

MIXING TIME AND EIGENVALUES OF THE ABELIAN SANDPILE MARKOV CHAIN

DANIEL C. JERISON, LIONEL LEVINE, AND JOHN PIKE

ABSTRACT. The abelian sandpile model defines a Markov chain whose states are integer-valued functions on the vertices of a simple connected graph G . By viewing this chain as a (nonreversible) random walk on an abelian group, we give a formula for its eigenvalues and eigenvectors in terms of ‘multiplicative harmonic functions’ on the vertices of G . We show that the spectral gap of the sandpile chain is within a constant factor of the length of the shortest non-integer vector in the dual Laplacian lattice, while the mixing time is at most a constant times the smoothing parameter of the Laplacian lattice. We find a surprising inverse relationship between the spectral gap of the sandpile chain and that of simple random walk on G : If the latter has a sufficiently large spectral gap, then the former has a *small* gap! In the case where G is the complete graph on n vertices, we show that the sandpile chain exhibits cutoff at time $\frac{1}{4\pi^2} n^3 \log n$.

1. INTRODUCTION

Let $G = (V, E)$ be a simple connected graph with n vertices, one of which is designated the *sink* $s \in V$. A *sandpile* on G is a function

$$\sigma : V \setminus \{s\} \rightarrow \mathbb{N}$$

from the nonsink vertices to the nonnegative integers. In the abelian sandpile model [4, 11], certain sandpiles are designated as *stable*, and any sandpile can be *stabilized* by a sequence of local moves called *topplings*. (For the precise definitions see Section 2.) Associated to the pair (G, s) is a Markov chain whose states are the stable sandpiles. To advance the chain one time step, we choose a vertex v uniformly at random, increase $\sigma(v)$ by one, and stabilize. We think of $\sigma(v)$ as the number of sand grains at vertex v . Increasing it corresponds to dropping a single grain of sand on the pile. Topplings redistribute sand, and the role of the sink is to collect extra sand that falls off the pile.

This Markov chain $(\sigma_t)_{t \in \mathbb{N}}$ can be viewed as a random walk on a finite abelian group, the *sandpile group*, and consequently its stationary distribution is uniform

Date: November 2, 2015.

2010 Mathematics Subject Classification. 60J10; 82C20; 05C50.

Key words and phrases. abelian sandpile model, chip-firing, Laplacian lattice, mixing time, multiplicative harmonic function, pseudoinverse, sandpile group, smoothing parameter, spectral gap.

D.C. Jerison and J. Pike were supported in part by NSF grant DMS-0739164.

L. Levine was supported by NSF grant [DMS-1455272](#) and a Sloan Fellowship.

on the recurrent states. How long does it take to get close to this uniform distribution? We seek to answer this question in terms of graph-theoretic properties of G and algebraic properties of its Laplacian. What graph properties control the mixing time? In other words, how much randomly added sand is enough to make the sandpile ‘forget’ its initial state?

As we will see in Section 2, the characters of the sandpile group are indexed by *multiplicative harmonic functions* on G . These are functions

$$h : V \rightarrow \mathbb{C}^*$$

satisfying $h(s) = 1$ and

$$h(v)^{\deg(v)} = \prod_{w \sim v} h(w)$$

for all $v \in V$, where $\deg(v)$ is the number of edges incident to v , and we write $w \sim v$ if $\{v, w\} \in E$. There are finitely many such functions, and they form an abelian group whose order is the number of spanning trees of G . Associated to each such h is an eigenvalue of the sandpile chain,

$$\lambda_h = \frac{1}{n} \sum_{v \in V} h(v). \quad (1)$$

Our first result relates the spectral gap of the sandpile chain to the shortest vector in a lattice. Both gap and lattice come in two flavors: continuous time and discrete time. Writing \mathcal{H} for the set of multiplicative harmonic functions, the continuous time gap is defined as

$$\gamma_c = \min\{1 - \operatorname{Re}(\lambda_h) : h \in \mathcal{H}, h \neq 1\}$$

and the discrete time gap is defined as

$$\gamma_d = \min\{1 - |\lambda_h| : h \in \mathcal{H}, h \neq 1\}.$$

At times it will be useful to enumerate the vertices as $V = \{v_1, \dots, v_n\}$. In this situation we always take v_n to be the sink vertex. The *full Laplacian* of G is the $n \times n$ matrix

$$\bar{\Delta}(i, j) = \begin{cases} \deg(v_i), & i = j \\ -1, & v_i \sim v_j \\ 0, & \text{else.} \end{cases} \quad (2)$$

The *reduced Laplacian*, which we will denote by Δ , is the submatrix of $\bar{\Delta}$ omitting the last row and column (corresponding to the sink). Note that Δ is invertible, but $\bar{\Delta}$ is not, since its rows sum to zero.

The lattice that controls the continuous time gap is $\Delta^{-1}\mathbb{Z}^{n-1}$, the integer span of the columns of Δ^{-1} . To define the analogue for the discrete time gap, we consider the *Moore-Penrose pseudoinverse* $\bar{\Delta}^+$ of the full Laplacian. We also define:

- \mathbb{R}_0^n : the subspace of vectors in \mathbb{R}^n whose coordinates sum to zero,
- \mathbb{Z}_0^n : the set of vectors in \mathbb{R}_0^n with integer coordinates,
- \mathbb{W}^n : the orthogonal projection of \mathbb{Z}^n onto \mathbb{R}_0^n .

Theorem 1.1. *The spectral gap of the continuous time sandpile chain satisfies*

$$8 \frac{\|y\|_2^2}{n} \leq \gamma_c \leq 2\pi^2 \frac{\|y\|_2^2}{n},$$

where y is a vector of minimal Euclidean length in $(\Delta^{-1}\mathbb{Z}^{n-1}) \setminus \mathbb{Z}^{n-1}$.

The spectral gap of the discrete time sandpile chain satisfies

$$8 \frac{\|z\|_2^2}{n} \leq \gamma_d \leq 2\pi^2 \frac{\|z\|_2^2}{n},$$

where z is a vector of minimal Euclidean length in $(\bar{\Delta}^+\mathbb{Z}_0^n) \setminus \mathbb{W}^n$.

The upper and lower bounds in Theorem 1.1 match up to a constant factor. In Theorem 2.11 below, we will also prove a simple lower bound for the discrete time gap, namely

$$\gamma_d \geq \frac{8}{d_*^2 n} \tag{3}$$

where d_* is the penultimate term in the nondecreasing degree sequence. The reason for the appearance of the second-largest degree is that γ_d (unlike γ_c) does not depend on the choice of sink, which can be moved to the vertex of largest degree.

The standard bound of the mixing time in terms of spectral gap is $t_{\text{mix}} = O(\gamma_d^{-1} \log |\mathcal{G}|)$ where \mathcal{G} is the sandpile group. The size of this group—the number of spanning trees of G —can be exponential in $n \log d_*$, so the factor of $\log |\mathcal{G}|$ can be significant. The resulting upper bound on the mixing time obtained from (3) is often far from the truth. To improve it, our next result gives an upper bound for the L^2 mixing time of the sandpile chain in terms of a lattice invariant called the *smoothing parameter*. Let Λ be a lattice in \mathbb{R}^m , and let $V = \text{Span}(\Lambda) \subseteq \mathbb{R}^m$. Denote the dual lattice by $\Lambda^* = \{x \in V : \langle x, y \rangle \in \mathbb{Z} \text{ for all } y \in \Lambda\}$. For $s > 0$, the function

$$f_\Lambda(s) = \sum_{x \in \Lambda^* \setminus \{0\}} e^{-\pi s^2 \|x\|_2^2}$$

is continuous and strictly decreasing, with a limit of ∞ as $s \rightarrow 0$ and a limit of 0 as $s \rightarrow \infty$. For $\varepsilon > 0$, the smoothing parameter of Λ is defined as $\eta_\varepsilon(\Lambda) := f_\Lambda^{-1}(\varepsilon)$.

Theorem 1.2. *Fix $\varepsilon > 0$ and let σ be any recurrent state of the sandpile chain. Let H_σ^t and P_σ^t denote the distributions at time t of the continuous and discrete time sandpile chains, respectively, started from σ . The L^2 distances from U , the uniform distribution on recurrent states, satisfy:*

$$\begin{aligned} \|H_\sigma^t - U\|_2^2 &\leq \varepsilon \text{ for all } t \geq \frac{\pi}{16} n \cdot \eta_\varepsilon^2(\Delta\mathbb{Z}^{n-1}), \\ \|P_\sigma^t - U\|_2^2 &\leq \varepsilon \text{ for all } t \geq \frac{\pi}{16} n \cdot \eta_\varepsilon^2(\bar{\Delta}\mathbb{Z}^n). \end{aligned}$$

The smoothing parameter η_ε can be bounded in terms of n and d_* , allowing us to show in Theorem 4.3 that every sandpile chain has mixing time of order at most $d_*^2 n \log n$. In the case of the complete graph K_n , we use eigenfunctions

to provide a matching lower bound. The upper and lower bounds taken together demonstrate that the sandpile chain on K_n exhibits total variation cutoff.

Theorem 1.3. *Let P_σ^t be the distribution of the discrete time sandpile chain on K_n after t steps started at a fixed recurrent state σ , and let U be the uniform distribution on recurrent states. For any $c \geq 5/4$,*

$$\begin{aligned} \|P_\sigma^t - U\|_{\text{TV}} &\leq e^{-c} \text{ for all } t \geq \frac{1}{4\pi^2} n^3 \log n + cn^3, \\ \|P_\sigma^t - U\|_{\text{TV}} &\geq 1 - e^{-35c} \text{ for all } 0 \leq t \leq \frac{1}{4\pi^2} n^3 \log n - cn^3. \end{aligned}$$

While the lower bound is specific to K_n , the upper bound holds for *any* graph with n vertices. Therefore, as the underlying graph varies with a fixed number of vertices, the sandpile chain on the complete graph mixes asymptotically slowest: For any graph sequence G_n such that G_n has n vertices,

$$\limsup_{n \rightarrow \infty} \frac{t_{\text{mix}}(G_n)}{t_{\text{mix}}(K_n)} \leq 1.$$

1.1. An inverse relationship for spectral gaps. Given the central role of the graph Laplacian $\bar{\Delta}$ in Theorems 1.1 and 1.2, one might ask how the eigenvalues $\{\lambda_h : h \in \mathcal{H}\}$ of the sandpile chain relate to the eigenvalues of the Laplacian itself. In particular, how does mixing of the sandpile chain on G relate to mixing of the simple random walk on G ? The sandpile chain typically has an exponentially larger state space, so there is not necessarily a simple relationship. Indeed, certain local features of G that have little effect on the Laplacian spectrum can decrease the spectral gap of the sandpile chain. An example is the existence of two vertices x, y with a common neighborhood, which enables the multiplicative harmonic function

$$h(x) = e^{2\pi i/d}, \quad h(y) = e^{-2\pi i/d}, \quad h(z) = 1 \text{ for } z \neq x, y,$$

where $d = \deg(x) = \deg(y)$. The corresponding eigenvalue λ_h is close to 1, so the spectral gap of the sandpile chain is rather small whenever two such vertices exist.

Nevertheless there is a curious inverse relationship between the spectral gap of the sandpile chain and the spectral gap of the Laplacian. According to our next result, if the simple random walk on G mixes sufficiently quickly, then the sandpile chain on G mixes *slowly*!

Theorem 1.4. *The spectral gap of the discrete time sandpile chain satisfies*

$$\gamma_d \leq \frac{4\pi^2}{\beta_1^2 n}$$

where $0 = \beta_0 < \beta_1 \leq \dots \leq \beta_{n-1}$ are the eigenvalues of the full Laplacian $\bar{\Delta}$.

For a bounded degree expander graph, this upper bound on γ_d matches the lower bound (3) up to a constant factor.

1.2. Related work. We were drawn to the topic of sandpile mixing times in part by its intrinsic mathematical interest and in part by questions arising in statistical physics. Many questions about the sandpile model, even if they appear unrelated to mixing, seem to lead inevitably to the study of its mixing time. For instance, the failure of the ‘density conjecture’ [15, 24, 32] and the distinction between the stationary and threshold states [27] are consequences of slow mixing. Our characterization of the eigenvalues of the sandpile chain may also help explain the recent finding of ‘memory on two time scales’ for the sandpile dynamics on the 2-dimensional square grid [34].

The significance of the characters of \mathcal{G} for the sandpile model was first remarked by Dhar, Ruelle, Sen and Verma [12], who in particular analyzed the sandpile group of the square grid graph with sink at the boundary. Our multiplicative harmonic functions h are related to the toppling invariants Q of [12] by $h = e^{2\pi i Q}$.

In the combinatorics literature the abelian sandpile model is known as ‘chip-firing’ [8]. For an alternative construction of the sandpile group, using flows on the edges instead of functions on the vertices, see [3]; also [7] and [19, Ch. 14]. The geometry of the Laplacian lattice $\bar{\Delta}\mathbb{Z}^n$ is studied in [1, 31] in connection with the combinatorial Riemann-Roch theorem of Baker and Norine [5]. We do not know of a direct connection between that theorem and sandpile mixing times, but the central role of the Laplacian lattice in both is suggestive.

The pseudoinverse $\bar{\Delta}^+$ has appeared before in the context of sandpiles: It is used by Björner, Lovász and Shor [8] to bound the number of topplings until a configuration stabilizes; see [20] for a recent improvement. In addition, the pseudoinverse is a crucial ingredient in the ‘energy pairing’ of Baker and Shokrieh [6].

The smoothing parameter was introduced by Micciancio and Regev [29] in the context of lattice-based cryptography. Our interest lies in results that relate the smoothing parameter to other lattice invariants, many of which have natural interpretations in the setting of the sandpile chain. Relationships of this kind have been found by [30, 17, 36].

1.3. Outline. After formally defining the sandpile chain and recalling how it can be expressed as a random walk on the sandpile group, Section 2 applies the classical eigenvalue formulas and mixing bounds for random walk on a group, giving the characterization (1) of eigenvalues in terms of multiplicative harmonic functions. Although the sandpile chain is nonreversible in general, the usual mixing bounds hold thanks to the orthogonality of the characters of the sandpile group. Our first application of (1) is the lower bound (3) on the spectral gap γ_d . Section 2 concludes by showing how one can sometimes exploit local features of G , which we call ‘gadgets,’ to infer an upper bound on γ_d . In many cases, this upper bound matches the lower bound from (3) up to a constant factor.

The next two sections are the heart of the paper. Section 3 proves a correspondence between multiplicative harmonic functions and equivalence classes of vectors in the dual Laplacian lattices $\Delta^{-1}\mathbb{Z}^{n-1}$ and $\bar{\Delta}^+\mathbb{Z}_0^n$, leading to Theorems 1.1 and 1.4. Section 4 proves Theorem 1.2 and thereby obtains sharp bounds on mixing time in terms of the number of vertices and maximum degree of G .

The final Section 5 treats several examples: cycles, complete and complete bipartite graphs, the torus, and ‘rooted sums.’ For the complete graph, we show a lower bound on mixing time that matches the upper bound from Section 4, proving Theorem 1.3. The examples are collected in one place to improve the flow of the paper, but some readers will want to look at them before or in parallel to reading the proofs in Sections 3 and 4.

1.4. Notation. Throughout the paper, $\mathbb{N} = \{0, 1, 2, \dots\}$ and \log is the natural logarithm. We denote the usual inner product and Euclidean norm on \mathbb{R}^n by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|_2$. We use standard Landau notation: $f(n) = o(g(n))$ if $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$; $f(n) = O(g(n))$ if $f(n) \leq Cg(n)$ for some $C < \infty$ and all n ; $f(n) = \Omega(g(n))$ if $f(n) \geq cg(n)$ for some $c > 0$ and all n ; and $f(n) = \Theta(g(n))$ if $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$. At times we write $f(n) \ll g(n)$ to mean $f(n) = o(g(n))$.

2. THE SANDPILE CHAIN

We begin by formally defining the sandpile chain. Let $G = (V, E)$ be a simple connected graph with finite vertex set $V = \{v_1, \dots, v_n\}$. We call v_n the *sink* and write $s = v_n$. A sandpile is a collection of indistinguishable *chips* distributed amongst the non-sink vertices $\tilde{V} = V \setminus \{s\}$, and thus can be represented by a function $\eta : \tilde{V} \rightarrow \mathbb{N}$. The *configuration* η corresponds to the sandpile with $\eta(v)$ chips at vertex v . We say that η is *stable at the vertex* $v \in \tilde{V}$ if $\eta(v) < \deg(v)$, and say that η is *stable* if it is stable at each non-sink vertex. If the number of chips at v is greater than or equal to its degree, then the vertex is allowed to *topple*, sending one chip to each of its neighbors. This leads to the new configuration η' with $\eta'(v) = \eta(v) - \deg(v)$ and $\eta'(u) = \eta(u) + 1$ if $\{u, v\} \in E$, and $\eta'(u) = \eta(u)$ otherwise. Toppling v may cause other vertices to become unstable, which can lead to further topplings, and any chip that falls into the sink is gone forever. Since we are assuming that G is connected, the presence of the sink ensures that one can reach a stable sandpile from any initial configuration (in finitely many steps) by successively performing topplings at unstable sites. An easy argument shows that the final stable configuration, which we denote by η° , does not depend on the order in which the topplings are carried out, hence the appellation abelian sandpile [11].

Define the sum of two configurations by $(\sigma + \eta)(v) = \sigma(v) + \eta(v)$. The abelian property shows that if we restrict our attention to the set of stable configurations $\mathcal{S} = \{\eta \in \mathbb{N}^{\tilde{V}} : \eta(v) < \deg(v) \text{ for all } v \in \tilde{V}\}$, then the operation of addition followed by stabilization, $\eta \oplus \sigma = (\eta + \sigma)^\circ$, makes \mathcal{S} into a commutative monoid. The identity is the empty configuration $\iota \equiv 0$. In light of this semigroup structure, it is natural to consider random walks on \mathcal{S} : If μ is a probability on \mathcal{S} , then beginning with some initial state η_0 , define $\eta_{t+1} = \eta_t \oplus \sigma_{t+1}$ where $\sigma_1, \sigma_2, \dots$ are drawn independently from μ . A natural candidate for μ is the uniform distribution on the configurations $\delta_v(u) = 1\{u = v\}$ as v ranges over V . (Note that $\delta_s = \iota$.) In words, at each time step we add a chip to a random vertex and stabilize.

