AUSTRALASIAN JOURNAL OF COMBINATORICS
Volume **68(3)** (2017), Pages 330–345

A chip-firing variation and a Markov chain with uniform stationary distribution

Dave Perkins*

Mathematics Department
Hamilton College
Clinton, NY 13323
U.S.A.
dperkins@hamilton.edu

P. Mark Kayll†

Department of Mathematical Sciences
University of Montana
Missoula, MT 59812
U.S.A.
mark.kayll@umontana.edu

Abstract

We continue our study of burn-off chip-firing games on graphs initiated in [*Discrete Math. Theor. Comput. Sci.* **15** (2013), no. 1, 121–132]. Here we introduce randomness by choosing each successive seed uniformly from among all possible nodes. The resulting stochastic process—a Markov chain $(X_n)_{n\geq 0}$ with state space the set $\mathcal{R}$ of relaxed legal chip configurations $C\colon V \to \mathbb{N}$ on a connected graph $G=(V,E)$—has the property that, with high probability, each state appears equiproportionally in a long sequence of burn-off games. This follows from our main result that (X_n) has a doubly stochastic transition matrix. As tools supporting our main proofs, we establish several properties of the chip addition operator $E_{(\cdot)}(\cdot)\colon V\times\mathcal{R}\to\mathcal{R}$ that may be of independent interest. For example, if $V=\{v_1,\ldots,v_n\}$, then $C\in\mathcal{R}$ if and only if C is fixed by the composition $E_{v_1}\circ\cdots\circ E_{v_n}$; this property eventually yields the irreducibility of (X_n).

* Part of this work appears in the author's PhD dissertation [29]

† This work was partially supported by a grant from the Simons Foundation (#279367 to Mark Kayll).

1 Introduction

This article continues our study in [23] of a variant of chip-firing games—burn-off games—in which each iteration simulates the loss of energy from a complex system. Our point of departure is to introduce randomness to burn-off games by choosing each successive 'seed' uniformly at random from among all possible nodes. The resulting stochastic process turns out to have the property that, with high probability, each 'relaxed legal configuration' appears equiproportionally in a long sequence of burn-off games. This follows from our main result (Theorem 3.2) that the transition matrix for the underlying Markov chain is doubly stochastic. The uniform visitation of states (configurations) is at once expected (as the seed choices are uniform) and surprising (as it does not depend on the underlying graph).

We are not the first authors to bring a stochastic component to chip-firing games; see, e.g., [33], which studies a Markov chain different from the one in the present article. However, as far as we can tell, when we introduced (in [29]) the chain studied here, it was novel. By now, it is an example of the Markov chain on the so-called 'recurrent states' of a more abstract process called 'abelian networks', examined, e.g., in [10, Section 3.2]. A similar continuous process was studied in [21], where, interestingly, the uniform measure on the 'allowed configurations' also made an appearance. We also point the reader to [31], which—though it makes no explicit mention of chip firing—we cannot help feeling ought to be related to the present work through the shared thread of spanning tree enumeration; cf. [2–5, 20, 23].

The remainder of this introduction is an abbreviated version of the one in [23], from which we borrow heavily.

Chip-firing games attracted attention in part thanks to their ability to mimic the behaviour of complex real-world systems of individuals that interact primarily with their neighbours. For example, these games have been used to predict the frequencies of earthquakes [1], the volume of avalanches in sandpiles [14], the extent of traffic jams [28], and the intensity of bursts of brain activity [34]. (Some recent, more theoretical, work on these games appears, e.g., in [2, 3, 18–20, 26, 27, 30]; [25] provides a wonderfully succinct overview.) Because such behaviours typically involve the loss of energy (heat loss via friction, for example), we introduced in [29] a variant in which a chip is lost from the game each time a node fires. We called such variants 'burn-off games' (and will carefully define them in Section 1.1).

The published account of [29] started with [22] and continued in [23], the latter providing a foundation for the present article. In [23], following a characterization of the 'relaxed legal' configurations for general (connected) graphs and an enumeration of the 'legal' ones for complete graphs, we re-established the by now well-known connection between chip-firing games and spanning tree enumeration, specifically, a link between our burn-off games and Cayley's tree enumeration theorem [12] (or, e.g., [11]). Other articles exposing this connection are (nonexhaustively) listed at the end of the second paragraph above.

The present paper begins (§1.1) with the definition of a burn-off game (and the supporting terminology) and a handful of results imported from our earlier article [23] and from the survey [20] by Holroyd et al. Section 1.2 also contains one new (to us) result—Proposition 1.9—that greatly simplified two of our original proofs in Section 2. There, we establish three essential properties of the 'chip addition operator'—defined in (1.3) below—which sets the stage for Section 3, where we draw our main conclusions about the Markov chain underlying a sequence of burn-off games.

Notation and terminology

We mainly follow usual graph theory conventions as found, for instance, in [11]. A graph theory reference that addresses chip firing is [17]. For probability background, see the classic [16].

1.1 Description of a burn-off game

A *chip-firing game* plays on a connected graph $G = (V, E)$ with no loops or parallel edges. We initiate a game by distributing 'chips' onto the nodes; i.e., each $v \in V$ begins with a nonnegative number $C(v)$ of chips, and we call the function $C\colon V \to \mathbb{N}$ a *configuration*.

