

THE NON-BACKTRACKING SPECTRUM OF THE UNIVERSAL COVER OF A GRAPH

OMER ANGEL, JOEL FRIEDMAN, AND SHLOMO HOORY

ABSTRACT. A non-backtracking walk on a graph, H , is a directed path of directed edges of H such that no edge is the inverse of its preceding edge. Non-backtracking walks of a given length can be counted using the non-backtracking adjacency matrix, B , indexed by H 's directed edges and related to Ihara's Zeta function.

We show how to determine B 's spectrum in the case where H is a tree covering a finite graph. We show that when H is not regular, this spectrum can have positive measure in the complex plane, unlike the regular case. We show that outside of B 's spectrum, the corresponding Green function has "periodic decay ratios." The existence of such a "ratio system" can be effectively checked, and is equivalent to being outside the spectrum.

We also prove that the spectral radius of the non-backtracking walk operator on the tree covering a finite graph is exactly $\sqrt{\text{gr}}$, where gr is the growth rate of the tree. This further motivates the definition of the graph theoretical Riemann hypothesis proposed by Stark and Terras [ST].

Finally, we give experimental evidence that for a fixed, finite graph, H , a random lift of large degree has non-backtracking new spectrum near that of H 's universal cover. This suggests a new generalization of Alon's second eigenvalue conjecture.

1. INTRODUCTION

Let H be a finite, simple (i.e., without self-loops or multiple edges), connected graph. The main result of this paper is to describe the spectrum of operators on the universal cover, \tilde{H} , of H , that are "lifts" from operators on H . Let us briefly describe this result, saving precise definitions and statements for later sections.

A *cover* of H is a graph homomorphism $\pi : G \rightarrow H$ such that π induces a bijection between the edges incident on x and those incident on $\pi(x)$. Note that G may be infinite. The *universal cover* \tilde{H} of H is the unique (up to isomorphism) cover of H that is a tree. A cover of H may be thought of as a graph with the same local structure as H . The universal cover is in several senses an approximation to typical large covers of H .

A *connected local operator* on H is a matrix M indexed by H 's vertices, such that if $x \neq y$ then $M_{x,y} \neq 0$ iff $x \sim y$, where $x \sim y$ means that x is adjacent to y ; M may be viewed as an operator, in that for a function f on H we have a function Mf defined by $(Mf)_x = \sum_{y \sim x} M_{x,y} f_y$. Using the bijection on neighbourhoods induced by π , any such operator extends in a canonical way to a bounded operator \tilde{M} on $\ell^2(\tilde{H})$. (See Section 2 for more detailed definitions.) One of our main results is an algebraic description of the spectrum of \tilde{M} .

Date: November 2007.

This research was done while OA and SH were at UBC; OA held a PIMS postdoctoral fellowship.

Theorem 1.1. *Fix a finite graph, H , and symmetric, connected, local operator, M , on H . There is an explicit (in terms of H, λ) set, S , of polynomial equations in variables $\{r_e\}_{e \in E(H)}$ taking values in $\mathbb{C} \cup \{\infty\}$, such that $\widetilde{M} - \lambda I$ is invertible iff S has a solution with norm $\alpha(r) < 1$, where $\alpha(r)$ is the largest eigenvalue of a matrix whose entries are given via H as polynomials in the $\{r_e\}$.*

See Theorem 1.4 below for a statement including the algebraic representation. This result gives an algorithm for finding the spectrum of \widetilde{H} , and shows that it is determined by some algebraic curves.

A significant and motivating case of Theorem 1.1 is the non-backtracking spectrum of a graph. A non-backtracking walk on a graph, H , is a directed path in H such that no edge is the inverse of the preceding edge. The non-backtracking matrix of the graph allows counting such paths in the same way that the adjacency matrix of a graph is used to count paths in the graph. Non-backtracking walks appear in numerous applications and estimating their number is a fundamental problem. The matrix B is indexed by directed edges of H , with $B_{e,f} = 1$ iff (e, f) form a non-backtracking path of length 2. This may be seen as an asymmetric, local operator on the edge graph of H , and in this way extended to an operator on \widetilde{H} . This operator is asymmetric, but we can still use Theorem 1.4 to characterize \widetilde{B} 's spectrum. Note that while B and \widetilde{B} are real operators, they are not self-adjoint and so their spectra may be complex.

Theorem 1.2. *Given H, λ , there is an explicit set of polynomial equations in variables $\{r_e\}_{e \in E(H)}$ taking values in $\mathbb{C} \cup \{\infty\}$, such that $\widetilde{B} - \lambda I$ is invertible iff there is a solution with norm $\alpha(r) < 1$, where $\alpha(r)$ is explicitly given in terms of H .*

As an example, we consider the case where H is K_4 with one edge deleted. Already in that case, the spectrum of the non-backtracking path operator on \widetilde{B} has a non-trivial structure, see Figure 1. We are also interested in the relation between the non-backtracking spectrum of \widetilde{H} and that of a typical large but finite cover of H . Figure 2 shows the spectrum of a randomly chosen cover of this H . This and other numerical experiments give strong indication that the spectra converge to the spectrum of the universal cover (determined rigorously in this paper.)

1.1. Motivation. Non-backtracking walks seem to be an essential part of various aspects of graph theory, including the Broder-Shamir trace method ([BS87, Fri91, Fri93, Fri]), Ihara zeta functions of graphs ([Has89, KS00, ST]), Bollobas' question concerning girth and average degree ([AHL02, Hoo05]), and decoding of low density linear codes ([Gal63, Sho04, RU]). Given H , there is a matrix, $B = B(H)$, whose powers count non-backtracking walks of a given length. B is indexed by the directed edges of H , and is analogous to the adjacency matrix for paths (see below). Many interesting quantities concerning non-backtracking walks can be expressed in terms of spectral properties of B .

Several applications of walks in a graph H , require understanding the spectrum of the universal cover of H , denoted \widetilde{H} . Similarly, there are connections between non-backtracking walks on H and the spectrum of the non-backtracking walk operator on \widetilde{H} .




FIGURE 1. The non-backtracking spectrum of the universal cover of K_4 minus an edge.




FIGURE 2. The non-backtracking spectrum of a random 300-lift of K_4 minus an edge. The shaded area is the non-backtracking spectrum of its universal cover; the circled points are the spectrum for the base graph.

For regular graphs, i.e., graphs in which each vertex has the same degree, there is a direct and simple relationship between the spectrum of B and that of the adjacency matrix, $A = A(H)$. This fact renders many questions quite easy for regular graphs, masking formidable difficulties that may occur for non-regular graphs. In the case

of a d -regular graph H , the universal cover is a d -regular tree, and the spectrum of $\tilde{A} = A(\tilde{H})$ is known to be the interval

$$\sigma(\tilde{A}) = \left[-2\sqrt{d-1}, 2\sqrt{d-1}\right].$$

In this case, the spectrum of $\tilde{B} = B(\tilde{H})$ is the union of a circle and two real intervals in the complex plane. Numerical experiments, such as those of Stark and Terras (see [ST] and the previous papers in that series), suggest that \tilde{B} , for general (non-regular) H , can be two dimensional. As an example, when H is K_4 minus an edge, the non-backtracking spectrum of \tilde{H} is depicted in Figure 1.

The main goal of this paper is to give a method for finding the spectrum of \tilde{B} for any finite graph, H . This will prove the above noted two-dimensionality. A consequence of this method is that $\rho(\tilde{B}) = \sqrt{\text{gr}}$, i.e., the spectral radius of \tilde{B} is the square root of the growth rate of \tilde{H} . This further motivates the definition of the graph theoretical Riemann hypothesis proposed by Stark and Terras [ST].

Alon's conjecture. Alon's conjecture states that for any $d \geq 3$ and any $\varepsilon > 0$, the probability that a random d -regular graph on n vertices has second largest eigenvalue $\lambda_2 \leq 2\sqrt{d-1} + \varepsilon$ tends to 1 as $n \rightarrow \infty$. This was proven for various notions of "random graph," i.e., various probability distributions on d -regular graphs on n vertices, in [Fri] (see also [BS87, Fri91]). A main tool used to this end is the analysis of non-backtracking walks in random, regular graphs.

As we shall see (and see [Fri]), for finite d -regular graphs each $\mu \in \sigma(A)$ corresponds to two $\lambda \in \sigma(B)$ given by the solutions of $\lambda^2 - \mu\lambda + d - 1 = 0$. In particular, since $\sigma(A(H)) \subset [-d, d]$, the non-backtracking spectrum is contained in the union of the circle $\{z : |z| = \sqrt{d-1}\}$ and the two real segments $[1-d, -1]$ and $[1, d-1]$. This simple relation yields a form of Alon's conjecture in terms of the non-backtracking spectrum, namely that the spectral radius, $\rho(B)$, of B is at most $\sqrt{d-1} + \varepsilon$.

In [Fri03], Alon's conjecture is relativized, meaning that for any graph, H (the "base graph"), we consider a random lift, $\pi: G \rightarrow H$, of degree n ; eigenpairs of $A(H)$ pull back to eigenpairs of $A(G)$, we call such eigenpairs of $A(G)$ *old* eigenpairs of π , while eigenpairs of $A(G)$ not arising from $A(H)$ we call *new* eigenpairs of π ; we conjecture that for any fixed H and $\varepsilon > 0$, most random lifts have all new eigenvalues of absolute value bounded above by ε plus the spectral radius of $A(\tilde{H})$; and a weaker new eigenvalue bound is proved via the Broder-Shamir technique (Alon's original conjecture is the case where H has one vertex). The full conjecture might be provable with a refined Broder-Shamir method (as in [Fri91, Fri]) that seems to require a translation between the adjacency matrix spectrum and non-backtracking spectrum. However, for irregular graphs we are aware of no such translation. This makes it difficult to apply the trace method in the general case.

Computer experiments indicate that the new (i.e., not arising from the base) non-backtracking spectrum of large random lifts is distributed in some fixed region, as seen in Figure 2. We would like to prove that this region is the non-backtracking spectrum of the universal cover of the base graph, shown in Figure 1. This suggests a generalization of the relative Alon conjecture, roughly stating that the new non-backtracking spectrum of a large random lift is near the non-backtracking spectrum of the universal cover; for non-regular graphs, unlike the case of regular graphs, it is

not clear how this relates to the usual relative Alon conjecture regarding adjacency matrix spectrum. This might allow us to apply the trace methods of [BS87, Fri93, Fri] in order to prove some analog of Alon's conjecture for lifts of general base graphs (perhaps first for the non-backtracking spectrum and then, ultimately, on the adjacency matrix spectrum).

The relativized Alon conjecture raises the questions of how to determine the non-backtracking spectrum of the universal cover of a finite graph (answered in this paper) and how can one translate spectral bounds between adjacency matrices and non-backtracking adjacency matrices in the irregular case.

The Ihara Zeta function. For a finite graph H , the eigenvalues of $B(H)$ are the reciprocals of the poles of the *Ihara zeta function* of the graph. This zeta function is a graph theoretic analogue of the number theoretic Riemann zeta function. For further details, we refer the reader to Hashimoto [Has89], Kotani and Sunada [KS00], and to Stark and Terras [ST].

Bounds for graphs with large girth. A third motivation is the result of Alon, Hoory, and Linial [AHL02], giving a lower bound on the number of vertices in a graph with a specified girth and average degree. They answer affirmatively a longstanding open question by Bollobás. Their result is basically a lower bound on the Perron eigenvalue of B in terms of the average degree. In further work, Hoory [Hoo05], gave a similar bound on the spectral radius of H 's universal cover, $\rho(A(\tilde{H}))$.

Further questions raised by these works concern the relation between \tilde{H} 's growth rate $\text{gr} = \rho(B)$, and the spectral radii of the remaining three operators A , \tilde{A} , and \tilde{B} . In this paper we partially answer this question by proving that $\rho(\tilde{B}) = \sqrt{\text{gr}}$.

