Relationships between cycles spaces, gain graphs, coverings, path homology, and graph curvature

Mark Kempton* Florentin Münch† Shing-Tung Yau‡

September 12, 2017

Abstract

We prove a homology vanishing theorem for graphs with positive Bakry-Émery curvature, analogous to a result of Bochner on manifolds [2]. Specifically, we prove that if a graph has positive curvature at every vertex, then the first path homology group is trivial. We prove this by first showing that a graph with positive curvature can have no non-trivial infinite cover preserving 3-cycles and 4-cycles, and then by giving a combinatorial interpretation of the first path homology in terms of the cycle space of a graph. We also relate cycle spaces of graphs to gain graphs with abelian gain group, and relate these to coverings of graphs.

1 Introduction

In recent years there has been growing interest in applying tools and ideas from continuous geometry to discrete setting on graphs. One of the principal developments in this area concerns curvature for graphs. Numerous notions of curvature have been put forward for graphs [5, 16]. An important and very general notion of curvature for graphs has been defined via various formulas due to Bakry and Émery, which is called the Bakry-Émery curvature of a graph (see [1, 17, 14]).

In addition, there are various notions of homology and cohomology for graphs. Recent work has introduced one such theory called the path homology [7]. Path homology has been shown to be a non-trivial homology theory for graphs which satisfies nice functorial properties, namely the Künneth formula holds for graph products [8].

In this paper, we prove an important connection between these two notions, namely we prove a homology vanishing theorem for graphs with positive Bakry-Émery curvature. Homology vanishing theorems are ubiquitous in continuous geometry, and give important structural information about manifolds. Our vanishing theorem is analogous to a result of Bochner on manifolds ([2]).

hom_van **Theorem 1.1.** *If G has positive Bakry-Émery curvature at every vertex, then its first path homology group is trivial.*

Our proof is executed by developing relationships between curvature and homology with other well-studied combinatorial aspects of graphs. Particularly, we look at graph coverings, gain graphs, and the cycle space of a graph. One of the main contributions of this paper is to develop connections between each of these disparate areas of graph theory. In particular, we provide a combinatorial interpretation of the first path homology group. In addition, one of our main contributions is the notion of a Γ-circuit generator of a graph for a group Γ. This is a set of cycles $\mathcal{B}$ for which, if a gain function is balanced on $\mathcal{B}$, it is balanced on the entire graph. Such circuit generators can be viewed as a new kind of cycle basis. See Section 2 for more detailed definitions.

*Center of Mathematical Sciences and Applications, Harvard University, Cambridge MA, USA. mkempton@cmsa.fas.harvard.edu

†Universität Potsdam, Potsdam, Germany. chmeunch@uni-potsdam.de

‡Harvard University, Cambridge MA, USA. yau@math.harvard.edu

1.1 Organization and main results

The remainder of this paper will be organized as follows. In Section 2, we will give the technical preliminaries, including definitions and known results concerning gain graphs, cycle bases, covers of graphs, graph homology, and graph curvature. In addition to relevant known results, we present some new lemmas that will be useful later.

In Section 3, we connect the Bakry-Émery curvature to coverings of graph. We prove

Theorem 1.2. *If G is a graph with positive curvature, then G has no infinite covering that preserves 3-cycles and 4-cycles.*

Section 4 explores the relationship between gain graphs and covers of graphs. In particular, we prove

Theorem 1.3. *There is a 1-1 correspondence between non-trivial covers of a graph preserving a set of cycles and gain graphs that are balanced on those cycles.*

Section 5 investigates the relationships between cycle spaces and gain graphs. In particular we prove

Theorem 1.4. *A collection of circuits $\mathcal{B}$ of a graph is a $\mathbb{F}$-circuit generator if and only if $\mathcal{B}$ is an $\mathbb{F}$-cycle basis.*

In addition, we give a complete determinantal characterization of Γ circuit generators for Abelian groups Γ.

Theorem 1.5. *For an Abelian group Γ, a collection of cycles $\mathcal{B}$ is a Γ-circuit generator if and only if $g^{\det \mathcal{B}} \neq e_\Gamma$ for all $g \in \Gamma \setminus \{e_\Gamma\}$.*

Finally, in Section 6, we prove a relationship between the path homology of graphs and cycle spaces.

Theorem 1.6. *The first homology group $H_1(G, \mathbb{F})$ for a field $\mathbb{F}$ is isomorphic to the $\mathbb{F}$-cycle space modulo the space generated by all 3- and 4-cycles.*

With these results, we will be able to prove Theorem 1.1.

2 Preliminaries

sec:prelim

2.1 Gain graphs and circuit generators

Let $G = (V, E)$ an undirected graph. We will denote by $\vec{E}$ the set that contains two directed arcs, one in each direction, for each edge in E. Let Γ be a group. A *gain graph* is a triple (G, ϕ, Γ) where $\phi : \vec{E} \to \Gamma$ is a map satisfying $\phi(xy) = \phi(yx)^{-1}$ for all edges $(xy) \in \vec{E}$. The map ϕ is called the *gain function* of the gain graph. Denote by $\Phi(G, \Gamma)$ the set of all gain functions from the graph G to the group Γ.

Gain graphs have also been referred to as *voltage graphs*, and are special cases of *biased graphs* (see [18]). When Γ is a group of invertible linear transformations, they are also called *connection graphs* [3], and the map ϕ can be considered as a connection corresponding to a vector bundle on the graph [11].

For the most part, we will take terminology about gain graphs and biased graphs from [18]. A circuit or simple cycle C is a simple closed walk $(x_1, \ldots, x_n)$. Throughout the paper, we will refer to circuits of length 3 as "triangles," and circuits of length 4 as "squares." We write $\mathcal{S}(G)$ for the set of all circuits in G. A *theta graph* is the union of three internally disjoint simple paths that have the same two distinct endpoint vertices. A *biased graph* is a pair $(G, \mathcal{B})$ where $\mathcal{B} \subset \mathcal{S}(G)$ is a set of distinguished circuits, called *balanced circuits*, that form a *linear subclass* of circuits, that is, $\mathcal{B}$ has the property that if any two circuits of a theta graph are in $\mathcal{B}$, then so is the third. We say $\mathcal{B} \subset \mathcal{S}(G)$ is a *cyclomatic circuit set* if the cardinality of $\mathcal{B}$ equals the cyclomatic number of G, $|E| - |V| + 1$.

Gain graphs are biased graphs in a natural way. Define define the *order* of a circuit $C = (x_1, \ldots, x_n)$ under a gain function ϕ via

$$o_\phi(C) := \inf \left\{ r > 0 : \left[\phi(x_1 x_2) \ldots \phi(x_{n-1} x_n) \phi(x_n x_1)\right]^r = e_\Gamma \right\}.$$

We say the gain function ϕ is *balanced* on a circuit $C = (x_1, \ldots, x_n)$ if $o_\phi(C) = 1$. Denote by $\mathcal{B}(\phi)$ the set of balanced circuits. Then $(G, \mathcal{B}(\phi))$ defines a biased graph.

Definition 2.1 (Γ-circuit generator)**.** We say a set $\mathcal{B}$ of circuits is a *strong* Γ*-circuit generator* if for all $\phi \in \Phi(G, \Gamma)$, ϕ balanced on $\mathcal{B}$ implies that ϕ is balanced on the entire graph G. We say a set $\mathcal{B}$ of circuits is a *weak* Γ*-circuit generator* if for all $\phi \in \Phi(G, \Gamma)$, ϕ balanced on $\mathcal{B}$ implies that all circuits of G have finite order under ϕ.

Definition 2.2 (Canonical gain graph)**.** Let $G = (V, E)$ be a graph and let $\mathcal{B}$ be a set of circuits. We define the group $\Gamma(G, \mathcal{B})$ via the presentation

$$\Gamma(G, \mathcal{B}) = \langle \vec{E} | \mathcal{B} \rangle,$$

i.e., $\Gamma(G, \mathcal{B})$ is generated by the oriented edges and each circuit $C = (x_1, \ldots, x_n) \in \mathcal{B}$ gives a relation $(x_1x_2) \ldots (x_{n-1}x_n)(x_nx_1) = e_\Gamma$ where we identify $(xy) = (yx)^{-1} \in \vec{E}$. There is a natural gain function $\phi_\mathcal{B}$ given by the natural mapping of $\vec{E}$ into $\Gamma(G, \mathcal{B})$, and the corresponding gain graph $(G, \phi_\mathcal{B}, \Gamma(G, \mathcal{B}))$ is called the *canonical gain graph* associated with the biased graph $(G, \mathcal{B})$.