To each $v \in V$ we assign a *critical number*. Early chip-firing games (e.g. [4–7]) typically set the critical number of a node equal to its degree, but our variant tweaks this, as we explain in the next paragraph. Should $C(v)$ exceed the critical number of v, we say that v can *fire*. A configuration in which no node can fire is *relaxed*, and to start a chip-firing game, we add a chip to some selected node (called a *seed*) in a relaxed configuration. This may allow the seed to fire. A node that fires sends a chip to each of its neighbours, which may allow a new node to fire. The game proceeds in this way until the resulting configuration is relaxed. The *length* of a game equals the number of node firings in passing from an unrelaxed configuration to a relaxed one.

These rules may trigger games of infinite length, as Björner et al. [7] observed. To eliminate this possibility, we modified the rules: in our *burn-off game* variant, the critical number of each $v \in V$ is $\deg_G(v) + 1$, and when v fires, one of its chips is lost from the game. Formally, when v *fires*, C is modified to a *successor* configuration C' such that

$$C'(u) = \begin{cases} C(v) - \deg_G(v) - 1 & \text{if } u = v, \\ C(u) + 1 & \text{if } uv \in E(G), \\ C(u) & \text{if } v \neq u \not\sim v. \end{cases} \tag{1.1}$$

In [23], we noted that chip firing according to this rule is equivalent to the 'dollar game' of Biggs [5] in the particular case when his 'government' node is adjacent to every other node in the underlying graph; see also [4]. More recently, we noticed that burn-off chip firing on (undirected) graphs is a special case of conventional chip firing on directed graphs (initiated in [6]). Given one of our graphs G, replace each edge by a pair of oppositely oriented arcs, and add a new node s, to which every node of G is joined by an arc headed to s. In the language of the recent survey [20], the resulting digraph is 'bidirected' with s as its 'global sink' and so falls within the class of digraphs considered there and elsewhere. A still more general (and abstract) process—abelian networks—is introduced and studied in the important sequence of articles [8–10].

In a graph G with configuration C, we say that a node v is *critical* when $C(v) = \deg_G(v)$, i.e., when v will fire with the addition of a single chip. It is sometimes convenient to denote such an addition algebraically, so we write $\mathbf{1}_v$ for the configuration with one chip on v and zero chips on every other node. The action of passing from C to $C + \mathbf{1}_v$ is called *seeding* C at v or sometimes just *seeding* v; we do not require v to be critical before doing this. A *supercritical* node v is one for which $C(v) > \deg_G(v)$, i.e., one that must fire before a game in this state can end. During a burn-off game, it is possible for any of a number of nodes to become supercritical, but the decision of which of these to fire first has no bearing on the final relaxed configuration. This property of chip firing, first established in [13], was also proved by several other people, including this article's first author; see, e.g., [7], [15], or [29]. A formal statement of this result appears as Lemma 1.4 below.

We arrive at the notion of a 'legal' configuration by typifying those that are encountered in a long game sequence. Let us start by calling a configuration *supercritical* if every node is supercritical. As in [23], we follow [1] and focus on the configurations that can result from relaxing supercritical ones. To understand these, first consider what happens when a burn-off game is played in reverse. Looking at (1.1), we see that to *reverse-fire* a node v (each of whose neighbours u necessarily satisfies $C'(u) \geq 1$) means to modify C' to a configuration C such that

$$C(u) = \begin{cases} C'(v) + \deg_G(v) + 1 & \text{if } u = v, \\ C'(u) - 1 & \text{if } uv \in E(G), \\ C'(u) & \text{if } v \neq u \not\sim v. \end{cases}$$

The requirement that $C'(u) \geq 1$ for neighbours u of v ensures that C is a configuration. Now a configuration C is *legal* if there exists a reverse-firing sequence starting with C and ending with a supercritical configuration.

1.2 Preliminaries

Legal configurations

Our first three tools concern legal configurations. They were proved originally in [29], followed by published proofs in [23]. The first characterizes the relaxed legal configurations on general graphs G. In its statement, N_G denotes the 'earlier neighbour' set; i.e., given an ordering $(w_1, \ldots, w_n)$ of $V(G)$, define $N_G(w_i) := \{w_j \colon w_i w_j \in E(G) \text{ and } j < i\}$.

Lemma 1.1 ([29]) *A relaxed configuration $C \colon V \to \mathbb{N}$ is legal if and only if it is possible to label V as $w_1, \ldots, w_n$ so that*

$$C(w_i) \geq |N_G(w_i)| \text{ for } 1 \leq i \leq n. \tag{1.2}$$

We next restate our algorithm for determining the legality of a given configuration.

Algorithm 1.2 ([29])

INPUT:	a graph $G = (V, E)$ and a chip configuration $C\colon V \to \mathbb{N}$ on G
OUTPUT:	an answer to the question 'Is C legal?'
(1)	Let $G^* = G$.
(2)	If $C(v) < \deg_{G^*}(v)$ for all $v \in V(G^*)$, then stop. Output 'No.'
(3)	Choose any $v \in V(G^*)$ with $C(v) \geq \deg_{G^*}(v)$.
(4)	Delete v and all incident edges from G^* to create a graph G^-.
(5)	If $V(G^-) = \varnothing$, then stop. Output 'Yes.'
(6)	Let $G^* = G^-$ and go to step 2.