LDPC error correcting codes. A final motivation is the analysis of a decoding algorithm for a family of error correcting codes, known as low density parity check (LDPC) codes. This decoding algorithm, which is receiving a lot of attention at present, was first suggested by Gallager [Gal63]. LDPC codes are binary linear codes of length n , described by a set of highly sparse linear equations over $GF(2)$. It is convenient to represent such a code by a bipartite graph, where the two sides correspond to variables and equations. Each equation is adjacent to the variables occurring in it. There is a natural decoding algorithm for such codes based on message passing. At each iteration of the algorithm, a message is sent along each directed edge of the graph (alternating between variables to equations and equations to variables rounds). We refrain from further description of this algorithm, and refer the reader to a survey by Shokrollahi [Sho04], or a book by Richardson and Urbanke [RU].

The connection of this algorithm to non-backtracking walks is in the fact that the k -th generation message computed at edge e depends only on the initial value at the endpoints of all length k non-backtracking walks starting at e . It seems that to analyze such an algorithm, one needs a non-backtracking analog of expansion. For example, the Mixing Lemma for expanders states that the number of edges $e(S, T)$ between any two vertex sets S, T in a d -regular graph satisfies $|e(S, T) - d|S||T|/N| \leq \lambda_2 \sqrt{|S||T|}$, where λ_2 is the second largest eigenvalue in absolute value

(see [HLW] for a survey on expander graphs). It would be helpful to develop a non-backtracking analogue of the expander Mixing Lemma.

1.2. Main results. In the following unless otherwise specified, all graphs will be *simple*, meaning without self-loops or multiple edges, and connected; in Section 9 we explain our conventions and generalizations to graphs that are not simple. Throughout this paper, if G is a graph, then $V(G)$ denotes G 's vertices, and $E(G)$ denotes G 's *directed edges*, which (for G simple) we define as the union of $\{(u, v), (v, u)\}$ over all u, v joined by an edge of G ; for $e = (u, v) \in E(G)$ we set $e^{-1} = (v, u)$; we write $u \sim v$ to denote that there is an edge joining u and v . (All graphs in this paper are locally finite, meaning that each vertex is incident upon a finite number of edges.)

Let $\pi: T \rightarrow H$ be the universal cover of a finite graph, H . For an operator, M , on $l^2(V(T))$, we set $M_{uv} = \langle \delta_u, M\delta_v \rangle$ for $u, v \in V(T)$ (where δ denotes the Dirac delta function); we say that M is

- (1) *symmetric* if $\overline{M}^* = M$;
- (2) *local* if $M_{uv} = 0$ for u, v of distance greater than one;
- (3) *connected* if $M_{uv} \neq 0$ for u, v of distance one;
- (4) *a pullback (from π or, abusively, H)* if M_{uv} depends only on $\pi(u)$ and $\pi(v)$ assuming $u \sim v$.

The main result of this paper gives a way to compute the spectrum of local, symmetric operators on T pulled back from H . If M has an inverse, we define the *Green function for M* to be $G(u, v) = \langle \delta_u, M^{-1}\delta_v \rangle$. If an operator A on $l^2(V(H))$ pulls back to π^*A on T , then determining the spectrum of π^*A amounts to determining for which λ the operator $M = M_\lambda = \pi^*A - I\lambda$ has a Green function.

The key idea in this paper to determine the Green function, G , of M , is that for any edge, $e = (u, v)$, of $E(T)$, the ratio $r_e = G(u, v)/G(u, u)$ should depend only on $\pi(e)$. This surprised us at first, although the intuitive explanation is simple. If $\pi(e) = \pi(e')$ for $e' = (u', v') \in E(T)$, then there is an automorphism of T taking e to e' ; using this automorphism we can “graft” appropriate linear combinations of pieces of $G(u, \cdot)$ and $G(u', \cdot)$ onto $G(u', \cdot)$; if $r_e \neq r_{e'}$, then this “grafting” process will yield more than one Green function, which is impossible (see Lemma 5.3).

There are numerous technical difficulties in turning this intuition into precise theorems. One is that an example of Figà-Talamanca and Steger (see Section 8) shows that $G(u, u)$ may vanish, forcing us to work with the more delicate situation of zero and infinite ratios. Another is that knowing M and a ratio system, r , while it is easy to determine G , it does not seem easy to know if M is invertible, at least when H is infinite.

To make a precise statement, if $a, b \in \mathbb{C}$ with at least one of a, b not vanishing, we interpret b/a to be ∞ if $a = 0$ and the usual ratio otherwise; in this case we say the *generalized ratio*, b/a , exists. If $u, v, w \in V(T)$ are vertices of a tree, T , we write $[u, w]$ for the unique path (a sequence alternating between vertices and edges) from u to w in T , and say that v is between u and w if v lies in $[u, w]$.

Definition 1.3. A *ratio system* for M is a function $r: E(H) \rightarrow \mathbb{C} \cup \{\infty\}$ satisfying the following:

(a) If $r_e \neq 0$ for some $e = (u, v)$ then

$$0 = m_{vv} + \frac{m_{e^{-1}}}{r_e} + \sum_{e' \text{ s.t. } e \rightarrow e'} m_{e'} r_{e'}, \quad (1)$$

with the convention that $c/\infty = 0$ for any c , and where $e \rightarrow e'$ means that $e' \neq e^{-1}$ but (e, e') forms a path of length two (i.e., e 's head is the tail of e').

(b) If $r_e = 0$ then there is some f with $e \rightarrow f$ and $r_f = \infty$.

(c) If $r_f = \infty$ and $e \rightarrow f$, then $r_e = 0$ and $r_{f^{-1}} \neq 0$.

Here condition (a) comes from the fact that for fixed u we have $MG(u, \cdot) = \delta_u$; conditions (b) and (c) are needed when the ratios become zero or infinite, i.e., when $G(u, u) = 0$ for some u . Note that there is no assumed relationship between r_e and $r_{e^{-1}}$.

For a ratio system we define r 's *decay rate*, $\alpha(r)$, to be the Perron eigenvalue of the matrix R whose rows and columns are indexed by the directed edges of H with nonzero (but possibly infinite) ratios, defined by:

$$R_{e,f} = \begin{cases} \sum_{e' \text{ s.t. } e \rightarrow e' \rightarrow f} |m_f/m_{e^{-1}}|^2 & r_f = \infty, \\ |r_f|^2 & r_f \neq \infty \text{ and } e \rightarrow f, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Theorem 1.4. Let $\pi: T \rightarrow H$ be a universal cover of a finite graph, H , and let M be a pulledback, symmetric, local operator on $l^2(V(T))$. Then M has a bounded inverse if and only if M has a ratio system $r: E(H) \rightarrow \mathbb{C} \cup \{\infty\}$ with $\alpha(r) < 1$.

Furthermore, for an edge (u, v) of type e and any w such that $u \in [w, v]$, we have $G(w, v)/G(w, u) = r_e$ whenever this generalized ratio is defined.

Theorem 1.4 has numerous applications, including to adjacency matrices and Laplacians. The application to non-backtracking spectrum of a graph, G , derives from its connection to the local operator

$$Q_\lambda = Q - \lambda A + \lambda^2 I$$

where $A = A_G$ is the adjacency matrix of G and $Q = Q_G$ is the diagonal matrix with entries $Q_{vv} = d_v - 1$ where d_v is the degree of v . Specifically we shall prove the following theorem for graphs, G , of bounded degree, regarding the spectrum, $\sigma(B)$, of $B = B(G)$.

Theorem 1.5. Let G be a graph with degrees bounded by d_{\max} , then

$$\sigma(B(G)) = \{\pm 1\} \cup \{\lambda : Q_\lambda \text{ is not invertible}\}.$$

For $\lambda = 0$, Q_λ is symmetric, and otherwise

$$\lambda^{-1} Q_\lambda = \lambda^{-1} Q - A + \lambda I$$

is (real) symmetric, and so Theorem 1.4 applies to Q_λ .

The following provides an interesting contrast to Theorem 1.4.

Theorem 1.6. Let $\pi: T \rightarrow H$ be as in Theorem 1.4. If r is a ratio system with $\alpha(r) = 1$, then M is not invertible.

Calculations with K_4 minus an edge, as in Figure 1, show that a given M can have more than one ratio system. Theorem 1.6 says that if there is a ratio system with $\alpha = 1$, then there are none with $\alpha < 1$.

The rest of this paper is divided as follows. In Section 2 we give some preliminary remarks. In Section 3 we discuss the non-backtracking spectra of trees and prove Theorem 1.5. In Section 4 we bound the non-backtracking spectrum in an annulus in terms of the growth rate of the tree. In Section 5 we discuss the Green function and relative kernels (that make precise the “grafting” idea discussed earlier showing that Green’s function has “periodic ratios”). In Section 6 we prove Theorem 1.4 in the case of a ratio system with no zero or infinite ratios. In Section 7 we prove Theorem 1.4 in full generality. In Section 8 we give some interesting examples illustrating various aspects of Theorem 1.4. In Section 9 we show how to extend our theory to much more general base graphs (or pregraphs), H . In Section 10 we make some remarks, mentioning a few of the many questions left open by this paper.

Acknowledgments. We wish to thank Chen Greif, Audry Terras, Alex Gamburd, and Bálint Virág for illuminating discussions.

2. PRELIMINARIES

A directed edge $e' = (u', v')$ of a graph, H , can follow edge $e = (u, v)$, denoted $e \rightarrow e'$, if $u' = v$ and $e' \neq e^{-1}$. A non-backtracking walk on H is a sequence of directed edges e_1, e_2, \dots, e_r such that $e_k \rightarrow e_{k+1}$ for all k .

If H is a graph and $v \in V(H)$, the *1-neighbourhood* of v is the subgraph consisting of all undirected edges incident upon v . A graph homomorphism $\nu: H' \rightarrow H$ is a *covering map* (see [Fri93]) if for each $v' \in V(H')$, ν gives a bijection of the edges of the 1-neighbourhood of v' with those of v (note that this definition is correct even when the graphs have self-loops; see Section 9). H' is a *cover* of H . The *type* of a vertex $v \in H'$, is defined as the vertex $\nu(v) \in H$. If $\nu(v') = v$ for vertices v', v , we say that v' is *above* v (or v is *below* v'). These terms are similarly defined for edges.

Note that a cover of a cover is a cover. The *universal cover* of H , denoted \tilde{H} , is the (unique up to isomorphism) cover of H that is also a cover of every other cover of H . The universal cover may be constructed as follows. Fix some (arbitrary) vertex $v_0 \in H$. The vertices of \tilde{H} are the finite non-backtracking walks on H starting at v_0 . Two vertices of \tilde{H} are adjacent if one extends the other by a single step. The covering map $\pi: V(\tilde{H}) \rightarrow V(H)$ maps a walk to its terminal vertex in H . It is easy to see that \tilde{H} is a tree, hence it is also referred to as the covering tree of H . The universality of \tilde{H} can be stated as follows: If $\nu: H' \rightarrow H$ is any covering map, and $v' \in V(H')$ is of type v_0 , then there is a unique covering map $\mu: \tilde{H} \rightarrow H'$ with $\pi = \nu\mu$ and $\mu(\phi) = v'$, where ϕ is the empty path.)

If the covering map $\nu: H' \rightarrow H$ is n -to-1 for some finite n , then we also say that H' is a *lifting*, or an *n -lifting*, of H . Every cover is either a lifting or is ∞ -to-1. Equivalently, the graph H' , which we denote H^n , can be characterized as follows: the vertex set of H^n consists of n copies of H ’s vertices: $V(H^n) = V(H) \times [n]$, with the covering map $\pi((v, i)) = v$. The edges in H^n are given by a perfect matching between $\pi^{-1}(v)$ and $\pi^{-1}(u)$ for every undirected edge $(u, v) \in E(H)$. This

characterization gives rise to a natural distribution on n -lifts of H , by picking the perfect matchings independently and uniformly at random.