In [19], it is shown that the canonical gain graph satisfies a universal property with respect to gain graphs balanced on $\mathcal{B}$.

rop:universal **Proposition 2.3** (Theorem 2.1 of [19])**.** *Given any gain graph (G, ψ, Γ) such that ψ is balanced on $\mathcal{B}$, then there exists a homomorphism $h : \Gamma(G, \mathcal{B}) \to \Gamma$ such that $\psi = h \circ \phi_\mathcal{B}$ as defined above.*

Definition 2.4 (Combinatorial circuit generator)**.** We say, $\mathcal{B} \subset \mathcal{S}$ is a *combinatorial circuit generator* if the linear subclass generated by $\mathcal{B}$ is the entire set of circuits $\mathcal{S}(G)$.

For convenience in talking about combinatorial circuit generators and linear subclasses of circuits, we define the following operation on circuits. For circuits C_1, C_2 belonging to the same theta graph, we define $C_1 \oplus C_2$ to be the third cycle in the theta graph. Then we may describe a linear subclass as a set of circuits that is closed under the operation $\oplus$.

rop:products **Proposition 2.5.** *Let Γ_1, Γ_2 be groups, let G be a finite graph and let $\mathcal{B} \subset \mathcal{S}(G)$ be a set of cycles. T.f.a.e.*

1. *$\mathcal{B}$ is a $\Gamma_1 \times \Gamma_2$ circuit generator.*
2. *$\mathcal{B}$ is a Γ_1 circuit generator and a Γ_2 circuit generator.*

Moreover if $\Gamma_1 \subset \Gamma_2$ is a subgroup and if $\mathcal{B}$ is a Γ_2 circuit generator, then $\mathcal{B}$ is also a Γ_1 circuit generator.

Proof. The fact that the circuit generator property is inherited to subgroups directly follows from the definition. This also proves $1 \Rightarrow 2$. We now prove $2 \Rightarrow 1$ indirectly. Suppose $\mathcal{B}$ is not a $\Gamma_1 \times \Gamma_2$ circuit generator. Then, there exists $C = (x_1, \ldots, x_n) \in \mathcal{S}(G)$ and $\phi \in \Phi(G, \Gamma_1 \times \Gamma_2)$ s.t. ϕ is balanced on $\mathcal{B}$ but not on C, i.e., $\phi(x_1x_2) \ldots \phi(x_{n-1}x_n)\phi(x_nx_1) = (g_1, g_2) \in \Gamma_1 \times \Gamma_2$ with $g_k \neq e_{\Gamma_k}$ for some $k \in \{1, 2\}$. Define $\phi_i \in \Phi(G, \Gamma_i)$ s.t. $\phi(e) = (\phi_1(e), \phi_2(e))$ Now, ϕ_i is balanced on $\mathcal{B}$ for $i = 1, 2$ and ϕ_k is not balanced on C. Therefore, $\mathcal{B}$ is not a Γ_k-circuit generator. This is a contradiction and finishes the proof. □

We prove some facts about circuit generators that will be useful to us later on.

elyGenerated **Proposition 2.6.** *Let Γ_0 be a group, let G be a finite graph and let $\mathcal{B} \subset \mathcal{S}(G)$ be a set of cycles. T.f.a.e.:*

1. *$\mathcal{B}$ is a Γ_0-circuit generator.*
2. *$\mathcal{B}$ is a Γ-circuit generator for all finitely generated subgroups of Γ_0.*

Proof. $1 \Rightarrow 2$ follows from Proposition 2.5.

$2 \Rightarrow 1$: Suppose $\mathcal{B}$ is not a Γ_0-circuit generator. Then there exists $\phi \in \Phi(G, \Gamma_0)$ that is balanced on $\mathcal{B}$ but not on some $C \in \mathcal{S}(G)$. Let Γ be the group generated by all $\phi(e)$ for all $e \in \vec{E}$. Then, Γ is a finitely generated subgroup of Γ_0 and $\phi \in \Phi(G, \Gamma)$ is balanced on $\mathcal{B}$ but not on C. Therefore, $\mathcal{B}$ is not a Γ-circuit generator. This finishes the proof. □

2.2 Cycle bases and fields

If $\mathbb{F}$ is a field with additive group Γ, then we write $\Phi(G,\mathbb{F}) := \Phi(G,\Gamma)$ being a $\mathbb{F}$-vector space.

The $\mathbb{F}$-cycle space is defined by

$$\mathcal{C}(G,\mathbb{F}) := \left\{ \phi \in \Phi(G,\mathbb{F}) : \sum_{y:(xy)\in\vec{E}} \phi(xy) = 0 \text{ for all } x \in V \right\}.$$

Every circuit can be identified with a cycle $\phi \in \mathcal{C}(G,\mathbb{F})$ s.t. $\phi(\vec{E}) \subset \{0,\pm 1\}$ and such that $\phi^{-1}(\{1\})$ is a simple connected cycle. Therefore, we have $\mathcal{S} \hookrightarrow \mathcal{C}(G,\mathbb{F})$ allowing us to add circuits and to multiply them with scalars. For a cycle $C \in \mathcal{C}(G,\mathbb{F})$ and a gain function $\phi \in \Phi(G,\mathbb{F})$ define

$$\phi(C) := \sum_{e\in\vec{E}} \phi(e)C(e).$$

Definition 2.7 (F-cycle basis)**.** A subset $\mathcal{B} \subset \mathcal{C}(G,\mathbb{F})$ is called an $\mathbb{F}$-*cycle basis* of G if $\mathcal{B}$ is a vector space basis of $\mathcal{C}(G,\mathbb{F})$. In the literature, a $\mathbb{Q}$-cycle basis is also called a *directed* cycle basis and an $\mathbb{F}_2$-cycle basis is also called an *undirected* cycle basis.

According to [13], we define the determinant of a set of cycles. Let r be the cyclomatic number of G and let $\mathcal{B} \subset \mathcal{S}(G)$ be of size r. Corresponding to [13, Definition 22], consider the matrix $M(\mathcal{B},\mathbb{F})$ over the field $\mathbb{F}$ with the incidence vectors of $\mathcal{B}$ as columns. Let $M(\mathcal{B},\mathbb{F},T)$be the $r \times r$ submatrix that arises when deleting the arcs of the spanning T tree of G. Remark that $M(\mathcal{B},\mathbb{F})$ consists only of the entries 0 and ± 1. Now write

$$\det \mathcal{B} := |\det M(\mathcal{B},\mathbb{Q},T)|.$$

It is shown in [13] that $\det \mathcal{B}$ does not depend on the choice of the spanning tree T. The following theorem is a simple generalization of the characterization of directed and undirected cycle basis via determinants (see [13]).

thmCycleDet **Theorem 2.8.** *A set $\mathcal{B} \subset \mathcal{C}(G,\mathbb{F})$ is an $\mathbb{F}$-basis if and only if $\det \mathcal{B} \not\equiv 0 \mod \chi(\mathbb{F})$ where $\chi(\mathbb{F})$ denotes the characteristic of the field $\mathbb{F}$.*

Proof. It is easy to see that $\mathcal{B}$ is a $\mathbb{F}$-vector space basis of $\mathcal{C}(G,\mathbb{F})$ if and only if $M(\mathcal{B},\mathbb{F},T)$ is invertible as a matrix over $\mathbb{F}$ for a given spanning tree T of G. This holds true if and only if $\det M(\mathcal{B},\mathbb{F},T) \neq 0$ which is equivalent to

$$\det M(\mathcal{B},\mathbb{Q},T) \not\equiv 0 \mod \chi(\mathbb{F}).$$

This directly implies the theorem. □

Besides the $\mathbb{F}$ cycle basis, there are also stricter notions of cycle basis. The following definitions can be found in [13].

Definition 2.9 (Integral cycle basis)**.** A set $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{C}(G,\mathbb{Q})$ of oriented circuits of a graph G is an *integral cycle basis* of G if every oriented circuit C of G can be written as an integer linear combination of circuits in $\mathcal{B}$, i.e., there exist $\lambda_i \in \mathbb{Z}$ s.t.

$$C = \lambda_i C_i + \ldots + \lambda_r C_r.$$

Proposition 2.10 ([12])**.** *Let $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{C}(G,\mathbb{Q})$ be a set of oriented circuits of a graph G. T.f.a.e.:*

1. *$\mathcal{B}$ is an integral basis.*
2. *$\det \mathcal{B} = 1$.*

Definition 2.11 (Totally unimodular cycle basis)**.** A $\mathbb{Q}$-cycle basis $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{C}(G,\mathbb{Q})$ of a graph G is a *totally unimodular cycle basis* of G if its cycle matrix $M(\mathcal{B},\mathbb{Q})$ is totally unimodular, i.e., each sub determinant is either 0 or ± 1.