The following basic result establishes that containing a legal configuration is an inherited property for graphs.

Lemma 1.3 ([29]) *For a configuration $C\colon V(G) \to \mathbb{N}$ and a subgraph H of G, if C is legal on G, then $C|_{V(H)}$ is legal on H.*

Firing order independence and the abelian property

The next result actually holds in the more general setting of chip firing on digraphs, where the existence of a global sink is not even required. For our purposes, we need only this version for burn-off games on connected graphs G (but connectivity is not needed for its proof).

Lemma 1.4 ([13, 15]) *If $C_0, C_1, \ldots, C_\ell$ is a sequence of configurations on G, each a successor of the one before and $C'_0, C'_1, \ldots, C'_k$ is another such sequence such that $C'_0 = C_0$ and both of C_ℓ, C'_k are relaxed, then $k = \ell$, each node fires the same number of times in both sequences, and $C'_\ell = C_\ell$.*

Thus, in a burn-off game, the nodes can be fired in any order without affecting the length or final configuration of the game. As noted in Section 1.1, Lemma 1.4 has appeared in several works; [20] contains a particularly succinct proof.

In concert with Corollary 1.10 below, the result implies that starting a burn-off game in a configuration C leads unambiguously to a unique relaxed configuration, which we call the *relaxation* of C and denote by C° (the notation here and in (1.3) borrowed from [20]). Now the seed/relax action of a burn-off game can be written in terms of the *chip addition operator* E_v, defined as the map on configurations C that seeds a node v and then allows the system to relax; i.e.,

$$E_v(C) := (C + \mathbf{1}_v)^\circ. \tag{1.3}$$

It turns out that in burn-off games (and also some more general chip-firing settings), the chip addition operators commute. This 'abelian property' gave rise to the name 'abelian sandpile model' for a system introduced by Dhar [13] equivalent to chip-firing games. The abelian property admits the following strengthening, a consequence of which we use in the proof of Proposition 1.9.

Proposition 1.5 ([20]) *Applying a sequence of chip addition operators to a configuration yields the same result as adding all the associated chips simultaneously, then relaxing.*

Given two configurations C, C', we call C *reachable* from C' if there exists a configuration D such that $C = (C' + D)^{\circ}$. Notice that in this case, C is necessarily relaxed. Proposition 1.5 shows that the reachability of C from C' can also be viewed as the possibility of passing from C' to C via a sequence of individual chip addition operators, i.e., via a sequence of seeding nodes and firing supercritical ones. A moment's reflection shows that we can strengthen Proposition 1.5 still further.

Corollary 1.6 *Applying a sequence of chip addition operators to a configuration yields the same result as any sequence of seeding all the associated nodes and firing all the eventual supercritical ones, including sequences interspersing the seeding and firing operations.*

Recurrence

A relaxed configuration is *recurrent* if it is reachable from every configuration. As noted by Holroyd et al. [20], several notions of 'recurrent' appear in the chip-firing/abelian sandpile literature, and their Lemma 2.17 establishes the equivalence of the various definitions for digraphs with a global sink. Phrased in our language, the slice of their lemma needed in Section 3 is

Lemma 1.7 ([20]) *In a burn-off game on a connected graph G, a configuration C is recurrent if and only if G contains a node v such that C is reachable from $E_v(C)$.*

For the remainder of this article, $\mathcal{R}$ denotes the set of all relaxed legal configurations on a given connected graph. Using Lemma 1.7, we shall establish in Proposition 3.1 that configurations $C \in \mathcal{R}$ are recurrent; and indeed, the converse is also true. Though we do not make use of this equivalence, we note it here for completeness.

Proposition 1.8 *In a burn-off game on a connected graph G, a configuration C is recurrent if and only if $C \in \mathcal{R}$.*

As we just noted, the sufficiency in Proposition 1.8 is addressed by Proposition 3.1 below; see [24, Proposition 1.1] for a proof of the necessity.

In our proofs of Propositions 2.2 and 2.3, we find one further characterization of recurrence to be particularly fruitful. For a graph G with subgraph F and a node $x \in V(G)$, we denote by $\Gamma_F(x)$ the set of neighbours of x lying in $V(F)$.

Proposition 1.9 *In a burn-off game on a connected graph with $V = \{v_1, \ldots, v_n\}$, a configuration C lies in $\mathcal{R}$ if and only if*

$$E_{v_n} \circ E_{v_{n-1}} \circ \cdots \circ E_{v_2} \circ E_{v_1}(C) = C. \tag{1.4}$$

In a disguised form, Proposition 1.9 first appeared in [5, Lemma 3.6]; we include our own proof here for completeness.

Proof. To see the sufficiency of the condition (1.4) for $C \in \mathcal{R}$, first note that $E_{v_n} \circ \cdots \circ E_{v_1}(C)$ is, by definition, relaxed. To verify that C is also legal, apply Corollary 1.6 to obtain a deletion sequence—per Algorithm 1.2—confirming that legality. We omit the details and turn to the necessity of (1.4) for $C \in \mathcal{R}$.