Now define the *saturated configuration* by $\eta_*(v) = \deg(v) - 1$ for each $v \in \tilde{V}$. Our assumptions ensure that from any initial state, the random walk on \mathcal{S} driven by the uniform distribution on $\{\delta_v\}_{v \in V}$ will visit η_* in finitely many steps with full probability. (If $M = \sum_{v \in \tilde{V}} (\deg(v) - 1)$, then the probability of visiting η_* within M steps from any stable configuration is at least n^{-M} .) The random walk will thus eventually be absorbed by the communicating class $\mathcal{G} = \{\eta \in \mathcal{S} : \eta = \eta_* \oplus \sigma \text{ for some } \sigma \in \mathcal{S}\}$. Accordingly, the configurations in \mathcal{G} are called *recurrent*. In the language of semigroups, $\mathcal{G} = \eta_* \oplus \mathcal{S}$ is the minimal ideal of \mathcal{S} . To see that this is so, suppose that \mathcal{I} is an ideal of \mathcal{S} (that is, $\mathcal{I} \oplus \mathcal{S} \subseteq \mathcal{I}$) and let $\sigma \in \mathcal{I}$. Define $\bar{\sigma}(v) = \deg(v) - 1 - \sigma(v)$. Then $\bar{\sigma} \in \mathcal{S}$, so $\eta_* = \sigma + \bar{\sigma} = \sigma \oplus \bar{\sigma} \in \mathcal{I} \oplus \mathcal{S} \subseteq \mathcal{I}$, hence $\mathcal{G} \subseteq \mathcal{I}$. As the minimal ideal of a commutative semigroup, \mathcal{G} is a nonempty abelian group under \oplus . (Briefly, for any $a \in \mathcal{G}$, $\mathcal{G} \subseteq a \oplus \mathcal{G} \subseteq \mathcal{G}$ since \mathcal{G} is the minimal ideal, hence $a \oplus x = b$ has a solution in \mathcal{G} for all $a, b \in \mathcal{G}$.) The article [2] contains an excellent exposition of this perspective.

Because the random walk will eventually end up in the *sandpile group* \mathcal{G} anyway, it makes sense to restrict the state space to \mathcal{G} to begin with. Rather than thinking of this Markov chain in terms of \mathcal{S} acting on \mathcal{G} , it is more convenient to consider it as a random walk on \mathcal{G} so that we may draw on a rich existing theory. To this end, let id denote the identity in \mathcal{G} (which is not equal to ι in general). Then for each $v \in V$, $\sigma_v := \delta_v \oplus id \in \mathcal{G}$ since \mathcal{G} is an ideal of \mathcal{S} containing id . Also, $\sigma_v \oplus \eta = \delta_v \oplus id \oplus \eta = \delta_v \oplus \eta$ for all $\eta \in \mathcal{G}$. The process of successively adding chips to random vertices and stabilizing can thus be represented as the random walk on \mathcal{G} driven by the uniform distribution on $S = \{\sigma_v\}_{v \in V}$.¹ Since \mathcal{G} is an abelian group generated by S , we can conclude, for example, that the chain is irreducible with uniform stationary distribution and that the characters of \mathcal{G} form an orthonormal basis of eigenfunctions for the transition operator [33].

Though the foregoing is all very nice in theory, even computing the identity element of the sandpile group of a specific graph in these abstract terms is typically quite involved, so it is useful to establish a more concrete realization of \mathcal{G} . Recall that the reduced Laplacian Δ of G is the $(n-1) \times (n-1)$ submatrix of the full Laplacian $\bar{\Delta}$ formed by deleting the n th row and column (corresponding to the sink vertex). We claim that

$$\mathcal{G} \cong \mathbb{Z}^{n-1} / \Delta \mathbb{Z}^{n-1},$$

an isomorphism being given by $\eta \mapsto (\eta(v_1), \dots, \eta(v_{n-1})) + \Delta \mathbb{Z}^{n-1}$. Note that this implies that $|\mathcal{G}| = \det(\Delta)$, which is equal to the number of spanning trees in G by the matrix-tree theorem [35]. The interpretation is that $z \in \mathbb{Z}^{n-1}$ corresponds to the configuration having z_i chips at vertex v_i , where we are allowing vertices to have a negative number of chips—a hole or a debt, say. For $x \in \mathbb{Z}^{n-1}$, adding Δx to z corresponds to performing $-x_i$ topplings at each vertex v_i . (A negative toppling means that the vertex takes one chip from each of its neighbors.) Isomorphism is established by showing that each coset contains exactly one vector corresponding

¹More precisely, the pushforward of the uniform distribution on V under the map $v \mapsto \sigma_v$. Lemma 4.6 in [5] implies that the σ_v 's are distinct as elements of \mathcal{G} if and only if G is 2-edge connected.

to a recurrent configuration. One way to see this is to observe that the lattice $\Delta\mathbb{Z}^{n-1}$ contains points with arbitrarily large smallest coordinate, so from any $z \in \mathbb{Z}^{n-1}$, there is an $x \in \Delta\mathbb{Z}^{n-1}$ with $(x+z)_i \geq \deg(v_i)$ for all $1 \leq i \leq n-1$. Then use the fact that stabilizing such a configuration amounts to adding some $y \in \Delta\mathbb{Z}^{n-1}$ and results in a unique recurrent configuration. We refer the reader to [21] for a more detailed account.

In this view, the sandpile chain is a random walk on $\Gamma := \mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1}$ driven by the uniform distribution on $\{e_1, \dots, e_{n-1}, \mathbf{0}\}$, where e_i is the vector with a one in the i th coordinate and zeros elsewhere. Geometrically, we are performing a random walk on the positive orthant of \mathbb{Z}^{n-1} by taking steps of unit length in a direction chosen uniformly at random from the standard basis vectors of \mathbb{Z}^{n-1} (with a holding probability of $1/n$), but we are concerned only with our relative location within cells of the lattice $\Delta\mathbb{Z}^{n-1}$.

2.1. Spectral properties. Since the sandpile chain is a random walk on a finite abelian group Γ , we can find the eigenvalues and eigenfunctions of the associated transition matrix in terms of the *characters* of Γ , that is, elements of the dual group

$$\widehat{\Gamma} := \text{Hom}(\Gamma, \mathbb{T})$$

where \mathbb{T} is the set of complex numbers of modulus 1. We emphasize that the sandpile chain is *not* reversible in general, as our generating set $\{e_1, \dots, e_{n-1}, \mathbf{0}\}$ is generally not closed under negation modulo $\Delta\mathbb{Z}^{n-1}$. Nevertheless we will see that the usual bounds on mixing still hold due to orthogonality of the characters.

Our starting point is the following well-known lemma, which is particularly simple in our case thanks to the fact that all irreducible representations of an abelian group are one-dimensional. See [13] for the general (nonabelian) case.

Lemma 2.1. *If μ is a probability on a finite abelian group A , then the transition matrix for the associated random walk has an orthonormal basis of eigenfunctions consisting of the characters of A . The eigenvalue corresponding to a character χ is given by the evaluation of the Fourier transform*

$$\widehat{\mu}(\chi) := \sum_{a \in A} \mu(a)\chi(a).$$

Proof. Let $Q(x, y) = \mu(yx^{-1})$ denote the transition matrix. For any character χ , we have

$$\begin{aligned} (Q\chi)(x) &= \sum_{y \in A} Q(x, y)\chi(y) = \sum_{y \in A} \mu(yx^{-1})\chi(y) \\ &= \sum_{z \in A} \mu(z)\chi(zx) = \chi(x) \sum_{z \in A} \mu(z)\chi(z), \end{aligned}$$

hence χ is an eigenfunction with eigenvalue $\widehat{\mu}(\chi)$. The result follows by observing that there are $|A|$ characters and they are orthonormal with respect to the standard inner product $(f, g) = \frac{1}{|A|} \sum_{a \in A} f(a)\overline{g(a)}$ on \mathbb{C}^A . \square

To apply the preceding lemma, we are going to express the characters of the abstract group Γ in terms of functions defined on the vertices of our graph G . These are the *multiplicative harmonic functions*

$$h : V \rightarrow \mathbb{T}$$

satisfying $h(s) = 1$ and

$$h(v)^{\deg(v)} = \prod_{u \sim v} h(u) \quad (4)$$

for all $v \in V$. We will refer to (4) as the ‘geometric mean value property’ at vertex v . We pause to make two remarks about this definition.

Remark 2.2. We could equivalently take the codomain of h to be all nonzero complex numbers. The geometric mean value property implies $\bar{\Delta}(\log |h|) \equiv 0$, so $\log |h|$ is a constant function. Since $h(s) = 1$ it follows that h takes values on the unit circle \mathbb{T} .

Remark 2.3. The identity

$$\prod_{v \in V} h(v)^{\deg(v)} = \prod_{v \in V} \prod_{u \sim v} h(u) \quad (5)$$

holds for *any* function h . Therefore if the geometric mean value property holds at every vertex but one, then in fact it holds at every vertex.

Denote by \mathcal{H} the set of all multiplicative harmonic functions on G . This set is nonempty since it contains the constant function 1. One readily checks that it is an abelian group under pointwise multiplication.

For each $h \in \mathcal{H}$, define $\chi'_h : \mathbb{Z}^{n-1} \rightarrow \mathbb{T}$ by

$$\chi'_h(z) = \prod_{j=1}^{n-1} h(v_j)^{z_j}, \quad (6)$$

and define $\chi_h : \mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1} \rightarrow \mathbb{T}$ by $\chi_h(z + \Delta\mathbb{Z}^{n-1}) = \chi'_h(z)$. The ensuing proof shows that χ_h is well-defined.

Lemma 2.4. *The characters of Γ are precisely $\{\chi_h\}_{h \in \mathcal{H}}$.*

Proof. From (6), each function χ'_h is a homomorphism. For each standard basis vector $e_j \in \mathbb{Z}^{n-1}$,

$$\chi'_h(\Delta e_j) = h(v_j)^{\deg(v_j)} \prod_{v_k \sim v_j} h(v_k)^{-1} = 1, \quad (7)$$

using the geometric mean value property and $h(s) = 1$. Therefore χ'_h is constant on cosets of $\Delta\mathbb{Z}^{n-1}$ and thus descends to $\chi_h \in \widehat{\Gamma}$.

Conversely, every $\chi \in \widehat{\Gamma}$ lifts to a homomorphism $\chi' : \mathbb{Z}^{n-1} \rightarrow \mathbb{T}$ that is identically 1 on $\Delta\mathbb{Z}^{n-1}$. Define $h : V \rightarrow \mathbb{T}$ by $h(v_j) = \chi'(e_j)$ for $j \leq n-1$ and $h(s) = 1$. Since χ' is a homomorphism, it satisfies equation (6), so $\chi = \chi_h$. By equation (7), h satisfies the geometric mean value property at every non-sink vertex, implying that $h \in \mathcal{H}$. \square

Observe that the map $h \mapsto \chi_h$ gives an isomorphism between \mathcal{H} and $\widehat{\Gamma}$. Since Γ is a finite abelian group, we have $\mathcal{H} \cong \widehat{\Gamma} \cong \Gamma \cong \mathcal{G}$. Thus examining the multiplicative harmonic functions on G can give insights about the structure of the corresponding sandpile group. For example, the following proposition gives another way to see that the sandpile group of an undirected graph does not depend on the choice of sink.

Proposition 2.5. *If \mathcal{H}_v denotes the group of multiplicative harmonic functions on G with sink at v , then the map $\phi : \mathcal{H}_u \rightarrow \mathcal{H}_w$ given by $\phi(h)(v) = h(w)^{-1}h(v)$ is an isomorphism.*

Returning to the language of sandpiles, we see that the characters of \mathcal{G} are the functions $\{f_h : \mathcal{G} \rightarrow \mathbb{T}\}_{h \in \mathcal{H}}$, where

$$f_h(\eta) = \prod_{v \in \widetilde{V}} h(v)^{\eta(v)}.$$

Thus Lemma 2.1 implies the following.

Theorem 2.6. *An orthonormal basis of eigenfunctions for the transition matrix of the sandpile chain on G is $\{f_h\}_{h \in \mathcal{H}}$. The eigenvalue associated with f_h is*

$$\lambda_h := \frac{1}{n} \sum_{v \in V} h(v).$$

Proof. Letting p denote the uniform distribution on $\{\sigma_v\}_{v \in V}$, we see that the eigenvalue associated with f_h is

$$\widehat{p}(f_h) = \sum_{\eta \in \mathcal{G}} p(\eta) f_h(\eta) = \sum_{v \in V} \frac{1}{n} f_h(\sigma_v) = \frac{1}{n} \sum_{v \in V} h(v). \quad \square$$

Before proceeding to our main topic of mixing times, we remark on two ways to extend the above analysis.

Remark 2.7. Theorem 2.6 holds for an arbitrary chip-addition distribution μ , with the only change being that f_h then has eigenvalue $\lambda_h = \sum_{v \in V} \mu(v)h(v)$. For example, we may choose to add chips at a uniform *nonsink* vertex, in which case the eigenvalue associated with f_h is

$$\widetilde{\lambda}_h = \frac{1}{n-1} \sum_{v \in \widetilde{V}} h(v).$$

Since the only multiplicative harmonic function that is constant on \widetilde{V} is $h \equiv 1$ (as can be seen by applying the geometric mean value property at a neighbor of the sink), we have $|\widetilde{\lambda}_h| < 1$ for all nontrivial $h \in \mathcal{H}$. Therefore, this ‘non-lazy’ version of the sandpile chain is aperiodic despite the lack of holding probabilities.

Remark 2.8. Suppose that G is a directed multigraph in which every vertex has a directed path to the sink. In this setting there is a Laplacian $\overline{\Delta}$ analogous to (2) although it is no longer a symmetric matrix: its off-diagonal entries are $\overline{\Delta}_{uv} = -e_{uv}$ where e_{uv} is the number of edges directed from u to v , and its diagonal

entries are $\bar{\Delta}_{uu} = d_u - e_{uu}$ where d_u is the outdegree of vertex u . The sandpile group is naturally isomorphic to $\mathbb{Z}^{n-1}/\Delta_s^\top \mathbb{Z}^{n-1}$, where Δ_s^\top is the transpose of the submatrix omitting the row and column of $\bar{\Delta}$ corresponding to the sink s . Unlike in the undirected case, this group may depend on the choice of s (but see [14, Theorem 2.10]). If we define

$$\mathcal{H}_s := \left\{ h : V \rightarrow \mathbb{T} \text{ such that } h(s) = 1 \text{ and } h(u)^{d_u} = \prod_{v \in V} h(v)^{e_{uv}} \text{ for all } u \neq s \right\}$$

then nearly all of this section carries over to the directed case. The exceptions are (5) and Proposition 2.5, which hold only when G is Eulerian.

2.2. Mixing times. Since the sandpile chain is an irreducible and aperiodic random walk on a finite group, the law of the chain at time t approaches the uniform distribution on \mathcal{G} as $t \rightarrow \infty$. Our interest is in the rate of convergence. To avoid trivialities, we assume throughout that G is not a tree, so that $|\mathcal{G}| > 1$.

The metrics we consider are the L^2 distance

$$\|\mu - \nu\|_2 = \left(|\mathcal{G}| \sum_{g \in \mathcal{G}} |\mu(g) - \nu(g)|^2 \right)^{\frac{1}{2}}$$

and the total variation distance

$$\|\mu - \nu\|_{\text{TV}} = \frac{1}{2} \sum_{g \in \mathcal{G}} |\mu(g) - \nu(g)| = \max_{A \subseteq \mathcal{G}} (\mu(A) - \nu(A)).$$

Note that Cauchy-Schwarz immediately implies $\|\mu - \nu\|_{\text{TV}} \leq \frac{1}{2} \|\mu - \nu\|_2$.

We always assume that the chain is started from a deterministic state $\sigma \in \mathcal{G}$. As a random walk on a group, the distance to stationarity after t steps under either of these metrics is independent of σ , so without loss of generality we take $\sigma = id$ henceforth. Writing P_{id}^t for the distribution of the sandpile chain at time t and U for the uniform distribution on \mathcal{G} , the following lemma (proved in [33]) shows that $\|P_{id}^t - U\|_2$ is completely determined by the eigenvalues from Theorem 2.6.

Lemma 2.9. *Let Q_{id}^t be the t -step distribution of a random walk on a finite abelian group A started at the identity and driven by a probability measure μ , and let π be the uniform distribution on A . Then*

$$\|Q_{id}^t - \pi\|_2^2 = \sum_{\chi \neq 1} |\hat{\mu}(\chi)|^{2t}$$

where the sum is over all nontrivial characters.

For the sandpile chain, this says that

$$\|P_{id}^t - U\|_2^2 = \sum_{h \in \mathcal{H} \setminus \{1\}} |\lambda_h|^{2t}. \quad (8)$$

Lemma 2.9 is a special case of the famous Fourier bound from [13]. It also follows by an eigenfunction expansion exactly as in the standard proof of the spectral bound for reversible chains (see [26, Ch. 12]). Though sandpile chains are not reversible in general, many of the same arguments still apply because

of the orthogonality of the eigenfunctions. Probabilistically, this orthogonality is equivalent to the statement that the transition operator commutes with its time reversal.

Total variation convergence rates can also be estimated in terms of spectral information. To make this precise, we introduce some terminology. Let $\lambda_* = \max_{h \in \mathcal{H} \setminus \{1\}} |\lambda_h|$ denote the size of the subdominant eigenvalue of P . The (discrete time) *spectral gap* is defined as $\gamma_d = 1 - \lambda_*$, and the *relaxation time* as $t_{\text{rel}} = \gamma_d^{-1}$. The *mixing time* is $t_{\text{mix}}(\varepsilon) = \min\{t \in \mathbb{N} : \|P_{id}^t - U\|_{\text{TV}} \leq \varepsilon\}$. The following proposition is standard for reversible Markov chains. It holds in our context as well using the above definition of relaxation time.

Proposition 2.10. *The mixing time of the sandpile chain satisfies*

$$\log\left(\frac{1}{2\varepsilon}\right)(t_{\text{rel}} - 1) \leq t_{\text{mix}}(\varepsilon) \leq \left\lceil \log\left(\frac{|\mathcal{G}|^{\frac{1}{2}}}{2\varepsilon}\right) t_{\text{rel}} \right\rceil.$$

Proof. The upper bound follows from (8) and the L^2 bound on total variation by bounding the summands with λ_*^{2t} . Arithmetic manipulations then give the mixing time bound; see [26, Ch. 12] for details.

The lower bound is Theorem 12.4 from [26]. Reversibility is used there only to conclude that the constant function is orthogonal to nontrivial eigenfunctions in the inner product weighted by the stationary distribution. However, the same statement holds also for nonreversible chains since the stationary distribution, as a left eigenfunction with eigenvalue 1, is orthogonal to all nontrivial right eigenfunctions under the standard inner product. \square

Our next result gives a universal lower bound on the spectral gap γ_d (and thus an upper bound on t_{rel}) in terms of the number of vertices n and the second largest degree d_* of the underlying graph.

Theorem 2.11. *Suppose that G has degree sequence $d_1 \leq \dots \leq d_{n-1} \leq d_n$ and set $d_* = d_{n-1}$. Then the spectral gap of the (discrete time) sandpile chain on G satisfies*

$$\gamma_d \geq \frac{8}{d_*^2 n}.$$

The proof uses the following inequality.

