

ABELIAN NETWORKS II. HALTING ON ALL INPUTS

BENJAMIN BOND AND LIONEL LEVINE

ABSTRACT. Abelian networks are systems of communicating automata satisfying a local commutativity condition. We show that a finite irreducible abelian network halts on all inputs if and only if all eigenvalues of its production matrix lie in the open unit disk.

1. INTRODUCTION

Automata networks in general are nondeterministic: the same input can produce many different outputs depending on the order of events at different nodes of the network. However, there is a sizable class of automata networks, proposed by Dhar in [Dha99] and termed *abelian networks* in [BL13], for which the output is uniquely determined by the input.

If we take the view that an abelian network is a kind of asynchronous computer, then one of the most fundamental questions is its *halting problem*. A countably infinite abelian network can emulate a Turing machine with infinite tape [Cai15], so the question of whether a given abelian network halts on a given input is undecidable. However, for a *finite* abelian network it is decidable (below we give an argument using Dickson’s lemma) and its computational complexity is an interesting question.

The first halting criterion in an abelian network that we know of is due to Tardos [Tar88], who found an efficient check for finite termination of chip-firing on an undirected graph. Chip-firing belongs to a subclass of abelian networks called *toppling networks*, which can be described informally as follows. Let L be an integer square matrix with positive diagonal entries nonpositive off-diagonal entries. Each vertex v has a number of *chips* c_v and is allowed to *topple* if $c_v \geq L_{vv}$; the result of toppling v is that v loses L_{vv} chips and each other vertex $u \neq v$ gains $-L_{uv}$ chips. Halting means that there exists a finite sequence of topplings after which $c_v < L_{vv}$ for all v .

The Tardos halting criterion applies when L is the Laplacian of an undirected (or Eulerian directed) graph. The Laplacian of a general directed graph is more difficult because there exist directed graphs with chip configurations c that require exponentially many topplings to halt. Björner, Lovász and Shor [BLS91] remarked,

Date: August 4, 2015.

2010 Mathematics Subject Classification. 68Q10, 37B15, 20M14, 20M35, 05C50,

Key words and phrases. abelian distributed processors, asynchronous computation, automata network, chip-firing, commutative monoid action, Dickson’s lemma, least action principle, M-matrix, sandpile, torsor.

“one can ask for a characterization of those digraphs and initial chip configurations that guarantee finite termination.” This is a difficult problem: Björner and Lovász [BL92] gave an algorithm that takes exponential time in the worst case. Polynomial time algorithms are available for Eulerian and coEulerian graphs, but the problem is NP-complete for a general directed multigraph [FL15].

The halting problem for abelian networks is at least as hard as the problem posed by Björner, Lovász and Shor. In this paper we address the softer question of characterizing which finite abelian networks \mathcal{N} halt on *all* inputs. We associate to \mathcal{N} a *production matrix* P whose entries are nonnegative rational numbers, and show that \mathcal{N} halts on all inputs if and only if the Perron-Frobenius eigenvalue of P is strictly less than 1. This result generalizes an unpublished theorem of Gabrielov [Gab94], who considered the case of a toppling network.

Outline. Section 2 contains mathematical background on commutative monoid actions, Dickson’s lemma and toppling matrices. Theorem 2.5 is of some independent interest: it gives a general mechanism for how torsors (free transitive actions) of abelian groups arise from monoid actions.

After reviewing the definition of an abelian network in Section 3, we define two basic algebraic objects associated to an abelian network, the *total kernel* and *production matrix*, in Section 4. We then examine two ways of decomposing an abelian network: local components and strong components. The former have a smaller state space, the latter a smaller alphabet.

In Section 5 we prove our criterion for halting on all inputs.

In Section 6 we define the *sandpiling* $\mathcal{S}(\mathcal{N})$ of an abelian network \mathcal{N} , and show that \mathcal{N} halts on all inputs if and only if $\mathcal{S}(\mathcal{N})$ does.

We conclude in Section 7 with a short discussion of further research directions.

2. MATHEMATICAL BACKGROUND

2.1. Commutative monoid actions. Any finite commutative monoid M contains an abelian group whose identity element is the minimal idempotent of M . Every monoid action of M induces a corresponding group action. The main result of this section is Theorem 2.5 relating these two actions. We have not seen this theorem stated explicitly in the literature, but some of the lemmas in this section are well-known in the semigroup community. They trace their origins to the work of Green [Gre51] and Schützenberger [Sch57]; see [Gri01, Ste10] for modern treatments. We include their short proofs here in order to highlight the beauty and simplicity of the commutative case. For refinements of some of the lemmas below, and extensions to a certain class of infinite semigroups (π -regular semigroups) see [Gri07].

Let M be a commutative monoid, that is, a set equipped with a commutative and associative operation $(m, m') \mapsto mm'$ with an identity element $\epsilon \in M$ satisfying $\epsilon m = m$ for all $m \in M$. Let

$$\begin{aligned} \mu : M \times X &\rightarrow X \\ (m, x) &\mapsto mx \end{aligned}$$

be a monoid action of M on a set X ; that is, $\epsilon x = x$ and $m(m'x) = (mm')x$ for all $m, m' \in M$ and all $x \in X$. We say that μ is *irreducible* if there does not exist a partition of X into nonempty subsets X_1 and X_2 such that $MX_1 \subset X_1$ and $MX_2 \subset X_2$. Consider the relation \sim on X defined by

$$x \sim x' : \quad \exists m, m' \in M \text{ such that } mx = m'x' \quad (1)$$

Lemma 2.1. \sim is an equivalence relation, and μ is irreducible if and only if \sim has just one equivalence class.

