cover of B_2, and therefore a Cayley bigraph.

$\square$

Although we shall not need it, we mention that the idea in this last proof can be extended from $G \to B_2$ and $K \to B_2$ to an arbitrary number of étale maps, $L_i \to B_2$, and gives the following interesting fact.

Theorem 5.10 *For any étale bigraphs $L_1, \ldots, L_k$, there are covering maps $L_i' \to L_i$ and a Cayley bigraph, S, such that each L_i' is a subgraph of S that is dense (i.e., L_i' has the same vertex set as S).*

5.4 The Proof of Theorem 5.3

We finish this section with the proof of Theorem 5.3.

Claim (1) of the theorem is a simple base change argument: if L is a subgraph of a Cayley bigraph, G, and $L' \to B_2$ is any étale bigraph, let $L'' = L' \times_{B_2} G$. Then, by base change (see Subsection 5.2), $L'' \to G$ is étale and $L'' \to L'$ is a covering map. Then (L, L') satisfies the SHNC iff (L, L'') does. Hence L is universal for the SHNC iff (L, L'') satisfies the SHNC for all étale bigraphs, L'', whose colouring map factor as $L'' \to G \to B_2$.

Claim (2) is an immediate consequence of Corollary 5.7, with G, K, K_1, K_2 respectively replaced by B_2, G, L, L', noting that $\mathrm{Aut}(G/B_2) = \mathcal{G}$.

$\square$

6 ρ-kernels

In this section we introduce a collection of sheaves that are central to our proof of the SHNC. They are called ρ-kernels. Before defining them, we motivate their construction by showing how their study is connected to the SHNC. First we need to set some notation on Cayley graphs.

6.1 Sheaves on Cayley graphs

Let $G = \mathrm{Cayley}(\mathcal{G}; g_1, g_2)$ be a Cayley bigraph on a group, $\mathcal{G}$. Recall that since our generators act on the left, e.g., the colour 1 edges are of the form $(g, g_1 g)$, the Galois group of G is $\mathcal{G}$ acting on the right. Now we define a right action of $\mathcal{G}$ on sheaves on G. We shall state this in slightly more general terms. This is completely straightforward and mildly tedious, but convenient in this section and vital to Section 7.

Definition 6.1 *We say that a group, $\mathcal{G}$,* acts *on a digraph, G,* on the right*, if associated to each $g \in \mathcal{G}$ is an isomorphism π_g, of G such that $\pi_{g_1 g_2} = \pi_{g_2}\pi_{g_1}$ for all $g_1, g_2 \in G$. We will identify g with π_g if no confusion can arise. If L is a subgraph of G, we write Lg for the image of L under g (i.e., under π_g); similarly if $P \in V_G \amalg E_G$, Pg denotes the image of P under g.*

Of course, if G is a Cayley bigraph on a group, $\mathcal{G}$, then $\mathcal{G}$ acts on G on the right.

Theorem 6.2 *Let a group, $\mathcal{G}$, act on a digraph, G, on the right. Then each element of $\mathcal{G}$ acts naturally as a functor on sheaves, via the association $g \mapsto \pi^*_{g^{-1}}$, such that*

1. *$\mathcal{G}$ acts on the right, i.e., if we write $\mathcal{F}g$ for $\pi^*_{g^{-1}}\mathcal{F}$ for any sheaf, $\mathcal{F}$, on G, then for any $g_1, g_2 \in \mathcal{G}$ we have $\mathcal{F}g_1g_2 = (\mathcal{F}g_1)g_2$, and similarly with the sheaf $\mathcal{F}$ replaced by a morphism of sheaves;*

2. *for each sheaf, $\mathcal{F}$, on G, any $g \in \mathcal{G}$, and any $P \in V_G \amalg E_G$ we have*

$$(\mathcal{F}g)(P) = \mathcal{F}(Pg^{-1});$$

   and

3. for each subgraph $L \subset G$ and field, $\mathbb{F}$, we have

$$\underline{\mathbb{F}}_L\, g = \underline{\mathbb{F}}_{Lg}.$$

In Section 7 it will be important to use the fact that for each $g \in \mathcal{G}$, $\pi^*_{g^{-1}}$ is a *functor*, i.e., it acts (compatibly) on morphisms of sheaves as well as on sheaves.

Proof For item (1), we recall that for any $u\colon G' \to G$, u^* is a functor on sheaves, and for composable morphisms of digraphs, u_1, u_2, we have $(u_1u_2)^* = u_2^*u_1^*$; hence, since $\mathcal{G}$ acts on the right, for any $g_1, g_2 \in \mathcal{G}$ we have

$$\pi^*_{g_1^{-1}}\pi^*_{g_2^{-1}} = (\pi_{g_2^{-1}}\pi_{g_1^{-1}})^* = (\pi_{g_1^{-1}g_2^{-1}})^* = \pi^*_{(g_2g_1)^{-1}}$$

and so $g \mapsto \pi^*_{g^{-1}}$ is defines an action on sheaves and morphisms of sheaves that acts on the right.

Item (2) follows immediately from the definition of the pullback. Item (3) follows since for all $P \in V_G \amalg E_G$ and $L \subset G$ and $g \in \mathcal{G}$ we have

$$(\underline{\mathbb{F}}_L\, g)(P) = \underline{\mathbb{F}}_L(Pg^{-1});$$

but $Pg^{-1} \in L$ iff $P \in Lg$, so

$$\underline{\mathbb{F}}_{Lg}(P) = \underline{\mathbb{F}}_L(Pg^{-1}) = (\underline{\mathbb{F}}_L\, g)(P).$$

Hence $\underline{\mathbb{F}}_{Lg} = \underline{\mathbb{F}}_L\, g$.

□

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

$$\mathcal{F}\mathcal{G} = \bigoplus_{g \in \mathcal{G}} \mathcal{F}g.$$

In particular, for $L \subset G$, if we set

$$L\mathcal{G} = \coprod_{g \in \mathcal{G}} Lg$$

(akin to the notation in Theorem 5.3) then

$$\underline{\mathbb{F}}_L\, \mathcal{G} \simeq \underline{\mathbb{F}}_{L\mathcal{G}}.$$

6.2 Kernels and the SHNC

The following theorem summarizes our approach to the SHNC.

Theorem 6.3 *Let L be a subgraph of a Cayley bigraph, G. Assume there is an exact sequence*

$$0 \to \mathcal{K} \to \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^{\rho(L)} \to 0 \tag{12}$$

such that $\mathrm{m.e.}(\mathcal{K}) = 0$*. Then L is universal for the SHNC.*

Proof According to Theorem 5.3, it suffices to show that for each étale $u\colon L'' \to G$ we have that (L, L'') satisfies the SHNC. Tensoring equation (12) with $\mathbb{F}_{L''} = u_! \underline{\mathbb{F}}$ gives

$$0 \to \mathcal{K} \otimes \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_{L''}^{\rho(L)} \to 0. \tag{13}$$

Since $\mathrm{m.e.}(\mathcal{K}) = 0$ and $u\colon L'' \to G$ is étale, we have

$$\mathrm{m.e.}(\mathcal{K} \otimes \underline{\mathbb{F}}_{L''}) = 0$$

(this, as we mentioned earlier, is an example of a "contagious vanishing theorem" proven in [Fri]). Since the maximum excess is a first quasi-Betti number, this and equation (13) implies that

$$\mathrm{m.e.}(\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''}) \le \mathrm{m.e.}(\underline{\mathbb{F}}_{L''}^{\rho(L)}) = \rho(L)\rho(L'').$$

But

$$\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''} \simeq \underline{\mathbb{F}}_{(L\mathcal{G}) \times_G L''} \simeq \underline{\mathbb{F}}_{L \times_{B_2} L''}$$

(using equation (10) and Corollary 5.7), and so we have

$$\rho(L \times_{B_2} L'') = \mathrm{m.e.}(\underline{\mathbb{F}}_{L \times_{B_2} L''}) = \mathrm{m.e.}(\underline{\mathbb{F}}_L \mathcal{G} \otimes \underline{\mathbb{F}}_{L''}) \le \rho(L)\rho(L'').$$

$\square$

6.3 Definition and Existence of ρ-Kernels and k-th Power Kernels

We begin with some notation to describe the kernels we introduce here and study throughout the rest of this paper.

Let G be a Cayley bigraph on a group, $\mathcal{G}$. For any integer $k \ge 0$, let $\mathbb{F}^{k \times \mathcal{G}}$ be the set of $k \times |\mathcal{G}|$ matrices with entries $m_{ig} \in \mathbb{F}$ indexed over $i = 1, \ldots, k$

and $g \in \mathcal{G}$. If $M \in \mathbb{F}^{k \times \mathcal{G}}$, then we can view M as a map from $\mathbb{F}^{\mathcal{G}}$ to $\mathbb{F}^k$. Then M gives rise to a morphism of constant sheaves

$$\underline{M} \colon \underline{\mathbb{F}}^{\mathcal{G}} \to \underline{\mathbb{F}}^k.$$

For any $L \subset G$ and $g \in \mathcal{G}$, we have an inclusion $\underline{\mathbb{F}}_{Lg} \to \underline{\mathbb{F}}$, which gives us an inclusion

$$\underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}} \mathcal{G} \simeq \underline{\mathbb{F}}^{\mathcal{G}}.$$

Thus we get a monomorphism

$$\iota_{L\mathcal{G}} \colon \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^{\mathcal{G}},$$

and, for any $M \in \mathbb{F}^{k \times \mathcal{G}}$, a composite morphism

$$\underline{M} \iota_{L\mathcal{G}} \colon \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^k.$$

We shall often write ι instead of $\iota_{L\mathcal{G}}$, since the subscript $L\mathcal{G}$ can be inferred from the source (even if two different ι's are involved).

Definition 6.4 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be a field. For any integer $k \geq 0$, we say that $M \in \mathbb{F}^{k \times \mathcal{G}}$ is* L-surjective *if the map, $\underline{M} \iota_{L\mathcal{G}}$ is surjective. If so, we that its kernel, $\mathcal{K} = \mathcal{K}_M(L, G, \mathcal{G})$, is a k-th power kernel for $(L, G, \mathcal{G})$; if, in addition, $k = \rho(L)$, we also say that $\mathcal{K}$ is a* ρ-kernel for $(L, G, \mathcal{G})$.

Note that when kernels are defined in category theory, i.e., for a category with a zero morphism, then a kernel is defined only up to (unique) isomorphism. However, for sheaves on a graph, we can define the kernel of a morphism $\mathcal{F}_1 \to \mathcal{F}_2$ uniquely, as the subsheaf of $\mathcal{F}_1$ that is the kernel. Hence we will speak of *the* kernel of a morphism, or *its* kernel, for convenience; when we say "a kernel" we shall mean the category theory notion, i.e., any morphism $\mathcal{K} \to \mathcal{F}_1$ that is the equalizer of $\mathcal{F}_1 \to \mathcal{F}_2$ and the zero morphism.