Lemma 2.12. *Suppose $0 \leq r \leq 2\pi$. Then $\cos(x) \leq 1 - cx^2$ for all $|x| \leq r$, where $c = \frac{1 - \cos(r)}{r^2}$.*

Proof. Let $f(x) = [1 - \cos(x)]/x^2$, with $f(0) = 1/2$ so that f is continuous on \mathbb{R} . We will show that f is decreasing on the interval $[0, 2\pi]$. The inequality $f(x) \geq f(r)$ for $0 \leq x \leq r \leq 2\pi$ rearranges into the desired inequality $\cos(x) \leq 1 - cx^2$. Since f is even, the inequality also holds when $-2\pi \leq -r \leq x \leq 0$.

We have $f'(x) = g(x)/x^3$, where $g(x) = x \sin(x) - 2(1 - \cos(x))$. Thus, it suffices to show that g is negative on the interval $(0, 2\pi)$. Note that $g(0) = g(2\pi) = 0$. As well, $g'(x) = x \cos(x) - \sin(x)$ and $g''(x) = -x \sin(x)$. This means g' is strictly

decreasing for $x \in [0, \pi]$ and strictly increasing for $x \in [\pi, 2\pi]$. Since $g'(0) = 0$ and $g'(2\pi) = 2\pi$, there is $x_0 \in (\pi, 2\pi)$ such that $g'(x_0) = 0$ and $g'(x) < 0$ for $x \in (0, x_0)$ and $g'(x) > 0$ for $x \in (x_0, 2\pi]$. Therefore g is strictly decreasing for $x \in [0, x_0]$ and strictly increasing for $x \in [x_0, 2\pi]$. Because $g(0) = g(2\pi) = 0$, we conclude that $g(x) < 0$ for $x \in (0, 2\pi)$, as desired. \square

Proof of Theorem 2.11. We note at the outset that Proposition 2.5 implies that the moduli of the eigenvalues are invariant under change of sink, so we may assume without loss of generality that s is located at a vertex of maximum degree. Accordingly, $\deg(v) \leq d_*$ for all $v \in \tilde{V}$.

Now fix $h \in \mathcal{H}$ and suppose that there exists an arc $C_{ab} = \{e^{i\theta} : 2\pi a \leq \theta \leq 2\pi b\}$ with $0 < b - a < 1/d_*$ such that $h(v) \in C_{ab}$ for every $v \in V$. We will show this implies $h \equiv 1$. Write $h(v) = e^{2\pi i g(v)}$, where $g : V \rightarrow [a, b]$. Note that for all $v \in V$,

$$\bar{\Delta}g(v) = \sum_{w \sim v} (g(v) - g(w))$$

is an integer, by the geometric mean value property of h at v . On the other hand, since $g(V) \subseteq [a, b]$,

$$|\bar{\Delta}g(v)| \leq \sum_{w \sim v} |g(v) - g(w)| \leq d_*(b - a) < 1$$

for all $v \in \tilde{V}$. Since the left side is an integer it must be zero, so g is harmonic on \tilde{V} in the usual sense. Uniqueness of harmonic extensions implies $g \equiv g(s)$ and thus $h \equiv h(s) = 1$, as desired.

For any fixed $h \in \mathcal{H} \setminus \{1\}$ write $\lambda_h = r e^{i\theta}$ with $r \geq 0$. The preceding argument shows that $h(V)$ is not contained in any segment of the unit circle having arc length less than $2\pi/d_*$. For each $v \in V$, let $A(v)$ be the unique angle $-\pi < \phi \leq \pi$ such that $h(v) = e^{i(\theta+\phi)}$. Let $\phi_1 = \min_{v \in V} A(v)$ and $\phi_2 = \max_{v \in V} A(v)$, so that $\phi_2 - \phi_1 \geq 2\pi/d_*$. We have

$$|\lambda_h| = \frac{1}{n} \sum_{v \in V} \cos(A(v)) \leq \frac{1}{n} [n - 2 + \cos(\phi_1) + \cos(\phi_2)].$$

Using the identity

$$\cos(\phi_1) + \cos(\phi_2) = 2 \cos\left(\frac{\phi_2 + \phi_1}{2}\right) \cos\left(\frac{\phi_2 - \phi_1}{2}\right) \leq 2 \cos\left(\frac{\phi_2 - \phi_1}{2}\right),$$

along with $\pi/d_* \leq (\phi_2 - \phi_1)/2 < \pi$ and the fact that cosine is decreasing on $[0, \pi]$, we have

$$|\lambda_h| \leq \frac{1}{n} \left[n - 2 + 2 \cos\left(\frac{\pi}{d_*}\right) \right].$$

Recalling our assumption that G is not a tree, $\pi/d_* \leq \pi/2$, so Lemma 2.12 gives

$$|\lambda_h| \leq \frac{1}{n} \left[n - 2 + 2 \left(1 - \frac{4}{\pi^2} \left(\frac{\pi}{d_*} \right)^2 \right) \right] = 1 - \frac{8}{d_*^2 n}. \quad \square$$

Note that the number of stable configurations, and thus the order of the sandpile group, is at most d_*^{n-1} , so it follows from Theorem 2.11 and Proposition 2.10 that $t_{\text{mix}}(\varepsilon) = O((d_*^2 \log d_*)n^2)$. In Section 4, we will show that this bound can be improved to $O(d_*^2 n \log n)$.

2.3. Gadgets. In many cases, one can determine the order of the relaxation time t_{rel} of the sandpile chain by constructing just a single multiplicative harmonic function. For bounded degree graphs, Theorem 2.11 shows that $t_{\text{rel}} = O(n)$. Thus to show that $t_{\text{rel}} = \Theta(n)$, it suffices to find an $h \in \mathcal{H}$ with $|\lambda_h| \geq 1 - C/n$ for some constant C . Since the eigenvalues are all of the form $\lambda_h = \frac{1}{n} \sum_{v \in V} h(v)$, we see that large eigenvalues correspond to ‘nearly constant’ multiplicative harmonic functions. In particular, if there is an $h \in \mathcal{H}$ which is constant on $U \subseteq V$ with $|V \setminus U| = m$, then $|\lambda_h| \geq (n - 2m)/n$.

For example, consider the m -fold Sierpinski gasket graph $SG_2(m)$ defined recursively as follows: $SG_2(0)$ is the triangle K_3 and $SG_2(m)$ for $m \geq 1$ is obtained by gluing three copies of $SG_2(m-1)$ to obtain a triangle with center cut out. We take the ‘topmost’ vertex as the sink (Figure 1).




FIGURE 1. Sierpiński Gasket Graphs

It is easy to see that $SG_2(m)$ has $|V(m)| = \frac{1}{2}(3^{m+1} + 3)$ vertices, and it is known [9] that the number of spanning trees is $|\mathcal{G}(SG_2(m))| = 2^{\alpha(m)} 3^{\beta(m)} 5^{\gamma(m)}$ with $\alpha(m) = \frac{1}{2}(3^m - 1)$, $\beta(m) = \frac{1}{4}(3^{m+1} + 2m + 1)$, $\gamma(m) = \frac{1}{4}(3^m - 2m - 1)$. To the best of the authors’ knowledge, the invariant factor decomposition of $\mathcal{G}(SG_2(m))$ is an open problem.

Each $SG_2(m)$ contains a copy of $SG_2(1)$ in the lower right corner. Referring to Figure 1, define h by $h(a) = h(b) = h(c) = -1$ where a, b, c are the three inner vertices of this copy of $SG_2(1)$; and $h(v) = 1$ for all other vertices v of $SG_2(m)$. Then h is multiplicative harmonic with associated eigenvalue $\lambda_h = 1 - 6/|V(m)|$, hence the relaxation time is $\Omega(|V(m)|)$. Since $SG_2(m)$ has bounded degree, we conclude that $t_{\text{rel}} = \Theta(|V(m)|)$.

Because large eigenvalues so often arise in this fashion, it is useful to introduce the following definition.

Definition 2.13. Let $G' = (V', E')$ be a vertex induced subgraph of G , and let h' be a function from V' to \mathbb{T} . We call (G', h') a *gadget* of size $m > 0$ in G if:

- (1) h' satisfies the geometric mean value property (with respect to G') at every $v \in V'$.
- (2) The *interior* $\text{int}(V') = \{v \in V' : h'(v) \neq 1\}$ has size m .
- (3) The *boundary* $\partial V' = V' \setminus \text{int}(V')$ contains the subgraph boundary $\{v' \in V' : \{v', u\} \in E \text{ for some } u \in V \setminus V'\}$.

Often we refer to the gadget (G', h') simply as G' for convenience. For example, the Sierpiński gasket graph $SG_2(m)$ has a gadget $SG_2(1)$ of size $|\{a, b, c\}| = 3$.

Proposition 2.14. *If (G', h') is a gadget of size m in G , then $\gamma_d \leq \frac{2m}{n}$.*

Proof. Since γ_d is independent of the choice of sink vertex, we may assume that $s \notin \text{int}(V')$. Define $h : V \rightarrow \mathbb{T}$ by $h \equiv h'$ on V' and $h \equiv 1$ on $V \setminus V'$. It is easily checked that $h \in \mathcal{H}(G)$ and $|\lambda_h| \geq \text{Re}(\lambda_h) \geq 1 - \frac{2m}{n}$. \square

The next example shows that a gadget can have size as small as 2. Suppose that $v_1, v_2 \in \tilde{V}$ have common neighborhood $N = \{u \in V : u \sim v_1\} = \{u \in V : u \sim v_2\}$ with $d = |N| > 1$. Then the induced subgraph with vertex set $V' = \text{int}(V') \cup \partial V'$, $\text{int}(V') = \{v_1, v_2\}$, $\partial V' = N$, is a gadget of size 2 in G . Indeed, if ω is any nontrivial d th root of unity, then the function $h : V \rightarrow \mathbb{T}$ given by

$$h(v) = \begin{cases} \omega, & v = v_1 \\ \omega^{-1}, & v = v_2 \\ 1, & \text{else} \end{cases}$$

is multiplicative harmonic on G . Taking $\omega = e^{2\pi i/d}$, the eigenvalue corresponding to h is $\lambda_h = 1 - \frac{2}{n} \left(1 - \cos\left(\frac{2\pi}{d}\right)\right)$. Thus, in any graph on n vertices, two of which have the same neighborhood of size d , the spectral gap of the discrete time sandpile chain has order at most $1/(d^2n)$.

The common neighborhood gadget can also be understood from a more probabilistic perspective. The eigenfunction corresponding to h gives information about the difference mod d between the number of chips at v_1 and v_2 . For the chain to equilibrate, it must run long enough for this mod d difference to randomize. Since v_1 and v_2 have all neighbors in common, the mod d difference is invariant under toppling; it changes only when a chip is added at v_1 or v_2 .

Small gadgets can drastically affect the mixing time of the sandpile chain. For example, the sandpile chain on the cycle C_n mixes completely after a single step (see Section 5); but if we add just two extra vertices u, w and $2d$ extra edges $\{u, v_j\}$ and $\{w, v_j\}$, $j = 1, \dots, d$ for some $d \geq 2$, then the relaxation time of the sandpile chain becomes $\Omega(d^2n)$.

The following theorem provides one way to rule out the presence of a small gadget.

Theorem 2.15. *If G has girth g , then all gadgets in G have size at least $g/2$.*

Proof. We note at the outset that the relevant definitions preclude the existence of gadgets of size 1, so the theorem is vacuously true when $g \leq 4$. When $g \geq 5$, assume for the sake of contradiction that $G = (V, E)$ is a counterexample having vertex set of minimum size. Let (G', h') be a gadget in G whose interior W has

size $|W| < g/2$, and let h be the extension of h' to V by setting $h \equiv 1$ on $V \setminus V'$, so that h satisfies the geometric mean value property at every $v \in V$. Any vertex of degree 1 could be deleted without affecting the girth of G or the value of h at the other vertices, so by minimality, every vertex of G has degree at least 2.

We now define a new graph K on the vertex set W . We draw an edge between the distinct vertices $a, b \in W$ if $\{a, b\} \in E$ or if there exists $c \in V \setminus W$ with $\{a, c\}, \{b, c\} \in E$. Since any cycle in K would give rise to a cycle in G of length at most $2|W|$, the graph K has no cycles. Therefore there is $a \in W$ such that $\deg_K(a) = 0$ or 1.

Since $\deg_G(a) \geq 2$, the vertex a has at least one G -neighbor $v \in V \setminus W$. If v has no other neighbors in W , then $1 = h(v)^{\deg(v)} = \prod_{w \sim v} h(w) = h(a)$, a contradiction. Therefore v is adjacent to some $b \in W$, which must be the unique K -neighbor of a . The edge $\{a, b\} \notin E$ because G has no triangles. Thus a has no G -neighbors in W , and it must have another G -neighbor $v' \in V \setminus W$. By the reasoning used for v , v' is also adjacent to b . But now the edges $\{a, v\}, \{v, b\}, \{b, v'\}, \{v', a\}$ form an illegal 4-cycle in G . This contradiction proves the theorem. \square

3. DUAL LATTICES

In Section 2 we saw that the sandpile chain is a random walk on the lattice quotient $\Gamma = \mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1}$. Its character group $\widehat{\Gamma}$ is naturally isomorphic to

$$\widehat{\Gamma} \cong (\Delta^{-1}\mathbb{Z}^{n-1})/\mathbb{Z}^{n-1}.$$

We refer to $\Delta^{-1}\mathbb{Z}^{n-1}$ as the *dual Laplacian lattice* because it equals $(\Delta\mathbb{Z}^{n-1})^* = \{x \in \mathbb{R}^{n-1} : \langle x, y \rangle \in \mathbb{Z} \text{ for all } y \in \Delta\mathbb{Z}^{n-1}\}$. Since the group \mathcal{H} of multiplicative harmonic functions is isomorphic to $\widehat{\Gamma}$ (see Lemma 2.4), each $h \in \mathcal{H}$ can be identified with an equivalence class $x_h + \mathbb{Z}^{n-1} \subseteq \Delta^{-1}\mathbb{Z}^{n-1}$ of dual lattice vectors.

Given $h \in \mathcal{H}$, choose $x_h \in \Delta^{-1}\mathbb{Z}^{n-1}$ of minimal Euclidean length that corresponds to h . The first main result in this section, Theorem 3.4, relates the length $\|x_h\|_2$ to the eigenvalue $\lambda_h = \frac{1}{n} \sum_{v \in V} h(v)$ of the sandpile chain. Specifically, the length $\|x_h\|_2$ determines the gap $1 - \operatorname{Re}(\lambda_h)$ up to a constant factor. This property lets us translate information about the lengths of dual lattice vectors into information about the convergence of the *continuous time* sandpile chain.

What about the discrete time sandpile chain? As it turns out, we get parallel results to the continuous case if we use a slightly different dual lattice, constructed using the *pseudoinverse* of the full Laplacian matrix $\bar{\Delta}$. The pseudoinverse construction leads to a quick proof of Theorem 1.4, which states that if the spectral gap of $\bar{\Delta}$ is large, then the spectral gap of the discrete time sandpile chain is small.

This section is organized as follows. Section 3.1 provides preliminary details about the discrete and continuous time sandpile chains. Sections 3.2 and 3.3 construct the two dual lattices and prove the correspondence between dual lattice vector lengths and sandpile chain eigenvalues. Section 3.4 proves the inverse relationship (Theorem 1.4) and uses it to determine the order of the sandpile chain spectral gap for families of expander graphs.

3.1. Discrete and continuous time. In this subsection we compare the discrete time sandpile chain with its continuous time analogue. In discrete time, the mixing properties are essentially independent of the choice of sink vertex, while in continuous time, the choice of sink can affect the mixing by up to a factor of n .

Recall that the transition matrix P for the discrete time sandpile chain has eigenvalues $\{\lambda_h : h \in \mathcal{H}\}$, where $\lambda_h = \frac{1}{n} \sum_{v \in V} h(v)$. The continuous time chain, which proceeds by dropping chips according to a rate 1 Poisson process, has kernel

$$H^t(\eta, \sigma) = \sum_{k=0}^{\infty} e^{-t} \frac{t^k}{k!} P^k(\eta, \sigma) = e^{-t(I-P)}(\eta, \sigma).$$

The eigenvalues of H^t are $\{e^{-t(1-\lambda_h)} : h \in \mathcal{H}\}$.

If the chains are started from the identity configuration, then after time t the L^2 distances from the uniform distribution U are

$$\begin{aligned} \|P_{id}^t - U\|_2^2 &= \sum_{h \in \mathcal{H} \setminus \{1\}} |\lambda_h|^{2t}, \\ \|H_{id}^t - U\|_2^2 &= \sum_{h \in \mathcal{H} \setminus \{1\}} \left| e^{-t(1-\lambda_h)} \right|^2 = \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1-\operatorname{Re}(\lambda_h))}. \end{aligned}$$

The mixing of the discrete time chain is controlled by the eigenvalues with modulus close to 1, while the mixing of the continuous time chain is controlled by the eigenvalues with real part close to 1. This motivates the definitions for the discrete and continuous time spectral gaps given in Section 1:

$$\begin{aligned} \gamma_d &= \min\{1 - |\lambda_h| : h \in \mathcal{H}, h \neq 1\}, \\ \gamma_c &= \min\{1 - \operatorname{Re}(\lambda_h) : h \in \mathcal{H}, h \neq 1\}. \end{aligned}$$

Proposition 3.1. *Every eigenvalue λ_h of the discrete time sandpile chain satisfies*

$$1 - |\lambda_h| \leq 1 - \operatorname{Re}(\lambda_h) \leq n(1 - |\lambda_h|).$$

Hence $\gamma_d \leq \gamma_c \leq n\gamma_d$.

Proof. The lower bound is immediate. For the upper bound, since $h(s) = 1$, we can write

$$\lambda_h = \frac{1}{n} [1 + (n-1)z],$$

where $z = \frac{1}{n-1} \sum_{v \in \tilde{V}} h(v)$ satisfies $|z| \leq 1$. Therefore

$$1 - \operatorname{Re}(\lambda_h) = \frac{n-1}{n} [1 - \operatorname{Re}(z)]$$

and

$$|\lambda_h|^2 = \frac{1}{n^2} [1 + 2(n-1)\operatorname{Re}(z) + (n-1)^2|z|^2].$$

Using that $|z|^2 \leq 1$,

$$1 - |\lambda_h|^2 \geq \frac{2(n-1)}{n^2} [1 - \operatorname{Re}(z)] = \frac{2}{n} [1 - \operatorname{Re}(\lambda_h)].$$

Hence

$$1 - \operatorname{Re}(\lambda_h) \leq \frac{n}{2}(1 + |\lambda_h|)(1 - |\lambda_h|) \leq n(1 - |\lambda_h|). \quad \square$$

We have seen that the magnitudes $|\lambda_h|$, and thus the values of γ_d and $\|P_{id}^k - U\|_2$, do not depend on the location of the sink. By contrast, the choice of sink can affect the values of γ_c and $\|H_{id}^t - U\|_2$. Proposition 3.1 shows that the value of γ_c cannot vary by more than a factor of n . In Section 5, we will see two examples where moving the sink changes γ_c by a factor of roughly $n/2$.