It will suffice to show that when $E_{v_n} \circ \cdots \circ E_{v_1}$ is applied to C, every node $x \in V$ fires exactly once. For then, the final chip count on x will be its initial chip count $C(x)$, plus the $\deg_G(x)$-many chips received from its neighbours, plus one for the seed chip associated with E_x (in the sequence $E_{v_n} \circ \cdots \circ E_{v_1}$), and less the $(\deg_G(x) + 1)$-many chips lost when x itself fires. That is, x will contain

$$C(x) + \deg_G(x) + 1 - (\deg_G(x) + 1) = C(x)$$

chips under $E_{v_n} \circ \cdots \circ E_{v_1}(C)$, which shows that $E_{v_n} \circ \cdots \circ E_{v_1}$ fixes C. So we shall complete the proof by establishing the following two claims.

Claim 1. When $E_{v_n} \circ \cdots \circ E_{v_1}$ is applied to C, every $x \in V$ fires at least once.

Claim 2. When $E_{v_n} \circ \cdots \circ E_{v_1}$ is applied to C, every $x \in V$ fires at most once.

Proof of Claim 1. By Corollary 1.6, we may reorder the sequence in

$$E_{v_n} \circ \cdots \circ E_{v_1} \tag{1.5}$$

without changing its function. Without loss of generality, let us suppose that it is already in the order of a deletion sequence establishing the legality of C, per Lemma 1.1. For an $x \in V$, we consider the moment when it is seeded in (1.5); say $x = v_i$ and $v_1, \ldots, v_{i-1}$ were seeded prior to x. Before any further relaxing, let $F = \{v_1, \ldots, v_{i-1}\}$ and $H = \{v_i, v_{i+1}, \ldots, v_n\}$. We may now see by induction on i that each node in $F \cup \{x\}$ can be fired as soon as it is seeded in the sequence (1.5). First note that because of our arranged seeding order and the legality of C (cf. (1.2)), we have

$$C(v_i) \geq \deg_H(v_i) \tag{1.6}$$

for $i = 1, 2, \ldots, n$. Now when $i = 1$, so that $F = \varnothing$, the relation (1.6) gives $C(x) \geq \deg_G(x)$, which implies that x can fire if it is the first seeded node; this establishes the basis for our induction. Now fix $i > 1$, and assume that each node in F can be fired at the moment it is seeded. For convenience, we assume that these nodes do indeed fire when they can. By this hypothesis, aside from the chips on $x = v_i$ guaranteed by (1.6), x will also have gained (at least) one chip from each of its neighbours in F. Thus its chip count at the moment of seeding is at least $\deg_H(x) + |\Gamma_F(x)| = \deg_G(x)$, and so seeding x makes it supercritical. This completes the induction and hence the proof of Claim 1. □

Proof of Claim 2. By contradiction. If the claim is false, we may consider a node $w \in V$ as the first one to fire a second time. Let there be c chips on w at the outset (prior to the global seeding of (1.5)). As C is relaxed, we know that

$$c = C(w) \leq \deg_G(w). \tag{1.7}$$

Let k denote the number of neighbours of w that have fired once before w fires for the first time. Just before w fires for the first time, it contains $c + k + 1 > \deg_G(w)$ chips (accounting for the seed chip). For it to fire again, w must regain the threshold and so must gain at least

$$t := \deg_G(w) + 1 - (c + k + 1 - (\deg_G(w) + 1))$$

chips from its neighbours. But

$$t = (\deg_G(w) - k) + (\deg_G(w) - c) + 1 > \deg_G(w) - k,$$

by (1.7). Now $\deg_G(w) - k$ is precisely the number of neighbours of w that did not fire before w fired for the first time. So before w can gain $t > \deg_G(w) - k$ chips, one of its neighbours must fire for a second time. This contradicts our assumption that w is the first node to fire a second time. □

As we noted before the statements of the claims, their veracity suffices to establish Proposition 1.9. □

The paper [35] contains an idea similar to the one in our proof of Claim 2. As burn-off games seed one node at a time, Claim 2 leads to the following

Corollary 1.10 ([29]) *During a burn-off game that starts with a relaxed configuration, no node fires more than once.*

Corollary 1.10 formalizes our remark before (1.1) concerning the finiteness of burn-off games. In addition, it shows that for a given graph $G = (V, E)$, the length ℓ of a burn-off game on G satisfies $0 \leq \ell \leq |V|$. In our main application of this result—in the proof of Proposition 2.1—the configurations are legal as well as relaxed, but it is worth noting that legality is not needed to establish the corollary. We close this section with a lemma taking aim at the distribution of chips on the ℓ nodes that fire in a game of length two or more. It leans on Lemma 1.4 and Corollary 1.10 and is invoked in the proof of Proposition 2.1.

Lemma 1.11 ([29]) *Let C be a relaxed configuration on a connected graph G. Consider a burn-off game of length at least two that results in the configuration C'. Let H be the subgraph induced on the nodes that fire during the game. After relaxation, only the seed v will be critical in H.*

Proof. The seed v must fire because the game length is at least two. When v fires, it loses all its chips. Each of its neighbours in H fires, sending a chip back to v. By Corollary 1.10, none of these nodes may fire a second time. Thus, at the end of the game, $C'(v) = \deg_H(v)$; that is, v is critical in H.