We define 0-forms to be functions on the vertices of a graph, and 1-forms to be functions on directed edges of a graph (with no constraint relating $f(e)$ and $f(e^{-1})$.) Denote by $A = A(H)$ the adjacency operator acting on the Hilbert space of 0-forms, i.e., $(Af)(v) = \sum_{u \sim v} f(u)$ with multiple edges taken into account.) For graphs with minimal degree at least two, we consider the non-backtracking adjacency operator, $B = B(H)$, which acts on 1-forms. B 's action is given by

$$(Bf)(e) = \sum_{e' \text{ s.t. } e \rightarrow e'} f(e').$$

The definitions of A, B are valid for any locally finite graph, and in particular for \tilde{H} . When the degrees in a graph are globally bounded, A and B may be viewed as bounded operators on $l^2(V)$ and $l^2(E)$, respectively. In the finite case, B is the adjacency matrix of the directed graph with an edge from e_1 to e_2 iff $e_2 \rightarrow e_1$. Notice that in the literature B is occasionally defined as the adjoint of our definition.

The *non-backtracking spectrum* of a graph H is defined to be the spectrum of the corresponding operator B (as opposed to the spectrum of H which is the spectrum of A .) Note that A is self-adjoint, hence any graph's spectrum is real. However, the non-backtracking spectrum generally contains complex numbers. $\tilde{A} = A(\tilde{H})$ and $\tilde{B} = B(\tilde{H})$ denote the corresponding operators on \tilde{H} .

We denote by $\rho(u, v)$ the graph distance in \tilde{H} . We let gr denote the growth rate of the universal cover \tilde{H} , i.e.

$$\text{gr} = \lim_{r \rightarrow \infty} |\text{Ball}(v, r)|^{1/r},$$

where $\text{Ball}(v, r)$ is the radius r ball around some (arbitrary) vertex $v \in V(\tilde{H})$; for connected H the limit is independent of the choice of v . It is straightforward to see gr is the Perron-Frobenius eigenvalue of B .

3. NON-BACKTRACKING SPECTRUM BASICS

We start with some background from spectral theory. Let T be a bounded linear operator acting on a Hilbert space \mathcal{H} (in this paper we work with $l^2(X)$, where X will be one of $V, \tilde{V} = V(\tilde{H}), E, \text{ or } \tilde{E} = E(\tilde{H})$). Recall that the spectrum of T is

$$\sigma(T) = \{\lambda | T - \lambda I \text{ does not have a bounded inverse}\},$$

and its spectral radius is

$$\rho(T) = \max\{|\lambda| : \lambda \in \sigma(T)\}$$

(no confusion should arise regarding this ρ and the distance, ρ , that involves two arguments). If \mathcal{H} is finite dimensional, then $\lambda \in \sigma(T)$ iff there is a nonzero f satisfying $Tf = \lambda f$. In the general case we have [Rud91, Chapter 12]:

Theorem 3.1. *For a bounded linear operator $T: \mathcal{H} \rightarrow \mathcal{H}$, we have $\lambda \in \sigma(T)$ iff one of the following three possibilities hold:*

- (i) *there exists a nonzero f such that $(T - \lambda I)f = 0$,*
- (ii) *there exists a nonzero f such that $(T^* - \bar{\lambda}I)f = 0$,*
- (iii) *for every $\varepsilon > 0$ exists a nonzero f such that $\|(T - \lambda I)f\| < \varepsilon \|f\|$.*

In case (iii), a sequence of f 's with $\varepsilon \rightarrow 0$ is referred to as *approximate eigenfunctions*; case (i) is a special case of (iii).

Definition 3.2. A bounded linear operator T is called *symmetric* (sometimes *real symmetric* in the literature) if $\overline{T^*} = T$.

If T is symmetric, then case (ii) of Theorem 3.1 is equivalent to case (i), and so λ is in the spectrum iff case (iii) holds, i.e., we have a sequence of approximate eigenfunctions for λ .

While the operator B is not symmetric, to identify its spectrum it suffices to check only case (iii). To see this, define the unitary operator $U: l^2(E) \rightarrow l^2(E)$ by its action $(Uf)_e = f_{e^{-1}}$. Note that $U^{-1} = U$. Since B is real, it is easy to see that $U^{-1}BU = B^T = B^*$. Consequently, if f is an eigenfunction of B^* with eigenvalue $\overline{\lambda}$, then Uf is an eigenfunction of B with eigenvalue λ . Thus case (ii) of Theorem 3.1 holds for the same λ as case (i), and it suffices to check case (iii).

At this point we are ready to prove Theorem 1.5. Before doing so we make a few remarks on the theorem.

Note that if G is a d -regular graph, $Q = (d-1)I$ and so $Q_\lambda = (\lambda^2 + d - 1)I - \lambda A$ is singular iff $\lambda^2 - \mu\lambda + d - 1 = 0$ for some $\mu \in \sigma(A)$. Theorem 1.5 is known (in more precision) for finite graphs:

Theorem 3.3 (Bass [Bas92], Kotani and Sunada [KS00]). *For any finite graph*

$$\det(B - \lambda I) = (\lambda^2 - 1)^{|E|-|V|} \cdot \det(Q_\lambda).$$

Note 3.4. *A version of Theorem 3.3 holds for graphs with multiple edges and self-loops. Each half-loop (see Section 9) contributes a factor of $\lambda - 1$, rather than $\lambda^2 - 1$.*

We thank Chen Greif for the following remark:

Note 3.5. *Consider the $2n \times 2n$ matrix $X = \begin{pmatrix} A & -Q \\ I & 0 \end{pmatrix}$. One can easily verify by row elimination that $\det(Q_\lambda) = \det(X - \lambda I)$. This observation has practical importance for calculating the non-backtracking spectrum of specific graphs, since eigenvalue calculation is a standard feature in most numeric/symbolic matrix manipulation programs.*

Proof of Theorem 1.5. If G is finite, then Theorem 3.3 implies our result, hence we may restrict ourselves to infinite G . The result follows from the following three statements:

- (1) we have $\pm 1 \in \sigma(B)$;
- (2) if $\lambda \neq \pm 1$ and Q_λ is not invertible, then neither is $B - \lambda I$;
- (3) if $\lambda \neq \pm 1$ and $B - \lambda I$ is not invertible, then neither is Q_λ .

To prove (1), we explicitly construct approximate eigenfunctions. G contains some infinite path of directed edges $\{e_i\}_{i \geq 0}$. For any k , define the function

$$f = \sum_{i=1}^k \delta_{e_i} - \sum_{i=0}^{k-1} \delta_{e_i^{-1}},$$

i.e., f is 1 on e_1, \dots, e_k , f is -1 on $e_0^{-1}, \dots, e_{k-1}^{-1}$, and zero on all other edges. One can easily check that

$$(B - I)f = \delta_{e_0} - \delta_{e_0^{-1}} - \delta_{e_k} + \delta_{e_k^{-1}}.$$

Thus $\|(B - I)f\| = \sqrt{4}$, while $\|f\| = \sqrt{2k}$, and so taking $k \rightarrow \infty$ gives approximate eigenfunctions for $\lambda = 1$.

$\lambda = -1$ is similar, with

$$f = \sum_{i=1}^k (-1)^i \delta_{e_i} + \sum_{i=0}^{k-1} (-1)^i \delta_{e_{i-1}}.$$

To prove (2), assume Q_λ is not invertible. Since Q_λ is symmetric, as noted above it has approximate (perhaps exact) eigenfunctions. Thus for every ε there is an f such that $\|f\| = 1$ and $\|Q_\lambda f\| < \varepsilon$. Note that f is defined on the vertices of G . From f we construct a 1-form, g , given by $g_{uv} = f_u - \lambda f_v$. Then:

$$\begin{aligned} ((B - \lambda I)g)_{uv} &= \sum_{\substack{w \sim v \\ w \neq u}} g_{vw} - \lambda g_{uv} = \sum_{w \sim v} g_{vw} - g_{vu} - \lambda g_{uv} \\ &= \sum_{w \sim v} (f_v - \lambda f_w) + (\lambda^2 - 1)f_v \\ &= (d_v - 1)f_v - \lambda(Af)_v + \lambda^2 f_v = (Q_\lambda f)_v. \end{aligned}$$

Therefore,

$$\|(B - \lambda I)g\|^2 \leq d_{\max} \cdot \|Q_\lambda f\|^2 < \varepsilon^2 d_{\max}.$$

It remains to show that $\|g\|$ is bounded from below. To this end, note that $\begin{pmatrix} g_{uv} \\ g_{vu} \end{pmatrix} = \begin{pmatrix} 1 & -\lambda \\ -\lambda & 1 \end{pmatrix} \begin{pmatrix} f_u \\ f_v \end{pmatrix}$ is an invertible linear transformation for $\lambda \neq \pm 1$. Thus for $\lambda \neq \pm 1$, for some $c = c(\lambda) > 0$

$$|g_{uv}|^2 + |g_{vu}|^2 \geq c(|f_u|^2 + |f_v|^2).$$

Summing over all edges we find $\|g\|^2 \geq c$.

To prove (3), let S_o, S_i be the outbound and inbound sum operators from 1-forms to 0-forms, defined by $(S_o h)_v = \sum_{u \sim v} h_{vu}$, and $(S_i h)_v = \sum_{u \sim v} h_{uv}$. For any v , and any 1-form h :

$$\begin{aligned} (S_i B h)_v &= \sum_{u \sim v} (B h)_{uv} = \sum_{\substack{u, w \sim v \\ w \neq u}} h_{vw} = (d_v - 1)(S_o h)_v \\ (S_o B h)_v &= \sum_{u \sim v} (B h)_{vu} = \sum_{\substack{w \sim u \sim v \\ w \neq v}} h_{uw} = (A S_o h - S_i h)_v. \end{aligned}$$

And so we have

$$S_i B = Q S_o, \quad \text{and} \quad S_o B = A S_o - S_i.$$

Assume $B - \lambda I$ is not invertible for some $\lambda \neq \pm 1$. As before, since B is real, for any $\varepsilon > 0$ there is a function g such that $\|g\| = 1$, and $\|(B - \lambda I)g\| < \varepsilon$. We claim that $f = S_o g$ is an approximate eigenfunction of Q_λ . Indeed:

$$\begin{aligned} \|Q_\lambda f\| &= \|(\lambda^2 I - \lambda A + Q)S_o g\| \\ &= \|(\lambda^2 S_o - \lambda S_o B - \lambda S_i + S_i B)g\| \\ &= \|(S_i - \lambda S_o)(B - \lambda I)g\| \\ &\leq (1 + |\lambda|)\sqrt{d_{\max}} \varepsilon \end{aligned}$$

(since $\|S_o\| = \|S_i\| = \sqrt{d_{\max}}$). Furthermore, for $\lambda \neq \pm 1$ we shall bound $\|f\|$ from below. For any edge vu ,

$$(S_o g)_v = (B g)_{uv} + g_{vu} - \lambda g_{uv} + \lambda g_{uv};$$

notice that S_i^* is the ‘‘head map,’’ $(S_i^* h)_{uv} = h_v$ for all $h \in L^2(V)$, and similarly S_o^* is the ‘‘tail map.’’ So we may write the last equation as $S_i^* S_o g = (B - \lambda I)g + (\lambda I + U)g$ (recall $(Ug)_{uv} = g_{vu}$.) Since $\|g\| = 1$, $U^2 = 1$ and $\|U\| = 1$, we have

$$\begin{aligned} |\lambda^2 - 1| &= \|(\lambda^2 - 1)g\| = \|(\lambda - U)(\lambda + U)g\| \\ &\leq (|\lambda| + 1)\|(\lambda + U)g\|. \end{aligned}$$

Since clearly $\|h\|^2 \geq \|S_i^* h\|^2 / d_{\max}$ for any $h \in L^2(V)$, we have

$$\begin{aligned} \sqrt{d_{\max}}\|f\| &\geq \|S_i^* f\| = \|S_i^* S_o g\| \geq \|(\lambda I + U)g\| - \|(B - \lambda I)g\| \\ &\geq \frac{|\lambda^2 - 1|}{|\lambda| + 1} - \|(B - \lambda I)g\|. \end{aligned}$$

Since $\|(B - \lambda I)g\| < \varepsilon$, it follows that for ε small enough $\|f\|$ is bounded from below. \square

Note 3.6. *It is possible to prove Theorem 3.3 by a similar argument, with cycles taking the place of the infinite path, resulting in exact eigenfunctions; the eigenspace of 1 is spanned by functions which are 1 on edges of an oriented cycle and -1 on the reverse cycle. The eigenspace of -1 is similar but with alternating signs along cycles of even length.*

The following theorem restricts the non-backtracking spectrum of G to the union of an annulus in the complex plane, and two real intervals. For finite graphs, the theorem is due to Kotani and Sunada [KS00].