Note that we could also define k-th power kernels when $\underline{M} \iota_{L\mathcal{G}}$ is not surjective, as the element of the derived category (see [GM03]) as a single shift of the mapping cone of $\underline{M} \iota_{L\mathcal{G}}$; we shall not pursue this here.

The important point to notice is that if $k \leq \rho(L)$, "most" matrices $M \in \mathbb{F}^{k \times \mathcal{G}}$ are L-surjective. We now demonstrate this, in a rather explicit fashion.

Definition 6.5 *We say that $M \in \mathbb{F}^{k\times\mathcal{G}}$ is* totally linearly independent *(or just* totally independent*) if every subset $\mathcal{G}'$ of $\mathcal{G}$ of size k we have $\{m^g\}_{g\in\mathcal{G}'}$ is linearly independent, where m^g denotes the column of M corresponding to $g \in \mathcal{G}$.*

Lemma 6.6 *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Then the number of vertices of L and the number of edges of either colour in L are all at least $\rho(L)$.*

Proof Adding vertices and edges to a graph does not decrease its reduced cyclicity (i.e., its ρ). So if P is an edge of colour 2, let L' be L union all vertices of G and all edges of colour 1. Then $\rho(L') \geq \rho(L)$ and L' has the same number of edges of colour 2 as L. But if we discard the edges of colour 2 from L' we are left with a union of cycles, for which $\rho = 0$, and discarding one edge decreases ρ by at most one (given equation (6)). Hence the number of edges of colour 2 in L' is at least $\rho(L)$, and so the same is true of the number of colour 2 edges in L.

Similarly L must have at least $\rho(L)$ edges of colour 1. Finally, since each vertex of L is the head of at most one edge of colour 1, the number of vertices is also at least $\rho(L)$.

□

Now we wish to describe ρ-kernels, both as a kernel of a sheaf morphism and, alternatively, by explicitly giving their values and restrictions.

Definition 6.7 *Fix a subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$. Fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ for an integer $k \geq 0$. For a subset $T \subset \mathcal{G}$, the T*-free subspace of $\ker(M)$ *we mean the set*

$$\mathrm{Free}_T = \mathrm{Free}_T(M) = \{\vec{a} \in \ker(M) \mid \forall g \notin T,\ a_g = 0\}.$$

A free subspace of $\ker(M)$ *is a subspace that is T-free for some $T \subset \mathcal{G}$. For $P \in V_G \amalg E_G$, we set*

$$\mathcal{G}_L(P) = \{g \in \mathcal{G} \mid P \in Lg\}.$$

In the above definition, if $M \in \mathbb{F}^{k\times\mathcal{G}}$ is totally independent, then for all $T \subset \mathcal{G}$ we have

$$\dim(\mathrm{Free}_T) = \max(0, |T| - k). \tag{14}$$

Lemma 6.8 *Let L be a subgraph of a Cayley bigraph, G, on a graph $\mathcal{G}$. Let $M \in \mathbb{F}^{k\times\mathcal{G}}$ be totally independent, for some $k \le \rho(L)$. Then*

$$\underline{M}\iota \colon \underline{\mathbb{F}}_L\mathcal{G} \to (\underline{\mathbb{F}}_G)^k$$

is surjective. Furthermore, if $\mathcal{K}_M$ denotes its kernel, then for each $P \in V_G \amalg E_G$ we have

$$\mathcal{K}_M(P) = \mathrm{Free}_{\mathcal{G}_L(P)}(M),$$

in the notation of Definition 6.7, and the restriction maps for $\mathcal{K}_M$ are the inclusions. In particular,

$$\dim\Big(\mathcal{K}_M(P)\Big) = n_P - \rho(L),$$

where $n_P = |\mathcal{G}_L(P)|$. (We shall sometimes write $\mathcal{K}_M$ as $\mathcal{K}_M(L)$ or $\mathcal{K}_M(L, G, \mathcal{G})$ to emphasize $\mathcal{K}_M$'s dependence upon L, G, and $\mathcal{G}$.)

Proof For each $g \in \mathcal{G}$ we have $\underline{\mathbb{F}}_L g = \underline{\mathbb{F}}_{Lg}$. Hence for each $P \in V_G \amalg E_G$, we have

$$(\underline{\mathbb{F}}_L g)(P) = (\underline{\mathbb{F}}_{Lg})(P) = \begin{cases} \mathbb{F} & \text{if } P \in Lg \\ 0 & \text{if } P \notin Lg \end{cases}$$

Hence

$$(\underline{\mathbb{F}}_L\mathcal{G})(P) \simeq \bigoplus_{g\in\mathcal{G}_L(P)} \mathbb{F},$$

and the image of $\underline{M}\iota$ in $(\mathbb{F}_G)^k$ at P is the span of the subcollection of the n_P columns of M corresponding to the elements of $\mathcal{G}_L(P) \subset \mathcal{G}$. Since G is a Cayley bigraph, n_P is either the number of vertices, edges of colour 1, or edges of colour 2 in L. By Lemma 6.6 we have $n_P \ge \rho(L)$, and hence this subcollection of n_P vectors in M spans $\mathbb{F}^k$. Hence $\underline{M}\iota$ is surjective at P, and its kernel, $\mathrm{Free}_{\mathcal{G}_L(P)}$, is of dimension $n_P - k$. The restriction maps on $\underline{\mathbb{F}}_L\mathcal{G}$ are, component by component, those of the individual $\underline{\mathbb{F}}_L g$ over all $g \in \mathcal{G}$, and those are just inclusions; since $\mathcal{K}$ is a subsheaf of $\underline{\mathbb{F}}_L\mathcal{G}$, we have that $\mathcal{K}$ inherits those restriction maps.

□

Note that it is easy to see, even with $\mathcal{G} = \mathbb{Z}/3\mathbb{Z}$ and L consisting of five edges, that there need not be any graph theoretic surjections $L\mathcal{G} \to G^{\rho(L)}$, where $G^{\rho(L)}$ is $\rho(L)$ disjoint copies of G; so in passing from the graphs $L\mathcal{G}$ and $G^{\rho(L)}$ to the sheaves $\underline{\mathbb{F}}_L\mathcal{G}$ and $\mathbb{F}^{\rho(L)}$, there exists a surjection of sheaves that does not arise from any surjection of graphs. So an added benefit of working with sheaves (aside from using them to form kernels useful in studying the SHNC) is that sheaves give "additional surjections" that don't exist in graph theory.

7 Symmetry and Algebra of the Excess

In this section we make some general observations about the maximum excess of k-th power kernels. The main observation is that given $(L, G, \mathcal{G})$ as usual, the maximum excess of $\mathcal{K}_M(L)$ for generic $M \in \mathbb{F}^{k \times \mathcal{G}}$ is divisible by $|\mathcal{G}|$, where by "generic" we mean for M in some subset of $\mathbb{F}^{k \times \mathcal{G}}$ that contains a nonempty, Zariski open subset of $\mathbb{F}^{k \times \mathcal{G}}$. Let us outline this argument.

First, in Subsection 7.1, we will show that for any $M \in \mathbb{F}^{k \times \mathcal{G}}$ and $g \in \mathcal{G}$ we have $\mathcal{K}_M(L)g \simeq \mathcal{K}_{Mg}(L)$ where Mg is obtained by an appropriate action of g on the columns of M. This means that if $\mathcal{F}$ is the maximal (or minimal) excess maximizer for $\mathcal{K}_M(L)$, then $\mathcal{F}g$ is isomorphic to the maximal (or, respectively, minimal) excess maximizer for $\mathcal{K}_{Mg}(L)$. It may be helpful, albeit somewhat fanciful, to understand this symmetry via two "observers" looking at the exact sequence

$$0 \to \mathcal{K}_M \to \underline{\mathbb{F}}_L \mathcal{G} \xrightarrow{M\iota} \underline{\mathbb{F}}^k \to 0,$$

one who examines this at $P \in V_G \amalg E_G$, and the other at Pg, for g fixed and P varying; for example, $\underline{\mathbb{F}}_L \mathcal{G}$ "looks" the same to both observers, except that its summands appear permuted from one observer to the other.

In Subsection 7.2 we discuss the generic maximum excess of $\mathcal{K}_M = \mathcal{K}_M(L)$ with L fixed and $M \in \mathbb{F}^{k \times \mathcal{G}}$ variable (and $G, \mathcal{G}, k$ fixed). The key to this discussion is considering what we call "dimension profiles," which we now define.

Definition 7.1 *By a* dimension profile *on a bigraph, G, we mean a function*

$$n \colon V_G \amalg E_G \to \mathbb{Z}_{\geq 0}.$$

For any such n, we set

$$\chi(n) = \sum_{v \in V_G} n(v) - \sum_{e \in E_G} n(e); \qquad |n| = \sum_{P \in V_G \amalg E_G} n(P).$$

Any sheaf, $\mathcal{F}$, on G determines a dimension profile, $\dim(\mathcal{F})$*, as the function $P \mapsto \dim(\mathcal{F}(P))$. For any dimension profile, n, of a Cayley bigraph, G, on a group, $\mathcal{G}$, any subgraph, $L \subset G$, any field, $\mathbb{F}$, and any $k \geq 0$, let*

$$\mathcal{M}(n) = \mathcal{M}(n, L, G, \mathcal{G}, \mathbb{F}, k)$$

be the set of $M \in \mathbb{F}^{k \times \mathcal{G}}$ for which $\mathcal{K}_M = \mathcal{K}_M(L, G, \mathcal{G})$ exists (i.e., M is L-surjective) and has a subsheaf, $\mathcal{F}$, with $\dim(\mathcal{F}) = n$.

We easily see that for all n, $\mathcal{M}(n)$ is a constructible subset of $\mathbb{F}^{k\times\mathcal{G}}$. Let $\mathcal{N}$ be the set of n for which $\mathcal{M}(n)$ is generic, i.e., contains a Zariski open subset of $\mathbb{F}^{k\times\mathcal{G}}$; since $\mathcal{M}(n)$ is constructible, it is equivalent to say that its Zariski closure is $\mathbb{F}^{k\times\mathcal{G}}$. The generic (in M) value of the maximum excess of $\mathcal{K}_M$ is the largest value of $-\chi(n)$ among those $n \in \mathcal{N}$; let $\mathcal{N}'$ be the subset of $n \in \mathcal{N}$ which attain this largest $-\chi(n)$ value. Since $\mathcal{G}$ is finite, for any $n \in \mathcal{N}$ there is a generic set of M such that $Mg \in \mathcal{M}(n)$ for all $g \in \mathcal{G}$. Hence, if $n \in \mathcal{N}'$ is chosen with $|n|$ at a maximum value (or minimum value), then by the uniqueness of the maximum (or minimum) maximizer of the excess, the symmetry $\mathcal{K}_M g \simeq \mathcal{K}_{Mg}$ implies that $n(P) = n(Pg^{-1})$ for all $P \in V_G \amalg E_G$ and $g \in \mathcal{G}$. Hence the generic maximum excess of $\mathcal{K}_M$, which equals $-\chi(n)$ for any $n \in \mathcal{N}'$, is divisible by $\mathcal{G}$.