3.2. Dual lattice: Continuous time. In this subsection, we define the first of two dual lattices and describe the correspondence between multiplicative harmonic functions and dual lattice vectors. Then we show the relationship between lengths of dual lattice vectors and eigenvalues of the continuous time sandpile chain.

Proposition 3.2. *The map from $\Delta^{-1}\mathbb{Z}^{n-1} \rightarrow \mathcal{H}$ given by*

$$x = (x_1, \dots, x_{n-1}) \mapsto h(v_j) = \begin{cases} e^{2\pi i x_j} & \text{if } 1 \leq j \leq n-1, \\ 1 & \text{if } j = n, \end{cases} \quad (9)$$

is a surjective homomorphism with kernel \mathbb{Z}^{n-1} . Therefore $\mathcal{H} \cong \Delta^{-1}\mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}$.

We treat x as a column vector; the notation $x = (x_1, \dots, x_{n-1})$ is purely for convenience.

Proof. Suppose $x \in \mathbb{R}^{n-1}$ maps to h by (9). We claim that $h \in \mathcal{H}$ if and only if $\Delta x \in \mathbb{Z}^{n-1}$. Indeed, h satisfies the geometric mean value property at v_j if and only if

$$1 = h(v_j)^{\deg(v_j)} \prod_{v_k \sim v_j} h(v_k)^{-1} = e^{2\pi i (\Delta x)_j},$$

where the second equality holds because $h(s) = 1$. Therefore $\Delta x \in \mathbb{Z}^{n-1}$ if and only if h satisfies the geometric mean value property at every non-sink vertex (which implies the geometric mean value property at the sink).

It follows that (9) defines a surjective homomorphism from $\Delta^{-1}\mathbb{Z}^{n-1}$ to \mathcal{H} . It is immediate from the definition that the kernel is \mathbb{Z}^{n-1} . \square

Remark 3.3. The isomorphism $\mathcal{H} \cong \Delta^{-1}\mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}$ is the dual version of the isomorphism $\mathcal{G} \cong \mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1}$. Although finite abelian groups are non-canonically isomorphic to their duals, in this case there is a natural map from $\mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1}$ to $\Delta^{-1}\mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}$, namely multiplication by Δ^{-1} . The corresponding map from \mathcal{G} to \mathcal{H} can be described as follows: Given a sandpile configuration η viewed as an element of \mathbb{Z}^{n-1} , let $x = \Delta^{-1}\eta$. Set $h(v_j) = e^{2\pi i x_j}$ for all $1 \leq j \leq n-1$ and $h(s) = 1$. If η and η' are equivalent configurations (that is, $\eta - \eta' \in \Delta\mathbb{Z}^{n-1}$), then the resulting functions h, h' will be equal.

When G is a directed graph, Δ is not symmetric, and $\mathcal{H} \cong \Delta^{-1}\mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}$ whereas $\mathcal{G} \cong \mathbb{Z}^{n-1}/\Delta^T\mathbb{Z}^{n-1}$; these groups are still isomorphic, but not naturally.

Fix $h \in \mathcal{H}$, and choose $|x_j| \leq 1/2$ such that $h(v_j) = e^{2\pi i x_j}$ for all $1 \leq j \leq n$. (In particular, $x_n = 0$.) The next theorem quantifies the following simple idea: If the eigenvalue

$$\lambda_h = \frac{1}{n} \sum_{j=1}^n e^{2\pi i x_j}$$

is close to 1, then the vector (x_1, \dots, x_{n-1}) should be close to 0; more precisely, the squared length $\|x\|_2^2$ is within a constant factor of $1 - \operatorname{Re}(\lambda_h)$.

Theorem 3.4. *Given $h \in \mathcal{H}$, choose $x \in \Delta^{-1}\mathbb{Z}^{n-1}$ of minimal Euclidean length such that $x \mapsto h$ via (9). Then*

$$8 \frac{\|x\|_2^2}{n} \leq 1 - \operatorname{Re}(\lambda_h) \leq 2\pi^2 \frac{\|x\|_2^2}{n}.$$

Proof. We know that $|x_j| \leq 1/2$ for all $1 \leq j \leq n-1$, so $|2\pi x_j| \leq \pi$. We will use the inequality $1 - t^2/2 \leq \cos(t) \leq 1 - 2t^2/\pi^2$ for all $|t| \leq \pi$, where the upper bound is Lemma 2.12. Since

$$1 - \operatorname{Re}(\lambda_h) = \frac{1}{n} \sum_{j=1}^n [1 - \operatorname{Re}(h(v_j))] = \frac{1}{n} \sum_{j=1}^{n-1} [1 - \cos(2\pi x_j)],$$

it follows that

$$\frac{1}{n} \sum_{j=1}^{n-1} \frac{2}{\pi^2} (2\pi x_j)^2 \leq 1 - \operatorname{Re}(\lambda_h) \leq \frac{1}{n} \sum_{j=1}^{n-1} \frac{1}{2} (2\pi x_j)^2,$$

which is the desired result. \square

We can bound the continuous time spectral gap γ_c by minimizing Theorem 3.4 over all $h \neq 1$. Since $h \equiv 1$ corresponds to vectors $x \in \mathbb{Z}^{n-1}$, we immediately obtain the first statement in Theorem 1.1:

Corollary 3.5. *The spectral gap of the continuous time sandpile chain satisfies*

$$8 \frac{\|y\|_2^2}{n} \leq \gamma_c \leq 2\pi^2 \frac{\|y\|_2^2}{n},$$

where y is a vector of minimal Euclidean length in $(\Delta^{-1}\mathbb{Z}^{n-1}) \setminus \mathbb{Z}^{n-1}$.

There are cases in which the shortest vector in $\Delta^{-1}\mathbb{Z}^{n-1} \setminus \mathbb{Z}^{n-1}$ is much longer than the shortest nonzero vector in $\Delta^{-1}\mathbb{Z}^{n-1}$. For example, if G is a cycle on n vertices, then the length of the shortest vector in $\Delta^{-1}\mathbb{Z}^{n-1} \setminus \mathbb{Z}^{n-1}$ has order \sqrt{n} . But since $\mathbb{Z}^{n-1} \subseteq \Delta^{-1}\mathbb{Z}^{n-1}$, the lattice $\Delta^{-1}\mathbb{Z}^{n-1}$ contains vectors of length 1.

3.3. Dual lattice: Discrete time. To analyze the discrete time sandpile chain using $\mathcal{H} \cong \Delta^{-1}\mathbb{Z}^{n-1}/\mathbb{Z}^{n-1}$, we could argue as follows. The eigenvalue λ_h is close to the unit circle when the values $h(v_j)$ are close to each other. This happens when the equivalence class in $\Delta^{-1}\mathbb{Z}^{n-1}$ associated with h contains a vector x close to the line $\{c\mathbf{1} : c \in \mathbb{R}\}$, where $\mathbf{1}$ denotes the all-ones vector. Instead of pursuing this approach, we start over with a different dual lattice construction that puts the sink on an equal footing with the other vertices. We will show that

the analogue of $\Delta^{-1}\mathbb{Z}^{n-1}$ in this new construction is naturally associated with the discrete time sandpile chain in the same way that $\Delta^{-1}\mathbb{Z}^{n-1}$ is associated with the continuous time sandpile chain. Although the construction is somewhat involved, the reward—an easy proof of Theorem 1.4—makes it worthwhile.

We first provide a sink-independent analogue of the group $\mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1}$. Recall that $\bar{\Delta}$ is the full n -dimensional Laplacian matrix of G . If $x \in \mathbb{R}^n$, then $\bar{\Delta}x \in \mathbb{R}_0^n$. It can be checked that under the map from \mathbb{Z}^{n-1} to \mathbb{Z}_0^n that sends $(x_1, \dots, x_{n-1}) \mapsto (x_1, \dots, x_{n-1}, -\sum_{j=1}^{n-1} x_j)$, the image of $\Delta\mathbb{Z}^{n-1}$ is exactly $\bar{\Delta}\mathbb{Z}^n$. Hence $\mathbb{Z}^{n-1}/\Delta\mathbb{Z}^{n-1} \cong \mathbb{Z}_0^n/\bar{\Delta}\mathbb{Z}^n$. This isomorphism is used in [21] to prove that the sandpile group is independent of the choice of sink.

The analogue to $\Delta^{-1}\mathbb{Z}^{n-1}$ is the dual lattice $(\bar{\Delta}\mathbb{Z}^n)^* = \{x \in \mathbb{R}_0^n : \langle x, y \rangle \in \mathbb{Z} \text{ for all } y \in \bar{\Delta}\mathbb{Z}^n\}$. We will see below that $(\bar{\Delta}\mathbb{Z}^n)^* = \bar{\Delta}^+\mathbb{Z}_0^n$, where $\bar{\Delta}^+$ is the Moore-Penrose pseudoinverse of $\bar{\Delta}$, an n -dimensional symmetric matrix which we now define.

Because the graph G is connected, $\bar{\Delta}$ has eigenvalue 0 with multiplicity 1, and its kernel is $\ker \bar{\Delta} = \{c\mathbf{1} : c \in \mathbb{R}\}$. The orthogonal subspace to $\ker \bar{\Delta}$ is \mathbb{R}_0^n . Since $\bar{\Delta}$ is symmetric, \mathbb{R}_0^n is an invariant subspace for the map $x \mapsto \bar{\Delta}x$. Indeed, this map is invertible when restricted to \mathbb{R}_0^n . Any vector in \mathbb{R}^n can be written uniquely as $y + c\mathbf{1}$ for $y \in \mathbb{R}_0^n$ and $c \in \mathbb{R}$. The matrix $\bar{\Delta}^+$ is defined by $\bar{\Delta}^+(y + c\mathbf{1}) = x$, where x is the unique vector in \mathbb{R}_0^n such that $\bar{\Delta}x = y$. Thus the maps $x \mapsto \bar{\Delta}x$ and $y \mapsto \bar{\Delta}^+y$ are inverses on \mathbb{R}_0^n . Also, $\ker \bar{\Delta}^+ = \ker \bar{\Delta}$, and the maps $x \mapsto \bar{\Delta}^+\bar{\Delta}x$ and $y \mapsto \bar{\Delta}\bar{\Delta}^+y$ on \mathbb{R}^n are both the same as orthogonal projection onto \mathbb{R}_0^n . If the eigenvalues of $\bar{\Delta}$ are $0 = \beta_0 < \beta_1 \leq \beta_2 \leq \dots \leq \beta_{n-1}$, then the eigenvalues of $\bar{\Delta}^+$ are $0 < \beta_{n-1}^{-1} \leq \beta_{n-2}^{-1} \leq \dots \leq \beta_1^{-1}$.

To see why $(\bar{\Delta}\mathbb{Z}^n)^* = \bar{\Delta}^+\mathbb{Z}_0^n$, suppose first that $x = \bar{\Delta}^+a$ for some $a \in \mathbb{Z}_0^n$. For any $b \in \mathbb{Z}^n$, $\langle x, \bar{\Delta}b \rangle = \langle \bar{\Delta}\bar{\Delta}^+a, b \rangle = \langle a, b \rangle \in \mathbb{Z}$. Thus $x \in (\bar{\Delta}\mathbb{Z}^n)^*$. Conversely, if $x \in (\bar{\Delta}\mathbb{Z}^n)^*$, then $x \in \mathbb{R}_0^n$ and $\langle \bar{\Delta}x, e_j \rangle = \langle x, \bar{\Delta}e_j \rangle \in \mathbb{Z}$ for all $1 \leq j \leq n$. This means $\bar{\Delta}x \in \mathbb{Z}_0^n$, so $x \in \bar{\Delta}^+\mathbb{Z}_0^n$.

Here is the parallel to Proposition 3.2.

Proposition 3.6. *The map from $\bar{\Delta}^+\mathbb{Z}_0^n \rightarrow \mathcal{H}$ given by*

$$x = (x_1, \dots, x_n) \mapsto h(v_j) = e^{2\pi i(x_j - x_n)} \quad (10)$$

is a surjective homomorphism with kernel \mathbb{W}^n . Therefore $\mathcal{H} \cong \bar{\Delta}^+\mathbb{Z}_0^n/\mathbb{W}^n$.

Proof. We first observe that the function $h(v_j) = e^{2\pi i y_j}$ is in \mathcal{H} if and only if the vector $y = (y_1, \dots, y_n)$ satisfies $y_n \in \mathbb{Z}$ and $\bar{\Delta}y \in \mathbb{Z}_0^n$. This is because the geometric mean value property at v_j is the statement

$$1 = h(v_j)^{\deg(v_j)} \prod_{v_k \sim v_j} h(v_k)^{-1} = e^{2\pi i(\bar{\Delta}y)_j}.$$

Note that $\bar{\Delta}y \in \mathbb{Z}_0^n$ if and only if $\bar{\Delta}y \in \mathbb{Z}^n$ since the columns of $\bar{\Delta}$ sum to zero. The condition $y_n \in \mathbb{Z}$ is equivalent to $h(s) = 1$.

Let $x \in \bar{\Delta}^+\mathbb{Z}_0^n$ and define h by (10). We write $h(v_j) = e^{2\pi i y_j}$, where $y = x - x_n\mathbf{1}$. Then $y_n = 0$ and $\bar{\Delta}y = \bar{\Delta}x \in \bar{\Delta}\bar{\Delta}^+\mathbb{Z}_0^n = \mathbb{Z}_0^n$, implying that $h \in \mathcal{H}$.

It is clear that the map (10) is a homomorphism. To see that it is surjective, take any $h \in \mathcal{H}$. Let $h(v_j) = e^{2\pi i y_j}$, so that $\bar{\Delta}y \in \mathbb{Z}_0^n$. Let $c = \frac{1}{n}(y_1 + \dots + y_n)$ and define $x = y - c\mathbf{1}$. Then $x \in \mathbb{R}_0^n$ and $\bar{\Delta}x = \bar{\Delta}y \in \mathbb{Z}_0^n$, meaning that $x = \bar{\Delta}^+ \bar{\Delta}x \in \bar{\Delta}^+ \mathbb{Z}_0^n$. In addition, $x_j - x_n = y_j - y_n$ and so $e^{2\pi i(x_j - x_n)} = h(v_j)/h(v_n) = h(v_j)$ for each j . Thus $x \mapsto h$ via (10).

For the kernel of (10), first note that $\mathbb{W}^n = \bar{\Delta}^+ \bar{\Delta} \mathbb{Z}^n \subseteq \bar{\Delta}^+ \mathbb{Z}_0^n$. A vector $x \in \bar{\Delta}^+ \mathbb{Z}_0^n$ maps to $h \equiv 1$ via (10) if and only if $x_j - x_n \in \mathbb{Z}$ for all j . If $x \in \mathbb{W}^n$, we can write $x = y - c\mathbf{1}$ for some $y \in \mathbb{Z}^n$ and $c \in \mathbb{R}$. Then each $x_j - x_n = y_j - y_n \in \mathbb{Z}$. Conversely, if $x_j - x_n \in \mathbb{Z}$ for all j , let $y = x - x_n \mathbf{1} \in \mathbb{Z}^n$. Since $x \in \mathbb{R}_0^n$ is a translate of y by a multiple of $\mathbf{1}$, x is the orthogonal projection of y onto \mathbb{R}_0^n , so $x \in \mathbb{W}^n$. This proves that \mathbb{W}^n is the kernel of (10). \square

Remark 3.7. As in the other dual lattice construction, multiplication by $\bar{\Delta}^+$ gives a natural map from $\mathbb{Z}_0^n / \bar{\Delta} \mathbb{Z}^n$ (which is isomorphic to \mathcal{G}) to $\bar{\Delta}^+ \mathbb{Z}_0^n / \bar{\Delta}^+ \bar{\Delta} \mathbb{Z}^n = \bar{\Delta}^+ \mathbb{Z}_0^n / \mathbb{W}^n$ (which is isomorphic to \mathcal{H}). The corresponding map from \mathcal{G} to \mathcal{H} is the same one described in the remark following Proposition 3.2.

Suppose $h \in \mathcal{H}$ is given. Every x that maps to h via (10) arises by the following recipe: Choose $y \in \mathbb{R}^n$ such that $h(v_j) = e^{2\pi i y_j}$ for all j ; set $c = \frac{1}{n}(y_1 + \dots + y_n)$; and let $x = y - c\mathbf{1}$. Then

$$\|x\|_2^2 = \sum_{j=1}^n (y_j - c)^2,$$

so x is short when the values y_j are close to each other. This can only happen if the eigenvalue $\lambda_h = \frac{1}{n} \sum_{j=1}^n h(v_j)$ is close to the unit circle. The next theorem, which is parallel to Theorem 3.4, makes this intuition precise.

Theorem 3.8. *Given $h \in \mathcal{H}$, choose $x \in \bar{\Delta}^+ \mathbb{Z}_0^n$ of minimal Euclidean length such that $x \mapsto h$ via (10). Then*

$$8 \frac{\|x\|_2^2}{n} \leq 1 - |\lambda_h| \leq 2\pi^2 \frac{\|x\|_2^2}{n}.$$

Proof. The upper bound is straightforward: Since $\cos(t) \geq 1 - t^2/2$,

$$\begin{aligned} |\lambda_h| &= \left| \frac{1}{n} \sum_{j=1}^n e^{2\pi i(x_j - x_n)} \right| = \frac{1}{n} \left| \sum_{j=1}^n e^{2\pi i x_j} \right| \geq \frac{1}{n} \operatorname{Re} \left[\sum_{j=1}^n e^{2\pi i x_j} \right] \\ &= \frac{1}{n} \sum_{j=1}^n \cos(2\pi x_j) \geq 1 - \frac{1}{n} \sum_{j=1}^n 2\pi^2 x_j^2 = 1 - 2\pi^2 \frac{\|x\|_2^2}{n}. \end{aligned}$$

For the lower bound, let $r = |\lambda_h|$ and write $\lambda_h = r e^{2\pi i \theta}$. We will construct a vector $x_0 \in \bar{\Delta}^+ \mathbb{Z}_0^n$ that maps to h via (10) for which $8\|x_0\|_2^2/n \leq 1 - r$. Since $\|x\|_2 \leq \|x_0\|_2$, the result will follow.