Theorem 3.7. *Let $d_{\min} \geq 2$ and d_{\max} be the minimal and maximal degrees of some (finite or infinite) graph G . Then*

$$\sigma(B) \subset \left\{ \lambda \in \mathbb{C} : \sqrt{d_{\min} - 1} \leq |\lambda| \leq \sqrt{d_{\max} - 1} \right\} \cup \left\{ \lambda \in \mathbb{R} : 1 \leq |\lambda| \leq d_{\max} - 1 \right\}.$$

Proof. We first prove that

$$\sigma(B) \subset \{ \lambda \in \mathbb{C} : 1 \leq |\lambda| \leq d_{\max} - 1 \}.$$

Fix some vertex v . Clearly UB acts on functions supported on edges directed out of v independently of its action on all other edges. On these d_v edges, the action of UB is given by the matrix $J - I$ where J is the all ones matrix. Since $\sigma(J - I) = \{-1, d_v - 1\}$ and since $J - I$ (being symmetric) is orthonormally diagonalizable, we easily see that

$$\frac{\|(J - I)f\|}{\|f\|} \in [1, d_v - 1]$$

for any function f supported on these edges; therefore, since U is unitary, for any 1-form, f , we have

$$\frac{\|Bf\|}{\|f\|} = \frac{\|UBf\|}{\|f\|} \in [1, d_{\max} - 1].$$

If $\lambda \in \sigma(B)$, then we have for any ε an approximate eigenfunction, f , with $\|f\| = 1$ and $\|(B - \lambda)f\| \leq \varepsilon$. So

$$|\lambda| = \|\lambda f\| \leq \|\lambda f - Bf\| + \|Bf\| \leq \varepsilon + d_{\max} - 1.$$

Taking $\varepsilon \rightarrow 0$ yields $|\lambda| \leq d_{\max} - 1$. Similarly

$$|\lambda| = \|\lambda f\| \geq \|Bf\| - \|Bf - \lambda f\| \geq 1 - \varepsilon,$$

and we conclude $\|\lambda\| \geq 1$.

It remains to prove that the non-real $\lambda \in \sigma(B)$ satisfy $\sqrt{d_{\min} - 1} \leq |\lambda| \leq \sqrt{d_{\max} - 1}$. Given $f \in L^2(V)$, define

$$\alpha = \frac{\langle f, Af \rangle}{\|f\|^2}, \quad \beta = \frac{\langle f, Qf \rangle}{\|f\|^2},$$

and note that $\langle f, Q_\lambda f \rangle = \|f\|^2 \cdot (\lambda^2 - \alpha\lambda + \beta)$. Further, note that since A, Q are real symmetric operators, that α, β are real. From the definition of A, Q , we clearly have $|\alpha| \leq d_{\max}$ and $\beta \in [d_{\min} - 1, d_{\max} - 1]$.

If f is an eigenfunction then $\lambda^2 - \alpha\lambda + \beta = 0$. If Q_λ has only approximate eigenfunctions, f_ε , then for those we have $\alpha = \alpha_\varepsilon, \beta = \beta_\varepsilon$ satisfying

$$|\lambda^2 - \alpha_\varepsilon\lambda + \beta_\varepsilon| = |\langle f, Q_\lambda f \rangle| \leq \varepsilon$$

and as before, $|\alpha_\varepsilon| \leq d_{\max}$ and $\beta_\varepsilon \in [d_{\min} - 1, d_{\max} + 1]$. By passing to a subsequence we may assume that α_ε and β_ε have limits, α, β , as $\varepsilon \rightarrow 0$. We find that, again, $\lambda^2 - \alpha\lambda + \beta = 0$ with real α, β satisfying $\beta \in [d_{\min} - 1, d_{\max} - 1]$.

Solving the quadratic equation we get $\lambda = \alpha/2 \pm \sqrt{\alpha^2/4 - \beta}$. If λ is not real then the square root is of a negative number and

$$|\lambda|^2 = (\alpha/2)^2 + \beta - \alpha^2/4 = \beta,$$

and the claim follows. \square

4. NON-BACKTRACKING SPECTRUM OF TREES

The non-backtracking spectrum of a tree has a fourfold symmetry that is not hard to rigorously prove:

Proposition 4.1. *For any bipartite graph, G , $\sigma(B(G))$ is invariant under reflection in the real and imaginary axis.*

As we will be primarily concerned with the non-backtracking spectrum of trees, note that this result applies whenever G is a tree.

Proof. Since (for any graph) B is a real operator, $\sigma(B)$ is invariant under complex conjugation.

For a bipartite graph G , let X, Y be the two sides of the graph. Consider the unitary operator N on 1-forms that negates f_{uv} for $u \in X$ and leaves f_{uv} unchanged for $u \in Y$. Clearly $N^{-1} = N$ and it is easy to see that $N^{-1}BN = -B$. Thus $\sigma(B) = -\sigma(B)$. \square

The next result improves on the upper bound of Theorem 3.7 for sufficiently nice trees. A graph has *uniformly bounded growth*, $\overline{\text{gr}}$, if

$$\overline{\text{gr}} = \limsup M_r^{1/r}$$

is finite, where

$$M_r = \sup_v |\text{Ball}(v, r)|.$$

Theorem 4.2. *Assume a tree without leaves has uniformly bounded growth $\overline{\text{gr}}$, then*

$$\sigma(B) \subset \{\lambda : 1 \leq |\lambda| \leq \sqrt{\overline{\text{gr}}}\}.$$

Note that for graphs with uniformly bounded degrees, $M_r = O((d_{\max} - 1)^r)$, and so $\overline{\text{gr}}$ is finite. For universal covers of finite graphs, $\overline{\text{gr}} = \text{gr} = \lim_{r \rightarrow \infty} |\text{Ball}(v, r)|^{1/r}$ for any fixed v . We shall later see that for universal covers of finite graphs, Theorem 4.2 is sharp and in fact $\{\lambda : |\lambda| = \sqrt{\overline{\text{gr}}}\} \subset \sigma(B)$.

Proof. Denote $e \rightarrow_k e'$ if there is a non-backtracking path $e \rightarrow e_1 \rightarrow \dots \rightarrow e_{k-1} \rightarrow e'$. Note that if such a path exists, then it is unique and its length is determined by e, e' . By definition,

$$\|B^k\| = \sup_{\|f\|=\|g\|=1} |\langle g, B^k f \rangle|.$$

For f, g with $\|f\| = \|g\| = 1$, and any collection of coefficients $a(e, e') > 0$, we have by Cauchy-Schwarz

$$\begin{aligned} |\langle g, B^k f \rangle| &= \left| \sum_{e' \rightarrow_k e} f(e)g(e') \right| \\ &\leq \sum_{e' \rightarrow_k e} a(e, e') |f(e)|^2 + a(e, e')^{-1} |g(e')|^2 \\ &\leq \left(\sup_e \sum_{e' \rightarrow_k e} a(e, e') \right) \|f\|^2 + \left(\sup_{e'} \sum_{e' \rightarrow_k e} a(e, e')^{-1} \right) \|g\|^2. \end{aligned}$$

It remains to choose the coefficients $a(e, e')$ so that the sums above are small. To this end, fix some (arbitrary) vertex in the tree to be the root. The unique path from e' to e must descend toward the root some number of steps i , and then ascend $k - i$ steps on a different branch to reach e . Fix $a(e, e') = \overline{\text{gr}}^{(i-k/2)}$.

For any $\varepsilon > 0$ there is some $c = c(\varepsilon)$, such that the total number of non-backtracking paths of length ℓ from any vertex is bounded by $c \cdot ((1 + \varepsilon)\overline{\text{gr}})^\ell$. Since there is always a unique path toward the root, the number of paths starting at e' with i, k as above is at most $c \cdot ((1 + \varepsilon)\overline{\text{gr}})^{k-i}$. Consequently

$$\begin{aligned} \sup_e \sum_{e' \rightarrow_k e} a(e, e') &\leq \sum_{i=0}^k c((1 + \varepsilon)\overline{\text{gr}})^{k-i} \overline{\text{gr}}^{(i-k/2)} \\ &\leq c(k+1)(1 + \varepsilon)^k \overline{\text{gr}}^{k/2}. \end{aligned}$$

Similarly, the number of paths of type i ending at e is bounded by $c_\delta((1 + \varepsilon)\overline{\text{gr}})^i$, and hence also

$$\sup_{e'} \sum_{e' \rightarrow_k e} a(e, e')^{-1} \leq c(k+1)(1 + \varepsilon)^k \overline{\text{gr}}^{k/2}.$$

Combining the bounds we find

$$\rho(B) = \lim (\|B^k\|)^{1/k} \leq (1 + \varepsilon)\sqrt{\text{gr}}.$$

Since ε was arbitrary, this proves the claim. \square

5. THE GREEN FUNCTION AND THE RELATIVE KERNELS

Let $\pi: T \rightarrow H$ and M as in Theorem 1.4. By deleting edges $e = \{u, v\}$ from $E(H)$ for which $m_{uv} = 0$, we may assume M is connected, i.e., $m_{uv} \neq 0$ whenever $(u, v) \in E(V)$. If $e = (u, v)$ is an edge of T or H , we write m_e for m_{uv} .

Assume that M is invertible. Then there is a unique *Green function*, $G = G_M$, defined by the requirement that $MG(x, \cdot) = \delta_x(\cdot)$, or equivalently,

$$G(u, v) = (M^{-1}\delta_u)(v).$$

Notice that the symmetry $M^* = \overline{M}$ implies the symmetry of the Green function,

$$\begin{aligned} G(u, v) &= \langle M^{-1}\delta_u, \delta_v \rangle = \langle \delta_u, (M^*)^{-1}\delta_v \rangle \\ &= \overline{\langle (\overline{M})^{-1}\delta_v, \delta_u \rangle} = \langle M^{-1}\delta_v, \delta_u \rangle = G(v, u). \end{aligned}$$

Proposition 5.1. *Let $\pi: T \rightarrow H$, M , and G , be as above. Let σ be an automorphism of T for which $\sigma^*M = M\sigma^*$, where σ^* is the pullback (i.e., $\sigma^*f = f \circ \sigma$). Then $G(x, y) = G(\sigma x, \sigma y)$ for all $x, y \in V(T)$.*

Proof. We have $\delta_{\sigma z} = \sigma^*\delta_z$ for any $z \in V(T)$, and so

$$\begin{aligned} G(\sigma x, \sigma y) &= \langle M^{-1}\delta_{\sigma x}, \delta_{\sigma y} \rangle = \langle M^{-1}\sigma^*\delta_x, \sigma^*\delta_y \rangle \\ &= \langle (\sigma^*)^{-1}M^{-1}\sigma^*\delta_x, \delta_y \rangle = \langle M^{-1}\delta_x, \delta_y \rangle = G(x, y). \end{aligned} \quad \square$$

We introduce some further notions in order to derive additional properties of the Green function. For a directed edge $e = (u, v)$ of the covering tree, consider the subtree T_e consisting of v 's connected component in $T - e$, together with e itself. Thus T_e contains exactly those vertices contained in a non-backtracking path from u that begins with e . T_e is considered to be rooted at u , which is of degree 1 in it. For any vertex x , the unique neighbour closer to the root referred to as x 's parent; the other neighbours are x 's children.

Definition 5.2. Let T_e be a tree as above, $e = (u, v)$. The *relative kernel of M with respect to e* , denoted $\ker_e M$, is the subspace of $l^2(V(T))$ consisting of functions f supported on T_e for which Mf is zero on $T_e \setminus \{u\}$. The *relative null-kernel*, denoted $\ker_e^0 M$, is the subspace of the relative kernel consisting of those functions, f , for which $f(u) = 0$.