Now consider $u \in V(H) \smallsetminus v$. We know that $C(u) \leq \deg_G(u)$ as the game begins because C is relaxed. By Corollary 1.10, each node in H fires exactly once. Thus, each of the $\deg_H(u)$ neighbours of u in H fires and adds one chip to u. By Lemma 1.4, the firing order has no effect on C'. Thus,

$$C'(u) = C(u) + \deg_H(u) - (\deg_G(u) + 1) < \deg_H(u).$$

□

2 Key properties of the chip addition operator

With the lemmas of Section 1.2 in hand, we now prove three results that together will allow us to derive our main theorems in Section 3. Recall that $\mathcal{R}$ denotes the set of all relaxed legal configurations on a given connected graph $G = (V, E)$.

Proposition 2.1 *If $C \in \mathcal{R}$ is fixed, then the function $E_{(\cdot)}(C) \colon V \to \mathcal{R}$ mapping $v \mapsto E_v(C)$ is injective.*

For readers pondering the legality of $E_v(C)$, notice that this configuration can be reverse-fired to $C + \mathbf{1}_v$ by following (in reverse) the sequence of relaxing $C + \mathbf{1}_v$ to $(C + \mathbf{1}_v)^\circ$; then, because C is legal, so too is $C + \mathbf{1}_v$, and hence $E_v(C)$ can be reverse-fired to a supercritical configuration.

Proof. For a node $v \in V$, let us abbreviate $E_v(C)$ by R_v. We wish to show that for $u, v \in V$, with $u \neq v$, we have $R_u \neq R_v$. Suppose, for a contradiction, that $R_u = R_v$. If u is the seed and does not fire, then $R_u(u) = C(u) + 1$. If v is the seed and does not fire, then $R_v(u) = C(u)$. Thus, $R_u \neq R_v$, which contradicts our assumption. Thus, at least one of u, v must fire when chosen as the seed.

Suppose first (without loss of generality) that if G is seeded at u, then u fires, while if G is seeded at v, then v does not fire. Since v is not critical in G and we have assumed that $R_u = R_v$, then $C(u) = R_v(u) = R_u(u)$. If u is chosen as the seed, it fires and loses all of its chips. Because $R_u(u) = R_v(u)$, the node u must regain $\deg(u)$ chips. This happens only if every neighbour of u fires after u itself fires. Now for every neighbour of u to fire, at least one of those neighbours must itself be critical in G. Suppose that w is such a neighbour. In the game where v is the seed and does not fire, we thus know that w is critical in G. Because $R_u(w) = R_v(w)$, it must be true that w again becomes critical, after it fires, in the game where u is the seed. But this is impossible by Lemma 1.11.

Now we may suppose that both u and v fire if chosen as the seed. First, we show that if v is the seed (without loss of generality), then u must also fire. Suppose, again for a contradiction, that u does not fire, so $R_v(u) = \deg(u)$. If u and v are neighbours, then clearly u must fire if v is the seed; we therefore assume for the remainder of the proof that u and v are not neighbours. Then in the game where u is the seed, every neighbour of u must also fire so that $R_u(u) = \deg(u)$. Therefore, at least one neighbour w of u must be critical in G. By Lemma 1.11, we know that $R_u(w) < C(w)$. But $R_u = R_v$ implies that w must also fire in the game where v is the seed. Since u is critical in G, it too will fire when w fires. This contradicts our assumption. Thus, if either u or v is the seed, then the other node must fire during the game.

So in the game where v is the seed, at least one neighbour of v must also fire so that u will fire. Suppose that k neighbours of v fire. After v fires, it contains zero chips, so $R_v(v) = k$. Because $R_u = R_v$, we must have $R_u(v) = k$ as well. In the game where u is the seed, at least one neighbour of v must fire in order to allow v to fire. Say that w is such a neighbour. Now after v fires, it contains zero chips, so k of its neighbours must now fire to achieve $R_u(v) = k$. Since w has already fired in this game, it may not be one of these k neighbours (by Corollary 1.10). So at least $k+1$ neighbours of v fire in the game that starts with u as the seed.

Thus, we have k neighbours of v firing in the game where v is the seed and at least $k+1$ neighbours of v firing in the game where u is the seed. There is therefore a node z that is adjacent to v and that fires when u is the seed but not when v is the seed. The number of chips on z at the end of the game in which it does not fire must be at least $C(z)$. But the number of chips on z at the end of a game in which it does fire must be at most

$$C(z) + \deg(z) - (\deg(z) + 1) < C(z).$$

Thus, $R_u(z) < R_v(z)$, contradicting our assumption that $R_u = R_v$. □

Proposition 2.1 establishes that for a given relaxed legal configuration, the choice of different seeds necessarily results in games that relax to distinct configurations. Our next result addresses the origin of a relaxed legal configuration, arguing that each one may be generated by the choice of a seed in some other relaxed legal configuration.