Choose $y \in \mathbb{R}^n$ such that for all j , $h(v_j) = e^{2\pi i y_j}$ and $|y_j - \theta| \leq 1/2$. Then

$$\lambda_h = r e^{2\pi i \theta} = \frac{1}{n} \sum_{j=1}^n e^{2\pi i y_j}, \text{ which implies that } r = \frac{1}{n} \sum_{j=1}^n e^{2\pi i(y_j - \theta)}.$$

Since $\cos(t) \leq 1 - 2t^2/\pi^2$ for all $|t| \leq \pi$ (by Lemma 2.12),

$$|\lambda_h| = r = \operatorname{Re}(r) = \frac{1}{n} \sum_{j=1}^n \cos(2\pi(y_j - \theta)) \leq 1 - \frac{1}{n} \sum_{j=1}^n 8(y_j - \theta)^2. \quad (11)$$

It follows that

$$8 \frac{\|y - \theta \mathbf{1}\|_2^2}{n} \leq 1 - r.$$

Let x_0 be the orthogonal projection of y onto \mathbb{R}_0^n . Then $x_0 = y - c\mathbf{1}$, where $c = \frac{1}{n}(y_1 + \dots + y_n)$. The surjectivity argument in the proof of Proposition 3.6 shows that $x_0 \in \bar{\Delta}^+ \mathbb{Z}_0^n$ and that $x_0 \mapsto h$ via (10). Since x_0 is also the orthogonal projection of $y - \theta \mathbf{1}$ onto \mathbb{R}_0^n , $\|x_0\|_2 \leq \|y - \theta \mathbf{1}\|_2$. Hence

$$8 \frac{\|x\|_2^2}{n} \leq 8 \frac{\|x_0\|_2^2}{n} \leq 8 \frac{\|y - \theta \mathbf{1}\|_2^2}{n} \leq 1 - r. \quad \square$$

By minimizing Theorem 3.8 over all $h \neq 1$, noting that $h \equiv 1$ corresponds to vectors $x \in \mathbb{W}^n$, we obtain the second statement in Theorem 1.1:

Corollary 3.9. *The spectral gap γ_d of the discrete time sandpile chain satisfies*

$$8 \frac{\|z\|_2^2}{n} \leq \gamma_d \leq 2\pi^2 \frac{\|z\|_2^2}{n},$$

where z is a vector of minimal Euclidean length in $(\bar{\Delta}^+ \mathbb{Z}_0^n) \setminus \mathbb{W}^n$.

As with the remark following Corollary 3.5, the cycle on n vertices provides an example where the shortest vector in $\bar{\Delta}^+ \mathbb{Z}_0^n \setminus \mathbb{W}^n$ has length of order \sqrt{n} , while the shortest nonzero vector in $\bar{\Delta}^+ \mathbb{Z}_0^n$ has length slightly less than 1.

3.4. An inverse relationship. Let $0 = \beta_0 < \beta_1 \leq \dots \leq \beta_{n-1}$ be the eigenvalues of the full Laplacian matrix Δ . The least positive eigenvalue β_1 is called the spectral gap of the graph G , or sometimes the *algebraic connectivity* of G . When G is a d -regular graph, β_1/d is the spectral gap of the simple random walk on G .² Larger values of β_1 mean that G is ‘more connected,’ and in the d -regular case that the simple random walk on G mixes faster (setting aside issues of periodicity).

In this subsection we prove Theorem 1.4, which says that the spectral gap of the discrete time sandpile chain satisfies

$$\gamma_d \leq \frac{4\pi^2}{\beta_1^2 n}.$$

Surprisingly, large values of β_1 cause the sandpile chain to mix slowly! This inverse relationship is reminiscent of the bound of Björner, Lovász and Shor [8, Theorem 4.2] on the number of topplings until a configuration stabilizes, which has a factor of β_1 in the denominator and is also proved using the pseudoinverse $\bar{\Delta}^+$.

²The second-largest eigenvalue of the transition matrix is $1 - \beta_1/d$. This disagrees with our earlier definition of the spectral gap as the minimum *absolute* distance between a nontrivial eigenvalue and the unit circle, but it accords with standard usage for reversible Markov chains.

Proof of Theorem 1.4. By Corollary 3.9, γ_d is determined up to a constant factor by the length of the shortest vector in the dual lattice $\bar{\Delta}^+\mathbb{Z}_0^n$ that is not also in the sublattice \mathbb{W}^n . The eigenvalues of $\bar{\Delta}^+$ are $0 < \beta_{n-1}^{-1} \leq \beta_{n-2}^{-1} \leq \dots \leq \beta_1^{-1}$. Therefore a lower bound on β_1 gives an upper bound on the operator norm of $\bar{\Delta}^+$.

Since $\mathcal{H} \cong \bar{\Delta}^+\mathbb{Z}_0^n/\mathbb{W}^n$ by Proposition 3.6, we may assume that \mathbb{W}^n is a proper sublattice of $\bar{\Delta}^+\mathbb{Z}_0^n$ (otherwise we would have $|\mathcal{H}| = 1$ and γ_d itself would be meaningless). The lattice \mathbb{Z}_0^n is generated by the vectors $e_i - e_j$ with $1 \leq i < j \leq n$. Thus we can choose i, j such that $\bar{\Delta}^+(e_i - e_j) \notin \mathbb{W}^n$. We have

$$\|\bar{\Delta}^+(e_i - e_j)\|_2^2 \leq \frac{\|e_i - e_j\|_2^2}{\beta_1^2} = \frac{2}{\beta_1^2}.$$

By Corollary 3.9,

$$\gamma_d \leq \frac{2\pi^2}{n} \|\bar{\Delta}^+(e_i - e_j)\|_2^2 \leq \frac{4\pi^2}{\beta_1^2 n}. \quad \square$$

Later in this subsection we will consider a class of graphs for which the bound in Theorem 1.4 is accurate to a constant factor. In many cases, though, it can be far from sharp. For example, when G is the torus $\mathbb{Z}_m \times \mathbb{Z}_m$ (so that $n = m^2$), β_1 has order $1/n$.

The analogue of Theorem 1.4 for the continuous time spectral gap is harder to prove because we must use Δ^{-1} instead of $\bar{\Delta}^+$. If $0 < \rho_1 \leq \rho_2 \leq \dots \leq \rho_{n-1}$ are the eigenvalues of Δ , then the Cauchy interlacing theorem [23, Ch. 4] implies that $\beta_{j-1} \leq \rho_j \leq \beta_j$ for all $1 \leq j \leq n-1$. Therefore the eigenvalues $\rho_{n-1}^{-1}, \dots, \rho_2^{-1}$ of Δ^{-1} are all bounded above by β_1^{-1} , but ρ_1^{-1} may be much larger than β_1^{-1} . A parallel argument to the proof of Theorem 1.4 would break down because of this large eigenvalue. Nevertheless, with considerable effort we can show the following bound. We omit the proof for reasons of space.

Theorem 3.10. *Suppose the graph G has no vertices of degree 1. The spectral gap of the continuous time sandpile chain satisfies*

$$\gamma_c \leq \frac{4\pi^2 + O(1/n)}{\beta_1^2 n} \leq \frac{10\pi^2}{\beta_1^2 n}.$$

Combining Theorems 2.11 and 1.4 gives both upper and lower bounds on the discrete time spectral gap:

$$\frac{8}{d_*^2 n} \leq \gamma_d \leq \frac{4\pi^2}{\beta_1^2 n},$$

where d_* is the second-largest vertex degree. We now introduce a class of graphs for which the upper and lower bounds have the same order. These are, in our language, the expander graphs with bounded degree ratio.

Loosely, the graph G is an expander if it is sparse and if the simple random walk on G mixes quickly. Expanders have many applications in pure and applied mathematics as well as theoretical computer science. See the surveys [28, 22] for more information.

For the purposes of this paper, we will not need the requirement of sparsity. Let K be the transition matrix for the simple random walk on G , and let $L = I - K$

be the *random walk Laplacian*. Since K is reversible with respect to the stationary distribution $\pi(v) = \deg(v) / \sum_{w \in V} \deg(w)$, L has real eigenvalues $0 = \theta_0 < \theta_1 \leq \theta_2 \leq \dots \leq \theta_{n-1}$. The value θ_1 is called the spectral gap of L . When G is d -regular, $\bar{\Delta} = dL$ and so $\beta_1 = d\theta_1$.

Definition 3.11. Fix $\alpha > 0$, and let L be the random walk Laplacian matrix for the graph G . We say G is an α -*expander* if the spectral gap of L satisfies $\theta_1 \geq \alpha$.

The following lemma, whose proof is postponed to the end of this section, gives a relationship between β_1 and θ_1 when G is not regular.

Lemma 3.12. *Let G have minimum and maximum vertex degrees d_{\min} and d_{\max} . The eigenvalues $0 = \beta_0 < \beta_1 \leq \dots \leq \beta_{n-1}$ of $\bar{\Delta}$ and $0 = \theta_0 < \theta_1 \leq \dots \leq \theta_{n-1}$ of L are related by*

$$d_{\min}\theta_i \leq \beta_i \leq d_{\max}\theta_i, \quad 0 \leq i \leq n-1.$$

In particular, $\beta_1 \geq d_{\min}\theta_1$. Thus if G is an α -expander and the *degree ratio* d_{\max}/d_{\min} is bounded by a constant R , then $\beta_1 \geq (\alpha/R)d_{\max}$.

Proposition 3.13. *Suppose that for fixed $\alpha, R > 0$, the graph G is an α -expander and $d_{\max}/d_{\min} \leq R$. Then*

$$\frac{8}{d_{\max}^2 n} \leq \gamma_d \leq \frac{4\pi^2 R^2 / \alpha^2}{d_{\max}^2 n}.$$

Proof. The lower bound is Theorem 2.11. The upper bound is Theorem 1.4 combined with Lemma 3.12, as discussed above. \square

For expanders with bounded degree ratio, Proposition 3.13 determines the spectral gap γ_d of the discrete time sandpile chain up to a constant factor.

Note that we could replace d_{\max} with d_* in the statement of Proposition 3.13 to get a slightly stronger result with the exact same proof. This is useful in graphs where a single vertex has much higher degree than all the others.

We can prove an analogue of Proposition 3.13 for the continuous time spectral gap γ_c . The proof uses Theorem 3.10 when $d_{\min} \geq 2$ and a separate argument when $d_{\min} = 1$. The lower bound is the same, since $\gamma_c \geq \gamma_d$. In the upper bound, the numerator $4\pi^2 R^2 / \alpha^2$ is replaced by $1300R^3 / \alpha^2$.

Proposition 3.13 allows us to determine the order of the relaxation time for the sandpile chain on some well-known random graphs.

Proposition 3.14. *Fix $d \geq 3$, and let (G_n) be a sequence of random d -regular graphs with $|V(G_n)| = n \rightarrow \infty$.³ The spectral gaps $\gamma_d(n)$ for the discrete time sandpile chain satisfy $\gamma_d = \Theta(1/n)$ almost surely.*

Proposition 3.15. *Let $p = p(n)$ satisfy $0 \leq p \leq 1$ and $np/\log(n) \rightarrow \infty$ as $n \rightarrow \infty$. Let $G_n = G(n, p)$ be independently chosen Erdős-Rényi random graphs on n vertices with edge probability p . The spectral gaps $\gamma_d(n)$ for the discrete time sandpile chain satisfy $\gamma_d = \Theta(1/n^3 p^2)$ almost surely.*

³More precisely, assume that for each n in the sequence, the set of simple d -regular graphs on n vertices is nonempty. Choose G_n uniformly from this set, independently for each n .

Both these propositions are simple applications of Proposition 3.13. For Proposition 3.14, all we need is the well-known fact that random d -regular graphs are expanders. A famous result of Friedman [16] says that for every $\varepsilon > 0$, the random walk spectral gap satisfies $\theta_1 \geq 1 - 2\sqrt{d-1}/d - \varepsilon$ asymptotically almost surely.

For Proposition 3.15, we first note that in our regime, d_{\min}/np and d_{\max}/np converge to 1 almost surely as $n \rightarrow \infty$. (Use Bernstein's Inequality to show that the degree of each vertex concentrates around its mean $(n-1)p$, then take a union bound over all the vertices.) To prove expansion, Theorem 2 of [10] implies that because $np/\log(n) \rightarrow \infty$, almost surely $\theta_1 \geq 1 - o(1)$.

We conclude with the proof of Lemma 3.12.

Proof of Lemma 3.12. Define $Z = \sum_{j=1}^n \deg(v_j)$. The stationary distribution π of the simple random walk on G is $\pi(v_j) = \deg(v_j)/Z$. Consider these two inner products on the space $\mathbb{R}^V = \{f : V \rightarrow \mathbb{R}\}$:

$$\langle f, g \rangle_\pi = \sum_{j=1}^n f(v_j)g(v_j)\pi(v_j), \quad \langle f, g \rangle = \sum_{j=1}^n f(v_j)g(v_j).$$

The matrix $\bar{\Delta}$ is self-adjoint with respect to $\langle \cdot, \cdot \rangle$, while L is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\pi$. We have $\bar{\Delta} = DL$ where D is the diagonal matrix with entries $D(j, j) = \deg(v_j)$. The two inner products are related by

$$\langle f, Dg \rangle = \sum_{j=1}^n f(v_j) \deg(v_j)g(v_j) = Z \langle f, g \rangle_\pi.$$

In addition,

$$\frac{d_{\min}}{Z} \langle f, f \rangle \leq \sum_{j=1}^n \frac{\deg(v_j)}{Z} f(v_j)^2 = \langle f, f \rangle_\pi \leq \frac{d_{\max}}{Z} \langle f, f \rangle.$$

It follows that for $f \neq 0$,

$$\begin{aligned} d_{\min} \left(\frac{\langle f, Lf \rangle_\pi}{\langle f, f \rangle_\pi} \right) &= \frac{Z \langle f, Lf \rangle_\pi}{(Z/d_{\min}) \langle f, f \rangle_\pi} \leq \frac{\langle f, DLf \rangle}{\langle f, f \rangle} \\ &\leq \frac{Z \langle f, Lf \rangle_\pi}{(Z/d_{\max}) \langle f, f \rangle_\pi} = d_{\max} \left(\frac{\langle f, Lf \rangle_\pi}{\langle f, f \rangle_\pi} \right). \end{aligned} \quad (12)$$

Since L is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\pi$ and $\bar{\Delta} = DL$ is symmetric, the variational characterization of eigenvalues (see [23]) says that for each $0 \leq i \leq n-1$,

$$\beta_i = \min_{\substack{U \subseteq \mathbb{R}^V \\ \dim(\bar{U})=i+1}} \max_{f \in U, f \neq 0} \frac{\langle f, \bar{\Delta}f \rangle}{\langle f, f \rangle}. \quad (13)$$

Here the minimum is taken over all linear subspaces $U \subseteq \mathbb{R}^V$ of dimension $i+1$. Likewise,

$$\theta_i = \min_{\substack{U \subseteq \mathbb{R}^V \\ \dim(\bar{U})=i+1}} \max_{f \in U, f \neq 0} \frac{\langle f, Lf \rangle_\pi}{\langle f, f \rangle_\pi}. \quad (14)$$

Comparing the characterizations (13) and (14) and using the upper and lower bounds in (12), we conclude that

$$d_{\min}\theta_i \leq \beta_i \leq d_{\max}\theta_i. \quad \square$$

4. THE SMOOTHING PARAMETER

This section discusses the relationship between the mixing time of the sandpile chain and a lattice invariant called the *smoothing parameter*, which was introduced by Micciancio and Regev [29] in the context of lattice-based cryptography. Our main result, Theorem 4.1, is a slightly stronger version of Theorem 1.2. It provides an upper bound on the L^2 mixing time in terms of the smoothing parameter for the Laplacian lattice.

The starting point of the proof is the relationship between lengths of dual lattice vectors and eigenvalues of the sandpile chain developed in Sections 3.2 and 3.3. Previously we focused on the consequences for the spectral gap. Here we consider all the eigenvalues in order to control the L^2 distance from stationarity.

With Theorem 4.1 in hand, we can take advantage of existing results that relate the smoothing parameter to other lattice invariants. By this approach, we show that if d_* is the second-highest vertex degree in G , then the mixing times for both the discrete and continuous time sandpile chains are at most $\frac{1}{16}d_*^2n \log n$. The precise statement is Theorem 4.3 below.

We already knew from Theorem 2.11 that the relaxation time is at most $\frac{1}{8}d_*^2n$, so (by Proposition 2.10) the mixing time is at most $\frac{1}{16}(d_*^2 \log d_*)n^2$. The new bound of $\frac{1}{16}d_*^2n \log n$ is a significant improvement. When $d_* \gg \sqrt{n}$, Theorem 4.5 further improves the leading constant from $\frac{1}{16}$ to $\frac{1}{4\pi^2}$. We will see in Section 5 that this bound is sharp for the discrete time sandpile chain on the complete graph.

Let us begin by recalling the definition of the smoothing parameter from Section 1. Let Λ be a lattice in \mathbb{R}^m , and let $V = \text{Span}(\Lambda) \subseteq \mathbb{R}^m$. For $s > 0$, the function

$$f_\Lambda(s) = \sum_{x \in \Lambda^* \setminus \{\mathbf{0}\}} e^{-\pi s^2 \|x\|_2^2}$$

is continuous and strictly decreasing, with a limit of ∞ as $s \rightarrow 0$ and a limit of 0 as $s \rightarrow \infty$. For $\varepsilon > 0$, the smoothing parameter of Λ is defined as $\eta_\varepsilon(\Lambda) := f_\Lambda^{-1}(\varepsilon)$. Therefore,

$$\sum_{x \in \Lambda^* \setminus \{\mathbf{0}\}} e^{-\pi \eta_\varepsilon^2(\Lambda) \|x\|_2^2} = \varepsilon. \quad (15)$$

In [29] the smoothing parameter was defined only for lattices of full rank. The definition given above is a natural generalization. Indeed, suppose Λ has rank $k < m$. Fix a vector space isomorphism $\Phi : V \rightarrow \mathbb{R}^k$ that preserves the standard inner product. We see that $\Phi(\Lambda^*) = (\Phi(\Lambda))^*$, so $f_\Lambda(s) = f_{\Phi(\Lambda)}(s)$ and $\eta_\varepsilon(\Lambda) = \eta_\varepsilon(\Phi(\Lambda))$. Results that relate the smoothing parameter to other values invariant under Φ , such as lengths of vectors, carry over without change to the general setting.

The reason for the name ‘smoothing parameter’ is an alternative characterization that we will not use in our proofs. Let $\Lambda \subseteq \mathbb{R}^k$ be a full-rank lattice. Given $c \in \mathbb{R}^k$ and $r > 0$, push forward the density for a Gaussian on \mathbb{R}^k with mean c and

covariance matrix $\frac{r^2}{2\pi}\text{Id}$ by the quotient map $\mathbb{R}^k \rightarrow \mathcal{D} := \mathbb{R}^k/\Lambda$. It is shown in [29] that the L^∞ distance between this density and the uniform density on \mathcal{D} is independent of c , and equals ε precisely when $r = \eta_\varepsilon(\Lambda)$. Thus $\eta_\varepsilon(\Lambda)$ is the amount of scaling necessary for a Gaussian to be nearly uniformly distributed over \mathcal{D} , with an L^∞ error of ε .