We wish to remark that the generic maximum excess is not generally attained by all $M \in \mathbb{F}^{k\times\mathcal{G}}$. For example, our approach to the SHNC is based on the fact that the generic maximum excess of a ρ-kernel is zero, i.e., for $k = \rho(L)$ (see Theorem 9.3). However, if M is zero in one column, but totally independent in the others, then M will still be L-surjective provided that L has at least $\rho(L) + 1$ edges of each colour. (A simple example of such an L can be obtained by deleting one edge of each colour from $\mathrm{Cayley}(\mathbb{Z}/m\mathbb{Z}; 1, 1)$ with $m \geq 2$.) In such a situation, the fact that M has a column of zeros implies that $\mathcal{K}_M$ has $\mathbb{F}_L g$ as a subsheaf (more precisely, a direct summand) for some $g \in \mathcal{G}$, and hence the maximum excess of $\mathcal{K}_M$ will be at least $\rho(L)$. Hence any L that has at least $\rho(L) + 1$ edges of each colour, and for which $\rho(L) > 0$, has a ρ-kernel of positive maximum excess. Hence it is essential to study the maximum excess of $\mathcal{K}_M$ with some restrictions on M, i.e., requiring some special properties of M; in our case, these properties restrict M to some generic subset of $\mathbb{F}^{\rho(L)\times\mathcal{G}}$.

7.1 Symmetry of k-th Power Kernels

The point of this subsection is to establish the following symmetry of k-th power kernels.

Theorem 7.2 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be an arbitrary field. Let k be an arbitrary non-negative integer and $M \in \mathbb{F}^{k\times\mathcal{G}}$. Let Mg be the matrix (described earlier) whose g' column, for $g' \in \mathcal{G}$, is the $g'g^{-1}$ column of M. Then M is L-surjective iff Mg is L-surjective, and if so then $\mathcal{K}_M(L)g \simeq \mathcal{K}_{Mg}(L)$.*

Proof We begin our discussion of symmetry with a somewhat pedantic, but important, point. If $\mathcal{A}$ is a category in which finite direct sums exists, such as an additive category, and $\{A_s\}_{s\in S}$ is a family of objects in the category indexed upon a finite set, S, then their direct sum comes with projections

$$f_r \colon \bigoplus_{s\in S} A_s \to A_r$$

for each $r \in S$. If $\pi\colon S \to S$ is a permutation, then we have a "component permuting map," $P = P(\pi)$, given by

$$P(\pi)\left(\bigoplus_{s\in S} A_s\right) = \bigoplus_{s\in S} A_{\pi(s)},$$

The two direct sums in this last equation are isomorphic, but not equal (e.g., the direct sum on the right-hand-side has the projection f_r' whose target is $A_{\pi(r)}$, not to A_r, for each $r \in \mathcal{G}$). We shall need to keep the seemingly unimportant operator $P = P(\pi)$ in mind in order to make things precise for this subsection. If $A_\bullet$ is any direct sum indexed on $\mathcal{G}$, then we easily see $P(\pi_2)(P(\pi_1)A_\bullet) = P(\pi_2 \circ \pi_1)A_\bullet$.

Again, let $\mathbb{F}$ be a field, $k \geq 0$ an integer, L a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and $M \in \mathbb{F}^{k\times\mathcal{G}}$ that is L-surjective. We have an exact sequence.

$$0 \longrightarrow \mathcal{K} \longrightarrow \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g' \xrightarrow{\underline{M}\iota} \underline{\mathbb{F}}^k \longrightarrow 0. \tag{15}$$

For a $g \in \mathcal{G}$, applying $\pi^*_{g^{-1}}$, of Theorem 6.2, to this sequence gives an exact sequence:

$$0 \longrightarrow \mathcal{K}g \longrightarrow \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\pi^*_{g^{-1}}(\underline{M}\iota)} \underline{\mathbb{F}}^k g \longrightarrow 0. \tag{16}$$

We have $\underline{\mathbb{F}}g = \underline{\mathbb{F}}$ since $\underline{\mathbb{F}}$ is a constant sheaf (note the we mean that the two are equal, not merely isomorphic). Note that $\pi^*_{g^{-1}}$ acts on sheaves by renaming the vertices of G, so it acts on $\underline{M}\iota$ only by permuting sheaf inclusions $\underline{\mathbb{F}}_{Lg''} \to \underline{\mathbb{F}}$ for various values of g''; in other words,

$$\pi^*_{g^{-1}}(\underline{M}\iota_{\underline{\mathbb{F}}_L\mathcal{G}}) = \underline{M}\iota',$$

where ι' is ι with the source $\underline{\mathbb{F}}_L\mathcal{G}g$. Hence we may write equation (16) as

$$0 \longrightarrow \mathcal{K}g \xrightarrow{j} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\underline{M}\iota'} \underline{\mathbb{F}}^k \longrightarrow 0, \tag{17}$$

where j is an inclusion.

Also, we have

$$P(\pi_{g^{-1}})\left(\bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g\right) = \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, \pi_{g^{-1}}(g')g = \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g' = \underline{\mathbb{F}}_L\mathcal{G}.$$

Hence from equation (17) we get a sequence

$$0 \longrightarrow \mathcal{K}g \xrightarrow{j} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{P(\pi_{g^{-1}})} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{P(\pi_g)} \bigoplus_{g'\in\mathcal{G}} \underline{\mathbb{F}}_L\, g'g \xrightarrow{\underline{M}\iota'} \underline{\mathbb{F}}^k \longrightarrow 0,$$

and hence an exact sequence (since $P(\pi_g)$ and $P(\pi_{g^{-1}})$ are isomorphisms)

$$0 \longrightarrow \mathcal{K}g \xrightarrow{P(\pi_{g^{-1}})\circ j} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\underline{M}\iota' P(\pi_g)} \underline{\mathbb{F}}^k \longrightarrow 0, \tag{18}$$

with j being the inclusion in equation (17). But clearly

$$\underline{M}\iota' P(\pi_g) = \underline{M}P(\pi_g)\,\iota_{\underline{\mathbb{F}}_L\mathcal{G}} = \underline{M\pi_g}\,\iota_{\underline{\mathbb{F}}_L\mathcal{G}},$$

where π_g is viewed as operating vectors in $\mathbb{F}^{\mathcal{G}}$ sending $\alpha \in \mathbb{F}^{\mathcal{G}}$, viewed as a function $\alpha\colon \mathcal{G} \to \mathbb{F}$ to $\pi_g\alpha$ given by

$$g' \mapsto \alpha(\pi_g(g')) = \alpha(g'g).$$

Hence setting $Mg = M\pi_g$, we get a short exact sequence

$$0 \longrightarrow \mathcal{K}g \longrightarrow \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\underline{Mg}\,\iota} \underline{\mathbb{F}}^k \xrightarrow{0} . \tag{19}$$

Hence $\mathcal{K}_{Mg}(L)$ is, up to isomorphism, just $\mathcal{K}g$.

To complete the proof of the theorem, it remains to find that permutation that brings the columns of M to those of Mg. If $M = \{m_{i,g'}\}$ give M's entries, then for any $w \in \mathbb{F}^{\mathcal{G}}$, the i-th component of $(Mg)w = M(\pi_g w)$ is

$$(M(\pi_g w))_i = \sum_{g' \in \mathcal{G}} m_{i,g'}(\pi_g w)_{g'} = \sum_{g' \in \mathcal{G}} m_{i,g'} w_{g'g} = \sum_{g'' \in \mathcal{G}} m_{i,g''g^{-1}} w_{g''}.$$

Hence the i, g'' entry of Mg is $m_{i,g''g^{-1}}$, so the g'' column of Mg is the $g''g^{-1}$ of M.

□

We wish to make a few comments on equation (19) and how we derived it. First, kernels, in category theory, are defined only up to isomorphism; this is why we can "forget about" $P(\pi_{g^{-1}})j$ in equation (18); it is only important to know that this arrow gives an exact sequence there and in equation (19).

Note that the two actions of $g \in \mathcal{G}$ in equation (19) are right $\mathcal{G}$ actions on the exact sequence. To see this, first note that

$$(Mg)g' = (M\pi_g)\pi_{g'} = M\pi_{gg'} = Mgg'.$$

Then note that if we take the procedure for going from equation (15) to equation (19) and then do the same procedure with g replaced by g', then we easily see (paying close attention to the order of the P's, the π's, and the ι's) that we get the same equation as equation (15) with $\mathcal{K}$ replaced by $\mathcal{K}gg'$ and M replaced by Mgg'.

We wish to comment on something that seems a bit contradictory. The map $g \mapsto P(\pi_{g^{-1}})$ is a left action, and so it may seem strange that our forgotten monomorphism $P(\pi_{g^{-1}})j$ in equation (18) involves a left action. But note that if we apply g' to equation (19) we get

$$0 \longrightarrow \mathcal{K}gg' \xrightarrow{\pi^*_{g'^{-1}}(P(\pi_{g^{-1}})\circ j)} \underline{\mathbb{F}}_L\mathcal{G}g' \xrightarrow{\pi^*_{g'^{-1}}(\underline{M}\iota' P(\pi_g))} \underline{\mathbb{F}}^k \longrightarrow 0,$$

and hence an exact sequence

$$0 \longrightarrow \mathcal{K}gg' \xrightarrow{\alpha} \underline{\mathbb{F}}_L\mathcal{G} \xrightarrow{\beta} \underline{\mathbb{F}}^k \longrightarrow 0,$$

where

$$\begin{aligned} \alpha &= P(\pi_{g'^{-1}}) \circ \pi^*_{g'^{-1}}(P(\pi_{g^{-1}}) \circ j) = P(\pi_{g'^{-1}})P(\pi_{g^{-1}})j', \\ \beta &= \pi^*_{g'^{-1}}(\underline{M}\iota' P(\pi_g))P(\pi_{g'}), \end{aligned}$$

for an inclusion j'. Examining α we see that $P(\pi_{g'^{-1}})$ is applied to the left of $P(\pi_{g^{-1}})$, whose product equals $P(\pi_{(gg')^{-1}})$, so that g' appears to the right of g.

A similar remark applies for the column permuting rule taking M to Mg: $g \mapsto \pi_{g^{-1}}$ is a left action, not a right action. However, if $f\colon \mathcal{G} \to T$ is any function from $\mathcal{G}$ to a set, T, then defining a function fg via $(fg)(g') = f(g'g^{-1})$ defines a *right* action of $\mathcal{G}$ on functions from $\mathcal{G}$ to T; indeed, for $f\colon \mathcal{G} \to T$ and $g, g_1, g_2 \in \mathcal{G}$ we have

$$\Bigl((fg_1)g_2)\Bigr)(g) = (fg_1)(gg_2^{-1}) = f(gg_2^{-1}g_1^{-1}) = f\Bigl(g(g_1g_2)^{-1}\Bigr) = \Bigl(f(g_1g_2)\Bigr)(g).$$

So the left action $g \mapsto \pi_{g^{-1}}$ turns into a right action when it acts on the argument of a function.

We finish this subsection with a corollary of Theorem 7.2 that is our sole application of the theorem.