It is clear that either $\ker_e M = \ker_e^0 M$ or else $\ker_e M / \ker_e^0 M$ is one-dimensional.

Lemma 5.3. *Let M be invertible and let $e = (u, v)$.*

- (i) *If $G(u, u) \neq 0$, then $\dim(\ker_e M) = 1$.*
- (ii) *If $G(u, u) = G(u, v) = 0$, then $\dim(\ker_e M) = 1$.*
- (iii) *If $G(u, u) = 0$ and $G(u, v) \neq 0$, then $\dim(\ker_e^0 M) = 1$.*

This lemma shows that $\ker_e M$ is one-dimensional or two-dimensional, with the latter only possible in the case where $G(u, u) = 0$ and $G(u, v) \neq 0$. We do not know of any operator where $\ker_e M$ is two-dimensional, or whether such an example exists. There is a class of M 's where there are edges for which $G(u, u) = 0$ and $G(u, v) \neq 0$ (see Section 8.1).

Proof. Assume first $G(u, u) \neq 0$. Consider the function

$$f(w) = \begin{cases} G(u, w) & w \in T_e \\ 0 & w \notin T_e \end{cases}.$$

Mf differs from $MG(x, \cdot) = \delta_u$ only at u and outside T_e , and therefore $f \in \ker_e M$. Since $f(u) \neq 0$, the dimension of $\ker_e M$ is at least 1. To show that the dimension is not larger than one, suppose that $g \in \ker_e M$, and define \tilde{g} via

$$\tilde{g}(w) = \begin{cases} G(u, u)g(w) & \text{if } w \in V(T_e), \\ G(u, w)g(u) & \text{otherwise.} \end{cases}$$

Thus \tilde{g} is proportional to g on T_e and to $G(u, \cdot)$ outside. The two definitions coincide at u . It follows that $M\tilde{g}$ is supported on u , and so \tilde{g} is proportional to $G(u, \cdot)$ and hence g is proportional to f .

Assume from here on that $G(u, u) = 0$. Consider the restriction of the $G(u, \cdot)$ to the subtrees emerging from u . Since the Green function is unique, $G(u, \cdot)$ must vanish on all but one of these trees (otherwise the restriction of $G(u, \cdot)$ to each subtree gives another function with $Mf = \delta_u$). In particular, there is a unique neighbour x of u with $G(x, u) \neq 0$.

If $G(u, v) = 0$, then we repeat the argument above with f the restriction to T_e of $G(x, \cdot)$. As above, $f \in \ker_e$. If there is another $g \in \ker_e$, then let

$$\tilde{g}(w) = \begin{cases} G(x, u)g(w) & \text{if } w \in V(T_e), \\ G(x, w)g(u) & \text{otherwise.} \end{cases}$$

This time, $M\tilde{g}$ is supported on $\{x, u\}$, and so $\tilde{g} = c_1G(x, \cdot) + c_2G(u, \cdot)$. Applying this at x and u gives

$$G(x, x)g(u) = c_1G(x, x) + c_2G(u, x), \quad G(x, u)g(u) = c_1G(x, u) + c_2G(u, u)$$

and solving for c_1, c_2 shows $c_2 = 0$ (and $c_1 = g(u)$). Thus g is proportional to f , and $\dim(\ker_e) = 1$.

Finally, if $x = v$, then $G(u, \cdot)$ vanishes outside T_e , and so $f = G(u, \cdot) \in \ker_e^0 M$ implying the dimension is at least one. Conversely, for any $g \in \ker_e^0 M$ the function Mg is supported on u and hence must be proportional to f , implying the dimension is exactly one. \square

Corollary 5.4. *If $e = (u, v)$ and x, y are not in $V(T_e) \setminus \{u\}$, then the restrictions of $G(x, \cdot)$ and $G(y, \cdot)$ to T_e are proportional.*

Proof. The cases $G(u, u) \neq 0$ and $G(u, u) = G(u, v) = 0$ follow from the one-dimensionality of $\ker_e M$. The remaining case $G(u, u) = 0$ and $G(u, v) \neq 0$ follows from the fact mentioned in the above proof that $G(u, \cdot)$ is supported on T_e , so $G(u, x) = G(u, y) = 0$. Hence $G(x, u) = G(u, x) = 0$. Thus the restrictions to T_e of $G(x, \cdot)$ and $G(y, \cdot)$ are both in the one-dimensional $\ker_e^0 M$. \square

6. GREEN FUNCTIONS WITH FINITE PERIODIC RATIOS

Let $\pi: T \rightarrow H$ and M be as in Theorem 1.4.

Definition 6.1. A *finite ratio system* for M is a ratio system, r , whose values, r_e , are never 0 or ∞ ; in other words, (1) always holds.

As a step toward proving Theorem 1.4, we first discuss the simpler case where all ratios are finite.

Theorem 6.2. *The following are equivalent:*

- (1) M is invertible and $G(u, u) \neq 0$ for all u ,
- (2) there is a finite ratio system for M with $\alpha(r) < 1$.

We start the proof of Theorem 6.2 with the assumption that M is invertible and $G(u, u) \neq 0$ for all u , and prove that there is a ratio system for M with $\alpha(r) < 1$. Indeed, for any edge $e_0 \in H$, take some edge $e = (u, v)$ of type e_0 and define

$$r_{e_0} = \frac{G(u, v)}{G(u, u)}.$$

For a path $\vec{p} = \{e_i\}_{i=0}^n$ in H , it will be convenient to denote $r(\vec{p}) = \prod r_{e_i}$, i.e., the product of r_e over \vec{p} , with repeated edges taken into account.

For an edge e , we consider the function f on T_e defined by $f = \frac{G(u, \cdot)}{G(u, u)}$. As noted above, $f \in \ker_e M$. The following relates f to the finite ratio system.

Claim 6.3. *Consider a non-backtracking path \vec{p} in H beginning with e_0 . Lift the path to \tilde{H} , starting at u and let y be the terminal vertex. Then $f(y) = r(\vec{p})$.*

Proof. We induct on the path's length. If $n = 0$ the claim states $f(u) = 1$, which is true. If $n = 1$, the claim states $f(v) = r_e$ which is also true by the choice of r . For the inductive step, let $e' = (x, y)$ be an edge of type e_k in T_e , with x closer to the root, then it suffices to show that $f(y) = r_{e_k} f(x)$ (this is why we call r a finite ratio system).

Since the restriction of f to $T_{e'}$ is in $\ker_{e'} M$, by Lemma 5.3 it is proportional to $G(x, \cdot)$, and so

$$\frac{f(y)}{f(x)} = \frac{G(x, y)}{G(x, x)} = r_{e_k}. \quad \square$$

To check that r is indeed a finite ratio system, note that $(Mf)(v) = 0$. However,

$$(Mf)(v) = m_{vv}r_e + m_{e^{-1}} + \sum_{e \rightarrow e'} m_{e'}r_e r_{e'},$$

which is just (1).

It remains to show that $\alpha(r) < 1$. Consider $S_n = \sum_{\rho(x, u)=n} |f(x)|^2$, i.e., the contribution to $\|f(x)\|^2$ from all x at level n of T_e . The matrix R is constructed so that S_n is the square of the l^2 -norm of the row corresponding to e_0 in R^n (this is proved by induction using Claim 6.3).

By Perron-Frobenius theory, there is an e_0 and $c > 0$ such that $S_n \geq c\alpha(r)^n$ for n large (see e.g., [HJ85, Theorem 8.1.26]). Since $\sum_n S_n = \|f_n\|^2$ is finite, it must be that $\alpha(r) < 1$.

We now prove the second direction of Theorem 6.2. Given a ratio system r with $\alpha(r) < 1$, we will first construct the Green function for M , and then use it to construct the inverse of M . For a vertex $u \in T$ we define $f_u(v) = r(\vec{p}(u, v))$, where $\vec{p}(u, v)$ is (the projection to H of) the unique path from u to v . By convention, $f_u(u) = 1$. Note that $f_u \in l^2$, since for all $\varepsilon > 0$ we have $S_n = \sum_{\rho(x, u) = n} |f(x)|^2 = (\alpha(r) + \varepsilon)^n$ for large n , which converges if $\varepsilon < 1 - \alpha(r)$.

Claim 6.4. *We have $(Mf_u)(v) = 0$ whenever $v \neq u$. If e is (the type of) an edge incident to u then $(Mf_u)(u) = 0$ iff $r_e r_{e^{-1}} = 1$. In particular, either $r_e r_{e^{-1}} = 1$ for all edges, or for no edges.*

Proof. When $v \neq u$ the claim is immediate from (1).

Let e be an edges originating at u . We have $(Mf_u)(u) = m_{uu} + \sum_{e'=(u,v)} m_{e'} r_{e'}$, where the sum is over all edges originating at u . Multiplying this by $r_{e^{-1}}$ and subtracting (1) for e^{-1} gives

$$r_{e^{-1}}(Mf_u)(u) = m_e(r_e r_{e^{-1}} - 1).$$

If $r_e r_{e^{-1}} = 1$, then $r_{e^{-1}} \neq 0$ and therefor $(Mf_u)(u)$ must be 0. If $r_e r_{e^{-1}} \neq 1$, then (since we assume $m_e \neq 0$) $(Mf_u)(u)$ cannot be 0.

It follows that if e, e' are two edges incident on u , then $r_e r_{e^{-1}}$ and $r_{e'} r_{e'^{-1}}$ are either both 1 or both different from 1. Since H is connected, the same holds for any two edges. \square

We now show that $r_e r_{e^{-1}} \neq 1$ for all e . Assume the contrary. Take some non-backtracking cycle $\vec{p} \in H$, then also $r(\vec{p})r(\vec{p}^{-1}) = 1$, and so at least one of the two has absolute value at least one. Without loss of generality suppose $|r(\vec{p})| \geq 1$. However, in this case f_u cannot be in l^2 , since $|f_u| \geq 1$ at all endpoints of the path that is the lifting of \vec{p} repeated any number of times.

Thus we find that $(Mf_u)(u) \neq 0$ for all u . Consequently, there is a function g_u proportional to f_u such that $(Mg_u) = \delta_u$. Let us define $\tilde{G}(u, v) = g_u(v)$. Our next goal is to prove that \tilde{G} is the Green function for M . Let us establish some of its properties¹.

First, note that there are constants K_1, K_2 such that for any x, z and any y on the path $[x, z]$

$$K_1 \left| \tilde{G}(x, z) \right| \leq \left| \tilde{G}(x, y) \tilde{G}(y, z) \right| \leq K_2 \left| \tilde{G}(x, z) \right|. \quad (3)$$

This follows from the definition of \tilde{G} in terms of the finite ratio system, r , and the fact that $\tilde{G}(y, y)$ takes on only finitely many values as it depends only on y 's type.

The second property is that $\tilde{G}(x, \cdot)$ and $\tilde{G}(\cdot, x)$ are in l^2 . Moreover, multiplying a row or a column of \tilde{G} by a ‘‘slowly increasing function’’ (like that in the lemma that follows) still yields an l^2 function. Recall that we denote the distance between vertices x, y in T by $\rho(x, y)$.

Claim 6.5. *For any vertex u , let $h_1 = \rho(u, \cdot) \tilde{G}(u, \cdot)$ and $h_2 = \rho(u, \cdot) \tilde{G}(\cdot, u)$. There is some constant c independent of u such that $\|h_1\|_2 < c$ and $\|h_2\|_2 < c$.*

¹Note that since M is a symmetric operator, we expect \tilde{G} to be symmetric as well. While this fact has a direct proof, the following lemma is easier and suffices for our purposes. The symmetry of \tilde{G} will eventually follow by proving that \tilde{G} is the Green function of M .

Proof. Since there are finitely many types for u , it suffices to consider a single u . As before, let $S_n = \sum_{\rho(x,u)=n} |\tilde{G}(u, x)|^2$, so that for any $\varepsilon > 0$ we have $S_n = (\alpha(r) + \varepsilon)^n$ for large n . It follows that

$$\|h_1\|^2 = \sum_n n^2 S_n < \infty.$$

For h_2 , consider $S'_n = \sum_{\rho(x,u)=n} |\tilde{G}(x, u)|^2$. This equals the l^2 norm of a column of R^n , and therefore has the same asymptotics as S_n . \square

The next two claims imply that M is invertible.