Theorem 4.1. *Let H_{id}^t and P_{id}^t denote the distributions at time t of the continuous and discrete time sandpile chains, respectively, started from the identity, and let U denote the uniform distribution on \mathcal{G} . For any $\varepsilon > 0$,*

$$\|H_{id}^t - U\|_2^2 \leq \varepsilon \text{ for all } t \geq \frac{\pi}{16}n \cdot \eta_\varepsilon^2(\Delta\mathbb{Z}^{n-1}), \quad (16)$$

$$\max\{\|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2\} \leq \varepsilon \text{ for all } t \geq \frac{\pi}{16}n \cdot \eta_\varepsilon^2(\bar{\Delta}\mathbb{Z}^n). \quad (17)$$

Since $\|H_\sigma^t - U\|_2^2 = \|H_{id}^t - U\|_2^2$ and $\|P_\sigma^t - U\|_2^2 = \|P_{id}^t - U\|_2^2$ for any $\sigma \in \mathcal{G}$, Theorem 4.1 implies Theorem 1.2. The only difference is that the bound (17) applies to both the continuous and discrete time sandpile chains.

Proof of Theorem 4.1. The proofs of (16) and (17) are almost identical. For (16), start with the formula

$$\|H_{id}^t - U\|_2^2 = \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1 - \text{Re}(\lambda_h))}.$$

In Section 3.2 we saw that each $h \in \mathcal{H}$ corresponds to an equivalence class of vectors in the dual lattice $\Delta^{-1}\mathbb{Z}^{n-1}$. Let x_h be a member of the equivalence class corresponding to h , chosen to have minimal Euclidean length. By Theorem 3.4,

$$1 - \text{Re}(\lambda_h) \geq 8 \frac{\|x_h\|_2^2}{n}.$$

Let $W = \{x_h : h \in \mathcal{H}\} \subseteq \Delta^{-1}\mathbb{Z}^{n-1}$. Since $h \equiv 1$ corresponds to $x_h = \mathbf{0}$,

$$\|H_{id}^t - U\|_2^2 \leq \sum_{x \in W \setminus \{\mathbf{0}\}} e^{-\frac{16t}{n}\|x\|_2^2} \leq \sum_{x \in \Delta^{-1}\mathbb{Z}^{n-1} \setminus \{\mathbf{0}\}} e^{-\frac{16t}{n}\|x\|_2^2}.$$

Now using the lower bound on t in (16) along with (15),

$$\sum_{x \in \Delta^{-1}\mathbb{Z}^{n-1} \setminus \{\mathbf{0}\}} e^{-\frac{16t}{n}\|x\|_2^2} \leq \sum_{x \in \Delta^{-1}\mathbb{Z}^{n-1} \setminus \{\mathbf{0}\}} e^{-\pi\eta_\varepsilon^2(\Delta\mathbb{Z}^{n-1})\|x\|_2^2} = \varepsilon.$$

This proves (16).

To show (17), we have

$$\begin{aligned} \|H_{id}^t - U\|_2^2 &= \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1 - \text{Re}(\lambda_h))} \leq \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1 - |\lambda_h|)}, \\ \|P_{id}^t - U\|_2^2 &= \sum_{h \in \mathcal{H} \setminus \{1\}} |\lambda_h|^{2t} = \sum_{h \in \mathcal{H} \setminus \{1\}} e^{2t \log |\lambda_h|} \leq \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1 - |\lambda_h|)}. \end{aligned}$$

From there the proof is the same as above, using the dual lattice $\bar{\Delta}^+\mathbb{Z}_0^n$ in place of $\Delta^{-1}\mathbb{Z}^{n-1}$ and Theorem 3.8 in place of Theorem 3.4. \square

Theorem 4.4 will show that the leading constant $\pi/16$ in Theorem 4.1 can be improved under certain conditions.

To apply Theorem 4.1, we need to bound the smoothing parameter in terms of other lattice invariants. There are several such bounds in the literature. We will use the following result of [29].

Lemma 4.2 ([29], Lemma 3.3). *For any lattice Λ of rank k and any $\varepsilon > 0$,*

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\log(2k(1+1/\varepsilon))}{\pi}} \cdot \lambda_k(\Lambda),$$

where $\lambda_k(\Lambda)$ is the least real number r such that the closed Euclidean ball of radius r about the origin contains at least k linearly independent vectors in Λ .

(In [29] the lemma is stated for lattices of full rank. It extends to the general case by the previous remark about the isomorphism Φ .)

If we combine Theorem 4.1 with Lemma 4.2 and bound $\lambda_k(\Lambda)$ using the entries of the Laplacian matrix, we obtain the following bound on mixing time.

Theorem 4.3. *Let d_* be the second-highest vertex degree in G . For any $\varepsilon > 0$,*

$$\max \{ \|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2 \} \leq \varepsilon \text{ for all } t \geq \frac{1}{16}(d_*^2 + d_*)n \log[2(n-1)(1+1/\varepsilon)].$$

Proof. To bound $\lambda_{n-1}(\bar{\Delta}\mathbb{Z}^n)$, observe that $\bar{\Delta}\mathbb{Z}^n$ is generated by any $n-1$ columns of $\bar{\Delta}$. We omit the column corresponding to the highest-degree vertex. All the other columns satisfy $\|x\|_2^2 \leq d_*^2 + d_*$, implying that $\lambda_{n-1}^2(\bar{\Delta}\mathbb{Z}^n) \leq d_*^2 + d_*$.

Lemma 4.2 yields

$$\eta_\varepsilon^2(\bar{\Delta}\mathbb{Z}^n) \leq \frac{1}{\pi}(d_*^2 + d_*) \log[2(n-1)(1+1/\varepsilon)]. \quad (18)$$

The proof is finished by plugging this bound into (17). \square

Theorem 4.3 shows that every sandpile chain on a graph G with second-highest vertex degree d_* has L^2 mixing time at most $\frac{1}{16}d_*^2n \log n$ plus lower-order terms. When G is the complete graph, we will see below that the order $d_*^2n \log n$ is correct but the constant $\frac{1}{16}$ is not sharp. By contrast, when G is the cycle graph or the ‘triangle with tail’ (see Section 5), the order $d_*^2n \log n$ is not sharp.

We now state a modified version of Theorem 4.1 which improves the leading constant from $\pi/16$ to $1/4\pi$ at the cost of an extra term. As usual, η_ε denotes the smoothing parameter, and $|\mathcal{G}|$ the order of the sandpile group.

Theorem 4.4. *For any $\varepsilon > 0$,*

$$\|H_{id}^t - U\|_2^2 \leq 2\varepsilon \text{ for all } t \geq \frac{1}{4\pi}n \cdot \eta_\varepsilon^2(\Delta\mathbb{Z}^{n-1}) + \frac{\pi^2}{48}n \log\left(\frac{|\mathcal{G}|}{\varepsilon}\right),$$

$$\max \{ \|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2 \} \leq 2\varepsilon \text{ for all } t \geq \frac{1}{4\pi}n \cdot \eta_\varepsilon^2(\bar{\Delta}\mathbb{Z}^n) + \frac{\pi^2}{48}n \log\left(\frac{|\mathcal{G}|}{\varepsilon}\right).$$

When we combine the second statement of Theorem 4.4 with (18) and the inequality $\log(|\mathcal{G}|) \leq (n-1) \log(d_*)$, we obtain the following result.

Theorem 4.5. *Let d_* be as in Theorem 4.3. For any $c \geq 1$,*

$$\max \{ \|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2 \} \leq e^{-c} \text{ for all } t \geq \frac{1}{4\pi^2} d_*^2 n \log(n) + \frac{n^2}{4} \log(d_*) + c d_*^2 n.$$

We will show in Section 5 that the leading term $\frac{1}{4\pi^2} d_*^2 n \log n$ is sharp for the discrete time sandpile chain on the complete graph. In general, Theorem 4.5 is better than Theorem 4.3 when $d_* \gg \sqrt{n}$ and worse when $d_* \ll \sqrt{n}$.

Proof of Theorem 4.4. We start by proving the second statement of the theorem.

Fix $b > 0$; we will specify its value later. We partition the set \mathcal{H} of multiplicative harmonic functions into two subsets $\mathcal{H}_1, \mathcal{H}_2$ as follows. Given $h \in \mathcal{H}$, write $\lambda_h = r e^{2\pi i \theta}$ with $r \geq 0$. (If $r = 0$ then let $\theta = 0$.) Choose $y \in \mathbb{R}^n$ such that for all j , $h(v_j) = e^{2\pi i y_j}$ and $|y_j - \theta| \leq 1/2$. As in the proof of Theorem 3.8,

$$|\lambda_h| = \frac{1}{n} \sum_{j=1}^n \cos(2\pi(y_j - \theta)).$$

Now define

$$\begin{aligned} \mathcal{H}_1 &= \{h \in \mathcal{H} : \max_{1 \leq j \leq n} 2\pi|y_j - \theta| \leq b\}, \\ \mathcal{H}_2 &= \{h \in \mathcal{H} : \max_{1 \leq j \leq n} 2\pi|y_j - \theta| > b\}. \end{aligned}$$

Note that \mathcal{H}_2 is empty if $b \geq \pi$. The reason for the partition is that when $h \in \mathcal{H}_1$, the cosine approximation from Lemma 2.12 is improved; and when $h \in \mathcal{H}_2$, the existence of j for which $2\pi|y_j - \theta| > b$ makes $|\lambda_h|$ relatively far from 1.

For $h \in \mathcal{H}_1$, Lemma 2.12 yields

$$|\lambda_h| \leq \frac{1}{n} \sum_{j=1}^n \left[1 - \frac{1 - \cos(b)}{b^2} \cdot 4\pi^2 (y_j - \theta)^2 \right] \leq 1 - \frac{4\pi^2}{n} \left(\frac{1}{2} - \frac{b^2}{24} \right) \sum_{j=1}^n (y_j - \theta)^2,$$

using that $1 - \cos(b) \geq b^2/2 - b^4/24$ in the second inequality. Compare with (11). Let $x_h \in \bar{\Delta}^+ \mathbb{Z}_0^n$ be a dual lattice vector of minimal Euclidean length corresponding to h . Arguing as in the proof of Theorem 3.8, we conclude that

$$1 - |\lambda_h| \geq \frac{4\pi^2}{n} \left(\frac{1}{2} - \frac{b^2}{24} \right) \|x_h\|_2^2. \quad (19)$$

For $h \in \mathcal{H}_2$, there exists j such that $b < 2\pi|y_j - \theta| \leq \pi$, which gives the bound

$$|\lambda_h| \leq 1 - \frac{1}{n} + \frac{1}{n} \cos(b).$$

Since $b < \pi$, Lemma 2.12 says that $\cos(b) \leq 1 - 2b^2/\pi^2$, so

$$1 - |\lambda_h| \geq \frac{1}{n} \cdot \frac{2b^2}{\pi^2}. \quad (20)$$

We are now ready to bound $\max \{\|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2\}$. To begin,

$$\begin{aligned} \max \{\|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2\} &\leq \sum_{h \in \mathcal{H} \setminus \{1\}} e^{-2t(1-|\lambda_h|)} \\ &= \sum_{h \in \mathcal{H}_1 \setminus \{1\}} e^{-2t(1-|\lambda_h|)} + \sum_{h \in \mathcal{H}_2} e^{-2t(1-|\lambda_h|)}. \end{aligned}$$

For the first sum, let $W_1 = \{x_h : h \in \mathcal{H}_1\} \subseteq \bar{\Delta}^+ \mathbb{Z}^n$. By (19),

$$\sum_{h \in \mathcal{H}_1 \setminus \{1\}} e^{-2t(1-|\lambda_h|)} \leq \sum_{x \in W_1 \setminus \{0\}} e^{-2t \cdot \frac{4\pi^2}{n} \left(\frac{1}{2} - \frac{b^2}{24}\right) \|x\|_2^2}.$$

If we choose b so that

$$2t \cdot \frac{4\pi^2}{n} \left(\frac{1}{2} - \frac{b^2}{24}\right) \geq \pi \eta_\varepsilon^2(\bar{\Delta} \mathbb{Z}^n), \quad (21)$$

then this sum will be at most ε by the definition of the smoothing parameter.

The second sum is zero when $b \geq \pi$. When $b < \pi$, we use (20) and the inequality $|\mathcal{H}_2| \leq |\mathcal{G}|$ to obtain

$$\sum_{h \in \mathcal{H}_2} e^{-2t(1-|\lambda_h|)} \leq |\mathcal{G}| e^{-\frac{4tb^2}{\pi^2 n}} = e^{\log(|\mathcal{G}|) - \frac{4tb^2}{\pi^2 n}}.$$

This quantity will be at most ε provided that

$$\log |\mathcal{G}| - \frac{4tb^2}{\pi^2 n} \leq \log \varepsilon. \quad (22)$$

To prove the theorem, we find a value of b that satisfies both (21) and (22). Set

$$b = \left[\frac{\pi^2 n}{4t} \log \left(\frac{|\mathcal{G}|}{\varepsilon} \right) \right]^{1/2},$$

so that equality holds in (22). Then (21) holds if and only if t satisfies the lower bound in the second statement of Theorem 4.4. This completes the proof.

For the first statement of Theorem 4.4, we retain the same value of b . Given $h \in \mathcal{H}$, choose $x \in \Delta^{-1} \mathbb{Z}^{n-1}$ such that $h(v_j) = e^{2\pi i x_j}$ and $|x_j| \leq 1/2$ for all j . Define

$$\begin{aligned} \mathcal{H}_1 &= \{h \in \mathcal{H} : \max_{1 \leq j \leq n-1} 2\pi |x_j| \leq b\}, \\ \mathcal{H}_2 &= \{h \in \mathcal{H} : \max_{1 \leq j \leq n-1} 2\pi |x_j| > b\}. \end{aligned}$$

Since

$$1 - \operatorname{Re}(\lambda_h) = \frac{1}{n} \sum_{j=1}^{n-1} [1 - \cos(2\pi x_j)],$$

for any $h \in \mathcal{H}_1$ we can apply Lemma 2.12 to get

$$1 - \operatorname{Re}(\lambda_h) \geq \frac{4\pi^2}{n} \left(\frac{1}{2} - \frac{b^2}{24}\right) \|x\|_2^2$$

as in the previous argument. Likewise, for any $h \in \mathcal{H}_2$ we have

$$1 - \operatorname{Re}(\lambda_h) \geq \frac{1}{n}[1 - \cos(b)] \geq \frac{1}{n} \cdot \frac{2b^2}{\pi^2},$$

keeping in mind that \mathcal{H}_2 can only be nonempty when $b < \pi$.

With this preparation, we decompose

$$\|H_{id}^t - U\|_2^2 = \sum_{h \in \mathcal{H}_1 \setminus \{1\}} e^{-2t(1 - \operatorname{Re}(\lambda_h))} + \sum_{h \in \mathcal{H}_2} e^{-2t(1 - \operatorname{Re}(\lambda_h))}.$$

From here the proof is the same as before, substituting $\Delta\mathbb{Z}^{n-1}$ for $\bar{\Delta}\mathbb{Z}^n$. \square

The proof of Theorem 4.5 is elementary. No effort has been made to optimize the constants in the non-leading terms.

Proof of Theorem 4.5. The only graph on 3 vertices with nontrivial sandpile group is the cycle C_3 , for which the result can be checked directly. Therefore we may assume that $n \geq 4$.

Starting from the second statement of Theorem 4.4, we use (18) along with the inequality $\log|\mathcal{G}| \leq (n-1)\log(d_*)$. For any $\varepsilon > 0$, if

$$t \geq \frac{1}{4\pi^2}(d_*^2 + d_*)n \log[2(n-1)(1+1/\varepsilon)] + \frac{\pi^2}{48}n[(n-1)\log(d_*) + \log(1/\varepsilon)], \quad (23)$$

then $\max\{\|H_{id}^t - U\|_2^2, \|P_{id}^t - U\|_2^2\} \leq 2\varepsilon$. Since $\log(1/\varepsilon) < d_*^2 \log(2+2/\varepsilon)$, the right side of (23) is less than

$$\frac{1}{4\pi^2}(d_*^2 + d_*)n[\log(n) + \log(2+2/\varepsilon)] + \frac{\pi^2}{48}[n^2 \log(d_*) + d_*^2 n \log(2+2/\varepsilon)]. \quad (24)$$

Using the bound $d_* \log(n) \leq n \log(d_*)$ (since $2 \leq d_* < n$ and $n \geq 4$) along with $d_*^2 + d_* < 2d_*^2$, (24) is less than

$$\frac{1}{4\pi^2}d_*^2 n \log(n) + \left[\frac{1}{4\pi^2} + \frac{\pi^2}{48}\right]n^2 \log(d_*) + \left[\frac{2}{4\pi^2} + \frac{\pi^2}{48}\right]d_*^2 n \log(2+2/\varepsilon). \quad (25)$$

Since

$$\frac{1}{4\pi^2} + \frac{\pi^2}{48} < \frac{1}{4},$$

any $t \geq \frac{1}{4\pi^2}d_*^2 n \log(n) + \frac{1}{4}n^2 \log(d_*) + cd_*^2 n$ will be greater than (25) as long as

$$c \geq \left(\frac{2}{4\pi^2} + \frac{\pi^2}{48}\right) \log(2+2/\varepsilon). \quad (26)$$

It remains to show that (26) holds for $\varepsilon = \frac{1}{2}e^{-2c}$ whenever $c \geq 1$, as this implies that $\max\{\|H_{id}^t - U\|_2, \|P_{id}^t - U\|_2\} \leq \sqrt{2\varepsilon} = e^{-c}$. Using $e^{-2c} \leq 1$ followed by $c \geq 1$, we have

$$\log(2+2/\varepsilon) = \log\left(\frac{e^{-2c} + 2}{\frac{1}{2}e^{-2c}}\right) \leq \log(6) + 2c \leq \left(\frac{2}{4\pi^2} + \frac{\pi^2}{48}\right)^{-1} c,$$

and the proof is complete. \square

5. EXAMPLES

We conclude by illustrating our results in a variety of specific examples. We will use the notation $[n] = \{1, \dots, n\}$.

Cycle graph. If $G = C_n$ is an n -cycle with successive vertices $v_1, \dots, v_{n-1}, v_n = s$, then for each n th root of unity ω , the function defined by $h_\omega(v_k) = \omega^k$ is easily seen to be multiplicative harmonic. As there are n spanning trees in an n -cycle, this completely accounts for the characters of $\mathcal{G}(C_n)$. Moreover, this shows that $\mathcal{G}(C_n) \cong \mathcal{H}(C_n) \cong \mathbb{Z}/n\mathbb{Z}$.

Since $\sum_{k=1}^n \omega^k = 0$ whenever ω is a nontrivial n th root of unity, we see that $\lambda = 0$ is an eigenvalue of the discrete time sandpile chain with multiplicity $n - 1$. Thus (8) shows that the chain is stationary after a single step.