Corollary 7.3 *Let n be a dimension profile for Cayley bigraph, G, on a group, $\mathcal{G}$. Let L be a subgraph of G, let $\mathbb{F}$ be a field, and let $k \geq 0$ be an integer. Then for any $g \in \mathcal{G}$, we have*

$$\mathcal{M}(ng, L, G, \mathcal{G}, k) = \mathcal{M}(n, L, G, \mathcal{G}, k)g,$$

where ng is given by

$$(ng)(P) = n(Pg^{-1})$$

for all $P \in V_G \amalg E_G$.

(We easily check that the action $g \to ng$ in this corollary is a right action, similar to the above discussion of the action on functions from $\mathcal{G}$ to a set, T.)

Proof Let $g \in \mathcal{G}$ and $M \in \mathcal{M}(n)$. Then there exists an $\mathcal{F} \subset \mathcal{K}_M$ such that $\dim(\mathcal{F}) = n$. Then we have $\mathcal{F}g \subset \mathcal{K}_M g$ and we have $\dim(\mathcal{F}g) = \dim(\mathcal{F})g$, since

$$\dim\Bigl((\mathcal{F}g)(P)\Bigr) = \dim\Bigl(\mathcal{F}(Pg^{-1})\Bigr)$$

for all $P \in V_G \amalg E_G$. But we have an isomorphism $\iota_g\colon \mathcal{K}_M g \to \mathcal{K}_{Mg}$ of sheaves on G; so on the one hand we have $\iota_g \mathcal{F} \subset \mathcal{K}_{Mg}$, and on the other hand, since isomorphisms preserve the dimension profile, we have $Mg \in \mathcal{M}(n')$ where

$$n' = \dim(\iota_g \mathcal{F}g) = \dim(\mathcal{F}g) = ng.$$

Hence $M \in \mathcal{M}(n)$ implies that $Mg \in \mathcal{M}(ng)$. Applying this observation to M replaced with Mg and g replaced with g^{-1} (or simply reversing the argument in this proof) shows the converse. Hence $\mathcal{M}(n)g = \mathcal{M}(ng)$.

$\square$

7.2 Generic Maximum Excess

If $\mathbb{F}$ is a field and $r \geq 1$ an integer, then by a *generic* subset of $\mathbb{F}^r$ we mean a subset that contains

$$\{(x_1, \ldots, x_r) \in \mathbb{F}^r \mid p(x_1, \ldots, x_r) \neq 0\}$$

for some nonzero polynomial, p. Algebraic geometry and generic subsets are most commonly discussed (at least on the most basic level) under the assumption that $\mathbb{F}$ is algebraically closed. Under this situation, all generic sets are nonempty; this remains true if $\mathbb{F}$ is infinite, or if the polynomial, p, above is fixed and $\mathbb{F}$ is finite but sufficiently large.

In order to have a sensible definition of generic and to conform to the algebraic geometric literature, we will freely assume that $\mathbb{F}$ is algebraically closed. However, the theorems we obtain in this section and the next will be valid for any infinite field or "sufficiently large" finite field, $\mathbb{F}$, by applying these theorems to the algebraic closure of $\mathbb{F}$, finding the associated polynomials, p, to the generic sets of interest, and determining how large $\mathbb{F}$ needs to be so that the generic sets are nonempty. The reader may find it amusing to note that in all our discussion of generic sets and generic conditions, all that we ultimately care about is that certain of these generic sets are nonemtpy (e.g., that there is at least one ρ-kernel for $(L, G, \mathcal{G})$ with vanishing maximum excess).

Let us review some notation in algebraic geometry; see [Har77], Chapter 1, Section 1. Let us assume that $\mathbb{F}$ is algebraically closed. Let $\mathbb{A}^N = \mathbb{A}^N(\mathbb{F})$, where N is an integer or a set or a product thereof, denote affine N space over $\mathbb{F}$, i.e., the set $\mathbb{F}^N$, with its usual Zariski topology. (When we speak of topological notions on $\mathbb{F}^N$ we mean those of $\mathbb{A}^N(\mathbb{F})$; in the literature $\mathbb{A}^N(\mathbb{F})$ connotes $\mathbb{F}^N$ viewed as a topological space, or scheme, etc.) Recall that a *locally closed* set is the intersection of an open and closed set (i.e., a subset of $\mathbb{A}^N$ determined as the zeros of some polynomials and complement of the zeros of some other polynomials), and a *constructible* set on $\mathbb{A}^N$ amounts to a finite disjoint union of locally closed sets (see [Har77], Exercise II.3.18).

Lemma 7.4 *Let $\mathbb{F}$ be an algebraically closed field, $k \geq 0$ an integer, and L a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. For each $n\colon V_G \amalg E_G \to \mathbb{Z}_{\geq 0}$, $\mathcal{M}(n) = \mathcal{M}(n, L, G, \mathcal{G}, k)$ is a constructible set.*

Proof We introduce $|n|\,|\mathcal{G}|$ indeterminates as follows: for each $P \in V_P \amalg E_P$, and $i = 1, \ldots, n(P)$, let $x_{P,i}$ be a vector of indeterminates indexed on $\mathcal{G}$ (there are $|n|$ vector variables $x_{P,i}$, for a total of $|n|\,|\mathcal{G}|$ indeterminates). We note that $M \in \mathcal{M}(n)$ precisely when one can find a solution for M and $x_{P,i}$ to the conditions

1. M is L-surjective; i.e., for each $P \in V_G \amalg E_G$, $\mathbb{F}^k$ is spanned by the columns of M corresponding to the elements of $\mathcal{G}_P(L)$;
2. for all P and i we have that $x_{P,i}$ has zero components outside of $\mathcal{G}_L(P)$;
3. for all P and i, $Mx_{P,i} = 0$;
4. for all P, $x_{P,1}, \ldots, x_{P,n_P}$ are linearly independent;
5. for all $e \in E_G$ and all i we have that $x_{e,i}, x_{te,1}, x_{te,2}, \ldots, x_{te,n_{te}}$ are linearly dependent, and similarly with he replacing te.

The dependence or independence or spanning of vectors reduces to the vanishing or nonvanishing of determinants of the vectors' coordinates. Hence all the above equations give us a collection of polynomials $f_i \in \mathbb{F}[M, x]$ (polynomials in the entries of M and the $x_{P,i}$'s) and $\widetilde{f}_j \in \mathbb{F}[M, x]$ such that $M \in \mathcal{M}(n)$ iff for some x we have $(M, x) \in C$, where C is the set of (M, x) for which $f_i(M, x) = 0$ for all relevant i and $\widetilde{f}_j(M, x) \neq 0$ for all relevant j; hence C is constructible. But $M \in \mathcal{M}(n)$ iff $(M, x) \in C$ for some x; hence $\mathcal{M}(n)$ is the image of C under the projection

$$\mathbb{A}^{k \times \mathcal{G}} \times \mathbb{A}^{|n| \times \mathcal{G}} \to \mathbb{A}^{k \times \mathcal{G}}.$$

But any projection from an affine space to another by omitting some of the coordinates has the property that it takes constructible sets to constructible sets (see Exercise II.3.19 of [Har77] or Theorem 3.16 of [Har92], noting that such a projection is both regular and of finite type). Hence $\mathcal{M}(n)$, the image of C, is constructible.

$\square$

We recall that a generic subset, S, of some affine space, $\mathbb{F}^T$, is a subset that contains a nonemtpy Zariski open subset of the space; if S is constructible, then S is generic iff its Zariski closure is the affine space.

Next we claim that $\mathcal{M}(n)$ is generic in $\mathbb{F}^{k\times\mathcal{G}}$ for at least one n, provided that $k \le \rho(L)$, and that $\mathcal{M}(n) = \emptyset$ for all but finitely many n. Indeed, for any totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ we have that M is L-surjective (for $(L, G, \mathcal{G}, \mathbb{F}, k)$), and the zero sheaf, $\mathcal{Z}$, has $\dim(\mathcal{Z}) = 0$. Hence the Zariski closure of $\mathcal{M}(0)$ is $\mathbb{F}^{k\times\mathcal{G}}$. Furthermore, $\mathcal{K}_M(P)$, for any $P \in V_G \amalg E_G$, is of dimension at most $|\mathcal{G}| - k$; hence $\mathcal{M}(n) = \emptyset$ unless $|n(P)| \le |\mathcal{G}| - k$ for all $P \in V_G \amalg E_G$, and there are only finitely many such n.

Definition 7.5 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$, and let $\mathbb{F}$ be an algebraically closed field. Let $k \le \rho(L)$ be a non-negative integer. We say that $n\colon V_G \amalg E_G \to \mathbb{Z}_{\ge 0}$ is* generic for $(L, G, \mathcal{G}, \mathbb{F}, k)$ *if the Zariski closure of $\mathcal{M}(n)$ is $\mathbb{A}^{k\times\mathcal{G}}(\mathbb{F})$. We define the* generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ *to be the largest value of $-\chi(n)$ for which n is generic. We define n to be a* maximal profile *(respectively,* minimal profile*) of* $(L, G, \mathcal{G}, \mathbb{F}, k)$ *if n is generic, $-\chi(n)$ equals the generic maximum excess, and there is no $n' \ne n$ which is generic with $-\chi(n') = -\chi(n)$ and $n'(P) \ge n(P)$ (respectively $n'(P) \le n(P)$) for all $P \in V_G \amalg E_G$.*

Theorem 7.6 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and $k \le \rho(L)$ an integer. There is a unique maximal profile and a unique minimal profile for $(L, G, \mathcal{G}, \mathbb{F}, k)$. Furthermore, if n is either the maximal or minimal profile, and $P \in V_G \amalg E_G$, then $ng = n$ for all $g \in \mathcal{G}$ (in the notion of Corollary 7.3). In particular $-\chi(n)$ is divisible by $|\mathcal{G}|$.*

Actually, the proof below shows that the theorem is still true when $k > \rho(L)$, provided that L has at least k edges of each colour, so that a totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ is L-surjective.

Proof Let n_1, n_2 be two maximal profiles for $(L, G, \mathcal{G}, \mathbb{F}, k)$. Let us show that $n_1 = n_2$. Consider the subset, S, of $\mathbb{F}^{k\times\mathcal{G}}$, M, such that $M \in \mathcal{M}(n_i)$ for $i = 1, 2$ and $\mathrm{m.e.}(\mathcal{K}_M) = -\chi(n_1) = -\chi(n_2)$. Clearly

$$S = \mathcal{M}(n_1) \cap \mathcal{M}(n_2) \cap \bigcap_{n \text{ s.t. } -\chi(n) > -\chi(n_1)} \overline{\mathcal{M}(n)},$$

where $\overline{\mathcal{M}(n)}$ denotes the complement of $\mathcal{M}(n)$. But if $-\chi(n) > -\chi(n_1)$ then, by assumption, n is not generic, and hence S is the intersection of a finite number of generic subsets of $\mathbb{F}^{k\times\mathcal{G}}$; hence S is a generic subset of $\mathbb{F}^{k\times\mathcal{G}}$, as well. But any element, $M \in S$, has subsheaves $\mathcal{F}_1, \mathcal{F}_2$, of $\mathcal{K}_M$ which obtain the maximum excess of $\mathcal{K}_M$ and with $\dim(\mathcal{F}_i) = n_i$ for $i = 1, 2$. But then $\mathcal{F} = \mathcal{F}_1 + \mathcal{F}_2$ also achieves the maximum excess and has $\deg(\mathcal{F}) \geq n_i$ for $i = 1, 2$. Hence

$$S \subset \bigcup_{n \text{ s.t. } -\chi(n)=-\chi(n_1),\ n\geq N} \mathcal{M}(n),$$

where $N = \max(n_1, n_2)$. Since the union on the right-hand-side is a finite union of constructible sets, the closure of one of these sets is $\mathbb{F}^{k\times\mathcal{G}}$. Hence there is an n with $-\chi(n) = -\chi(n_1)$ and $n \geq N = \max(n_1, n_2)$ such that n is generic; but if $n_1 \neq n_2$, then n does not equal either of them and is at least as big as either, which contradicts the maximality of the n_i, $i = 1, 2$. Hence $n_1 = n_2$, and the maximal profile is unique.