Claim 6.6. *For every $h \in l^2(T)$ there exists $\theta \in l^2(T)$ such that $M\theta = h$.*

Claim 6.7. *The kernel of M is trivial.*

Proof of Claim 6.6. Given $h \in l^2(T)$, define θ as

$$\theta(v) = \sum_{u \in V(T)} \tilde{G}(u, v) h(u);$$

Claim 6.5 and Cauchy-Schwarz combined imply that the above sum converges. It also follows that $M\theta = h$ at each vertex. It remains to show that $\theta \in l^2(T)$, so we need to bound

$$\begin{aligned} \sum_v |\theta(v)|^2 &\leq \sum_{u,v,w} |h(u)h(w)\tilde{G}(u, v)\tilde{G}(w, v)| \\ &= \sum_{u,w} |h(u)h(w)H(u, w)|, \end{aligned} \tag{4}$$

where

$$H(u, w) = \sum_v |\tilde{G}(u, v)\tilde{G}(w, v)|.$$

The tree paths $[v, u]$ and $[v, w]$ separate at some point x , which is along the path $[u, w]$ (possibly, $x = u$ or $x = w$.) Separating the last sum according to the value of $x = x(v, u, w)$, using (3) and summing over v we get

$$\begin{aligned} H(u, w) &\leq c \sum_{\begin{array}{c} u-x-w \\ \downarrow \\ v \end{array}} |\tilde{G}(u, x)\tilde{G}(w, x)\tilde{G}(x, v)|^2 \\ &\leq c' \sum_{u-x-w} |\tilde{G}(u, x)\tilde{G}(w, x)|, \end{aligned} \tag{5}$$

where the first diagram means that we sum over all vertices u, w, v, x such that the subtree they span has the given structure (with x at the branching point in the first case, and along the path in the second). Note that it is possible to use (3) at this stage to eliminate x . We keep x to facilitate discussion in the next section. Inserting (5) in (4) gives

$$\sum_v |\theta(v)|^2 \leq c \sum_{u-x-w} |h(u)h(w)\tilde{G}(u, x)\tilde{G}(w, x)|.$$

This completes the proof of Theorem 6.2. Now we give a proof of Theorem 1.6 in the case of finite ratio systems.

Theorem 6.8. *If there is a finite ratio system for M with $\alpha(r) = 1$ then M is not invertible.*

Proof. Fix some edge $e = (u, v) \in \tilde{H}$. Using a finite ratio system r with $\alpha(r) = 1$ we consider the function

$$f_n(x) = \begin{cases} r(\vec{p}) & \text{if } x \in T_e \text{ and } \rho(x, u) = n, \\ 0 & \text{otherwise,} \end{cases}$$

and let $F_n = \sum_{k=0}^n f_k$. Thus $\|F_n\|^2 = \sum_{k=0}^n \|f_k\|^2$. Since MF_n is 0 on levels $1, \dots, n-1$ of the tree, it is supported on level 0 and levels $n, n+1$. Consequently $\|MF_n\|^2 = O(\|f_n\|^2)$.

Recall that $\|f_n\|^2$ is the sum of the row corresponding to e in R^n . Since the Perron eigenvalue of R is 1, this sum grows polynomially in n and we have $\|f_n\|^2 = P(n) + o(1)$ for some polynomial P . By choosing the edge e to be in the support of the Perron eigenfunction of R we can guarantee $P \neq 0$.

Assume $\deg P = k$ (which might be 0), with leading term cn^k . Then $\|F_n\|^2 = cn^{k+1}/(k+1) + O(n^k)$, while $\|MF_n\| = O(n^k)$. Thus $\|MF_n\|/\|F_n\| \rightarrow 0$ as $n \rightarrow \infty$, and so M cannot have a bounded inverse. \square

Corollary 6.9. *For any finite graph H , we have $\rho(\tilde{B}) = \sqrt{\text{gr}}$. Moreover, $\{\lambda : |\lambda| = \sqrt{\text{gr}}\} \subset \sigma(\tilde{B})$.*

Proof. By Theorem 1.5, it suffices to show that $M = Q_\lambda$ is invertible for $|\lambda| > \sqrt{\text{gr}}$, and is not invertible for $|\lambda| = \sqrt{\text{gr}}$. Since $m_{vv} = d_v - 1 + \lambda^2$ and $m_e = -\lambda$ for any e , it is straightforward to see that $r_e \equiv \lambda^{-1}$ satisfies (1). This finite ratio system has $\alpha(r) = \lambda^{-2}\text{gr}$. Theorem 6.2 shows that Q_λ is invertible whenever $|\lambda| > \sqrt{\text{gr}}$. However, if $|\lambda| = \sqrt{\text{gr}}$ then $\alpha(r) = 1$ and Theorem 6.8 implies that Q_λ is singular. \square

7. ZEROES AND INFINITIES

In this section we extend the results of the previous section to the case where $G(u, u)$ could vanish for some u , i.e., where ratios are allowed to be zero or infinity; we finish the proofs of Theorem 1.4 and 1.5. We discuss some examples where $G(u, u)$ vanishes for some u in Section 8.

We start with some remarks about Definition 1.3. Conditions (b) and (c) replace condition (a) in the situation of zero and infinite ratios. Condition (b) is easy to justify: $m_{e^{-1}}/r_e = \infty$, so there must be an infinite term in the sum as well. Condition (c) effectively states that the only way to have $r_f = \infty$ is as in condition (b). In particular, there cannot be two edges leaving a vertex with infinite ratios (since then their inverses both must and cannot have ratio zero).

As before, we need to extend r to non-backtracking paths, and would like $r(\vec{p})$ to be the product of the ratios along the edges of the path. In the presence of infinite ratios this needs clarification. We define $r(\vec{p})$ inductively. For paths of a single edge, $r(e) = r_e$, and for the an empty path, ε , $r(\varepsilon) = 1$.

For longer paths, note that an infinite ratio must be preceded by a zero ratio, so we must determine the value of $0 \cdot \infty$. For paths of length $m \geq 2$ we define $r(\vec{p})$

inductively as follows:

$$r(e_1, \dots, e_m) = \begin{cases} r_{e_1} r(e_2, \dots, e_m) & r_{e_1}, r_{e_2} \neq \infty, \\ -\frac{m_{e_2}}{m_{e_1-1}} r(e_3, \dots, e_m) & r_{e_1} = 0, r_{e_2} = \infty, \\ \infty & r_{e_1} = \infty. \end{cases} \quad (6)$$

In the case $r_{e_1} = 0$ and $r_{e_2} = \infty$, we effectively define “ $r_{e_1} r_{e_2}$ ” (literally zero times infinity) to be $r(e_1, e_2) = -m_{e_2}/m_{e_1-1}$, since the uniqueness of G implies that $G(v, v) = 0$ and $G(v, u) \neq 0$ implies that $G(v, w) \neq 0$ for exactly one other neighbour of u , and furthermore $m_{vw}G(v, w) + m_{vu}G(v, u) = 0$.

We remark that with this definition of $r(\vec{p})$, the proof of Theorem 6.8 carries over to prove Theorem 1.6.

We first assume that M is invertible, and find the promised ratio system $\{r_e\}$. Let G be the Green function. Let $e = (u, v)$ be some edge in T of type e_0 . We define r_e to be the generalized ratio $G(u, v)/G(u, u)$ as before, provided $G(u, v)$ and $G(u, u)$ do not both vanish. Otherwise, since $MG(u, \cdot)$ does not vanish at u , there must be a neighbour x of u with $G(u, x) \neq 0$. We then define $r_e = G(x, v)/G(x, u)$. In short there is some x , either $x = u$ or $x \notin T_e$ such that $G(u, x), G(x, v)$ are not both zero and

$$r_e = G(x, v)/G(x, u).$$

By Corollary 5.4, this ratio does not depend on our choice of x . As before, r_e depends only on the type of e , so this it is naturally equivalent to a function on edges of H .

Claim 7.1. *The function r defined above is a ratio system for M with $\alpha(r) < 1$.*

Proof. Throughout we assume that $e = (u, v)$ and that x is either u or outside T_e and $G(x, v)/G(x, u)$ is defined (and therefore equals r_e). In particular $x \neq v$ and so $MG(x, \cdot) = \delta_x$ vanishes at v . Thus we have

$$0 = m_{vv}G(x, v) + m_{e-1}G(x, u) + \sum_{\substack{w \sim v \\ w \neq u}} m_{vw}G(x, w). \quad (7)$$

To check (a), note that $r_e \neq 0$ implies $G(x, v) \neq 0$, and so the same x may be used to find r_f for $f = (v, w)$. Writing $G(x, w) = G(x, v)r_{vw}$ in (7) gives

$$0 = m_{vv} + m_{e-1} \frac{G(x, u)}{G(x, v)} + \sum_{\substack{w \sim v \\ w \neq u}} m_{vw} r_{vw}.$$

This is just (1), since $\frac{G(x, u)}{G(x, v)} = 1/r_e$ (meaning 0 if $r_e = \infty$).

To prove (b), suppose that $r_e = 0$. Then we have $G(x, u) \neq 0$ and $G(x, v) = 0$, and so (7) gives

$$\sum_{\substack{w \sim v \\ w \neq u}} m_{vw}G(x, w) = -m_{vu}G(x, u) \neq 0$$

Thus not all of $G(x, w)$ are 0. Let $w \sim v$ such that $G(x, w) \neq 0$, then the same x can be used to find r_f for $f = (v, w)$, and $r_{vw} = G(x, w)/G(x, v) = \infty$.

To verify (c), note that if $r_f = \infty$ for $f = (v, w)$, then by Corollary 5.4, we must have $G(x, v) = 0$ for all $x \notin V(T_f) \setminus \{v\}$. In particular $G(v, v) = 0$. It follows that $G(v, w) \neq 0$, so w can be used to find r_{f-1} and so $r_{f-1} = \frac{G(w, v)}{G(w, w)} \neq 0$. Finally, to

check that $r_e = 0$, suppose x is used to find r_e . Then $G(x, v) = 0$, for otherwise x could be used to find r_f , and r_f would be finite.

It remains to show that $\alpha(r) < 1$. As before, we show that the Perron eigenvalue of R satisfies $\alpha(r) < 1$ by proving that for all e with $r_e \neq 0$ the sum of row e in R^k tends to zero with k . Pick an edge $\tilde{e} = (u, v)$ of type e with $r_e \neq 0$. As above, fix some vertex $x \notin T_{\tilde{e}} \setminus \{u\}$ satisfying $G(x, v) \neq 0$. We claim that

$$\sum_{y \in T_{\tilde{e}}} |G(x, y)|^2 = |G(x, u)|^2 + |G(x, v)|^2 \sum_f \sum_{k=0}^{\infty} (R^k)_{ef} \quad (8)$$

This implies the required result, since the left-hand-side is finite. To see (8), for every $y \in T_{\tilde{e}}$ except for u , the path $[v, y]$ in T projects to some non-backtracking walk \vec{p}_y in H . This gives a bijection between such y and non-backtracking walks in H that begin with an edge e' such that $e \rightarrow e'$. We have $G(x, y) = G(x, v)r(\vec{p}_y)$.

The definition of R is such that

$$(R^k)_{ef} = \sum |r_{\vec{p}}|^2,$$

where the sum is over paths in H that begin with an edge e' such that $e \rightarrow e'$, end with edge f and, for some m , contain $k + m$ steps with m edges with $r_{e_i} = \infty$. The above identity follows. \square

Conversely, given a ratio system r with $\alpha(r) < 1$, we construct a candidate for the Green function $\tilde{G}(\cdot, \cdot)$ as in the proof of Theorem 6.2 and use it to construct M^{-1} . Our first goal is to find for each u a function f_u so that $Mf_u = c\delta_u$. We can then define $\tilde{G}(u, v) = c^{-1}f_u(v)$.