Definition 2.12 (Weakly fundamental cycle basis)**.** A set $\mathcal{B} = \{C_1, \ldots, C_r\}$ of circuits of a graph G is a *weakly fundamental cycle basis* of G if there exists some permutation σ such that for all $i = 2, \ldots, r$,

$$C_{\sigma(i)} \setminus (C_{\sigma(1)} \cup \ldots \cup C_{\sigma(i-1)}) \neq \emptyset.$$

Proposition 2.13 ([13])**.** *Let $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{C}(G, \mathbb{Q})$ be a set of oriented circuits of a graph G. T.f.a.e.:*

1. *$\mathcal{B}$ is a weakly fundamental cycle basis.*
2. *There exists a spanning tree T and a permutation of columns and rows such that $M(\mathcal{B}, \mathbb{Q}, T)$ is lower triangular.*

Definition 2.14 (Strictly fundamental cycle basis)**.** A set $\mathcal{B}$ of circuits of a graph G is a *strictly fundamental cycle basis* of G, if there exists some spanning tree $T \subseteq E$ such that $\mathcal{B} = \{C_e : e \in E \setminus T\}$, where C_e denotes the unique circuit in $T \cup \{e\}$.

Proposition 2.15 ([13])**.** *Let $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{C}(G, \mathbb{Q})$ be a set of oriented circuits of a graph G. T.f.a.e.:*

1. *$\mathcal{B}$ is a strictly fundamental cycle basis.*
2. *There exists a spanning tree T and a permutation of columns and rows such that $M(\mathcal{B}, \mathbb{Q}, T)$ is diagonal.*

2.3 Circuit preserving coverings

Let $G = (V, E)$ be a connected graph. If $\widetilde{G} = (\widetilde{V}, \widetilde{E})$ is a graph, and $\Psi : \widetilde{G} \to G$ is a surjective graph homomorphism such that Ψ is locally bijective (i.e., Ψ is bijective when restricted to the neighborhood of a single vertex), then the pair $(\widetilde{G}, \Psi)$ is called a *covering* of G.

Since Ψ is locally bijective, than it can be seen that $|\Psi^{-1}(x)|$ is constant for all vertices $x \in V$. If this constant value is m, we say that $(\widetilde{G}, \Psi)$ is a covering with m *sheets*, or is an m*-sheeted* covering. Here, m can be infinite.

We call a covering $(\widetilde{G}, \Psi)$ *trivial* if Ψ restricted to any connected component of $\widetilde{G}$ is a graph isomorphism. We say it is *non-trivial* otherwise, i.e., if there is at least one connected component of $\widetilde{G}$ on which Ψ is not one-to-one.

Let $\mathcal{B} \subset \mathcal{S}(G)$ be a set of circuits. We say a covering $(\widetilde{G}, \Psi)$ is a $\mathcal{B}$ *preserving covering* of G if for all circuits $C = (x_1, \ldots, x_n) \in \mathcal{B}$ and all $\widetilde{x}_1 \in \widetilde{V}$ with $\Psi(\widetilde{x}_1) = x_1$, there exist a circuit $\widetilde{C} = (\widetilde{x}_1, \ldots, \widetilde{x}_n) \in \widetilde{V}$ s.t. $\Psi(\widetilde{x}_k) = x_k$ for all k. Note in particular that every circuit in the pre-image of C has length equal to the length of C.

2.4 Path homology of graphs

In this section, we will give definitions for a homology theory on graphs that has been developed in recent years, called *path homology.* See [7, 6, 8]. This homology theory is most naturally described for directed graph, and the homology of an undirected graph is obtained by orienting each edge of an undirected graph in both possible directions.

For a directed graph $G = (V, E)$ (without self-loops) we start by defining an *elementary m-path* on V to be a sequence $i_0, ..., i_m$ of $m + 1$ vertices of V. For a field $\mathbb{F}$ we define the $\mathbb{F}$-linear space Λ_m to consist of all formal linear combinations of elementary m-paths with coefficients from $\mathbb{F}$. We identify an elementary m-path as an element of Λ_m denoted by $e_{i_0 \ldots i_m}$, and $\{e_{i_0 \ldots i_m} : i_0, ..., i_m \in V\}$ is a basis for Λ_m. Elements of Λ_m are call *m-paths*, and a typical m-path p can be written as

$$p = \sum_{i_0, \ldots, i_m \in V} a_{i_0 \ldots i_m} e_{i_0 \ldots i_m}, \quad a_{i_0 \ldots i_m} \in \mathbb{F}.$$

Note that Λ_0 is the set of all formal linear combinations of vertices in V.

We define the *boundary operator* $\partial : \Lambda_m \to \Lambda_{m-1}$ to be the $\mathbb{F}$-linear map that acts of elementary m-paths by

$$\partial e_{i_0 \ldots i_m} = \sum_{k=0}^{m} (-1)^k e_{i_0 \ldots \hat{i}_k \ldots i_m},$$

where $\hat{i}_k$ denotes the omission of index i_k.

For convenience, we define $\Lambda_{-1} = 0$ and $\partial : \Lambda_0 \to \Lambda_{-1}$ to be the zero map.

It can be checked that $\partial^2 = 0$, so that the Λ_m give a chain complex (see [8]). When it is important to make the distinction, we will use ∂_m to denote the boundary map on Λ_m, $\partial_m : \Lambda_m \to \Lambda_{m-1}$.

An elementary m-path $i_0...i_m$ is called *regular* if $i_k \neq i_{k+1}$ for all k, and is called *irregular* otherwise. Let I_m be the subspace of Λ_m spanned by all irregular m-paths, and define

$$\mathcal{R}_m = \Lambda_m / I_m.$$

The space $\mathcal{R}_m$ is isomorphic to the span of all regular m-paths, and the boundary map ∂ is naturally defined on $\mathcal{R}_m$, treating any irregular path resulting from applying ∂ as 0.

In the graph $G = (V, E)$, call an elementary m-path $i_0...i_m$ *allowed* if $i_k i_{k+1} \in E$ for all k. Define $\mathcal{A}_m$ to be the subspace of $\mathcal{R}_m$ given by

$$\mathcal{A}_m = \text{span}\{e_{i_0...i_m} : i_0...i_n \text{ is allowed}\}.$$

The boundary map ∂ on $\mathcal{A}_m$ is simply the restriction of the boundary map on $\mathcal{R}_n$, however, it can be the case that the boundary of an allowed m-path is not an allowed $(m-1)$-path. So we make one further restriction, and call an elementary m-path p *∂-invariant* if ∂p is allowed. We define

$$\Omega_m = \{p \in \mathcal{A}_n : \partial p \in \mathcal{A}_{m-1}\}.$$

Then it can be seen that $\partial \Omega_m \subseteq \Omega_{m-1}$. The Ω_m with the boundary map ∂ give us our chain complex of ∂-invariant allowed paths from which we will define our homology:

$$\cdots \Omega_m \xrightarrow{\partial} \Omega_{m-1} \xrightarrow{\partial} \cdots \to \Omega_1 \to \Omega_0 \to 0.$$

Observe that Ω_0 is the space of all formal linear combinations of vertices of G, and Ω_1 is space of all formal linear combinations of edges of G. We can now define the homology groups of this chain complex.

Definition 2.16 (path homology)**.** The *path homology groups* of the graph G over the field $\mathbb{F}$ are

$$H_n(G, \mathbb{F}) = Ker\ \partial|_{\Omega_n} / Im\ \partial|_{\Omega_{n+1}}.$$

A standard fact is that $dim\ H_0(G, \mathbb{F})$ counts the number of connected components of G ([7]).