We argue similarly for the minimal profile, replacing $\mathcal{F}_1+\mathcal{F}_2$ with $\mathcal{F}_1\cap\mathcal{F}_2$.

Let n be the maximal profile for $(L, G, \mathcal{G}, \mathbb{F}, k)$ (now known to be unique). Since $\mathcal{M}(n)$ is a generic subset of $\mathbb{F}^{k\times\mathcal{G}}$, so is $\mathcal{M}(ng) = \mathcal{M}(n)g$ for any $g \in \mathcal{G}$. But then ng is also a maximal profile, since clearly $\chi(n) = \chi(ng)$ and $|n| = |ng|$. Hence $n = ng$ for all $g \in \mathcal{G}$. The same is true of the minimal profile.

It follows that the maximal (or minimal) profile, n, is invariant under $\mathcal{G}$, and hence has the same value on all the vertices, on all the edges of colour 1, and on all the edges of colour 2. Hence $-\chi(n)$ is divisible by $|\mathcal{G}|$ for the maximal (or minimal) profile, and hence the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ is divisible by $|\mathcal{G}|$.

$\square$

8 Variability of k-th Power Kernels

The main goal of this section is to prove the following theorem.

Theorem 8.1 *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and let $k \leq \rho(L)$ be a positive integer. Then the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ is at most that of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$, and we have equality iff both excesses are zero.*

As a consequence we get the following theorem.

Theorem 8.2 *Let L be a subgraph of a Cayley bigraph, G, on a group $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field, and let $k \le \rho(L)$ be a positive integer. Let L' be obtained from L by removing a single edge. Then the generic maximum excess of $(L', G, \mathcal{G}, \mathbb{F}, k-1)$ is at least that of $(L, G, \mathcal{G}, \mathbb{F}, k)$.*

(As before, this theorem is also true if $k > \rho(L)$, provided that L has at least k edges of each colour, so that a totally independent $M \in \mathbb{F}^{k\times\mathcal{G}}$ is L-surjective.)

A second goal of this section is to establish some general relations between kernels $\mathcal{K} = \mathcal{K}_M(L)$ as M and L vary. We shall derive two interesting, short exact sequences. First we establish a short exact sequence

$$0 \to \mathcal{K}_M(L) \to \mathcal{K}_{M'}(L) \to \underline{\mathbb{F}} \to 0, \tag{20}$$

for any $M \in \mathbb{F}^{k\times\mathcal{G}}$ and M' obtained from M by deleting the last row. Second we establish a short exact sequence

$$0 \to \mathcal{K}_{M'}(L') \to \mathcal{K}_{M'}(L) \to \mathcal{E} \to 0, \tag{21}$$

with L, L' as in Theorem 8.2, $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ such that $\mathcal{K}_{M'}(L')$ exists (i.e., M' is L'-surjective), and $\mathcal{E}$ a sheaf with $\mathcal{E}(V) = 0$ and $\mathcal{E}(E)$ of dimension $|\mathcal{G}|$.

Equation (21) will be used along with Theorem 7.6 to show that Theorem 8.1 implies Theorem 8.2.

Theorem 8.1 will not be proven with short exact sequences, but rather with a careful analysis of the unique minimal maximizer of the excess of $\mathcal{K}_M(L)$ and of that of $\mathcal{K}_{M'}(L)$. The sequence in equation (20) gives a relationship between $\mathcal{K}_M(L)$ and $\mathcal{K}_{M'}(L)$, but we don't know how to directly use this to conclude anything interesting about the two sheaves, such as the result of Theorem 8.1.

At this point we will divide our discussion into subsections. In Subsection 8.1, we will discuss the exact sequences related to our proof. In Subsection 8.2 we give the main observation of how the maximum excess changes in passing to subsheaves, and give an intuitive reason why the generic maximum excess of $\mathcal{K}_{M'}(L)$, as above, should be strictly greater than that of $\mathcal{K}_M(L)$ provided that these numbers don't both vanish. In Subsection 8.3 we mimic the notation of Section 7 to include a discussion of $\mathcal{K}_{M'}(L)$ as above and make our arguments precise, finishing the proof of Theorem 8.1; this will easily yield Theorem 8.2.

8.1 Variability as Exact Sequences

Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. For any non-negative integer, $k \le \rho(L)$, we have that a generic $M \in \mathbb{F}^{k\times\mathcal{G}}$ gives rise to a short exact sequence

$$0 \to \mathcal{K}_M(L) \to \underline{\mathbb{F}_L\mathcal{G}} \to \underline{\mathbb{F}}^k \to 0. \tag{22}$$

First we considering the variance of this equation in M; in other words, fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ such that

$$\underline{M}\iota\colon \underline{\mathbb{F}_L\mathcal{G}} \to \underline{\mathbb{F}}^k$$

is surjective. Then we have an exact sequence given in equation (22); if $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ is M with its last row deleted, we have a similar exact sequence

$$0 \to \mathcal{K}_{M'}(L) \to \underline{\mathbb{F}_L\mathcal{G}} \to \underline{\mathbb{F}}^{k-1} \to 0. \tag{23}$$

Notice that this discussion and everything below will remain essentially the same if, more generally, M' is taken to be M followed by any surjective map $\mathbb{F}^k \to \mathbb{F}^{k-1}$. In any event, we get a digram:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] & \underline{\mathbb{F}_L\mathcal{G}} \arrow[r, "\underline{M\iota}"] \arrow[d, "\simeq"] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{K}_{M'}(L) \arrow[r] \arrow[u, dashed, <-] & \underline{\mathbb{F}_L\mathcal{G}} \arrow[r, "\underline{M'\iota}"] & \underline{\mathbb{F}}^{k-1} \arrow[r] & 0
\end{tikzcd}$$

where the dotted arrow from $\mathcal{K}_{M'}(L)$ to $\mathcal{K}_M(L)$ is inferred from the solid arrows; furthermore, given that the solid horizontal arrows consist of an isomorphism and epimorphism, we infer that the dotted arrow is a monomorphism.

We then complete the diagram to obtain a diagram

$$\begin{tikzcd}
 & & & 0 \arrow[d] & \\
 & 0 \arrow[d] & 0 \arrow[d] & \underline{\mathbb{F}} \arrow[d] & \\
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] \arrow[d] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "\underline{M\iota}"] \arrow[d] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d] & 0 \\
0 \arrow[r] & \mathcal{K}_{M'}(L) \arrow[r] \arrow[d] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r, "\underline{M'\iota}"] \arrow[d] & \underline{\mathbb{F}}^{k-1} \arrow[r] \arrow[d] & 0 \\
 & \mathcal{K}_{M'}(L)/\mathcal{K}_M(L) \arrow[d] & 0 & 0 & \\
 & 0 & & &
\end{tikzcd}$$

A simple diagram chase shows that the nonzero upper right sheaf, $\underline{\mathbb{F}}$, and the nonzero lower left sheaf, $\mathcal{K}_{M'}(L)/\mathcal{K}_M(L)$, are isomorphic. Hence we obtain the short exact sequence in equation (20).

An analogous exact sequence can be obtained by varying L in equation (22). Let $L' \subset L$ be a subgraph of L. Fix an $M \in \mathbb{F}^{k\times\mathcal{G}}$ that induces a surjection $\underline{\mathbb{F}}_{L'}\mathcal{G} \to \underline{\mathbb{F}}^k$. Then we get a diagram:

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{K}_M(L') \arrow[r] \arrow[d, dashed] & \underline{\mathbb{F}}_{L'}\mathcal{G} \arrow[r] \arrow[d] & \underline{\mathbb{F}}^k \arrow[r] \arrow[d, "\simeq"] & 0 \\
0 \arrow[r] & \mathcal{K}_M(L) \arrow[r] & \underline{\mathbb{F}}_L\mathcal{G} \arrow[r] & \underline{\mathbb{F}}^k \arrow[r] & 0
\end{tikzcd}$$

Since $\underline{\mathbb{F}}_{L'}\mathcal{G} \to \underline{\mathbb{F}}_L\mathcal{G}$ is an injection, and the last vertical arrow is an isomorphism, the inferred dotted arrow is an injection. We therefore add a bottom row to the diagram and infer from the 3×3 Lemma that

$$\mathcal{K}_M(L)/\mathcal{K}_M(L') \simeq (\underline{\mathbb{F}}_L/\underline{\mathbb{F}}_{L'})\mathcal{G}.$$

In particular, if L' is obtained from L by removing m edges, then we infer equation (21) (with M here replaced by M', and k implicit here replaced by $k-1$), where $\mathcal{E}$ is a sheaf with $\mathcal{E}(V) = 0$ and $\mathcal{E}(E)$ being of dimension $m|\mathcal{G}|$.

8.2 Maximum Excess and Subsheaves

The goal of this section is to explain the main idea we will use to prove Theorem 8.1; the formal proof will be given in the subsection after this one.

The following theorem gives a number of ways of demonstrating whether or not a sheaf and one of its subsheaves have the same maximum excess.

Theorem 8.3 *Let $\mathcal{F}' \subset \mathcal{F}$ be sheaves on a digraph, G. Let $U \subset \mathcal{F}(V)$ be the minimal maximizer of the excess of $\mathcal{F}$, and let $U' \subset \mathcal{F}'(V)$ be the minimal maximizer of the excess of $\mathcal{F}'$. Then*

$$\text{m.e.}(\mathcal{F}') \leq \text{m.e.}(\mathcal{F}), \tag{24}$$

with equality iff $U = U'$ and

$$\Gamma_{\text{ht}}(\mathcal{F}, U') = \Gamma_{\text{ht}}(\mathcal{F}', U').$$

We already know equation (24) is true, since the maximum excess is a quasi-Betti number; the novelty of this theorem is that we have a simple condition to characterize when equality holds.