Proposition 2.2 *If $v \in V$ is fixed, then the function $E_v(\cdot)\colon \mathcal{R} \to \mathcal{R}$ mapping $C^* \mapsto E_v(C^*)$ is surjective.*

Proof. For $C \in \mathcal{R}$, we need a configuration $C^* \in \mathcal{R}$ such that $E_v(C^*) = C$. Let $v_n = v$ and $V = \{v_1, \ldots, v_n\}$. With

$$C^* := E_{v_{n-1}} \circ E_{v_{n-2}} \circ \cdots \circ E_{v_2} \circ E_{v_1}(C),$$

we have, by Proposition 1.9, that

$$E_v(C^*) = E_{v_n}(E_{v_{n-1}} \circ \cdots \circ E_{v_1}(C)) = C.$$

□

Our final result in this section establishes the key (cf. Lemma 1.7) to showing that each relaxed legal configuration is recurrent, which we formalize in the proof of Proposition 3.1.

Proposition 2.3 *If $C \in \mathcal{R}$ and v is critical in G, then C is reachable from $E_v(C)$.*

Proof. Let $v_1 = v$ and $V = \{v_1, \ldots, v_n\}$. As $C \in \mathcal{R}$, Proposition 1.9 applies, and we can thus restate (1.4) in the form

$$E_{v_n} \circ E_{v_{n-1}} \circ \cdots \circ E_{v_2}(E_v(C)) = C,$$

wherein the leftmost $(n-1)$ composition factors point out the chip operations leading from $E_v(C)$ back to C. □

3 Uniformity of the stationary distribution

Now we are prepared to present our main results. The state space of our Markov chain $(X_n)_{n\geq 0}$ is the set $\mathcal{R}$ of relaxed legal configurations on a connected graph G. Each transition is determined by seeding a node—via a uniform, random selection—and relaxing the resulting configuration; to be precise, given $X_n \in \mathcal{R}$, the next state is determined by choosing $v \in V$ uniformly at random and setting $X_{n+1} = E_v(X_n)$. Because V is finite and relaxed configurations C satisfy $C(x) \leq \deg_G(x)$ for every $x \in V$, so too is $\mathcal{R}$ finite; we take $r := |\mathcal{R}|$. Let D be the digraph representing the Markov chain and P its transition probability matrix.

Recall that a Markov chain is *irreducible* if all of its states communicate with one another. We proved Proposition 2.3 as a tool for deducing the following result.

Proposition 3.1 *The Markov chain (X_n) is irreducible.*

Proof. In chip-firing language, we need only argue that any given configuration $C \in \mathcal{R}$ is recurrent. With C being legal on G, this graph must contain a critical node v (cf. Algorithm 1.2), so Proposition 2.3 shows that C is reachable from $E_v(C)$. Therefore, Lemma 1.7 guarantees that C is recurrent. □

The next result—our main theorem—provided the motivation for proving Propositions 2.1 and 2.2.

Theorem 3.2 *The matrix P is doubly stochastic.*

PROOF. Starting in a state $C \in \mathcal{R}$, when we initiate a burn-off game on G by seeding one of its nodes v, our Markov chain transitions to $E_v(C)$, a uniquely determined relaxed legal configuration. Consequently, we see that

$$\mathrm{outdeg}_D(C) \leq |V| \text{ for all } C \in \mathcal{R}. \tag{3.1}$$

Summing over all relaxed legal configurations, we have

$$\sum_{C \in \mathcal{R}} \mathrm{outdeg}_D(C) \leq r|V|. \tag{3.2}$$

In Proposition 2.2 we saw that for any configuration $C \in \mathcal{R}$ and node v there is at least one configuration $C^* \in \mathcal{R}$ such that $E_v(C^*) = C$. In Proposition 2.1 we found that different choices of v correspond to different such configurations C^*. Together, these results imply that

$$\mathrm{indeg}_D(C) \geq |V| \text{ for all } C \in \mathcal{R}. \tag{3.3}$$

Summing over all relaxed legal configurations, we have

$$\sum_{C \in \mathcal{R}} \mathrm{indeg}_D(C) \geq r|V|. \tag{3.4}$$

Inequalities (3.2) and (3.4) together give

$$r|V| \geq \sum_{C \in \mathcal{R}} \mathrm{outdeg}_D(C) = \sum_{C \in \mathcal{R}} \mathrm{indeg}_D(C) \geq r|V|,$$

which establishes the identity

$$\sum_{C \in \mathcal{R}} \mathrm{outdeg}_D(C) = \sum_{C \in \mathcal{R}} \mathrm{indeg}_D(C) = r|V|. \tag{3.5}$$

Taking (3.5) along with (3.1) and (3.3) respectively, we conclude that

$$\mathrm{outdeg}_D(C) = |V| = \mathrm{indeg}_D(C) \text{ for all } C \in \mathcal{R}. \tag{3.6}$$

In pursuing (X_n), subsequent seeds are chosen uniformly at random, so every nonzero entry in P is $1/|V|$. Because $\mathrm{indeg}_D(C) = |V|$ for all $C \in \mathcal{R}$, the theorem follows. □

Corollary 3.3 *The stationary distribution of (X_n) is uniform.*

PROOF. It is well known (see, e.g., [16]) that for an irreducible Markov chain with a finite state space $\mathcal{R}$ and a doubly stochastic transition matrix, the (unique) stationary distribution is uniform (over $\mathcal{R}$). □

We denote the stationary distribution of (X_n) by $\pi = (\pi_C)_{C \in \mathcal{R}}$, which we have just proved to satisfy

$$\pi_C = \frac{1}{r} \text{ for all } C \in \mathcal{R}. \tag{3.7}$$

Limiting behaviour and future directions

Let us begin by reminding the reader of a few basic Markov chain notions. For each state $C \in \mathcal{R}$, the *mean recurrence time* μ_C is the expected number of transitions needed to return from the state C to itself. With our state space $\mathcal{R}$ being finite and (X_n) being irreducible, there are no issues with 'transience' or 'null-recurrency', so each μ_C is finite (i.e. (X_n) is 'positive recurrent'). For integers $m \geq 1$, the number of visits of (X_n) to state C during the first m transition epochs is denoted by $N_m(C)$. Recall also that the powers P^m of the transition matrix give the matrices of m-step transition probabilities for (X_n).