As a standard example of some of the interesting behavior of this homology, consider the directed 4-cycle pictured below.

$$\begin{tikzcd} y \arrow[r] & z \\ x \arrow[u] \arrow[r] & w \arrow[u] \end{tikzcd}$$

Note that the 2-paths e_{xyz} and e_{xwz} are allowed by not ∂-invariant, since the edge xz is missing from the graph. However, if we consider the linear combination $e_{xyz} - e_{xwz}$, then note

$$\partial(e_{xyz} - e_{xwz}) = e_{yz} - e_{xz} + e_{xy} - (e_{wz} - e_{xz} + e_{xw}) = e_{yz} + e_{xy} - e_{wz} + e_{xw} \in \Omega_1.$$

Thus $e_{xyz} - e_{xwz} \in \Omega_2$, and it turns out that $\Omega_2 = span\{e_{xyz} - e_{xwz}\}$. It can be seen that $Ker\ \partial_1$ is spanned by $e_{xy} + e_{yz} - e_{zw} - e_{xz}$, which is precisely $Im\ \partial_2$, so $H_1(G, \mathbb{F}) = 0$. For circuits of length more than 4, then $dim\ H_1(G, \mathbb{F}) = 1$. See [8] for details.

2.5 Curvature bounds in graphs

For a graph $G = (V, E)$, the *graph Laplacian* is the operator Δ on the space of functions $f : V \to \mathbb{R}$ given by

$$\Delta f(x) = \sum_{y \sim x} (f(y) - f(x)).$$

The *Bakry-Émery operators* are defined via

$$\Gamma(f,g) := \frac{1}{2}\left(\Delta(fg) - f\Delta g - g\Delta f\right)$$
$$\Gamma_2(f,g) := \frac{1}{2}\left(\Delta\Gamma(f,g) - \Gamma(f,\Delta g) - \Gamma(g,\Delta f)\right).$$

We write $\Gamma(f) := \Gamma(f,f)$ and $\Gamma_2(f) = \Gamma_2(f,f)$.

Definition 2.17 (Bakry-Émery Curvature)**.** A graph G is said to satisfy the *curvature dimension inequality* $CD(K,n)$ for some $K \in \mathbb{R}$ and $n \in (0,\infty]$ if for all f,

$$\Gamma_2(f) \geq \frac{1}{n}(\Delta f)^2 + K \cdot \Gamma(f).$$

Theorem 2.18 (Bonnet-Myers Theorem, Corollary 2.2 of [15])**.** *Let G be a graph satisfying $CD(K,\infty)$ for some $K > 0$, and with maximum degree Deg_{max}. Then*

$$diam(G) \leq \frac{2Deg_{max}}{K}.$$

3 Curvature and coverings

Theorem 3.1. *Suppose a finite graph satisfies $CD(K,\infty)$ for some $K > 0$. Then, there exists no infinite covering of G preserving all 3- and 4-cycles.*

Proof. Suppose there exists an infinite covering $(\widetilde{G},\Psi)$ of G preserving all 3- and 4- cycles. Then, $\widetilde{G}$ is locally isomorphic to G and thus satisfying the same curvature bound $CD(K,\infty)$. We observe that $\widetilde{G}$ has bounded vertex degree Deg. Now, Theorem 2.18 implies $diam(\widetilde{G}) \leq \frac{2Deg}{K}$ and thus finiteness of $\widetilde{G}$. This is a contradiction and therefore proves that there is no infinite covering of G preserving 3- and 4-cycles. $\square$

4 Coverings and circuit basis

For gain graphs, there is a natural construction of a covering of the graph that is derived from the gain function. This construction is given in [10] and is a variant on a construction from [9].

Definition 4.1. Let $G = (V,E)$ be a graph

1. Let ϕ be a gain function on G into group Γ. We define the *derived graph*, denoted $G^\phi = (V^\phi, E^\phi)$ as

$$V^\phi = V \times \Gamma$$
$$E^\phi = \left\{\{(u,g),(v,g\phi(uv))\} : uv \in E, g \in \Gamma\right\}.$$

2. For ϕ a gain function into the symmetric group S_n, $\sigma_{uv} = \phi(uv)$, we define the *permutation derived graph* $G^\sigma = (V^\sigma, E^\sigma)$ as

$$V^\sigma = V \times \{1,...,n\}$$
$$E^\sigma = \left\{\{(u,i),(v,\sigma_{uv}(i))\} : uv \in E, i = 1,...,n\right\}.$$

Note that, as pointed out in [10], a permutation derived graph is *not* simply a derived graph where the associated group is the symmetric group. Indeed, the latter would have $n!\cdot|V|$ vertices, while the permutation derived graph has $n|V|$ vertices. Also, in the definition, we allow $n = \infty$, and identify S_∞ with the infinite symmetric group $Sym(\mathbb{Z})$, in which case G^σ is an infinite graph.

There is a natural projection $\Psi : G^\phi \to G$ given by $\Psi((u,g)) = u$. It is clear, as noted in [10], that Ψ is a covering map, so that (G^ϕ,Ψ) is a covering of G with $|\Gamma|$ sheets. Similarly, G^σ is a covering with n sheets.

We give a slight variation of our definition of the order of a cycle in a gain graph when the group is the symmetric group S_n. For a gain graph where each edge has an associated permutation σ_{uv}, and for a circuit

C of the graph, we denote by σ_C the composition of all the permutations going around the cycle. For an element $i \in \{1, ..., n\}$ we denote

$$o_i(C) = \min\{k : \sigma_C^k(i) = i\}.$$

Note that for $i \neq j$, $o_i(C)$ may not be equal to $o_j(C)$, depending on the action of σ_C on i and j.

Lemma 4.2. *Let G be a gain graph with group S_m and corresponding permutation derived covering (G^σ, Ψ). Given a circuit C of length n of G, the pre-image $\Psi^{-1}(C)$ consists of a collection of vertex disjoint circuits $\{\widetilde{C}\}$, where if $\widetilde{C}$ contains the vertex (x_1, i) for some $i \in \{1, ..., m\}$, then the length of $\widetilde{C}$ is $o_i(C) \cdot n$. In the case $o_i(C)$ is infinite, $\Psi^{-1}(C)$ contains an infinite path.*

Proof. Let $C = (x_1, ..., x_n)$ be a circuit of G, and fix $i \in \{1, ..., m\}$. Set $i_1 = i$, and define $i_{j+1} = \sigma_{x_j x_{j+1}}(i_j)$ where the index j on the x's is taken $\pmod n$. Then since $x_j x_{j+1} \in E$ for all j, then $\{(x_j, i_j), (x_{j+1}, i_{j+1})\} \in E^\sigma$ for all j by definition. Now we ask, when (if ever) does the sequence $\widetilde{C} = ((x_1, i_1), (x_2, i_2), ...)$ return to its starting point at (x_1, i_1). Clearly, for this to be the case, the index j satisfies $j \equiv 1 \pmod n$, and every time j becomes $1 \pmod n$, the associated element gets mapped by σ_C. So the sequence comes back to (x_1, i_1) when the x_j have come back to x_1 $o_i(C)$ times. The sequence cannot intersect itself at any earlier point by minimality of $o_i(C)$. Therefore clearly $\widetilde{C}$ is a circuit of length $o_i(C) \cdot n$. If there are multiple distinct circuits in $\Psi^{-1}(C)$, it is clear that they are vertex disjoint by the definition of G^σ. □

A slight modification of the above proof can be applied to derived graphs (not permutation derived) for any group.

Lemma 4.3. *Let (G, ϕ, Γ) be a gain graph, with (G^ϕ, Ψ) the corresponding derived covering. Given a circuit C of length n of G, the pre-image $\Psi^{-1}(C)$ consists of a collection of vertex disjoint circuits $\{\widetilde{C}\}$ where each $\widetilde{C}$ is of length $o_\phi(C) \cdot n$. In the case $o_\phi(C)$ is infinite, $\Psi^{-1}(C)$ contains an infinite path.*

We remark that for the derived cover (G^ϕ, Ψ), every circuit in Ψ^{-1} has the same length, $o_\phi(C) \cdot n$, but for the permutation derived cover (G^σ, Ψ), $\Psi^{-1}(C)$ may contain circuits of different lengths, if the permutation σ_C has different orders for different elements of $\{1, ..., m\}$.

Corollary 4.4. *Let (G, ϕ, Γ) be a gain graph with (G^ϕ, Ψ) the associated derived covering, and if Γ is some permutation group, let (G^σ, Ψ) be the associated permutation derived covering. Then for both G^ϕ and G^σ, if $\mathcal{B}$ is a collection of circuits of G, then ϕ is balanced on $\mathcal{B}$ if and only if the covering preserves $\mathcal{B}$.*

Proof. For the permutation derived covering G^σ, a cycle C of length n is balanced if and only of σ_C is the identity, which holds if and only if $\sigma_C(i) = i$ for all $i \in \{1, ..., m\}$, in other words, $o_i(C) = 1$ for all i, so that $\Psi^{-1}(C)$ is a collection of vertex disjoint cycles of length n by Lemma 4.2. This is the definition of C being preserved under the cover.