Complete graph. Let $G = K_n$ be the complete graph on $n \geq 3$ vertices, $v_1, \dots, v_{n-1}, v_n = s$. Set $M = \left\{ z \in (\mathbb{Z}/n\mathbb{Z})^n : \sum_{j=1}^{n-1} z_j = z_n = 0 \right\}$ and $\omega = e^{2\pi i/n}$. For every $z \in M$, the function h_z defined by $h_z(v_j) = \omega^{z_j}$ is multiplicative harmonic since

$$\prod_{j=1}^n h_z(v_j) = \omega^{\sum_{j=1}^n z_j} = 1$$

and thus

$$\prod_{j \neq k} h_z(v_j) = h_z(v_k)^{-1} = h_z(v_k)^{n-1}$$

for all $k = 1, \dots, n$. Every element of M is uniquely determined by specifying the first $n - 2$ coordinates, so $|M| = n^{n-2}$. As this is equal to the number of spanning trees of K_n by Cayley's formula [25], we have $\mathcal{H} = \{h_z\}_{z \in M}$, so the eigenvalues of the discrete time sandpile chain on G are $\lambda_z = \frac{1}{n} \sum_{j=1}^n \omega^{z_j}$ as z ranges over M . This characterization of \mathcal{H} shows that K_n has sandpile group $(\mathbb{Z}/n\mathbb{Z})^{n-2}$.

By construction, no $z \in M \setminus \{\mathbf{0}\}$ can have all coordinates in $\{0, 1\}$ or $\{0, -1\}$, so $z_* = (1, -1, 0, \dots, 0)$ gives the maximum modulus of a nontrivial eigenvalue. (Any permutation of the first $n - 1$ coordinates of z_* or of $\pm(2, 1, \dots, 1, 0)$ also gives a nontrivial eigenvalue of maximum modulus.) The inequality $\cos(x) \leq 1 - \frac{27}{8\pi^2} x^2$ for $|x| \leq 2\pi/3$ (Lemma 2.12) gives

$$\lambda_{z_*} = \frac{1}{n} (n - 2 + \omega + \omega^{-1}) = 1 - \frac{2}{n} \left(1 - \cos \left(\frac{2\pi}{n} \right) \right) \leq 1 - \frac{27}{n^3}.$$

Using $\cos(x) \geq 1 - x^2/2$, we see that the relaxation time of the sandpile chain is

$$t_{\text{rel}}(K_n) = \Theta(n^3).$$

Now the standard bounds relating relaxation and mixing times (Proposition 2.10) imply the order of $t_{\text{mix}}(K_n)$ is between n^3 and $n^4 \log n$. Our next goal is to prove Theorem 1.3, which says that the truth is in between: $t_{\text{mix}}(K_n)$ has order $n^3 \log n$, and moreover the sandpile chain on K_n exhibits cutoff at time $\frac{1}{4\pi^2} n^3 \log n$. The upper bound follows easily from Theorem 4.5, which gives:

Proposition 5.1. *Let P be the transition operator for the discrete time sandpile chain on K_n and U the stationary distribution. For any $c \geq 5/4$, if $t \geq \frac{1}{4\pi^2}n^3 \log n + cn^3$, then $\|P_{id}^t - U\|_{\text{TV}} \leq e^{-c}$.*

To obtain a matching lower bound, we turn to the eigenfunctions. The basic idea is to find a *distinguishing statistic* φ for which the distance between the pushforward measures $P_{id}^t \circ \varphi^{-1}$ and $U \circ \varphi^{-1}$ can be bounded from below using moment methods. Specifically, we appeal to Proposition 7.8 in [26], which shows that for any $\varphi : \mathcal{G} \rightarrow \mathbb{R}$,

$$\|P_{id}^t - U\|_{\text{TV}} \geq 1 - \frac{4}{4 + R(t)} \text{ whenever } R(t) \leq \frac{2 \left(\mathbb{E}_{P_{id}^t}[\varphi] - \mathbb{E}_U[\varphi] \right)^2}{\text{Var}_{P_{id}^t}(\varphi) + \text{Var}_U(\varphi)}. \quad (27)$$

Natural candidates for φ are eigenfunctions of P corresponding to a large real eigenvalue λ . The reason for using such eigenfunctions is that if X_t is distributed as P_{id}^t and Y has the uniform distribution, then $\mathbb{E}[\varphi(X_t)] = \lambda^t \varphi(id)$ and $\mathbb{E}[\varphi(Y)] = 0$, so their difference is relatively large when t is not too big. When λ has high multiplicity, we can average over a basis of eigenfunctions to reduce the variances of $\varphi(X_t)$ and $\varphi(Y)$.

Proposition 5.2. *Let P and U be as in Proposition 5.1. For any $c \geq 0$, we have*

$$\|P_{id}^t - U\|_{\text{TV}} \geq 1 - \frac{100}{100 + e^{4\pi^2 c}} \text{ for all } 0 \leq t \leq \frac{1}{4\pi^2}n^3 \log n - cn^3.$$

Taken together, Propositions 5.1 and 5.2 show that the discrete time sandpile chain on K_n exhibits cutoff at time $\frac{1}{4\pi^2}n^3 \log n$ with cutoff window of size $O(n^3)$. Also, Proposition 5.2 implies the lower bound in Theorem 1.3 since for all $c \geq 5/4$,

$$\frac{100}{100 + e^{4\pi^2 c}} \leq \frac{100}{e^{4\pi^2 c}} \leq e^{-35c}.$$

Proof of Proposition 5.2. First note that the statement is true when $n = 3$ since the only integer t in the allowed range is $t = 0$, where the inequality can be checked directly. Therefore we may assume that $n \geq 4$.

To describe our choice of φ , write $D_n = \{(j, k) \in [n-1]^2 : j \neq k\}$ and set $\omega = e^{2\pi i/n}$. For $(j, k) \in D_n$, define

$$h_{j,k}(v) = \begin{cases} \omega, & v = v_j \\ \omega^{-1}, & v = v_k \\ 1, & \text{else} \end{cases}$$

It is clear that $h_{j,k} \in \mathcal{H}$, so the function $f_{j,k} : \mathcal{G} \rightarrow \mathbb{T}$ given by

$$f_{j,k}(\eta) = \prod_{v \in V} h_{j,k}(v)^{\eta(v)} = \omega^{\eta(v_j) - \eta(v_k)}$$

is an eigenfunction for the sandpile chain with eigenvalue

$$\lambda_1 = \frac{1}{n} \sum_{v \in V} h(v) = 1 - \frac{1}{n} \left[2 - 2 \cos \left(\frac{2\pi}{n} \right) \right].$$

As λ_1 does not depend on (j, k) , the function

$$\varphi(\eta) := \frac{1}{(n-1)(n-2)} \sum_{(j,k) \in D_n} f_{j,k}(\eta)$$

is also in the λ_1 -eigenspace. Moreover, φ is \mathbb{R} -valued because $f_{j,k} = \overline{f_{k,j}}$.

To apply (27), we first observe that since $\varphi(id) = 1$ and $P\varphi = \lambda_1\varphi$,

$$\mathbb{E}_{P_{id}^t}[\varphi] = \lambda_1^t \text{ and } \mathbb{E}_U[\varphi] = 0.$$

(The latter expectation used the fact that U is a left eigenfunction with eigenvalue $1 \neq \lambda_1$, so $\mathbb{E}_U[\varphi] = \langle U, \varphi \rangle = 0$.) So the numerator of (27) is

$$2 \left(\mathbb{E}_{P_{id}^t}[\varphi] - \mathbb{E}_U[\varphi] \right)^2 = 2\lambda_1^{2t}.$$

Next we need an upper bound on the denominator $\text{Var}_{P_{id}^t}(\varphi) + \text{Var}_U(\varphi)$. Since \mathcal{H} is closed under pointwise multiplication, we have $h_{j,k}h_{l,m} \in \mathcal{H}$ for all pairs $(j, k), (l, m) \in D_n$. The corresponding eigenfunction is

$$f_{j,k,l,m}(\eta) := \prod_{v \in V} h_{j,k}(v)^{\eta(v)} h_{l,m}(v)^{\eta(v)} = f_{j,k}(\eta) f_{l,m}(\eta).$$

The associated eigenvalue depends on the 4-tuple (j, k, l, m) . For example, when $(j, k) = (l, m)$, the product $h_{j,k}h_{l,m}$ sends v_j to ω^2 and v_k to ω^{-2} and all other vertices to 1, giving an eigenvalue of $1 - \frac{1}{n} \left[2 - 2 \cos \left(\frac{4\pi}{n} \right) \right]$. If $(k, j) = (l, m)$, then $h_{j,k}h_{l,m} \equiv 1$. If j, k, l are distinct and $k = m$, then $h_{j,k}h_{l,m}$ sends v_j and v_l to ω and v_k to ω^{-2} and all other vertices to 1, so the resulting eigenvalue is $1 - \frac{1}{n} \left[3 - 2\omega - \omega^{-2} \right]$.

Writing

$$(n-1)^2(n-2)^2\varphi^2 = \sum_{(j,k),(l,m) \in D_n} f_{j,k,l,m}$$

as a linear combination of eigenfunctions, we can count the number of $f_{j,k,l,m}$ corresponding to each eigenvalue. This information is given in Table 1.

Eigenvalue	Multiplicity in $(n-1)^2(n-2)^2\varphi^2$
$\lambda_0 = 1$	$(n-1)(n-2)$
$\lambda_1 = 1 - \frac{1}{n} \left[2 - 2 \cos \left(\frac{2\pi}{n} \right) \right]$	$2(n-1)(n-2)(n-3)$
$\lambda_2 = 1 - \frac{1}{n} \left[2 - 2 \cos \left(\frac{4\pi}{n} \right) \right]$	$(n-1)(n-2)$
$\lambda_3 = 1 - \frac{1}{n} \left[4 - 4 \cos \left(\frac{2\pi}{n} \right) \right]$	$(n-1)(n-2)(n-3)(n-4)$
$\lambda_4 = 1 - \frac{1}{n} \left[3 - 2\omega - \omega^{-2} \right]$	$(n-1)(n-2)(n-3)$
$\overline{\lambda_4} = 1 - \frac{1}{n} \left[3 - 2\omega^{-1} - \omega^2 \right]$	$(n-1)(n-2)(n-3)$

TABLE 1. Expansion of $(n-1)^2(n-2)^2\varphi^2$ in the eigenbasis for P .

It follows that

$$\begin{aligned} \text{Var}_{P_{id}^t}(\varphi) + \text{Var}_U(\varphi) &= \mathbb{E}_{P_{id}^t}[\varphi^2] + \mathbb{E}_U[\varphi^2] - \mathbb{E}_{P_{id}^t}[\varphi]^2 - \mathbb{E}_U[\varphi]^2 \\ &= \frac{2}{(n-1)(n-2)} + \frac{2(n-3)}{(n-1)(n-2)}\lambda_1^t + \frac{1}{(n-1)(n-2)}\lambda_2^t \\ &\quad + \frac{(n-3)(n-4)}{(n-1)(n-2)}\lambda_3^t + \frac{2(n-3)}{(n-1)(n-2)}\text{Re}(\lambda_4^t) - \lambda_1^{2t}. \end{aligned}$$

To simplify, we note that $\lambda_3 \leq \lambda_1^2$. Thus

$$\frac{(n-3)(n-4)}{(n-1)(n-2)}\lambda_3^t - \lambda_1^{2t} \leq 0.$$

In addition, $|\lambda_4| \leq \lambda_1$, so $\text{Re}(\lambda_4^t) \leq |\lambda_4|^t \leq \lambda_1^t$. This follows from the computation

$$\lambda_1^2 - |\lambda_4|^2 = \frac{2}{n} \left(1 - \frac{4}{n}\right) \left[1 - \cos\left(\frac{4\pi}{n}\right)\right] + \frac{4}{n^2} \left[\cos\left(\frac{2\pi}{n}\right) - \cos\left(\frac{6\pi}{n}\right)\right],$$

where all of the terms are nonnegative since $n \geq 4$. Finally, $\lambda_2 < \lambda_1$. Therefore

$$\text{Var}_{P_{id}^t}(\varphi) + \text{Var}_U(\varphi) \leq \frac{2}{(n-1)(n-2)} + \frac{4n-11}{(n-1)(n-2)}\lambda_1^t \leq \frac{6}{n^2} + \frac{5}{n}\lambda_1^t,$$

again using $n \geq 4$ in the final inequality. The function

$$R(t) = \frac{2\lambda_1^{2t}}{\frac{6}{n^2} + \frac{5}{n}\lambda_1^t}$$

thus satisfies the condition in (27).

Next we find a lower bound on λ_1^t . Since $\cos(x) \geq 1 - x^2/2$, we have $\lambda_1 \geq 1 - 4\pi^2/n^3$. Therefore

$$t \log(\lambda_1) \geq \left(\frac{1}{4\pi^2}n^3 \log(n) - cn^3\right) \log\left(1 - \frac{4\pi^2}{n^3}\right).$$

Since $n \geq 4$, $4\pi^2/n^3 \in [0, \pi^2/16]$. Now we use that $\log(1-x) \geq -x - x^2$ for $0 \leq x \leq \pi^2/16$. This is because the function $f(x) = x + x^2 + \log(1-x)$ satisfies $f(0) = 0$, $f(\pi^2/16) > 0$, and $f'(x) = x(1-2x)/(1-x)$, implying that f is increasing on $[0, 1/2]$ and decreasing on $[1/2, \pi^2/16]$. This yields

$$\begin{aligned} t \log(\lambda_1) &\geq \left(\frac{1}{4\pi^2}n^3 \log(n) - cn^3\right) \left(-\frac{4\pi^2}{n^3} - \frac{16\pi^4}{n^6}\right) \\ &\geq -\log(n) + 4\pi^2c - \frac{4\pi^2}{n^3} \log(n). \end{aligned}$$

Thus

$$n\lambda_1^t = e^{\log(n)+t \log(\lambda_1)} \geq e^{-4\pi^2 \log(n)/n^3} e^{4\pi^2c} \geq \alpha e^{4\pi^2c},$$

where $\alpha = e^{-4\pi^2 \log(4)/4^3}$. It follows that

$$R(t) = \frac{2(n\lambda_1^t)^2}{6 + 5(n\lambda_1^t)} \geq \frac{2(\alpha e^{4\pi^2c})^2}{6 + 5(\alpha e^{4\pi^2c})} \geq \frac{2\alpha^2 e^{8\pi^2c}}{(6 + 5\alpha)e^{4\pi^2c}} = \left(\frac{2\alpha^2}{6 + 5\alpha}\right) e^{4\pi^2c},$$

where the first inequality used that the function $x \mapsto 2x^2/(6+5x)$ is increasing for $x \geq 0$. Since $2\alpha^2/(6+5\alpha) > 0.04$, we conclude from (27) that

$$\|P_{id}^t - U\|_{\text{TV}} \geq 1 - \frac{4}{4 + 0.04e^{4\pi^2 c}} = 1 - \frac{100}{100 + e^{4\pi^2 c}}. \quad \square$$

Complete bipartite graph. Suppose that $G = K_{m,n}$ is the complete bipartite graph having vertices $\{u_1, \dots, u_m, v_1, \dots, v_n\}$ and edges $\{\{u_j, v_k\}\}_{j \in [m], k \in [n]}$. In this case, there are essentially two possible choices for the sink, u_m or v_n . We will assume the former as the other case then follows by interchanging the u 's and v 's (or appealing to Proposition 2.5). Any $h \in \mathcal{H}$ must satisfy $h(u_m) = 1$ and thus $1 = h(u_m)^n = \prod_{k=1}^n h(v_k) = h(u_j)^n$ for all $j = 1, \dots, m$. If we let $h(u_j)$ be arbitrary n th roots of unity for $j = 1, \dots, m-1$, then, writing $\rho = \prod_{j=1}^m h(u_j)$, we must have $h(v_k)^m = \rho$ for each $k = 1, \dots, n$. Letting $h(v_1), \dots, h(v_{n-1})$ be arbitrary m th roots of ρ , $h(v_n)$ is then determined by the condition $1 = \prod_{k=1}^n h(v_k)$. Any such function is multiplicative harmonic by construction, and there are $n^{m-1}m^{n-1}$ such functions corresponding to the different choices of n th roots of 1 for $h(u_1), \dots, h(u_{m-1})$ and m th roots of ρ for $h(v_1), \dots, h(v_{n-1})$. Since this is the number of spanning trees in $K_{m,n}$ [25], we have found all possible choices for h .

When $3 \leq m \leq n$, we claim that the discrete time sandpile chain has relaxation time $t_{\text{rel}} = \Theta(n^3)$. (The $m = 2$ case will be discussed shortly.) To see that this is so, observe that $d_* = n$, so Theorem 2.11 implies that $\gamma_d \geq 8/[n^2(m+n)] \geq 4/n^3$.

Conversely, the function

$$h(v) = \begin{cases} e^{\frac{2\pi i}{n}}, & v = u_1 \\ e^{-\frac{2\pi i}{n}}, & v = u_2 \\ 1, & \text{else} \end{cases}$$

is multiplicative harmonic and the corresponding eigenvalue has modulus

$$|\lambda_h| = \frac{1}{m+n} \left[m+n-2 + 2 \cos\left(\frac{2\pi}{n}\right) \right] = 1 - \frac{2}{m+n} \left(1 - \cos\left(\frac{2\pi}{n}\right) \right),$$

hence

$$\gamma_d \leq \frac{2}{m+n} \left(1 - \cos\left(\frac{2\pi}{n}\right) \right) \leq \frac{4\pi^2}{(m+n)n^2} \leq \frac{4\pi^2}{n^3}.$$

Torus. Let $G = \mathbb{Z}_m \times \mathbb{Z}_m$ be the m -by- m discrete torus (square grid with periodic boundary). Since G is 4-regular, each eigenvalue β of the full Laplacian $\bar{\Delta}$ arises from an eigenvalue λ of the transition matrix K for simple random walk on G , with $\beta = 4(1-\lambda)$. Since K is a tensor product of transition matrices for random walks on cycles, the eigenvalues of K are $\lambda_{j,k} = \frac{1}{2} \left[\cos\left(\frac{2\pi j}{m}\right) + \cos\left(\frac{2\pi k}{m}\right) \right]$ for $j, k \in [m]$, so the matrix-tree theorem gives

$$|\mathcal{G}| = \frac{1}{m^2} \prod_{(j,k) \in [m]^2 \setminus (m,m)} \left(4 - 2 \cos\left(\frac{2\pi j}{m}\right) - 2 \cos\left(\frac{2\pi k}{m}\right) \right).$$

Interpreting the log of the product as a Riemann sum shows that

$$\frac{1}{m^2} \log(|\mathcal{G}(\mathbb{Z}_m \times \mathbb{Z}_m)|) \rightarrow \frac{4\beta(2)}{\pi} \approx 1.1662 \quad \text{as } m \rightarrow \infty$$

where $\beta(2)$ is the Catalan constant [18].