Generally we take $f_u(x) = r(\vec{p})$ where \vec{p} is (the projection of) the path from u to x . However, if there is an edge e leaving u with $r_e = \infty$, then this results in some infinite values of f_u . In this case, roughly speaking, we “divide by ∞ ”. All finite values will become 0, and the infinite ones will become finite. We shall now make this more precise.

Let V_∞ denote the set of all u such that there is an edge $e = (u, v)$ with $r_e = \infty$:

$$V_\infty = \{u : \exists v, r_{uv} = \infty\}$$

As noted above, for $u \in V_\infty$, the v with $r_{uv} = \infty$ must be unique. Define u^* to be this neighbour, v ; otherwise, for $u \notin V_\infty$, let $u^* = u$. Then (recall $\pi: T \rightarrow H$ is the covering map) set

$$f_u(x) = \begin{cases} r(\vec{p}) & \text{if } u \notin V_\infty, \text{ where } \vec{p} = \pi([u, x]), \\ r(\vec{q}) & \text{if } u \in V_\infty \text{ and } x \in T_{(u, u^*)} \setminus \{u\}, \text{ where } \vec{q} = \pi([u^*, x]), \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

We further denote $Z(u) = \sum_{e=(u,v)} m_e r_e$ where the sum is over edges leaving u with $r_e \neq \infty$ (i.e., if one of the edges has $r_e = \infty$, then it is not included).

Claim 7.2. *For any u we have*

$$(Mf_u)(u) = \begin{cases} m_{uu^*} & u \in V_\infty, \\ m_{uu} + Z(u) & u \notin V_\infty. \end{cases} \quad (10)$$

Proof. If $u \in V_\infty$ then we have $r_{uu^*} = \infty$. Consequently f_u is supported on T_{uu^*} , $f_u(u) = 0$ and for $v \sim u$ we have $f_u(v)$ is 1 or 0 according to whether or not $v = u^*$. The claim follows. For $u \notin V_\infty$ we have $f_u(u) = 1$ and $f_u(v) = r_{uv}$ for all neighbours v of u , so the equality holds also in this case. \square

Define $V_0 = \{u : m_{uu} + Z(u) = 0\}$.

Claim 7.3. *If an edge $e = (u, v)$ has $r_e \notin \{0, \infty\}$ and $v \in V_0$, then $r_e r_{e^{-1}} = 1$ and $u \in V_0$.*

Proof. (1) can be written as $m_{vv} + m_{e^{-1}}/r_e + Z(v) - m_{e^{-1}}r_{e^{-1}} = 0$. If $v \in V_0$ this becomes $1/r_e - r_{e^{-1}} = 0$, or $r_e r_{e^{-1}} = 1$. Applying this to (1) for e^{-1} gives $m_{uu} + Z(u) = 0$, so $u \in V_0$ as well. \square

Claim 7.4. *If $u \in V_0$, some edge entering u has ratio in $\{0, \infty\}$, then one edge e entering u has $r_e = r_{e^{-1}} = \infty$, and all other edges entering u have ratio 0.*

Proof. If $r_{vu} = 0$ for some v , then by (b) there is an $e = (w, u)$ with $r_{e^{-1}} = \infty$. This implies by (c) that all edges e' entering u except for e have $r_{e'} = 0$. By (c), $r_e \neq 0$. Applying (a) to e reduces to $m_{uu} + m_{e^{-1}}/r_e + Z(u) = 0$. Since $u \in V_0$, we must have $r_e = \infty$.

If $r_e = \infty$ for some $e = (v, u)$, then we claim $r_{e^{-1}} = \infty$: if it were finite, applying (a) to e would give $m_{uu} + Z(u) - m_{e^{-1}}r_{e^{-1}} = 0$. This implies $r_{e^{-1}} = 0$, contradicting (c). Since $r_e = r_{e^{-1}} = \infty$, again by (c) all edges e' entering u except for e have $r_{e'} = 0$. \square

Claim 7.5. $V_0 \subset V_\infty$.

Proof. Suppose that $u \in V_0 \setminus V_\infty$. Since there are no edges with infinite ratio leaving u , by Claim 7.4 all edges entering u have finite nonzero ratios. By Claim 7.3 this implies that all of u 's neighbours are in V_0 , and by Claim 7.4 they are all in $V_0 \setminus V_\infty$.

It follows that if $V_0 \setminus V_\infty$ is nonempty then it must be the entire tree. If $V_0 \setminus V_\infty = T$ then for all edges $r_e r_{e^{-1}} = 1$. Then (as in Claim 6.4) taking a non-backtracking cycle, \vec{p} , in H gives $r(\vec{p})r(\vec{p}^{-1}) = 1$, and so $\alpha \geq 1$ — a contradiction. \square

Claim 7.6. *For any u , $(Mf_u)(u) \neq 0$.*

Proof. By Claim 7.2, either $u \in V_\infty$ and then $(Mf_u)(u) = m_{uu^*} \neq 0$, or $u \notin V_\infty$. In the latter case $u \notin V_0$, and therefore $(Mf_u)(u) = m_{uu} + Z(u) \neq 0$. \square

It follows that we may define

$$\tilde{G}(u, x) = \frac{f_u(x)}{(Mf_u)(u)}, \quad (11)$$

so that $M\tilde{G}(u, \cdot) = \delta_u$. It remains to prove the generalized form of Claim 6.6:

Claim 7.7. *Given a ratio system with $\alpha(r) < 1$, for every $h \in l^2(T)$ there exists a $\theta \in l^2(T)$ such that $M\theta = h$.*

The proof is very similar to the proof of Claim 6.6, and we only point out the differences.

Proof. Equation (3) only fails if $G(y, y) = 0$. This problem is easily overcome by shifting the splitting point from y to y^* , which is the unique neighbour of y satisfying $r_{yy^*} = \infty$. When $G(y, y) \neq 0$ we define y^* to be y . We have for suitable K_1, K_2 and any $y \in [x, z]$

$$K_1 |\tilde{G}(x, z)| \leq |\tilde{G}(x, y^*)\tilde{G}(y^*, z)| \leq K_2 |\tilde{G}(x, z)|. \quad (3')$$

This follows from the definition of \tilde{G} and the fact that y has only finitely many types.

As in the preceding section, we define $H(u, w) = \sum_v |\tilde{G}(u, v)\tilde{G}(w, v)|$. Using (3') gives the bound

$$H(u, w) \leq c \sum_{u-x-w} |\tilde{G}(u, x^*)\tilde{G}(x^*, w)|.$$

In all subsequent estimates x^* and z^* , respectively, take the place of x and z in all instances of \tilde{G} , but x and z remain unchanged in the “tree sum” notation and in ρ' .

The only difference is that in place of Claim 6.5 we need bounds of the form

$$\sum_x \rho(u, x)^2 \left(|\tilde{G}(u, x^*)|^2 + |\tilde{G}(x^*, u)|^2 \right) < c. \quad (12)$$

Such bounds are an easy consequence of Claim 6.5: The map $x \mapsto x^*$ maps no more than $d_{\max} + 1$ x 's to each x^* . Since also $|\rho(x^*, y^*) - \rho(x, y)| \leq 2$, replacing some of the variables by their starred versions introduces at most a constant factor into the bounds at each stage of the computation. \square

8. EXAMPLES

We wish to analyze two examples to highlight the necessity of considering ratio systems. Each of these examples demonstrates a general point. Then we give a third example, which seems to be the “smallest interesting example” of a graph with a non-regular universal cover. We describe the calculations needed to determine its non-backtracking spectrum. In each of these examples, the finite graph is a simple graph. In the next section we shall show how each of these can be viewed as a covering of an even smaller graph by using self-loops, multiple edges, or “passing” to a pregraph.

8.1. Example 1: Figà-Talamanca and Steger. [FTS94] contains various examples of graphs where $G(u, u) = 0$ for some u when M is the weighted adjacency matrix of the tree. Consider K_4 , the complete graph on four vertices, whose six edges are written as the disjoint union of three perfect matchings, E_1, E_2, E_3 . Let the edges in E_i have weight p_i , for some positive real p_1, p_2, p_3 . Without loss of generality $p_1 \geq p_2 \geq p_3$. By symmetry (i.e., using Proposition 5.1), for any ratio system r , the ratios r_e of all directed edges of E_i are equal. Call the three ratios r_1, r_2, r_3 , and consider the ratio system is $r_1 = \infty, r_2 = r_3 = 0$. We have $r(e_k, e_1) = -p_k/p_1$; the condition $\alpha(r) < 1$ amounts to

$$(p_2/p_1)^2 + (p_3/p_1)^2 < 1.$$

Thus the adjacency matrix for the tree is invertible with $G(u, u) = 0$ if $p_1^2 > p_2^2 + p_3^2$. (In [FTS94, Lemma 1.4] it is shown that if $p_1^2 \leq p_2^2 + p_3^2$ then the adjacency matrix is not invertible.)

It is of interest to note that the relative kernel with respect to an E_1 edge is only one-dimensional (when $p_1^2 > p_2^2 + p_3^2$).

8.2. Example 2: (2,3)-bipartite graphs. Consider the adjacency matrix of a (2,3)-bipartite graph, i.e., a bipartite graph with vertex set $V_2 \cup V_3$ such that each vertex in V_k is of degree k . Then consider $r_{32} = \infty$, $r_{23} = 0$, where e_{23} is an edge from V_2 to V_3 and e_{32} the opposite. This fails to be a generalized ratio for the adjacency matrix only at condition (c) of the definition. Furthermore, if we tried to omit condition (c), then this ratio system would have $\alpha < 1$: A vertex of degree 3 in the universal cover has two children so $r(e_{32}, e_{23}) = -1/2$. This gives rise to L^2 -functions, f , with $Mf = \delta_r$ for any vertex r of degree 2 (we can take f to have value 0 at all degree two vertices, and value $-(-2)^{-a}$ at all vertices at odd distance $2a - 1$ to r).

It is easy to see that the adjacency matrix is not invertible here; indeed, consider a long path in the tree, and a function supported on the degree two vertices of this path, alternating ± 1 ; this function, f , has Mf bounded but $\|f\|$ unbounded as the path length tends to infinity. This example shows that condition (c) in the definition of a ratio system cannot be omitted from Theorem 1.4.

8.3. Example 3: Case study of $P_{1,2,2}$. We consider here non-backtracking walks on lifts of the graph $P_{1,2,2}$ consisting of three paths between a pair of vertices u, u' of lengths 1, 2, 2. Thus u, u' have degree 3 and there are two other vertices, v, v' , of degree 2. This is the same as \widetilde{K}_4 with an edge removed.

To find the spectrum of $B(\widetilde{P}_{1,2,2})$ we need to determine invertibility of $M = Q_\lambda$. Let us look for ratio functions r_e that respect all the graph's symmetries, i.e. $r_{uv} = r_{uv'} = \dots$. There are only three distinct ratios:

$$x = r_{uv} \quad y = r_{uu'} \quad z = r_{vu}.$$

Thus y appears on two directed edges, and x, z on four each. The equations we need to satisfy are

$$\begin{aligned} (\lambda^2 + 1)x + \lambda(1 + xz) &= 0, \\ (\lambda^2 + 2)y + \lambda(1 + 2xy) &= 0, \\ (\lambda^2 + 2)z + \lambda(1 + xz + yz) &= 0. \end{aligned}$$

Using the first equation to eliminate z and the second to eliminate x yields:

$$(y + \lambda^{-1})((2\lambda^4 + 2\lambda^2 + 4)y^2 + \lambda^3(\lambda^2 + 1)y - \lambda^2(\lambda^2 + 1)) = 0.$$

This leads to the trivial solution $x = y = z = -\lambda^{-1}$ and two non-trivial solutions for x, y, z (or to be precise, two branches of one solution). The condition $\alpha(r) < 1$ reduces to $x^2 z^2 (1 + 2y^2) < 1$. A computer plot of the λ where none of the solutions has $\alpha(r) < 1$ is shown in Figure 1 (the shaded region).

9. SELF-LOOPS AND PREGRAPHS VIA FINITE TREE QUOTIENTS

We shall show that Theorem 1.4 can be generalized to universal covers of all of the following structures: graphs/pregraphs with multiple edges, self-loops, boundaries (with boundary conditions), edge and vertex weights, and edge lengths (see [Fri93]).