Proof Since $U' \subset \mathcal{F}(V)$ and

$$\Gamma_{\text{ht}}(\mathcal{F}', U') \subset \Gamma_{\text{ht}}(\mathcal{F}, U'),$$

we have that

$$\text{m.e.}(\mathcal{F}') = \text{excess}(\mathcal{F}', U') \leq \text{excess}(\mathcal{F}, U') \leq \text{m.e.}(\mathcal{F});$$

hence equality holds in equation (24) iff

$$\text{excess}(\mathcal{F}', U') = \text{excess}(\mathcal{F}, U') = \text{m.e.}(\mathcal{F}).$$

The first equality holds iff

$$\Gamma_{\text{ht}}(\mathcal{F}, U') = \Gamma_{\text{ht}}(\mathcal{F}', U').$$

The second equality holds iff U' is also a maximizer for $\mathcal{F}$. But since U is the minimal maximizer for $\mathcal{F}$, this implies that $U \subset U'$; but this means that $U \subset U' \subset \mathcal{F}'(V)$, so U is a maximizer for $\mathcal{F}'$, and hence $U' \subset U$ (since U' is the minimal maximizer for $\mathcal{F}'$). Hence $U = U'$.

□

Theorem 8.3 gives us a number of ways to conclude that equation (24) holds with strict inequality in certain situations. For example, consider a subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$, and consider a value, k for which

$$\text{m.e.}\Big(\mathcal{K}_M(L)\Big) > 0$$

for a generic $M \in \mathbb{F}^{k\times\mathcal{G}}$. Let M' be obtained from M by removing its bottom row, and consider the minimal maximizer, $U = U(M') \subset \mathcal{K}_{M'}(L)(V)$ of the excess of $\mathcal{K}_{M'}(L)$. We have that $\mathcal{K}_M \subset \mathcal{K}_{M'}$, and hence

$$\text{m.e.}\Big(\mathcal{K}_M(L)\Big) = \text{m.e.}\Big(\mathcal{K}_{M'}(L)\Big)$$

implies that $U(M')$, which is generically nonzero, lies entirely in $\mathcal{K}_M(L)(V)$. But if $w \in \mathcal{K}_{M'}(L)(V)$ is any nonzero vector supported on $v \in V_G$, then we may identify w with the corresponding element of $\mathcal{K}_{M'}(L)(v)$, and so

$$w \in (\underline{\mathbb{F}}_L\mathcal{G})(v) \simeq \bigoplus_{g\in\mathcal{G}_L(v)} \mathbb{F}_g,$$

where $\mathbb{F}_g$ denotes a copy of $\mathbb{F}$. In other words, w is given by its $\mathcal{G}$ components, which are (zero outside of $\mathcal{G}_L(v)$ and are) elements of $\mathbb{F}$. Hence, if we add a generic extra row to M' on the bottom, to form M, the row, $\vec{m} = (m_g)_{g\in\mathcal{G}}$ will (generically in $\vec{m}$) satisfy

$$\sum_{g\in\mathcal{G}} w_g m_g \neq 0. \tag{25}$$

Hence $w \notin \mathcal{K}_M(L)(V)$ generically, and therefore the minimal maximizers for $\mathcal{K}_M(L)$ and $\mathcal{K}_{M'}(L)$ will generically be different. Hence, by Theorem 8.3, we have

$$\text{m.e.}\Big(\mathcal{K}_{M'}(L)\Big) \geq 1 + \text{m.e.}\Big(\mathcal{K}_M(L)\Big)$$

for generic M (and M' obtained from M by deleting its bottom row). This argument will establish Theorem 8.1; all we need to do is to make this rigourous.

8.3 Proof of Theorems 8.1 and 8.2

In this subsection we precisely state the idea in the last subsection as Theorem 8.5 and use it to prove Theorem 8.1. We then easily conclude Theorem 8.2.

Let $\mathbb{F}$ be a field, $\mathcal{G}$ a group, and $k \geq 1$ an integer. If $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ and $\vec{m} \in \mathbb{F}^{\mathcal{G}}$, we define $\text{merge}(M', \vec{m})$ to be the element of $\mathbb{F}^{k\times\mathcal{G}}$ whose first $k-1$ rows consist of M' and whose k-th row consists of $\vec{m}$.

Definition 8.4 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Let $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ be a matrix, for some integer $k \geq 1$, that is L-surjective and whose kernel, $\mathcal{K}_{M'}(L)$, has nonzero maximum excess. Define the* redundancy *of M', denoted* $\mathrm{redund}(M')$*, to be the subset of $\mathbb{F}^{\mathcal{G}}$ consisting of $\vec{m}$ such that $M = \mathrm{merge}(M', \vec{m})$ is L-surjective, and such that*

$$\mathrm{m.e.}\Bigl(\mathcal{K}_{M'}(L)\Bigr) = \mathrm{m.e.}\Bigl(\mathcal{K}_{M}(L)\Bigr).$$

Theorem 8.5 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field and k a positive integer. Let $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ be a matrix that is L-surjective, and whose kernel, $\mathcal{K}_{M'}(L)$, has nonzero maximum excess. Then the redundancy of M' lies in a proper subspace of $\mathbb{F}^{\mathcal{G}}$.*

Proof This follows the argument of the last subsection. If U is the minimal maximizer of $\mathcal{K}_{M'}(L)$, then U is nonzero because the maximum excess is nonzero. Hence there exists a $w \in U$ supported at $v \in V_G$ with $w \neq 0$. So if $\mathrm{merge}(M', \vec{m})$ is L-surjective, we have $w \notin \mathcal{K}_M(L)$ if equation (25) holds. Since $w \neq 0$, equation (25) holds for all $\vec{m}$ outside of a proper subspace of $\mathbb{F}^{\mathcal{G}}$.

□

Proof of Theorem 8.1 If the generic maximum excesses were equal, then for a nonempty Zariski open subset, U, of $\mathbb{F}^{k\times\mathcal{G}}$, we would have for all $M \in U$ the maximum excess of $\mathcal{K}_M(L)$ is the same as that of $\mathcal{K}_{M'}(L)$, where M' is obtained from M by discarding its bottom row. Since U is nonempty and Zariski open, we have a polynomial, $p = p(M', \vec{m})$ such that

$$p(M', \vec{m}) \neq 0$$

implies that $(M', \vec{m}) \in U$. Write

$$p(M', \vec{m}) = \sum_{n \in \mathbb{Z}^{\mathcal{G}}_{\geq 0}} q_n(M')\vec{m}^n,$$

and fix any n such that $q_n \neq 0$. Then $q_n(M') \neq 0$ for all $M' \in U'$ for a nonempty Zariski open subset, U', of $\mathbb{F}^{(k-1)\times\mathcal{G}}$. For any fixed $M' \in U'$ we have $p(M', \vec{m})$ is a nonzero polynomial in $\vec{m}$; hence for fixed $M' \in U'$ we have that $(M', \vec{m}) \in U$ for a Zariski open subset of $\vec{m}$ in $\mathbb{F}^{\mathcal{G}}$.

On the other hand, assuming that the maximum excesses in Theorem 8.1 are not both zero, the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$ is positive. Hence $\mathcal{K}_{M'}(L)$ has positive maximum excess for all M' in some nonempty, Zariski open subset, U'', of $\mathbb{F}^{(k-1)\times\mathcal{G}}$. But by Theorem 8.5, for any $M' \in U''$ we have that $(M', \vec{m}) \notin U$ for $\vec{m}$ outside of a proper subspace of $\mathbb{F}^{\mathcal{G}}$. But U' and U'' must intersect (being two nonempty, Zariski open subsets of an irreducible variety), which gives a contradiction.

□

Proof of Theorem 8.2 Let the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, k)$ be m_k, that of $(L, G, \mathcal{G}, \mathbb{F}, k-1)$ be m_{k-1}, and that of $(L', G, \mathcal{G}, \mathbb{F}, k-1)$ be m'_{k-1}. Since $k \leq \rho(L)$ and hence $k-1 \leq \rho(L')$ (we can see $\rho(L') \geq \rho(L) - 1$ from equation (6)), we have that m_k, m_{k-1}, m'_{k-1} are all multiples of $|\mathcal{G}|$. The theorem is immediate if $m_k = 0$, so we may assume $m_k > 0$. In this case Theorem 8.1 implies that $m_{k-1} > m_k$, and since these numbers are both multiples of $|\mathcal{G}|$, we have

$$m_{k-1} \geq m_k + |\mathcal{G}|. \tag{26}$$

But the exact sequence in equation (21) shows that for any $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ we have

$$\text{m.e.}\Big(\mathcal{K}_{M'}(L')\Big) \geq \text{m.e.}\Big(\mathcal{K}_{M'}(L)\Big) - |\mathcal{G}|. \tag{27}$$

Let U, U', respectively, are the subsets of $M' \in \mathbb{F}^{(k-1)\times\mathcal{G}}$ at which $\mathcal{K}_{M'}(L), \mathcal{K}_{M'}(L')$, respectively, attain their generic value; hence U, U' are generic, and therefore so is $U \cap U'$. Then applying equation (27) to any $M' \in U \cap U'$ implies that

$$m'_{k-1} \geq m_{k-1} - |\mathcal{G}|.$$

Combining this with equation (26) gives $m'_{k-1} \geq m_k$, which proves the theorem.

□

9 Proof of the SHNC

In this section we prove the SHNC. At this point we have all the tools we need, except for one small detail.

Lemma 9.1 *Let L be an arbitrary étale bigraph with $\rho(L) > 0$. Then there exists an edge, $e \in E_L$, such that the graph, L', obtained by removing e from L has $\rho(L') = \rho(L) - 1$.*

Proof For each $F \in E_L$ let L_F denote the subgraph of L obtained by removing the edges in F from L. It is easy to see that for each $e \in E_L$ we have that $\rho(L_{\{e\}})$ is either $\rho(L)$ or $\rho(L)-1$; this can be seen from equation (6), since removing e from its connected component of L leaves h_1 the same or reduces it by one; alternatively, we can see this from the exact sequence

$$0 \to \underline{\mathbb{F}}_{L_{\{e\}}} \to \underline{\mathbb{F}}_L \to \underline{\mathbb{F}}_{L/L_{\{e\}}} \to 0,$$

using the fact that $L/L_{\{e\}}$ is (edge supported and) of maximum excess one.

For any two subgraphs, L', L'', of L we have an exact sequence

$$0 \to \underline{\mathbb{F}}_{L'\cap L''} \to \underline{\mathbb{F}}_{L'} \oplus \underline{\mathbb{F}}_{L''} \to \underline{\mathbb{F}}_{L'\cup L''} \to 0.$$

Hence

$$\rho(L' \cap L'') \geq \rho(L') + \rho(L'') - \rho(L' \cup L'').$$

Taking F', F'' to be disjoint subsets of E_L, we see that if $\rho(L_{F'}) = \rho(L_{F''}) = \rho(L)$, then setting $L' = L_{F'}$, $L'' = L_{F''}$ yields

$$\rho(L_{F'\cup F''}) \geq \rho(L),$$

and so $\rho(L_{F'\cup F''}) = \rho(L)$. Hence, if $\rho(L_F) = \rho(L)$ for all $F \subset E_L$ of size one, then by induction we can show this holds for $F \subset E_L$ of any size, which is impossible (since removing all the edges of a graph leaves it with $\rho = 0$). Hence there is at least one $e \in E_L$ for which $\rho(L_{\{e\}}) = \rho(L) - 1$.