A similar argument works for G^ϕ. □

The following theorem from [10] shows that permutation derived coverings account for all possible coverings of a graph.

Theorem 4.5 (Theorem 2 of [10])**.** *Let $(\widetilde{G}, \Psi)$ be a covering of G with m sheets. Then there is a gain function into the symmetric group S_m assigning each edge uv a permutation $\sigma_{uv} \in S_m$ on G such that the permutation derived graph G^σ is isomorphic to $\widetilde{G}$.*

We remark here that m need not be finite, in which case the associated symmetric group can be identified with $Sym(\mathbb{Z})$.

Theorem 4.6. *Let $G = (V, E)$ be a connected graph and let $\mathcal{B}$ be a set of circuits. T.f.a.e:*

1. *There exists no non-trivial covering of G preserving $\mathcal{B}$*
2. *The circuit set $\mathcal{B}$ is a strong Γ circuit generator of G for all groups Γ.*
3. *The circuit set $\mathcal{B}$ is a strong $Sym(\mathbb{Z})$ circuit generator of G.*
4. *The circuit set $\mathcal{B}$ is a strong $\Gamma(G, \mathcal{B})$ circuit generator of G.*

Proof. We begin by proving 1 $\implies$ 2 by contraposition. Suppose that there is some group Γ for which $\mathcal{B}$ is not a strong Γ circuit generator. That means there is a gain $\phi : \vec{E} \to \Gamma$ that is balanced on $\mathcal{B}$, but that is unbalanced on some circuit, say C of G. Then $o_\phi(C) > 1$. Consider the derived covering (G^ϕ, Ψ). By Corollary 4.4, this covering preserves $\mathcal{B}$. By Lemma 4.3, since $o_\phi(C) > 1$, $\Psi^{-1}(C)$ contains a circuit of length strictly larger than C. This implies that the covering is non-trivial.

That 2 $\implies$ 3 is immediate.

Now we prove 3 $\implies$ 1 by contraposition. Suppose that there is a non-trivial covering of $(\widetilde{G}, \Psi)$ of G that preserves $\mathcal{B}$, and suppose it has m sheets. Then by Theorem 4.5, this covering can be realized as a permutation derived covering (G^σ, Ψ) for some assignment $\sigma_{uv} \in S_m$ for all $uv \in E$. Since the covering preserves $\mathcal{B}$, then every circuit from $\mathcal{B}$ is balanced under this assignment by Corollary 4.4. We need to find some circuit of G, not in $\mathcal{B}$, that is unbalanced. Since (G^σ, Ψ) is a non-trivial covering, there exist two vertices, call them $\widetilde{x}_1$ and $\widetilde{x}_k$ in the same connected component of $\widetilde{G}$ such that $\Psi(\widetilde{x}_1) = \Psi(\widetilde{x}_k)$. Since these belong to the same component, there is a path connecting $\widetilde{x}_1$ to $\widetilde{x}_k$. Since the covering preserves edges, and $\Psi(\widetilde{x}_1) = \Psi(\widetilde{x}_k)$, the image of this path is a circuit of G. Call this circuit C, and suppose it has length n, and let $x_1 = \Psi(\widetilde{x}_1) = \Psi(\widetilde{x}_k)$. By Lemma 4.2, the pre-image $\widetilde{C}$ of C is a circuit of length $o_i(C) \cdot n$ for some i. But since the pre-image of x_1 has at least two vertices, then $\widetilde{C}$ has strictly more vertices than C, implying $o_i(C) > 1$. Thus the permutation σ_C cannot be the identity, so C is not balanced. So $\mathcal{B}$ is not a strong S_m circuit generator, for any m. Thus $\mathcal{B}$ is not a $Sym(\mathbb{Z})$ circuit generator either, by Proposition 2.6.

It is clear that 2 $\implies$ 4, and 4 $\implies$ 2 by Proposition 2.3. □

CoverWeakGen **Theorem 4.7.** *Let $G = (V, E)$ be a connected graph and let $\mathcal{B}$ be a set of circuits. T.f.a.e:*

1. *There exists no non-trivial infinite covering of G preserving $\mathcal{B}$*
2. *The circuit set $\mathcal{B}$ is a weak Γ circuit generator of G for all groups Γ.*
3. *The circuit set $\mathcal{B}$ is a weak $Sym(\mathbb{Z})$ circuit generator of G.*
4. *The circuit set $\mathcal{B}$ is a weak $\Gamma(G, \mathcal{B})$ circuit generator of G.*

Proof. The proof will follow the proof of Theorem 4.6. We begin by proving 1 $\implies$ 2 by contraposition. Suppose that there is some group Γ for which $\mathcal{B}$ is not a weak Γ circuit generator. That means there is a gain $\phi : \vec{E} \to \Gamma$ that is balanced on $\mathcal{B}$, but there is a circuit, say C of G, that has infinite order under ϕ (in particular, Γ is an infinite group). Consider the derived covering (G^ϕ, Ψ), which is an infinite cover of G. By Corollary 4.4, this covering preserves $\mathcal{B}$. By Lemma 4.3, since the order of C is infinite, $\Psi^{-1}(C)$ contains an infinite path. This implies that the covering is non-trivial.

That 2 $\implies$ 3 is immediate.

Now we prove 3 $\implies$ 1 by contraposition. Suppose that there is a non-trivial infinite covering of $(\widetilde{G}, \Psi)$ of G that preserves $\mathcal{B}$. Then by Theorem 4.5, this covering can be realized as a permutation derived covering (G^σ, Ψ) for some assignment $\sigma_{uv} \in Sym(\mathbb{Z})$ for all $uv \in E$. Since the covering preserves $\mathcal{B}$, then every circuit from $\mathcal{B}$ is balanced under this assignment by Corollary 4.4. We need to find some circuit C where σ_C has infinite order. Since (G^σ, Ψ) is a non-trivial infinite covering, there exist an infinite path containing an infinite sequence $\widetilde{x}_1, \widetilde{x}_2, ...$ of vertices, such that $\Psi(\widetilde{x}_k) = x_1 \in V$ for all k. Since the covering preserves edges, and $\Psi(\widetilde{x}_1) = \Psi(\widetilde{x}_k)$ for all k, the image of this path is a circuit of G. Call this circuit C, and suppose it has length n. By Lemma 4.2, the order of σ_C is infinite.

It is clear that 2 $\implies$ 4, and 4 $\implies$ 2 by Proposition 2.3. □

5 Cycle bases and gain graphs

sec:BasisGain **Theorem 5.1.** *Every weakly fundamental cycle basis is a combinatorial circuit generator.*

Proof. Let $\{C_1, \ldots, C_r\}$ be a weakly fundamental cycle basis. Due to the weakly fundamental property, we can assume without obstruction that C_r contains an edge $e = x_1x_2$ not contained in all other C_i. We assume by induction that the theorem holds true for all cyclomatic numbers smaller than r. Let C be a circuit. Due to induction, we can assume that C contains e since otherwise C can be be represented by $C_1, \dots, C_{r-1}$ and we can delete e and C_r from the graph and decrease the cyclomatic number. We aim to write

$$C = (\dots(C_r \oplus K_1) \oplus \dots) \oplus K_n \tag{1}$$

eq:weaklyFun

with $K_i \in \langle C_1, \dots, C_{r-1}\rangle$. This would prove the theorem since due to induction over the cyclomatic number, we can assume that $K_1, \dots, K_n$ are in the closure of $\mathcal{B}$ under $\oplus$.