At this point, it is convenient to state the following simple bound for random walks on abelian groups.

Lemma 5.3. *Suppose that A is an abelian group of order N which is generated by a set $S = \{s_1, \dots, s_n\}$, let π be the uniform distribution on A , and let Q_a^t be the distribution of the random walk driven by the uniform measure on S after t steps started at a . For any initial state a and any $\varepsilon > 0$,*

$$\|Q_a^t - \pi\|_{\text{TV}} \geq 1 - \varepsilon \quad \text{whenever } t \leq \log_2(\varepsilon N) - n.$$

Proof. Since A is abelian, the random walk at time t must be at an element of the form $as_1^{x_1} \cdots s_n^{x_n}$ where $x_1, \dots, x_n \in \mathbb{N}$ with $x_1 + \cdots + x_n = t$. Letting S_t denote the set of all elements of this form, we have $|S_t| \leq \binom{t+n-1}{t} \leq 2^{t+n-1}$. Since Q_a^t is supported on S_t ,

$$\|Q_a^t - \pi\|_{\text{TV}} \geq |Q_a^t(S_t) - \pi(S_t)| \geq 1 - \frac{2^{t+n-1}}{N} \geq 1 - \varepsilon$$

whenever $t \leq \log_2(\varepsilon N) - n$. □

We can now estimate the mixing time for the sandpile chain on the torus.

Proposition 5.4. *Let P be the transition operator for the discrete time sandpile chain on $\mathbb{Z}_m \times \mathbb{Z}_m$ and U the stationary distribution.*

For any $c \geq 2$,

$$\|P_{id}^t - U\|_{\text{TV}} \leq e^{-\frac{2}{5}c} \quad \text{for all } t \geq \frac{5}{2}m^2 \log(m) + cm^2.$$

There exists $m_0 > 0$ such that if $m \geq m_0$, then for any $c \geq 0$,

$$\|P_{id}^t - U\|_{\text{TV}} \geq 1 - 2^{-c} \quad \text{for all } t \leq 0.68m^2 - c.$$

Proof. The upper bound follows from Theorem 4.3 and the lower bound from Lemma 5.3 along with the fact that $|\mathcal{G}(\mathbb{Z}_m \times \mathbb{Z}_m)| \geq e^{1.166m^2}$ for m sufficiently large. □

Note that the preceding estimate gives bounds which differ only by a log factor of the number of vertices, and it did not require computing any of the approximately $e^{1.166m^2}$ multiplicative harmonic functions.

Continuous time; Location of the sink. So far in this section, all our examples have been in discrete time. Recall from Section 3.1 that the moduli of the eigenvalues (and thus the L^2 mixing time) of the discrete time sandpile chain do not depend on the location of the sink. However, changing the sink does affect the *arguments* of the eigenvalues and thus the L^2 mixing time for the continuous time sandpile chain. The next two examples illustrate this dependence.

Example (Triangle with tail). Let $G = (V, E)$, where $V = \{u, v, w, w_1, \dots, w_m\}$ and

$$E = \{\{u, v\}, \{u, w\}, \{v, w\}, \{w, w_1\}, \{w_1, w_2\}, \dots, \{w_{m-1}, w_m\}\}.$$

This graph is a triangle with a long ‘tail’ attached and has $n = m + 3$ vertices. Any multiplicative harmonic function h on G must satisfy $h(w) = h(w_1) = \dots = h(w_m)$. This can be seen by starting from w_m and working backwards.

Let \mathcal{H}_u be the group of multiplicative harmonic functions on G with the sink placed at u . Since G has three spanning trees, $|\mathcal{H}_u| = 3$. For $z \in \{1, e^{2\pi i/3}, e^{4\pi i/3}\}$, we define $h_z \in \mathcal{H}_u$ by $h_z(u) = 1$, $h_z(w) = h_z(w_1) = \dots = h_z(w_m) = z$, and $h_z(v) = z^2$. If $z = 1$ then $h_1 \equiv 1$ and the associated eigenvalue is 1. If z is one of the other cube roots of unity, the eigenvalue is $\lambda_{h_z} = (1 - \frac{3}{n})z$. Therefore the spectral gaps of the discrete and continuous time sandpile chains have different orders: $\gamma_d = \frac{3}{n}$ while $\gamma_c = \frac{3}{2}(1 - \frac{1}{n}) = \frac{n-1}{2}\gamma_d$. The L^2 distances from stationarity are also different:

$$\begin{aligned} \|P_{id}^t - U\|_2^2 &= 2 \left(1 - \frac{3}{n}\right)^{2t}, \\ \|H_{id}^t - U\|_2^2 &\leq 2e^{-2t}. \end{aligned}$$

If the sink is placed at v , then the eigenvalues are the same as when the sink is at u . By contrast, if the sink is placed at w or any of the w_j , then the multiplicative harmonic functions $\{h'_z : z^3 = 1\}$ are given by $h'_z(w) = h'_z(w_1) = \dots = h'_z(w_m) = 1$, $h'_z(v) = z$, and $h'_z(u) = z^2$. The associated eigenvalues are 1 and $1 - \frac{3}{n}$ with multiplicity 2. Hence $\gamma_c = \gamma_d = \frac{3}{n}$.

For intuition on this example, suppose η is a recurrent chip configuration on G . We take one step in the discrete time chain by adding a chip at a uniformly chosen vertex and then stabilizing. It can be shown that adding a chip at any of w, w_1, \dots, w_m leads to the same configuration η' , so $P(\eta, \eta') = \frac{n-2}{n}$. When the sink is at u or v , the transition matrix can be written as

$$P^{(u)} = P^{(v)} = \frac{1}{n} \begin{bmatrix} 1 & n-2 & 1 \\ 1 & 1 & n-2 \\ n-2 & 1 & 1 \end{bmatrix},$$

which is nearly periodic. When the sink is at w or one of the w_j , the transition matrix is

$$P^{(w)} = P^{(w_j)} = \frac{1}{n} \begin{bmatrix} n-2 & 1 & 1 \\ 1 & n-2 & 1 \\ 1 & 1 & n-2 \end{bmatrix}.$$

Both discrete time chains converge to stationarity at the same rate, but the continuous time chain associated with $P^{(u)} = P^{(v)}$ converges much faster than the continuous time chain associated with $P^{(w)} = P^{(w_j)}$.

Example (Unbalanced complete bipartite graph). Let $G = K_{2,m}$ be the complete bipartite graph where one part has two vertices. Here $V = \{u_1, u_2, v_1, \dots, v_m\}$ and $E = \{\{u_i, v_j\} : 1 \leq i \leq 2, 1 \leq j \leq m\}$. A full characterization of all the multiplicative harmonic functions was given previously.

When the sink is placed at u_1 , the multiplicative harmonic function $h \in \mathcal{H}_{u_1}$ given by $h(u_1) = 1$, $h(u_2) = e^{2\pi i(2/m)}$, and $h(v_j) = e^{2\pi i/m}$ for all j maximizes both $|\lambda_h|$ and $\operatorname{Re}(\lambda_h)$ over all $h \neq 1$. Writing $n = m + 2$, we compute

$$\begin{aligned}\gamma_d &= 1 - |\lambda_h| = \frac{2}{n} \left[1 - \cos \left(\frac{2\pi}{n-2} \right) \right] \approx \frac{4\pi^2}{n^3}, \\ \gamma_c &= 1 - \operatorname{Re}(\lambda_h) = \left[1 + \frac{2}{n} \cos \left(\frac{2\pi}{n-2} \right) \right] \left[1 - \cos \left(\frac{2\pi}{n-2} \right) \right] \approx \frac{2\pi^2}{n^2}, \\ \frac{\gamma_c}{\gamma_d} &= \frac{n}{2} + \cos \left(\frac{2\pi}{n-2} \right) \approx \frac{n}{2}.\end{aligned}$$

If instead the sink is placed at v_1 , then the function $h' \in \mathcal{H}_{v_1}$ given by $h'(u_1) = e^{-2\pi i/m}$, $h'(u_2) = e^{2\pi i/m}$, and $h'(v_j) = 1$ for all j maximizes both $|\lambda_{h'}|$ and $\operatorname{Re}(\lambda_{h'})$ over all $h' \neq 1$. In this case, $\lambda_{h'}$ is a positive real number and

$$\gamma_d = \gamma_c = \frac{2}{n} \left[1 - \cos \left(\frac{2\pi}{n-2} \right) \right] \approx \frac{4\pi^2}{n^3}.$$

Rooted sums. Let $G_1 = (V_1, E_1), \dots, G_k = (V_k, E_k)$ be simple connected graphs with distinguished vertices s_1, \dots, s_k , and let $G = (V, E)$ be any simple connected graph with $V = \{v_1, \dots, v_k\}$. Construct a new graph G' from the G_j 's by identifying each $s_j \in V_j$ with $v_j \in V$. That is, $G' = (V', E')$ with $V' = \bigcup_{j=1}^k V_j$ and $E' = \left(\bigcup_{j=1}^k E_k \right) \cup \{ \{s_i, s_j\} : \{v_i, v_j\} \in E \}$. Let \mathcal{H}_j denote the group of multiplicative harmonic functions on G_j with sink vertex s_j , \mathcal{H} the multiplicative harmonic functions on G with sink vertex v_k , and \mathcal{H}' those on G' with sink s_k .

Define the map $\psi : \mathcal{H} \times \mathcal{H}_1 \times \dots \times \mathcal{H}_k \rightarrow \mathcal{H}'$ by $\psi(h, h_1, \dots, h_k) = h'$ where $h' : V' \rightarrow \mathbb{T}$ is given by $h'(v) = h(v_j)h_j(v)$ for $v \in V_j$. One easily checks that ψ is an isomorphism: It is a well-defined injective homomorphism by routine verification, and it is surjective because every spanning tree in G' is formed by choosing a spanning tree from each of G, G_1, \dots, G_k . Writing $n_j = |V_j|$, $n = |V'| = \sum_{j=1}^k n_j$, we see that the eigenvalue for the discrete time sandpile chain on G' associated with $h' = \psi(h, h_1, \dots, h_k)$ is

$$\lambda_{h'} = \frac{1}{n} \sum_{v \in V'} h'(v) = \frac{1}{n} \sum_{j=1}^k h(v_j) \sum_{w \in V_j} h_j(w) = \frac{1}{n} \sum_{j=1}^k n_j h(v_j) \lambda_{h_j}.$$

Is a large spectral gap compatible with a large sandpile group? Our last example is a graph whose sandpile chain mixes relatively quickly. It is an instance of the rooted sum construction above: Let $G = P_k$ be a path on k vertices, and for $j \in [k]$ let G_j be a cycle on m_j vertices, with $2 < m_1 \leq \dots \leq m_k$. Since $\mathcal{H}(P_k)$ is trivial, the multiplicative harmonic functions on the rooted sum G' are given by $h'(v) = h_j(v)$ for $v \in G_j$ with $h_j \in \mathcal{H}(G_j)$. The sandpile group of G' is isomorphic to $(\mathbb{Z}/m_1\mathbb{Z}) \times \dots \times (\mathbb{Z}/m_k\mathbb{Z})$, and the largest nontrivial eigenvalue of the sandpile chain is

$$\lambda_* = 1 - \frac{m_1}{m_1 + \dots + m_k},$$

corresponding to choosing any nontrivial multiplicative harmonic function h_1 on the smallest cycle G_1 , and $h_2 \equiv \cdots \equiv h_k \equiv 1$.

When $m_1 = \cdots = m_k = m$ the rooted sum G' has mk vertices, sandpile group of order m^k , and sandpile spectral gap $\gamma_c = \gamma_d = 1/k$. Fixing k gives a sequence of graphs, indexed by m , with gap uniformly bounded away from 0.

The same construction with $k = n^\alpha$ and $m = n^{1-\alpha}$ for fixed $0 < \alpha < 1$ gives a graph with n vertices, sandpile group of order $e^{(1-\alpha)n^\alpha \log(n)}$, and sandpile spectral gap $1/n^\alpha \gg 1/d_*^2 n$. We conclude with two questions.

1. Does there exist a graph sequence G_n with $\gamma_d(G_n) > c > 0$ such that the size of the sandpile group $|\mathcal{G}(G_n)|$ grows faster than any power of n ?
2. Does there exist a graph sequence G_n with $\log |\mathcal{G}(G_n)| = \Omega(n \log d_*)$ such that $\gamma_d(G_n) \gg 1/d_*^2 n$?

ACKNOWLEDGMENTS

We thank Spencer Backman, Shirshendu Ganguly, Alexander Holroyd, Yuval Peres and Farbod Shokrieh for inspiring conversations.

REFERENCES

- [1] Omid Amini and Madhusudan Manjunath. Riemann-Roch for sub-lattices of the root lattice A_n . *Electronic J. Combin.*, 17(1):R124, 2010. [arXiv:1007.2454](#).
- [2] László Babai and Evelin Toumpakari. A structure theory of the sandpile monoid for directed graphs. 2010.
- [3] Roland Bacher, Pierre de la Harpe, and Tatiana Nagnibeda. The lattice of integral flows and the lattice of integral cuts on a finite graph. *Bull. Soc. Math. France*, 125(2):167–198, 1997. [MR1478029](#).
- [4] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Phys. Rev. A (3)*, 38(1):364–374, 1988.
- [5] Matthew Baker and Serguei Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215(2):766–788, 2007. [MR2355607](#).
- [6] Matthew Baker and Farbod Shokrieh. Chip-firing games, potential theory on graphs, and spanning trees. *J. Combin. Theory Ser. A*, 120(1):164–182, 2013. [MR2971705](#).
- [7] N. L. Biggs. Chip-firing and the critical group of a graph. *J. Algebraic Combin.*, 9(1):25–45, 1999. [MR1676732](#).
- [8] Anders Björner, László Lovász, and Peter W. Shor. Chip-firing games on graphs. *European J. Combin.*, 12(4):283–291, 1991. [MR1120415](#).
- [9] Shu-Chiuan Chang, Lung-Chi Chen, and Wei-Shih Yang. Spanning trees on the Sierpinski gasket. *J. Stat. Phys.*, 126(3):649–667, 2007. [MR2294471](#).
- [10] Fan Chung and Mary Radcliffe. On the spectra of general random graphs. *Electron. J. Combin.*, 18(1):Paper 215, 14, 2011. [MR2853072](#).
- [11] Deepak Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990. [MR1044086](#).

- [12] Deepak Dhar, Philippe Ruelle, Siddhartha Sen, and DN Verma. Algebraic aspects of abelian sandpile models. *Journal of Physics A: Mathematical and General*, 28(4):805–831, 1995. [arXiv:cond-mat/9408020](#).
- [13] Persi Diaconis. *Group representations in probability and statistics*. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 11. Institute of Mathematical Statistics, Hayward, CA, 1988. [MR964069](#).
- [14] Matthew Farrell and Lionel Levine. Coeulerian graphs. *Proceedings of the American Mathematical Society*, 2015, to appear. [arXiv:1502.04690](#).
- [15] Anne Fey, Lionel Levine, and David B. Wilson. Driving sandpiles to criticality and beyond. *Phys. Rev. Lett.*, 104:145703, Apr 2010. [arXiv:0912.3206](#).
- [16] Joel Friedman. A proof of Alon’s second eigenvalue conjecture and related problems. *Mem. Amer. Math. Soc.*, 195(910):viii+100, 2008. [MR2437174](#).
- [17] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 197–206. ACM, 2008.
- [18] M. L. Glasser and F. Y. Wu. On the entropy of spanning trees on a large triangular lattice. *Ramanujan J.*, 10(2):205–214, 2005. [MR2194523](#).
- [19] Chris Godsil and Gordon Royle. *Algebraic graph theory*, volume 207 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2001. [MR1829620](#).
- [20] Felix Goldberg. Chip-firing may be much faster than you think. *ArXiv e-prints*, November 2014. [arXiv:1411.1652](#).
- [21] Alexander E. Holroyd, Lionel Levine, Karola Mészáros, Yuval Peres, James Propp, and David B. Wilson. Chip-firing and rotor-routing on directed graphs. In *In and out of equilibrium. 2*, volume 60 of *Progr. Probab.*, pages 331–364. Birkhäuser, Basel, 2008. [MR2477390](#).
- [22] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561 (electronic), 2006. [MR2247919](#).
- [23] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, 1990. [MR1084815](#).
- [24] Hang-Hyun Jo and Hyeong-Chai Jeong. Comment on “driving sandpiles to criticality and beyond”. *Phys. Rev. Lett.*, 105:019601, Jun 2010.
- [25] Dieter Jungnickel. *Graphs, networks and algorithms*, volume 5 of *Algorithms and Computation in Mathematics*. Springer, Heidelberg, fourth edition, 2013. [MR2986014](#).
- [26] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. [MR2466937](#).
- [27] Lionel Levine. Threshold state and a conjecture of poghosyan, poghosyan, priezzhev and ruelle. *Comm. Math. Phys.*, 335(2):1003–1017, 2015. [arXiv:1402.3283](#).
- [28] Alexander Lubotzky. Expander graphs in pure and applied mathematics. *Bull. Amer. Math. Soc. (N.S.)*, 49(1):113–162, 2012. [MR2869010](#).
- [29] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302 (electronic),

2007. [MR2306293](#).
- [30] Chris Peikert. Limits on the hardness of lattice problems in l_p norms. *Comput. Complexity*, 17(2):300–351, 2008. [MR2417596](#).
- [31] David Perkinson, Jacob Perlman, and John Wilmes. Primer for the algebraic geometry of sandpiles. *Tropical and non-Archimedean geometry*, pages 211–256, 2011. [arXiv:1112.6163](#).
- [32] Su. S. Poghosyan, V. S. Poghosyan, V. B. Priezzhev, and P. Ruelle. Numerical study of the correspondence between the dissipative and fixed-energy abelian sandpile models. *Phys. Rev. E*, 84:066119, Dec 2011. [arXiv:1104.3548](#).
- [33] Laurent Saloff-Coste. Random walks on finite groups. In *Probability on discrete structures*, volume 110 of *Encyclopaedia Math. Sci.*, pages 263–346. Springer, Berlin, 2004. [MR2023654](#).
- [34] Andrey Sokolov, Andrew Melatos, Tien Kieu, and Rachel Webster. Memory on multiple time-scales in an abelian sandpile. *Physica A: Statistical Mechanics and its Applications*, 428:295–301, 2015.
- [35] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. [MR1676282](#).
- [36] Wei Wei, Chengliang Tian, and Xiaoyun Wang. New transference theorems on lattices possessing n^ϵ -unique shortest vectors. *Discrete Math.*, 315:144–155, 2014. [MR3130366](#).

DANIEL C. JERISON, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853.

E-mail address: jerison@math.cornell.edu

LIONEL LEVINE, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853.

URL: <http://www.math.cornell.edu/~levine>

JOHN PIKE, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853.

E-mail address: jpike@cornell.edu

URL: <http://www.math.cornell.edu/~jp999>