The common theme is that all these structure admit some sort of “universal cover,” T that is sufficiently symmetric that a mild generalization of Theorem 1.4 holds.

First, note that Definition 1.3 makes sense for any graph homomorphism $\pi: T \rightarrow H$, regardless of whether or not H is finite (and even whether or not π is a covering map).

Definition 9.1. A *tree quotient* is a tree, T , along with a group, \mathcal{G} , acting on it; we denote this T/\mathcal{G} . By $V(T/\mathcal{G})$ we mean the \mathcal{G} orbits of $V(T)$, and similarly for $E(T/\mathcal{G})$. The action of \mathcal{G} is cofinite if $V(T/\mathcal{G})$ and $E(T/\mathcal{G})$ are finite; alternatively we say T/\mathcal{G} is *finite*. By a symmetric local operator on T/\mathcal{G} we mean a symmetric local operator on T that is invariant under \mathcal{G} . By a ratio system on T we mean a ratio system as in Definition 1.3 with H there taken to be T ; by a ratio system on T/\mathcal{G} we mean a ratio system on T that is invariant under \mathcal{G} , i.e., r can be viewed as a function $r: E(T/\mathcal{G}) \rightarrow \mathbb{C} \cup \{\infty\}$.

Any finite graph and pregraph as described in [Fri93, FT04] (possibly with boundary, self-loops, multiple edges, vertex and edge weights, and edge lengths) admits a universal cover, $\pi: T \rightarrow G$; T is a tree (a graph) that inherits vertex and edge weights and lengths from G (boundary edges may be discarded with the Laplacian or adjacency matrix modified appropriately). By the universal property of T , for any two vertices in T above the same G vertex, T admits an automorphism from taking one vertex to the other. It follows that if \mathcal{G} is the group of all automorphisms of T over G , then T/\mathcal{G} is finite and has at most as many edges and vertices as G . Hence to generalize Theorem 1.4 to all the above situations it suffices to consider finite tree quotients T/\mathcal{G} .

Consider a symmetric, local operator, M , on a tree quotient, T/\mathcal{G} . The proof of Theorem 1.4 shows that a Green’s function, G , for M gives rise to generalized periodic ratios on T ; Proposition 5.1 shows that generalized periodic ratios also live on T/\mathcal{G} . Furthermore, given the periodic ratios $\{r_e\}$ from G , we can reconstruct G via (9) and (11).

Definition 9.2. A *classifying condition* for a tree quotient, T/\mathcal{G} , is a condition $C = C(r)$ or even $C = C(r, M)$ on generalized periodic ratios, r , such that $C(r, M)$ is satisfied iff r arises from the Green function, G , of an invertible, symmetric, local operator, M , on T/\mathcal{G} .

Given a tree quotient, we wish to find a classifying condition that is simple to check. Theorem 1.4 says that if $\pi: T \rightarrow H$ is a universal cover of a finite graph, H , and \mathcal{G} is the group of automorphisms of T over H , then “ $\alpha(r) < 1$ ” is a classifying condition for T/\mathcal{G} .

Our first task is to generalize the matrix R that describes α on a tree quotient, T/\mathcal{G} . For $e, f \in E(T)$ with $r_e, r_f \neq 0$, we define $R_{e,f}$ as in (2); for $a, b \in E(T/\mathcal{G})$ pick an e with $e \in a$ and set

$$R_{a,b} = \sum_{f \in b} R_{e,f}.$$

In this setting, our proof of Theorem 1.4 shows the following:

Theorem 9.3. *Let $C = C(r)$ be a condition on generalized periodic ratios of a tree quotient, T/\mathcal{G} , such that*

- (a) *if r arises from a Green function, then $C(r)$ is satisfied,*
- (b) *if r satisfies $C(r)$, then $(Mf_u)(u) \neq 0$ for all u (with f as in (9)), and*
- (c) *if r satisfied $C(r)$ then \tilde{G} is a bounded kernel, i.e., the map*

$$h(\cdot) \mapsto \sum_{u \in V(T)} \tilde{G}(u, \cdot)h(u),$$

is bounded in L^2 .

Then C is a classifying condition for T/\mathcal{G} . Furthermore,

- (1) *$(Mf_u)(u) \neq 0$ for all u as long as there can exist no sequence of distinct vertices, v_0, v_1, \dots in T for which $|r([v_0v_i])| \geq 1$ for all $i \geq 1$, and*
- (2) *any function $\tilde{G}: V(T) \times V(T) \rightarrow \mathbb{C}$ is a bounded kernel provided that it satisfies the bounds (3) and (12).*

Now let $\alpha(r)$ be the spectral radius of the operator R defined above as an operator on $l^2(V(T/\mathcal{G}))$. If \mathcal{G} acts cofinitely on T , then R is a finite real matrix with non-negative entries, and a periodic ratio arising from a Green function must satisfy $\alpha(r) < 1$ since the Green function is in l^2 , using the analogue of (8). Conversely, a generalized periodic ratio with $\alpha(r) < 1$ yields that

$$R + 2R^2 + 3R^3 + \dots$$

is a bounded operator, from which we conclude (12) and that there cannot exist distinct $v_0, v_1, \dots \in V(T)$ with $|r([v_0v_i])| \geq 1$; the cofinite action on \mathcal{G} on T shows that $((Mf_u)(u))^{-1}$ and r_e are uniformly bounded, yielding (3). We conclude the following theorem.

Theorem 9.4. *Let \mathcal{G} act cofinitely on T . The condition $\alpha(r) < 1$, defined above, is a classifying condition for T/\mathcal{G} .*

We do not know a classifying condition when \mathcal{G} does not act cofinitely on T .

We now demonstrate how the trees in the three examples in the previous section have smaller quotients, provided one is willing to pass to graphs or pregraphs that may have self-loops and multiple edges.

Example 1 in the previous section covers the graph consisting of one vertex with three half-loops, e_1, e_2, e_3 , of respective weights $p_1 \geq p_2 \geq p_3$. This is how it is presented in [FTS94]; they allow any number of half-loops with weights $p_1 \geq \dots \geq p_r$. The covering tree's adjacency matrix is invertible iff $p_1^2 > p_2^2 + \dots + p_r^2$. The r half-loop example is covered by any r -regular graph whose edges can be partitioned into r perfect matchings.

Example 2 covers the pregraph with two vertices, no self-loops, with one vertex of degree two, and one vertex of degree three.

Example 3 covers the graph $H = (V, E)$, where $V = \{u, v\}$, where there are two edges joining u and v and one half-loop at u .

10. CONCLUDING REMARKS AND FURTHER QUESTIONS

We have given a semi-algebraic characterization of invertible symmetric local operators on a finite tree quotient. It would be interesting to know what happens when the tree quotient is not finite and/or when the operator is not symmetric (with respect to l^2 of any measure on the tree vertices). For example, the non-backtracking *random* walk selects at each step a uniform neighbour of X_n different from X_{n-1} [ABLS, OW07].

Our characterization can be used, at least in principle, to compute the various spectra of trees, including adjacency and Laplacian, and, less obviously, non-backtracking. This may shed light on the questions mentioned in the introduction. One interesting question, especially for two-dimensional non-backtracking spectrum, is to know if there is a natural spectral measure there, analogous to the Wigner semi-circle law, as done by McKay for regular base graphs. Specifically, we conjecture that for any finite graph H there is a measure μ supported on the non-backtracking spectrum of the universal cover \tilde{H} , such that the uniform measure on $\sigma(H^n)$ converges in probability (in a suitable topology) to μ , where H^n is a random uniform n -lift of H (possibly with some corrections, say around ± 1).

The proof of Theorem 1.6 is based on explicitly constructing approximate eigenfunctions when there is a ratio system with $\alpha = 1$. It would be nice to be able to find a sequence of approximate eigenfunctions of $M - \lambda$ for general λ in M 's spectrum.

Finally we believe that for λ on the boundary of the spectrum of M (as before), $M - \lambda I$, has a ratio system, r , with $\alpha(r) = 1$. In particular, if $\{\lambda_n\}$ is a sequence of complex numbers with a finite limit, λ , and M is as usual, and r_n is a ratio system for $M - \lambda_n$, then r_n has a limit point, r , where we compactify $\mathbb{C} \cup \{\infty\}$ as usual. Is r necessarily a ratio system? In particular is it impossible to have $r_e = r_f = \infty$ for two edges, e, f , with the same tail? If this does happen, might it still be possible to define α for the limit point, r ?

REFERENCES

- [ABLS] N. Alon, I. Benjamini, E. Lubetzky, and S. Sodin. Non-backtracking random walks mix faster. to appear.
- [AHL02] N. Alon, S. Hoory, and N. Linial. The Moore bound for irregular graphs. *Graphs Combin.*, 18(1):53–57, 2002.
- [Bas92] H. Bass. The Ihara-Selberg zeta function of a tree lattice. *Internat. J. Math.*, 3(6):717–797, 1992.
- [BS87] A. Broder and E. Shamir. On the second eigenvalue of random regular graphs. In *The 28th Annual Symposium on Foundations of Computer Science*, pages 286–294, 1987.
- [Fri] J. Friedman. A proof of Alon's second eigenvalue conjecture. *Memoirs of the A.M.S.*, to appear.
- [Fri91] Joel Friedman. On the second eigenvalue and random walks in random d -regular graphs. *Combinatorica*, 11(4):331–362, 1991.
- [Fri93] Joel Friedman. Some geometric aspects of graphs and their eigenfunctions. *Duke Math. J.*, 69(3):487–525, 1993.
- [Fri03] J. Friedman. Relative expanders or weakly relatively Ramanujan graphs. *Duke Math. J.*, 118(1):19–35, 2003.
- [FT04] Joel Friedman and Jean-Pierre Tillich. Wave equation for graphs and the edge-based laplacian. *Pacific Journal of Mathematics*, 216(2):229–266, 2004.
- [FTS94] Alessandro Figà-Talamanca and Tim Steger. Harmonic analysis for anisotropic random walks on homogeneous trees. *Mem. Amer. Math. Soc.*, 110(531):xii+68, 1994.

- [Gal63] R. G. Gallager. *Low density parity check codes*. MIT Press, Cambridge, MA, 1963.
- [Has89] K. Hashimoto. Zeta functions of finite graphs and representations of p -adic groups. In *Automorphic forms and geometry of arithmetic varieties*, volume 15 of *Adv. Stud. Pure Math.*, pages 211–280. Academic Press, Boston, MA, 1989.
- [HJ85] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1985.
- [HLW] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. to appear.
- [Hoo05] S. Hoory. A lower bound on the spectral radius of the universal cover of a graph. *J. Combin. Theory Ser. B*, 93(1):33–43, 2005.
- [KS00] M. Kotani and T. Sunada. Zeta functions of finite graphs. *J. Math. Sci. Univ. Tokyo*, 7(1):7–25, 2000.
- [OW07] R. Ortner and W. Woess. Non-backtracking random walks and cogrowth of graphs. *Canad. J. Math.*, 59(4):828–844, 2007.
- [RU] T. Richardson and R. Urbanke. Modern coding theory. Draft of the book, <http://lthcwww.epfl.ch/papers/mct.ps>.
- [Rud91] W. Rudin. *Functional analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill Inc., New York, second edition, 1991.
- [Sho04] A. Shokrollahi. LDPC codes: an introduction. In *Coding, cryptography and combinatorics*, volume 23 of *Progr. Comput. Sci. Appl. Logic*, pages 85–110. Birkhäuser, Basel, 2004.
- [ST] H. Stark and A. Terras. Zeta functions of finite graphs and coverings, part iii. <http://math.ucsd.edu/~aterras/BrauerSiegel.pdf>.

(Omer Angel) MATH, UOFT
E-mail address: `angel@math.utoronto.ca`

(Joel Friedman) CS, UBC
E-mail address: `jf@cs.ubc.ca`

(Shlomo Hoory) HAIFA RESEARCH LAB., IBM
E-mail address: `shlomoh@il.ibm.com`