In our penultimate result, we summarize the conclusions about this stochastic process that can be drawn from Proposition 3.1, the relation (3.7), and the basic theory of finite Markov chains (see, e.g., [16] or [36]).

Theorem 3.4 *The entities μ_C, $N_m(C)$, and P satisfy the following conclusions:*

$$\mu_C = r \textit{ for all } C \in \mathcal{R};$$

$$\Pr\left\{\lim_{m\to\infty} \frac{N_m(C)}{m} = \frac{1}{r}\right\} = 1 \textit{ for all } C \in \mathcal{R} \textit{ (irrespective of the initial state)}; \quad (3.8)$$

and

$$\lim_{m\to\infty} \frac{1}{m} \sum_{k=1}^{m} P^m_{C'C} = \frac{1}{r} \textit{ for all pairs } C', C \in \mathcal{R}. \quad (3.9)$$

The statement (3.8) has the appealing interpretation that, with high probability, the long-term proportion of time that our Markov chain spends in any given state is equally spread across the states. In the general theory, (3.8) is often used to derive (3.9). This is a weaker conclusion than, say,

$$\lim_{m\to\infty} P^m_{C'C} = \frac{1}{r} \text{ for all pairs } C', C \in \mathcal{R}, \quad (3.10)$$

which would determine the individual limiting m-step transition probabilities as opposed to just their limiting mean. Though we have not been able to prove (3.10), we take some consolation in knowing at least that the Cesaro limit version (3.9) holds.

The barrier keeping us from (3.10) is our inability to determine the 'periodicity' of (X_n). Recall that a state s of a Markov chain (with transition matrix Q) has *period* $d \geq 1$ if $Q^n_{ss} > 0$ implies that d divides n and d is the largest positive integer with this property. Periodicity is a 'class property', and since our chain is irreducible, all states have the same period d. When $d > 1$, such a Markov chain is called *periodic* (with period d) and otherwise is *aperiodic*.

Our final result articulates the dichotomy inherent in the unknown periodicity of (X_n). Like Theorem 3.4, it interprets the well-known theory of finite Markov chains (see, e.g., [16] or [32]) in the context of our stochastic chip-firing process.

Theorem 3.5 *If (X_n) is aperiodic, then all pairs $C', C \in \mathcal{R}$ satisfy*

$$\lim_{m\to\infty} P^m_{C'C} = \frac{1}{r},$$

i.e., (3.10) *holds; if (X_n) is periodic with period d, then for all pairs $C', C \in \mathcal{R}$, there exists a nonnegative integer $k \le d-1$ such that*

$$\lim_{m\to\infty} P^{md+k}_{C'C} = \frac{d}{r} \quad \textit{and} \quad P^n_{C'C} = 0 \textit{ whenever } n \not\equiv k \pmod{d}. \tag{3.11}$$

Whether the limiting behaviour of $P^m_{C'C}$ is described by (3.10) or (3.11), uniformity in the nonzero limiting values is observed; however, we must leave it as an open problem to investigate the periodicity of the Markov chain (X_n).

We intend to present further results on burn-off games from [29] in a future paper, for which the present article will form a foundation.

Acknowledgements

The manuscript for this article was drafted while the second author was on sabbatical at the University of Otago in Dunedin, New Zealand. The author gratefully acknowledges the support of Otago's Department of Mathematics and Statistics. Both authors appreciate the care the referees exercised in reading and commenting on the manuscript.

References

[1] P. Bak and C. Tang, Earthquakes as a self-organized critical phenomenon, *J. Geophys. Res.* **94(B11)** (1989), 15635–15637.

[2] M. Baker and F. Shokrieh, Chip-firing games, potential theory on graphs, and spanning trees, *J. Combin. Theory Ser. A* **120(1)** (2013), 164–182.

[3] B. Benson, D. Chakrabarty and P. Tetali, G-parking functions, acyclic orientations and spanning trees, *Discrete Math.* **310(8)**, (2010), 1340–1353.

[4] N. Biggs and P. Winkler, Chip-firing and the chromatic polynomial, Research Report LSE-CDAM-97-03, Centre for Discrete and Applicable Mathematics, London School of Economics, London, U.K., 1997, 10pp.

[5] N.L. Biggs, Chip-firing and the critical group of a graph, *J. Algebraic Combin.* **9(1)** (1999), 25–45.

[6] A. Björner and L. Lovász, Chip-firing games on directed graphs, *J. Algebraic Combin.* **1(4)** (1992), 305–328.

[7] A. Björner, L. Lovász and P. W. Shor, Chip-firing games on graphs, *European J. Combin.* **12(4)** (1991), 283–291.