We inductively define K_k and

$$C_r^k := (\dots(C_r \oplus K_1) \oplus \dots) \oplus K_k$$

for all k. We start with $C_r^0 = C_r$. Write

$$C_r^k = (x_1, \dots, x_N).$$

Now, we can write

$$C = (x_1, \dots, x_K, y_1, \dots, y_L, x_M, \dots)$$

with $y_i \notin C_r^k$, and after x_M are some further elements of C that we need not specify. Define

$$K_{k+1} := (y_1, \dots, y_L, x_M, \dots, x_K).$$

Observe that K_{k+1} does not contain the edge x_1x_2, so K_{k+1} belongs to the closure of $\mathcal{B}$ under $\oplus$ by the induction hypothesis. Note also that C_r^k and K_{k+1} belong to a theta graph whose paths are given by $p_1 = x_K, x_{K+1}, \dots x_{M-1}, x_M$; $p_2 = x_K, y_1, \dots, y_L, x_M$; and $p_3 = x_M, \dots x_N, x_1, \dots, x_K$. Define

$$C_r^{k+1} := C_r^k \oplus K_{k+1} = (x_1, \dots, x_K, y_1, \dots, y_L, x_M, \dots, x_N).$$

Observe that C_r^{k+1} and C share a longer path than C_r^k and C do. Therefore, this process will terminate yielding (1). This finishes the proof. ☐

Theorem 5.2. *Every combinatorial circuit generator is a Γ-circuit generator for all groups Γ.*

Proof. Let $\mathcal{B}$ a combinatorial circuit generator and define inductively $\mathcal{B}_0 := \mathcal{B}$ and

$$\mathcal{B}_{k+1} := \mathcal{B}_k \cup (\mathcal{B}_k \oplus \mathcal{B}_k) = \mathcal{B}_k \cup \{C_1 \oplus C_2 : C_1, C_2 \in \mathcal{B}_k\}.$$

Since $\mathcal{B}$ is a combinatorial circuit generator, we have $\bigcup_{k\in\mathbb{N}} \mathcal{B}_k = \mathcal{S}(G)$. Suppose $\mathcal{B}$ is not a Γ-circuit generator for some group Γ. Let K be the minimal number such that there exists a circuit $C \in \mathcal{B}_k$ and $\phi \in \Phi(G, \Gamma)$ such that ϕ is balanced on $\mathcal{B}$ but not on C. Due to minimality, ϕ is also balanced on B_{k-1}. Since $C \in \mathcal{B}_k \setminus \mathcal{B}_{k-1}$, we can write $C = C_1 \oplus C_2$ for some $C_1, C_2 \in \mathcal{B}_{k-1}$, where C, C_1, and C_2 all belong to a single theta graph in G. Thus, we can write $C_i = (p_i, p_3)$ for some paths p_1, p_2, p_3. Moreover, $C = (p_1, p_2^{-1})$ where p_2^{-1} is the inverse of the path p_2. Due to induction, we have $\phi(p_i)\phi(p_3) = e_\Gamma$ and thus $\phi(p_1)\phi(p_2^{-1}) = e_\Gamma$ showing that ϕ is balanced on C. This is a contradiction and proves that $\mathcal{B}$ is a Γ-circuit generator. This finishes the proof. ☐

CircuitField

Theorem 5.3. *Let $\mathbb{F}$ be a field with additive group Γ. Let $\mathcal{B} = \{C_1, \dots, C_r\} \subset \mathcal{C}(G, \mathbb{F})$ be a cyclomatic circuit set of a graph G. T.f.a.e.:*

1. *$\mathcal{B}$ is a $\mathbb{F}-$cycle basis.*
2. *$\mathcal{B}$ is a Γ-circuit generator.*

Proof. $1 \Rightarrow 2$: We aim to show that every gain function ϕ is balanced on all cycles when assuming that ϕ is balanced on $\mathcal{B}$. Now, ϕ is balanced on C if and only $\phi(C) = 0$. Since ϕ is linear and $\phi(C) = 0$ for all $C \in \mathcal{B}$ due to assumption, we infer that $\phi(C) = 0$ for all $C \in \operatorname{span}(\mathcal{B}) = \mathcal{C}(G, \mathbb{F})$ since we assume that $\mathcal{B}$ is a $\mathbb{F}$-cycle basis of G.

$2 \Rightarrow 1$: We indirectly prove the claim. Assume $\mathcal{B}$ is not a $\mathbb{F}$-cycle basis. Then, there exists a basis $\widetilde{\mathcal{B}}$ and $C_0 \in \widetilde{\mathcal{B}}$ s.t. $\mathcal{B} \subset \operatorname{span}(\widetilde{\mathcal{B}} \setminus \{C_0\})$. The matrix $M(\widetilde{\mathcal{B}}, \mathbb{F})$ is a $r \times |E|$ matrix with full rank r. Hence, the multiplication with the gain functions $M(\widetilde{\mathcal{B}}, \mathbb{F}) : \Phi(G, \mathbb{F}) \to \mathbb{F}^{\widetilde{\mathcal{B}}}$ is surjective. In particular, there exists $\phi \in \Phi(G, \mathbb{F})$ s.t. for $C \in \widetilde{\mathcal{B}}$,

$$\phi(C) = \left[M(\widetilde{\mathcal{B}}, \mathbb{F})\phi\right](C) = \begin{cases} 1 & : C = C_0 \\ 0 & : C \in \widetilde{\mathcal{B}} \setminus \{C_0\} \end{cases}.$$

This implies $\phi(C) = 0$ for all $C \in \mathcal{B}$ and $\phi(C_0) = 1$ which proves that $\mathcal{B}$ is not a Γ circuit generator. This finishes the proof. ☐

cor:CurvBasis **Corollary 5.4.** *If a graph G satisfies $CD(K, \infty)$ for some $K > 0$, and if $\mathbb{F}$ is a field with characteristic 0, then there is an $\mathbb{F}$-cycle basis consisting of only 3- and 4-cycles.*

Proof. Let $\mathcal{B}$ be the set of 3- and 4-cycles of G. By Theorem 3.1, there is no infinite cover of G preserving preserving $\mathcal{B}$. Then take Γ to be the additive group of the field $\mathbb{F}$. By Theorem 4.7, the set $\mathcal{B}$ is a weak Γ-circuit generator of G. Since $\mathbb{F}$ has characteristic 0, no non-trivial element has finite order, so in fact $\mathcal{B}$ is a strong Γ-circuit generator of G. Then by Theorem 5.3, $\mathcal{B}$ contains a $\mathbb{F}$-cycle basis. ☐

Since all proper finitely generated subgroups of $(\mathbb{Q}, +)$ are isomorphic to $\mathbb{Z}$ we immediately obtain the following corollary by using Theorem 2.8.

itGenerators **Corollary 5.5.** *Let G be a finite graph and let $\mathcal{B} \subset \mathcal{S}(G)$ be a cyclomatic circuit set. T.f.a.e.:*

1. $\det \mathcal{B} \neq 0$.
2. $\mathcal{B}$ is a $\mathbb{Q}$-circuit generator.
3. $\mathcal{B}$ is a $\mathbb{Z}$-circuit generator.

cor:qnDetAbel **Corollary 5.6.** *Let $n \in \mathbb{N}$ and let q be a prime number. Let Γ be the cyclic group with q^n elements and let G be a graph with a cyclomatic circuit set $\mathcal{B} \subset \mathcal{S}(G)$. T.f.a.e.:*

1. $\det \mathcal{B} \not\equiv 0 \mod q$.
2. *$\mathcal{B}$ is a Γ circuit generator.*

Proof. The implication $2 \Rightarrow 1$ follows from Theorem 2.8, Theorem 5.3 and Proposition 2.5 since the additive group of $\mathbb{F}_q$ is a subgroup of Γ.

We next prove $1 \Rightarrow 2$. We say $\mathcal{B} = \{C_1, \ldots, C_r\}$. We canonically identify the elements of Γ with $\{0, \ldots, q^n - 1\} \subset \mathbb{Q}$ via a function $\eta : \Gamma \to \mathbb{Q}$. Since $\det \mathcal{B} \not\equiv 0 \mod q$, the matrix $M(\mathcal{B}, \mathbb{Q}, T)$ is invertible for a spanning tree T and every circuit $C \in \mathcal{C}(G, \mathbb{Q})$ can uniquely be written as

$$C = \frac{\lambda_1}{\det \mathcal{B}} C_1 + \ldots + \frac{\lambda_r}{\det \mathcal{B}} C_r$$

with integers $\lambda_1, \ldots, \lambda_r$. Suppose $\mathcal{B}$ is not a Γ circuit generator. Then there exists $\phi \in \Phi(G, \Gamma)$ that is balanced on $\mathcal{B}$ but not on some $C \in \mathcal{S}(G)$. Observe $\eta \circ \phi \in \Phi(G, \mathbb{F})$ and ϕ is balanced on a circuit C if and only if $(\eta \circ \phi)(C) \equiv 0 \mod q^n$. Therefore we can write, $(\eta \circ \phi)(C_i) = c_i q^n$ for integers c_i and $i = 1, \ldots, c_r$. This implies

$$(\eta \circ \phi)(C) \cdot \det \mathcal{B} = (\lambda_1 c_1 + \ldots + \lambda_r c_r) q^n \in q^n \mathbb{Z}.$$