□

Of course, one can give a purely graph theoretic proof of Lemma 9.1; we now sketch such a proof. From equation (6), it suffices to show that if L is connected with $h_1(L) \geq 2$ then we can remove an edge from L and reduce h_1 by one. We claim that it suffices to take any edge that remains after we repeatedly prune the leaves of L.

Definition 9.2 *Let L be a subgraph of a Cayley bigraph, G on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Then by the* generic maximum excess of the ρ-kernel of type $(L, G, \mathcal{G}, \mathbb{F})$ *we mean the generic maximum excess of $(L, G, \mathcal{G}, \mathbb{F}, \rho(L))$.*

Theorem 9.3 *Let L be a subgraph of a Cayley bigraph, G on a group, $\mathcal{G}$. Let $\mathbb{F}$ be an algebraically closed field. Then the generic maximum excess of the ρ-kernel of type $(L, G, \mathcal{G})$ is zero.*

Proof Fix G and $\mathcal{G}$ and let us prove the theorem for all L by induction on $\rho(L)$. The base case $\rho(L) = 0$ follows by definition, since the exact sequence

$$0 \to \mathcal{K} \to \underline{\mathbb{F}}_L \mathcal{G} \to \underline{\mathbb{F}}^0 \to 0$$

implies

$$\text{m.e.}(\mathcal{K}) \leq \text{m.e.}(\underline{\mathbb{F}}_L \mathcal{G}) = \sum_{g \in \mathcal{G}} \rho(Lg) = 0$$

(since $\rho(Lg) = \rho(L) = 0$ for all $g \in \mathcal{G}$). The inductive step of our induction on $\rho(L)$ is immediate from Theorem 8.2 applied to any L' obtained from L by removing a single edge so that $\rho(L') = \rho(L) - 1$; the existence of such an L' is given by Lemma 9.1.

□

Proof of Theorem 1.1, the SHNC By the graph theoretic reformulation of the SHNC, it suffices to show Theorem 3.1. By Theorem 5.2 it suffices to show that any subgraph, L, of a Cayley bigraph, G, on a group, $\mathcal{G}$, is universal for the SHNC. But by Theorem 9.3, there exists a ρ-kernel, $\mathcal{K} = \mathcal{K}_M(L)$ for $(L, G, \mathcal{G}, \mathbb{F})$ with vanishing maximum excess, for any algebraically closed $\mathbb{F}$. Hence we apply Theorem 6.3 to conclude that L is universal for the SHNC.

□

10 Concluding Remarks

We finish this paper with a few concluding remarks.

In this paper we have made no explicit reference to homology theories. In [Fri] we have used the twisted homology to prove that the maximum excess is a first quasi-Betti number; hence the theorems in this paper ostensibly rely on homology theories. However, we think it quite possible that one may able to prove that the maximum excess is a first quasi-Betti number directly, or give a direct proof of the inequalities we made use of in this paper. For example, if $\mathcal{F}' \to \mathcal{F}$ is a monomorphism, then since the maximum excess is a first quasi-Betti number we know that

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

But this inequality is clear from the subsheaf formulation of maximum excess in Theorem 4.6.

We remark that we first proved the SHNC using twisted homology theory, and then rewrote our proofs to use only maximum excess. In fact, twisted homology theory offers some additional intuition regarding the maximum excess. Twisted homology theory shows that (after pulling back appropriately, see [Fri]), the maximum excess can be interpreted as the dimension of a certain vector space of "twisted harmonic one-forms" of the sheaf. When this dimension is $d > 0$, one can impose d' linear conditions on the twisted harmonic one-forms and still have a $d - d'$ dimensional space of one-forms. This is how we view the variance in L of $\mathcal{K}_{M'}(L)$, as in the exact sequence of equation (21): to take a space of one-forms on $\mathcal{K}_{M'}(L)$ and obtain a one-form in $\mathcal{K}_{M'}(L')$, one has to impose $|\mathcal{G}|\,|E_L \setminus E_{L'}|$ conditions on the one-forms, namely the conditions that they vanish on the edges in $L\mathcal{G}$ that do not lie in $L'\mathcal{G}$. Of course, one has to pullback by an appropriate covering map to make this rigourous (see [Fri]), but all the edge counts and dimension counts scale appropriately under any covering.

The k-th power kernels in this paper are subsheaves of the constant sheaf $\underline{\mathbb{F}\mathcal{G}} \simeq \underline{\mathbb{F}}^{\mathcal{G}}$. We believe that subsheaves of constant sheaves satisfy some stronger properties than general sheaves, regarding their homological invariants and maximum excess. It would nice to study this further.

References

[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.

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

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

[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.

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

[GM03] Sergei I. Gelfand and Yuri I. Manin. *Methods of homological algebra*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, second edition, 2003.

[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.

[Har92] Joe Harris. *Algebraic geometry*, volume 133 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1992. A first course.

[How54] A. G. Howson. On the intersection of finitely generated free groups. *J. London Math. Soc.*, 29:428–434, 1954.

[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. Available as `http://arxiv.org/abs/math/0302009`.

[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.

[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.

[Neu56] Hanna Neumann. On the intersection of finitely generated free groups. *Publ. Math. Debrecen*, 4:186–189, 1956.

[Neu57] Hanna Neumann. On the intersection of finitely generated free groups. Addendum. *Publ. Math. Debrecen*, 5:128, 1957.

[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.

[Wis05] Daniel T. Wise. The coherence of one-relator groups with torsion and the Hanna Neumann conjecture. *Bull. London Math. Soc.*, 37(5):697–705, 2005.

A An Alternate View of ρ-kernels

In this section we give a direct proof that the vanishing generic maximum excess of ρ-kernels for all subgraphs, L, of any Cayley graph, G, implies the SHNC. We shall not use exact sequences. We shall require a few definitions, and some calculations to follow. While this gives some extra intuition about ρ-kernels, this section is not essential to the proof of the SHNC; we shall omit some of the easy but tedious graph theoretic details.

Definition A.1 *Let L be a subgraph of a Cayley bigraph, G, on a group, $\mathcal{G}$. As usual, for $P \in V_G \amalg E_G$, let $\mathcal{G}_L(P)$ be the set of $g \in \mathcal{G}$ such that Lg contains P. By a* vertex family *on $(L, G, \mathcal{G})$ we mean a function, $\mathcal{U}$, from V_G to $\mathcal{P}(\mathcal{G})$, the power set (i.e., set of subsets) of $\mathcal{G}$, such that for all $v \in V_G$ we have $\mathcal{U}(v) \subset \mathcal{G}_L(v)$. Similarly, an* edge family *on $(L, G, \mathcal{G})$ is a function $\mathcal{W} \colon E_G \to \mathcal{P}(\mathcal{G})$ such that $\mathcal{W}(e) \subset \mathcal{G}_L(e)$ for all $e \in E_G$. A vertex family, $\mathcal{U}$, and edge family, $\mathcal{W}$, are* compatible *if for all $e \in E_G$ we have*

$\mathcal{W}(e) \subset \mathcal{U}(te) \cap \mathcal{U}(he)$. *Given a vertex family,* $\mathcal{U}$, *the* induced edge family, $\mathcal{U}_E$, *of* $\mathcal{U}$ *is the edge family* $\mathcal{U}_E$ *given by*

$$\mathcal{U}_E(e) = \mathcal{U}(te) \cap \mathcal{U}(he) \cap \mathcal{G}_L(e).$$

The following lemmas motivate the above definitions; we omit their proofs, which are almost immediate.

Lemma A.2 *To each vertex family,* $\mathcal{U}$ *on* $(L, G, \mathcal{G})$, *and compatible edge family,* $\mathcal{W}$, *on* $(L, G, \mathcal{G})$, *there is a subgraph* $H \subset L \times_{B_2} G$ *determined via*

$$V_H = \{(vg^{-1}, v) \mid g \in \mathcal{U}(v)\}, \qquad E_H = \{(eg^{-1}, e) \mid g \in \mathcal{W}(e)\};$$

conversely, any subgraph $H \subset L \times_{B_2} G$ *arises from a unique vertex family,* $\mathcal{U}$, *and compatible edge family,* $\mathcal{W}$.

Lemma A.3 *For any vertex family,* $\mathcal{U}$, *on* $(L, G, \mathcal{G})$ *and compatible edge family,* $\mathcal{W}$, *we have*

$$\mathcal{W}(e) \subset \mathcal{U}_E(e).$$

In other words, $\mathcal{U}_E$ *is the "largest" edge family compatible with* $\mathcal{U}$.

Definition A.4 *Let* L *be a subgraph of a Cayley bigraph,* G, *on a group,* $\mathcal{G}$, *let* $M \in \mathbb{F}^{\rho(L)\times\mathcal{G}}$ *be totally linearly independent, and let* $\mathcal{K} = \mathcal{K}_M$ *be the resulting* ρ*-kernel. By a* straight *subspace of* $\mathcal{K}(V)$ *we mean a subspace*

$$U = \sum_{v \in V_G} U(v) \in \mathcal{K}(V),$$

such that for each $v \in V_G$, *we have*

$$U(v) = \mathrm{Free}_{\mathcal{U}(v)} \tag{28}$$

for some $\mathcal{U}(v) \subset \mathcal{G}$, *with notation as in Definition 6.7.*

Our goal for the rest of this section is to prove the following theorem.

Theorem A.5 *Let* L *be a subgraph of a Cayley graph,* G, *on a group,* $\mathcal{G}$. *The following conditions are equivalent:*

1. *for all* $L' \subset G$, *the SHNC holds for* (L, L');

2. *for every vertex family, $\mathcal{U}$, on $(L, G, \mathcal{G})$ we have*

$$\sum_{e\in E_G} |\mathcal{U}_E(e)|_{\rho(L)} \le \sum_{v\in V_G} |\mathcal{U}(v)|_{\rho(L)};$$

   and

3. *for some or any field, $\mathbb{F}$, and some or any totally independent $M \in \mathbb{F}^{\rho(L)\times\mathcal{G}}$, every straight subspace of $\mathcal{K}(V)$ with $\mathcal{K} = \mathcal{K}_M(L)$ has excess zero.*

We know by Theorem 5.2 that the SHNC holds iff it holds for all pairs (L, L') that are subgraphs of a Cayley graph, G. Hence Theorem 9.3, that implies condition (3) of this theorem, implies the SHNC.

Proof Conditions (2) and (3) are easily seen to be equivalent via equation (28) and equation (14).