[8] B. Bond and L. Levine, Abelian networks I: Foundations and examples, *SIAM J. Discrete Math.* **30(2)** (2016), 856–874.

[9] B. Bond and L. Levine, Abelian networks II: Halting on all inputs, *Selecta Math. (N.S.)* **22(1)** (2016), 319–340.

[10] B. Bond and L. Levine, Abelian networks III: The critical group, *J. Algebraic Combin.* **43(3)** (2016), 635–663.

[11] J.A. Bondy and U.S.R. Murty, *Graph theory*, volume 244 of *Graduate Texts in Mathematics*, Springer, New York, 2008.

[12] A. Cayley, A theorem on trees, In *The Collected Mathematical Papers of Arthur Cayley*, volume 13, pp. 26–28, Cambridge University Press, Cambridge, 1897, (*Quart. J. Math.* **23** (1889), 376–378).

[13] D. Dhar, Self-organized critical state of sandpile automaton models, *Phys. Rev. Lett.* **64(14)** (1990), 1613–1616.

[14] D. Dhar, P. Ruelle, S. Sen and D.-N. Verma, Algebraic aspects of abelian sandpile models, *J. Phys. A* **28(4)** (1995), 805–831.

[15] P. Diaconis and W. Fulton, A growth model, a game, an algebra, Lagrange inversion, and characteristic classes, *Rend. Sem. Mat. Univ. Politec. Torino* **49** (1991), no. 1 (1993), 95–119, Commutative algebra and algebraic geometry, II (Italian) (Turin, 1990).

[16] W. Feller, *An introduction to probability theory and its applications, Vol. I*, Third edition, John Wiley & Sons, Inc., New York-London-Sydney, 1968.

[17] C. Godsil and G. Royle, *Algebraic graph theory*, volume 207 of *Graduate Texts in Mathematics*, Springer-Verlag, New York, 2001.

[18] J. Guzmán and C. Klivans, Chip-firing and energy minimization on M-matrices, *J. Combin. Theory Ser. A* **132** (2015), 14–31.

[19] J. van den Heuvel, Algorithmic aspects of a chip-firing game, *Combin. Probab. Comput.* **10(6)** (2001), 505–529.

[20] A.E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp and D.B. Wilson, Chip-firing and rotor-routing on directed graphs, In *In and out of equilibrium, 2*, volume 60 of *Progr. Probab.*, pp. 331–364, Birkhäuser, Basel, 2008. Corrected in arXiv:0801.3306v4 [math.CO], 20 June 2013.

[21] W. Kager, H. Liu and R. Meester, Existence and uniqueness of the stationary measure in the continuous abelian sandpile, *Markov Process. Related Fields* **16(1)** (2010), 185–204.

[22] P.M. Kayll and D. Perkins, Combinatorial proof of an Abel-type identity, *J. Combin. Math. Combin. Comput.* **70** (2009), 33–40.

[23] P.M. Kayll and D. Perkins, A chip-firing variation and a new proof of Cayley's Formula, *Discrete Math. Theor. Comput. Sci.* **15(1)** (2013), 121–132.

[24] P.M. Kayll and D. Perkins, On lengths of burn-off chip-firing games, (submitted).

[25] L. Levine and J. Propp, What is ... a sandpile?, *Notices Amer. Math. Soc.* **57(8)** (2010), 976–979.

[26] M. Manjunath and B. Sturmfels, Monomials, binomials and Riemann-Roch, *J. Algebraic Combin.* **37(4)** (2013), 737–756.

[27] C. Merino, The chip-firing game, *Discrete Math.* **302(1-3)** (2005), 188–210.

[28] K. Nagel and M. Paczuski, Emergent traffic jams, *Phys. Rev. E* **51(4)** (1995), 2909–2918.

[29] D. Perkins, *Investigations of a chip-firing game*, Thesis (Ph.D.)–University of Montana, ProQuest LLC, Ann Arbor, MI, 2005.

[30] T.V. Pham and T.H.D. Phan, Lattices generated by chip firing game models: criteria and recognition algorithms, *European J. Combin.* **34(5)** (2013), 812–832.

[31] J.G. Propp and D.B. Wilson, How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph, *J. Algorithms* **27(2)** (1998), 170–217, 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996).

[32] S.M. Ross, *Stochastic processes*, Second edition, Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons, Inc., New York, 1996.

[33] M. Schulz, Addition of recurrent configurations in chip firing games: Finding minimal recurrent configurations with Markov chains, In S. Bandini, S. Manzoni, H. Umeo and G. Vizzari, editors, *Cellular Automata*, volume 6350 of *Lecture Notes in Computer Science*, pp. 209–218, Springer, Berlin, Heidelberg, 2010.

[34] D. Stassinopoulos and P. Bak, Democratic reinforcement: A principle for brain function, *Phys. Rev. E* **51(5)** (1995), 5033–5039.

[35] G. Tardos, Polynomial bound for a chip firing game on graphs, *SIAM J. Discrete Math.* **1(3)** (1988), 397–398.

[36] H.C. Tijms, *A first course in stochastic models*, John Wiley & Sons, Ltd., Chichester, 2003.

(Received 3 Apr 2016; revised 17 Mar 2017)