Since $\det \mathcal{B} \not\equiv 0 \mod q$ and since q is prime, we obtain $(\eta \circ \phi)(C) \equiv 0 \mod q^n$ which is equivalent to balance of ϕ on C. This contradicts the assumption that $\mathcal{B}$ is not a Γ circuit generator. This finishes the proof. ☐

Using the above theorem, we can fully characterize Γ circuit bases for Abelian groups Γ

Corollary 5.7. *Let Γ be an Abelian group and let $\mathcal{B} = \{C_1, ..., C_r\} \subset \mathcal{S}(G)$ be a cyclomatic circuit set of a graph G. T.f.a.e.:*

1. *$g^{\det \mathcal{B}} \neq e_\Gamma$ for all $g \in \Gamma \setminus \{e_\Gamma\}$*

2. *$\mathcal{B}$ is a Γ-circuit generator.*

Proof. We start proving $2 \Rightarrow 1$. First suppose $\det \mathcal{B} \neq 0$. Suppose $g^{\det \mathcal{B}} = e_\Gamma$ for some $g \in \Gamma \setminus \{e_\Gamma\}$. Then, we can assume without obstruction that g has prime order q and $\det \mathcal{B} \equiv 0 \mod q$. Theorem 2.8 and Theorem 5.3 imply that $\mathcal{B}$ is not a $\langle g \rangle$ circuit generator. Since $\langle g \rangle$ is a subgroup of Γ, Proposition 2.5 implies that $\mathcal{B}$ is not a Γ circuit generator. Now suppose $\det \mathcal{B} = 0$. If there exists $g \in \Gamma \setminus \{e_\Gamma\}$ with finite order, we can proceed as in the case above. Otherwise, there exists an element $g \in \Gamma$ with infinite order and therefore $(\mathbb{Z}, +) \cong \langle g \rangle \subset \Gamma$. Corollary 5.5 yields that $\mathcal{B}$ is not a $\langle g \rangle$ circuit basis since $\det \mathcal{B} = 0$. Since $\langle g \rangle$ is a subgroup of Γ, Proposition 2.5 also yields that $\mathcal{B}$ is not a Γ circuit basis. This finishes the proof of $2 \Rightarrow 1$.

We now prove $1 \Rightarrow 2$. W.l.o.g., Γ is not the one element group and therefore, we can assume $\det \mathcal{B} \neq 0$. Due to Proposition 2.6 we can assume without obstruction that Γ is finitely generated. Therefore, Γ is isomorphic to $\mathbb{Z}^n \oplus \mathbb{Z}_{q_1} \oplus \ldots \oplus \mathbb{Z}_{q_t}$ where $q_1, \ldots, q_t$ are powers of prime numbers and $\mathbb{Z}_q$ is the cyclic group with q elements. Due to Proposition 2.5 it suffices to show that $\mathcal{B}$ is a $\mathbb{Z}$ circuit generator and a $\mathbb{Z}_{q_i}$ circuit generator for $i = 1, \ldots, r$. Corollary 5.5 and $\det \mathcal{B} \neq 0$ implies that $\mathcal{B}$ is a $\mathbb{Z}$ circuit generator. Let $i \in \{1, \ldots r\}$. We know $q_i = p^n$ for some prime p and we know $\mathbb{Z}_{q_i}$ has order p. Therefore, assertion 1 of the theorem implies $\det \mathcal{B} \not\equiv 0 \mod p$. Now, Corollary 5.6 applied to $\mathbb{Z}_{q_i}$ yields that $\mathcal{B}$ is a $\mathbb{Z}_{q_i}$ circuit generator. Since i is arbitrary, this shows that $\mathcal{B}$ is a Γ circuit generator. This finishes the proof. □

6 Cycle basis and homology

BasisHomology

We will be dealing with undirected graphs $G = (V, E)$, so when considering the homology groups, we will view G as a directed graph in which each edge corresponds to two directed edges, one in each direction.

gyCycleSpace

Theorem 6.1. *Let $\mathbb{F}$ be a field with characteristic not equal to 2, and let $\mathcal{C}(G, \mathbb{F})$ denote the $\mathbb{F}$-cycle space of G. Let TS denote the subspace of $\mathcal{C}(G, \mathbb{F})$ that is generated by all simple triangles and squares of G. Then*

$$H_1(G, \mathbb{F}) \cong \mathcal{C}(G, \mathbb{F})/TS.$$

This section will be devoted to proving this theorem.

Recall that by definition,

$$H_1(G, \mathbb{F}) = Ker\ \partial_1 / Im\ \partial_2$$

where ∂_1 denotes the boundary operator on 1-paths and ∂_2 the boundary operator on 2-paths.

Observe that the space Ω_1 of 1-paths can be naturally identified with the space of functions from the edge set to the field $\mathbb{F}$; that is

$$\Omega_1 \cong \{\phi : \overrightarrow{E} \to \mathbb{F}\}.$$

This space can naturally be decomposed: define

$$\begin{aligned} \Omega_+ &= \{\phi \in \Omega_1 : \phi(xy) = \phi(yx) \text{ for all } x, y\} \\ \Omega_- &= \{\phi \in \Omega_1 : \phi(xy) = -\phi(yx) \text{ for all } x, y\}. \end{aligned}$$

Then it is clear that

$$\Omega_1 = \Omega_+ \oplus \Omega_-.$$

lem:ker

Lemma 6.2. *$Ker\ \partial_1 \cong \Omega_+ \oplus \mathcal{C}(G, \mathbb{F})$.*

Proof. The action of ∂_1 on Ω_1 can be given as

$$\partial_1 \phi = \sum_{x,y} \phi(xy)(e_y - e_x) = \sum_x \sum_y (\phi_{yx} - \phi_{xy}) e_x.$$

Therefore $\phi \in Ker\ \partial_1$ if and only if

$$\sum_{y \sim x} (\phi(yx) - \phi(xy)) = 0 \text{ for all } x.$$

In terms of the direct sum decomposition above, this then yields

$$\begin{aligned} Ker\ \partial_1 = &\{\phi : \phi(xy) = \phi(yx) \text{ for all } x, y\} \\ &\oplus \left\{ \phi : \phi(xy) = -\phi(yx) \text{ for all } x, y \text{ and } \sum_{y\sim x} \phi(xy) = 0 \text{ for all } x \right\}. \end{aligned}$$

The first term is clearly Ω_+ and the second is, by definition, the cycle space $\mathcal{C}(G, \mathbb{F})$. This gives the lemma. □

lem:im **Lemma 6.3.** *$Im\ \partial_2 \cong \Omega_+ \oplus TS$.*

Proof. First, any element of Ω_+ can be written

$$\sum_{x,y} \phi(xy)(e_{xy} + e_{yx}) = \partial_2 \left(\sum_{x,y} \phi(xy) e_{xyx} \right),$$

so $\Omega_+ \subset Im\ \partial_2$. It remains to show that the portion of $Im\ \partial_2$ that lies in Ω_- is equal to TS. That is, we must show that any element from $\mathcal{C}(G, \mathbb{F})$ is in $Im\ \partial_2$ if and only if it is in the space spanned by triangles and squares.

Suppose xyz is a triangle of G. Consider first any triangle, that is, some ϕ such that $\phi(xy) = \phi(yz) = \phi(zx) = -\phi(yx) = -\phi(zy) = -\phi(xz)$ and is 0 on all other edges. Note that

$$\begin{aligned} \partial_2(\phi(xy)(e_{xyz} - e_{zyx}) &= \phi(xy)e_{yz} - \phi(xy)e_{xz} + \phi(xy)e_{xy} - \phi(xy)e_{yx} + \phi_{xy}e_{zx} - \phi(xy)e_{yx} \\ &= \phi(xy)e_{xy} + \phi(yz)e_{yz} + \phi(zx)e_{zx} + \phi(yx)e_{yx} + \phi(zy)e_{zy} + \phi(xz)e_{xz} \end{aligned}$$

which is the triangle ϕ. So any triangle is contained in $Im\ \partial_2$.

Similarly if ϕ is an square $xyzw$ of G, i.e. $\phi(xy) = \phi(yz) = \phi(zw) = \phi(wx) = -\phi(yx) = -\phi(zy) = -\phi(wz) = -\phi(xw)$ and is 0 elsewhere, then in a similar way, it can be verified that

$$\partial_2(\phi(xy)(e_{xyz} - e_{xwz} - e_{zyx} + e_{zwx})) = \phi.$$

Therefore any square is in $Im\ \partial_2$ as well. It follows that the space $TS \subset Im\ \partial_2$.