If $\mathcal{U}$ is any vertex family on $(L, G, \mathcal{G})$, let the *positive set* of $\mathcal{U}$ consist of those $v \in V_G$ for which $|\mathcal{U}(v)| > \rho(L)$ and of those $e \in E_G$ for which $|\mathcal{U}_E(e)| > \rho(L)$. We easily see that the positive set forms a subgraph, L', of G, and that the pairs (Pg^{-1}, P) such that P is in the positive set and g lies in $\mathcal{U}(P)$ or $\mathcal{U}_E(P)$ (as is appropriate), forms a subgraph, H, of $L \times_{B_2} L'$. We see that

$$\begin{aligned} -\chi(H) &= \sum_{e\in E_{L'}} |\mathcal{U}_E(e)| - \sum_{v\in V_{L'}} |\mathcal{U}(v)| \\ &= \rho(L)\chi(L') + \sum_{e\in E_{L'}} \Big(|\mathcal{U}_E(e)| - \rho(L)\Big) - \sum_{v\in V_{L'}} \Big(|\mathcal{U}(v)| - \rho(L)\Big) \\ &= \rho(L)\chi(L') + \sum_{e\in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v\in V_G} |\mathcal{U}(v)|_{\rho(L)}. \end{aligned}$$

Hence we may write

$$-\chi(H) - \rho(L)\chi(L') = \sum_{e\in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v\in V_G} |\mathcal{U}(v)|_{\rho(L)}. \tag{29}$$

This equation is the main ingredient in the equivalence of conditions (1) and (3). Let us now state some graph theoretic lemmas that will firmly establish this equivalence.

Lemma A.6 *For any digraphs $H \subset G$, we have*

$$-\chi(H) \le \rho(G);$$

and equality holds if H consists of all connected components, X, of G with $h_1(X) > 0$ and any of those with $h_1(X) = 0$.

Proof The statement about equality is clear from the definition of ρ in equation (6). The inequality can be established graph theoretically by induction on the number of vertices and edges in G that are not in H. Alternatively, see the end of Section 4.

$\square$

Lemma A.7 *Let $\mathcal{U}$ be a vertex family for $(L, G, \mathcal{G})$, where L is a subgraph of a Cayley digraph, G, on a group, $\mathcal{G}$. Assume that*

$$\sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)} > 0. \tag{30}$$

Then there is a vertex family, $\mathcal{U}'$, which satisfies this inequality with $\mathcal{U}$ replaced by $\mathcal{U}'$, for which the positive set of $\mathcal{U}'$, L', satisfies $-\chi(L') = \rho(L')$.

Proof For any subgraph, $Y \subset G$ and vertex family $\mathcal{W}$ on $(L, G, \mathcal{G})$, set

$$f(\mathcal{W}, Y) = \sum_{e \in E_Y} |\mathcal{W}_E(e)|_{\rho(L)} - \sum_{v \in V_Y} |\mathcal{W}(v)|_{\rho(L)}.$$

Then clearly we have

$$f(\mathcal{U}, L') = \sum_{X \in \mathrm{conn}(L')} f(\mathcal{U}, X),$$

where $\mathrm{conn}(L')$ is the set of connected components of L'. But equation (30) says that $f(\mathcal{U}, G) > 0$, and clearly $f(\mathcal{U}, L') = f(\mathcal{U}, G)$. Hence we have $f(\mathcal{U}, X) > 0$ for some connected component, X, of L'; fix any such X.

We claim that $\rho(X) = -\chi(X)$. Since X is connected, this is true unless $\chi(X) = 1$; so it suffices to show that $\chi(X) = 1$ is impossible. If $\chi(X) = 1$, then by repeatedly pruning the leaves of X, i.e., deleting a vertex of degree one and its incident edge from X, we arrive at an isolated vertex. But if Y is any subgraph of G with a vertex, $v \in V_Y$, of degree one, and incident edge $e \in E_Y$, and if Y' is Y with v and e discarded, we claim that $f(\mathcal{U}, Y') \geq f(\mathcal{U}, Y)$; indeed, $\mathcal{U}_E(e) \subset \mathcal{U}(v)$, so

$$f(\mathcal{U}, Y') = f(\mathcal{U}, Y) + |\mathcal{U}(v)|_{\rho(L)} - |\mathcal{U}(e)|_{\rho(L)} \geq f(\mathcal{U}, Y).$$

Hence, by repeatedly pruning X we are left with X'' that is a single vertex with no edges, so $f(\mathcal{U}, X'') \geq f(\mathcal{U}, X) > 0$. But clearly $f(\mathcal{U}, X'') \leq 0$ for X''

consisting of a single vertex. Hence $\chi(X) = 1$ is impossible, and so $\chi(X) \leq 0$ and hence $\rho(X) = -\chi(X)$.

For any vertex family, $\mathcal{V}$ of $(L, G, \mathcal{G})$ and any subgraph $Y \in G$ define a vertex family $\mathcal{V}|_Y$ via

$$\mathcal{V}|_Y(v) = \begin{cases} \mathcal{V}(v) & \text{if } v \in V_Y, \\ \emptyset & \text{otherwise.} \end{cases}$$

for all $v \in V_G$. We easily see that

$$\mathcal{V}(e) \subset (\mathcal{V}|_L)_E(e) \tag{31}$$

for all $e \in E_L$. Hence

$$f(\mathcal{V}|_Y, G) = f(\mathcal{V}|_Y, Y) \geq f(\mathcal{V}, Y),$$

using equation (31). In particular, for $\mathcal{V} = \mathcal{U}$ and $Y = X$ we have

$$f(\mathcal{U}|_X, G) \geq f(\mathcal{U}, X) > 0.$$

So we take $\mathcal{U}' = \mathcal{U}|_X$ and let L' be its positive set. Then $f(\mathcal{U}', L') = f(\mathcal{U}', G) > 0$, and L' consists of X plus possibly some addition edges, so L' is connected and $\chi(L') \leq \chi(X) \leq 0$, so $\rho(L') = -\chi(L')$.

□

At this point condition (1) of Theorem A.5 easily implies condition (2). For if condition (2) does not hold, then for some $\mathcal{U}$, and with L' given as its positive set, we may assume $\rho(L') = -\chi(L')$ we have

$$\sum_{e \in E_G} |\mathcal{U}_E(e)|_{\rho(L)} - \sum_{v \in V_G} |\mathcal{U}(v)|_{\rho(L)} > 0$$

and hence

$$\rho(L \times_{B_2} L') \geq -\chi(H) > -\rho(L)\chi(L') = \rho(L)\rho(L').$$

Hence the SHNC is false on a pair of subgraphs of G.

It remains to show that condition (2) of Theorem A.5 implies condition (1). Again, we need some graph theoretic considerations.

Lemma A.8 *Assume the SHNC is false on a pair of subgraphs, (L, L'), of a Cayley bigraph G on a group $\mathcal{G}$. Then there is a subgraph, $L'' \subset L'$, such that*

1. *the SHNC is false on (L, L''),*
2. *L'' is connected,*
3. *$-\chi(L'') = \rho(L'')$, and*
4. *there is a subgraph, $H \subset L \times_{B_2} L''$ such that $-\chi(H) = \rho(L \times_{B_2} L'')$, and if $\mathcal{U}$ is the vertex family associated to H, then we have*

$$f(\mathcal{U}, G) > 0.$$

If condition (1) of Theorem A.5 is false, then the hypothesis of the above lemma holds; but item (4) of the lemma means that condition (2) of Theorem A.5 is false. Hence we conclude this section, and the proof of Theorem A.5 with the proof of the above lemma.

Proof So assume the SHNC is false on a pair of subgraphs, (L, L'). Similar to before, the SHNC must therefore be false on (L, L''), where L'' is some connected component of L. Fix such a connected component, L''.

We cannot have $\rho(L'') = 0$, for otherwise $\rho(L \times_{B_2} L'') = 0$ and the SHNC is not false on (L, L''). Hence we have $L'' \subset L'$ is connected and $\rho(L'') > 0$, whereupon we have $-\chi(L'') = \rho(L'')$.

Now let $L'' \subset L'$ be a minimal subgraph of L' (with respect to inclusion of subgraphs) with the properties that L'' is connected, $\rho(L'') > 0$, and the SHNC is false on (L, L''). We shall show that the lemma holds with this subgraph, L''; we have already established the first three items in the conclusion.

Then take any $H \subset L \times_{B_2} L''$ such that $-\chi(H) = \rho(L \times_{B_2} L'')$. We now make a number of claims regarding L'' that follow from the minimality of L''.

First, we claim that L'' has no leaves, i.e., no vertices of degree one. Otherwise, if v is a vertex of degree one, and e is its incident edge, then there are at least as many vertices in H over v as there are over v. So letting L''' be L'' with v and e discarded, we see that $\rho(L''') = \rho(L'')$; but if H' consists of the vertices and edges of H that do not lie over e or v, then H' is a subgraph of H and $\rho(H') = \rho(H)$ (since we obtain H' from H by discarding isolated vertices over v or vertices over v along with their single incident edges, lying over e). Hence the SHNC would fail also on (L, L'''), contradicting the minimality of L''.

Second, we claim that over each $e \in E_{L''}$ we there are at least $\rho(L)+1$ edges in H; if not, we delete e from L'', obtaining $L''' \subset L''$, and delete the at most $\rho(L)$ edges over e from H, obtaining H' that lies over L'''; this yields a strict subgraph, L''' of L'' such that

$$\begin{aligned}\rho(L \times_{B_2} L''') \geq -\chi(H') \geq -\chi(H) - \rho(L) = \rho(L \times_{B_2} L'') - \rho(L) \\ > \rho(L)\Big(\rho(L'') - 1\Big).\end{aligned}$$

But since L'' is pruned, we have $\rho(L''') = \rho(L'') - 1$. So, once again, we have the SHNC fails on (L, L''') for some a proper subgraph, L''', of L''; this contradicts the minimality of L''.

Third, we claim that over each $v \in V_{L''}$ there are at least $\rho(L)+1$ vertices in H. Indeed, if v is incident upon some edge, e, in L'', then e has at least $\rho(L)+1$ vertices in H above it, so v does as well. If v is isolated in L'', i.e., incident upon no edge, then L'' consists of only v, since L'' is connected; but this contradicts the fact that $\rho(L'') > 0$.

To H is associated a vertex family, $\mathcal{U}$, and an edge family, $\mathcal{W}$. According to the three claims established in the previous three paragraphs, we have

$$v \in V_{L''} \implies |\mathcal{U}(v)| \geq \rho(L)+1, \qquad e \in E_{L''} \implies |\mathcal{W}(e)| \geq \rho(L)+1.$$

Clearly also

$$v \notin V_{L''} \implies \mathcal{U}(v) = \emptyset, \qquad e \notin E_{L''} \implies \mathcal{W}(e) = \emptyset.$$

It follows that, as before

$$\begin{aligned}f(\mathcal{U}, G) \geq \sum_{e \in E_{L''}} |\mathcal{W}(e)|_{\rho(L)} - \sum_{v \in V_{L''}} |\mathcal{U}(e)|_{\rho(L)} \\ = -\chi(H) - \rho(L)(|E_{L''}| - |V_{L''}|) = \rho(L \times_{B_2} L'') - \rho(L)\rho(L''),\end{aligned}$$

using the fact that L'' is connected and $\rho(L'') > 0$ (so that $\rho(L'') = |E_{L''}| - |V_{L''}|$). Hence $f(\mathcal{U}, G) > 0$, which shows item (4) in the conclusion of the lemma.

□

□