Conversely, we must show $Im\ \partial_2 \subset \Omega_+ \oplus TS$. We already know $Im\ \partial_2 \subset Ker\ \partial_1 = \Omega_+ \oplus \mathcal{C}(G, \mathbb{F})$, so we will be done if we can show that any $\phi \in Im\ \partial_2 \cap \mathcal{C}(G, \mathbb{F})$ belongs to TS. Since $\phi \in Im\ \partial_2$ we can write

$$\begin{aligned} \phi &= \partial_2 \left(\sum_{xyz} a_{xyz} e_{xyz} \right) \\ &= \sum_{xyz} a_{xyz}(e_{yz} - e_{xz} + e_{xy}), \end{aligned}$$

where the sum is taken over ∂-invariant allowed paths xyz of G. Since Ω_2 consists only of ∂-invariant allowed elements, a_{xyz} is non-zero only for allowed paths xyz, implying that xy and yz are edges of G, and that either xz is an edge of G, or else the e_{xz} term cancels in the sum. We will therefore split the above sum into two parts,

$$\sum_{\substack{xyz \\ xz \in E(G)}} a_{xyz}(e_{yz} - e_{xz} + e_{xy}) + \sum_{\substack{xyz \\ xz \notin E(G)}} a_{xyz}(e_{yz} - e_{xz} + e_{xy}).$$

Since we are assuming $\phi \in \mathcal{C}(G, \mathbb{F})$ (in particular, $\phi(xz) = -\phi(zx)$) then it is clear that the first term above is a linear combination of triangles.

For the second term, since it is allowed, any e_{xz} term must cancel. Thus, for any xyz for which a_{xyz} is non-zero in the second sum, there must be some other allowed 2-path in which e_{xz} shows up as a term. Namely, there exists $w \neq y$ such that xwz is allowed in G, and the coefficient $a_{xwz} = -a_{xyz}$. It is clear then that the second sum is a linear combination of squares of G. Thus we have shown that $Im\ \partial_2 \cap \mathcal{C}(G, \mathbb{F}) \subset TS$, and we have shown the lemma. □

Lemma 6.2 and Lemma 6.3 together immediately give the proof of Theorem 6.1.

Theorem 6.1 along with Corollary 5.4 give the following corollary.

Corollary 6.4. *If G is a graph satisfying $CD(K,\infty)$ for some $K > 0$ and if $\mathbb{F}$ is a field with characteristic 0, then*

$$H_1(G, \mathbb{F}) = 0.$$

Of course, the converse of Corollary 6.4 does not hold. Indeed, it is well-known that in trees other than paths or the star on 4 vertices, there will typically be vertices with negative curvature (see, for instance [4]). However, all trees have trivial first homology [6].

It is natural to ask if the hypotheses of Corollary 6.4 can be weakened, but still obtain trivial first homology. Consider however the graph pictured below.

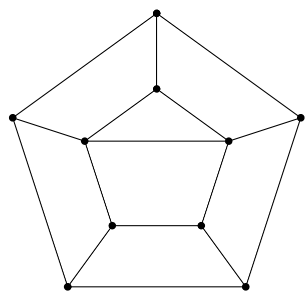

Computation shows that this graph has non-negative curvature at every vertex, and strictly positive curvature at some vertices. However $dim\ H_1(G, \mathbb{F}) = 1$, since the outer 5-cycle is not generated by any 3- or 4-cycles. This shows that we cannot weaken the hypothesis to simply non-negative curvature, even if there are some vertices with strictly positive curvature.

6.1 Clique Homology

A more commonly used notion of homology in graph theory is the *clique homology* coming from the *clique complex*, or *flag complex* of the graph. In this theory, the chain complex is

$$\cdots C_n \to C_{n-1} \to \cdots \to C_1 \to C_0 \to 0$$

where C_n is the space of all formal $\mathbb{F}$-linear combinations of n-cliques of the graph G. (Hence it is still the case that C_1 is all formal linear combinations of edges, and C_0 all formal linear combinations of vertices.) The boundary map ∂ of a clique is the sum of all its "faces," viewing the graph as a cell complex with an n-cell filling each n-clique. Then the clique homology groups are defined in the same way,

$$H_n^{\text{clique}}(G, \mathbb{F}) = Ker\ \partial|_{C_n} / Im\ \partial|_{C_{n+1}}.$$

Using techniques very similar to those used to prove Theorem 6.1, it is possible to prove an analogous theorem for clique homology.

Theorem 6.5. *Let $\mathbb{F}$ be a field with characteristic not equal to 2, and let $\mathcal{C}(G, \mathbb{F})$ denote the $\mathbb{F}$-cycle space of G. Let T denote the subspace of $\mathcal{C}(G, \mathbb{F})$ that is generated by all simple triangles of G. Then*

$$H_1^{clique}(G, \mathbb{F}) \cong \mathcal{C}(G, \mathbb{F})/T.$$

Hence for both the path and clique homology theories, the first homology group "counts" cycles of the graph, but there are certain types of cycles ignored depending on the theory; clique homology does not see triangles, and path homology sees neither triangles nor squares.

Observe in particular that the homology vanishing theorem for path homology, Corollary 6.4, does not hold for the clique homology (a simple 4-cycle being a counterexample). We take this as further evidence that the path homology is the more appropriate homology theory for graph theory.

References

BakryEmery85 [1] Dominique Bakry and Michel Émery. Diffusions hypercontractives. In *Séminaire de probabilités, XIX, 1983/84*, volume 1123 of *Lecture Notes in Math.*, pages 177–206. Springer, Berlin, 1985.

Bochner [2] S. Bochner. Curvature and Betti numbers. *Ann. of Math. (2)*, 49:379–390, 1948.

gZhaoKempton [3] Fan Chung, Wenbo Zhao, and Mark Kempton. Ranking and sparsifying a connection graph. *Internet Math.*, 10(1-2):87–115, 2014.

ing2016bakry [4] David Cushing, Shiping Liu, and Norbert Peyerimhoff. Bakry-Émery curvature functions of graphs. *arXiv preprint arXiv:1606.01496*, 2016.

Forman [5] Robin Forman. Bochner's method for cell complexes and combinatorial Ricci curvature. *Discrete Comput. Geom.*, 29(3):323–374, 2003.

Grigoryan2 [6] Alexander Grigor'yan, Yong Lin, Yuri Muranov, and Shing-Tung Yau. Homologies of path complexes and digraphs. *arXiv preprint: arXiv:1207.2834*, 2013.

Grigoryan1 [7] Alexander Grigor'yan, Yong Lin, Yuri Muranov, and Shing-Tung Yau. Cohomology of digraphs and (undirected) graphs. *Asian J. Math.*, 19(5):887–931, 2015.

Grigoryan3 [8] Alexander Grigor'yan, Yuri Muranov, and Shing-Tung Yau. Homologies of digraphs and künneth formulas. *preprint*, 2015.

Gross74 [9] Jonathan L. Gross. Voltage graphs. *Discrete Math.*, 9:239–246, 1974.

rossTucker77 [10] Jonathan L. Gross and Thomas W. Tucker. Generating all graph coverings by permutation voltage assignments. *Discrete Math.*, 18(3):273–283, 1977.

Kenyon [11] Richard Kenyon. Spanning forests and the vector bundle Laplacian. *Ann. Probab.*, 39(5):1983–2017, 2011.

n2003finding [12] Christian Liebchen. Finding short integral cycle bases for cyclic timetabling. In *European Symposium on Algorithms*, pages 715–726. Springer, 2003.

n2007classes [13] Christian Liebchen and Romeo Rizzi. Classes of cycle bases. *Discrete Applied Mathematics*, 155(3):337–355, 2007.

LinYau10 [14] Yong Lin and Shing-Tung Yau. Ricci curvature and eigenvalue estimate on locally finite graphs. *Math. Res. Lett.*, 17(2):343–356, 2010.

liu2016bakry [15] Shiping Liu, Florentin Münch, and Norbert Peyerimhoff. Bakry-emery curvature and diameter bounds on graphs. *arXiv preprint arXiv:1608.07778*, 2016.

ier2009ricci [16] Yann Ollivier. Ricci curvature of markov chains on metric spaces. *Journal of Functional Analysis*, 256(3):810–864, 2009.

998curvature [17] Michael Schmuckenschläger. Curvature of nonlocal markov generators. *Convex geometric analysis (Berkeley, CA, 1996)*, 34:189–197, 1998.

Zaslavsky [18] Thomas Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B*, 47(1):32–52, 1989.

Zaslavsky2 [19] Thomas Zaslavsky. Universal and topological gains for biased graphs. *Preprint*, 1989.