Proof. Clearly $x \sim x$, and $x \sim y$ implies $y \sim x$. To check transitivity, note that if $mx = m'x'$ and $m''x' = m'''x''$, then

$$m''mx = m''m'x' = m'm''x' = m'm'''x''$$

where the middle equality uses commutativity of M . Hence $x \sim x''$.

If X_1 is any equivalence class of \sim then $MX_1 \subset X_1$. Hence if μ is irreducible then \sim has just one equivalence class. Conversely, if \sim has just one equivalence class then for any partition of X into nonempty subsets X_1 and X_2 , we have $m_1x_1 = m_2x_2$ for some $x_1 \in X_1$, $x_2 \in X_2$ and $m_1, m_2 \in M$. If $MX_1 \subset X_1$ then $m_1x_1 \in X_1$ and hence $MX_2 \not\subset X_2$. \square

Assume now that M is finite. An *idempotent* is an element $f \in M$ such that $ff = f$. Among the powers of any element $m \in M$ is an idempotent: by the pigeonhole principle, $m^j = m^k$ for some $j < k$, and then $m^\ell = m^{2\ell}$ where ℓ is any integer multiple of $k - j$ such that $\ell \geq j$. Now consider the product of all idempotents in M :

$$e := \prod_{f \text{ idempotent}} f.$$

Note that e is again an idempotent since M is commutative. This *minimal idempotent* e is accessible from all of M : that is, $e \in mM$ for all $m \in M$.

Lemma 2.2. eM is an abelian group with identity element e .

Proof. Since e is an idempotent we have $e(em) = (ee)m = em$ for all $m \in M$, which verifies the identity axiom. Since e is accessible from all of M , we have $e \in (eme)M = (em)(eM)$ which verifies existence of inverses. \square

In particular, $e \neq \epsilon$ unless M itself is a group.

Lemma 2.3. (Recurrent Elements Of A Monoid Action) *Let M be a finite commutative monoid and $\mu : M \times X \rightarrow X$ an irreducible action. The following are equivalent for $x \in X$:*

- (1) $x \in My$ for all $y \in X$
- (2) $x \in M(mx)$ for all $m \in M$
- (3) $x \in mX$ for all $m \in M$
- (4) $x \in eX$
- (5) $x = ex$

Proof. (1) \Rightarrow (2): trivial.

(2) \Rightarrow (3): Since M is commutative, $M(mx) = (Mm)x = (mM)x \subset Mx$.

(3) \Rightarrow (4): trivial.

(4) \Rightarrow (5): If $x = ey$, then $ex = e(ey) = (ee)y = ey = x$.

(5) \Rightarrow (1): Here we use irreducibility. By Lemma 2.1, given $x, y \in X$ there exist $m, m' \in M$ such that $mx = m'y$. Let m'' be such that $m''m = e$. If $x = ex$, then $x = m''mx = m''m'y \in My$. \square

Any action of a finite commutative monoid

$$\mu : M \times X \rightarrow X$$

induces by restriction a corresponding group action

$$eM \times eX \rightarrow eX.$$

To see this, note that for any $m \in M$ and $x \in X$ we have $m(ex) = (me)x = (em)x = e(mx) \in eX$, so the action of M on X restricts to a monoid action of M on eX . Since $e(ex) = (ee)x = ex$, the element e acts by identity on eX . Since eM is a group with identity element e , it follows that $eM \times eX \rightarrow eX$ is a group action.

In fact slightly more is true. We say that $m \in M$ acts *invertibly* on a subset $Y \subset X$ if the map $y \mapsto my$ is a bijection $Y \rightarrow Y$.

Lemma 2.4. *Let M be a finite commutative monoid and $\mu : M \times X \rightarrow X$ a monoid action. Then every $m \in M$ acts invertibly on eX .*

Proof. For any $m \in M$ and $x \in X$ we have $(em)(ex) = (eme)x = (mee)x = (me)x = m(ex)$, so em and m have the same action on eX . Since $eM \times eX \rightarrow eX$ is a group action, em and hence m acts invertibly on eX . \square

We say that a monoid action $\mu : M \times X \rightarrow X$ is *faithful* if there do not exist distinct elements $m, m' \in M$ such that $mx = m'x$ for all $x \in X$. The next theorem shows that relatively weak properties of a monoid action (faithful and irreducible) imply stronger properties of the corresponding group action (free and transitive).

Let G be a group with identity element e . Recall that a group action $G \times Y \rightarrow Y$ is called *transitive* if $Gy = Y$ for all $y \in Y$, and is called *free* if for all $g \neq e$ there does not exist $y \in Y$ such that $gy = y$. If the action is both transitive and free, then for any two elements $y, y' \in Y$ there is a unique $g \in G$ such that $gy = y'$; in particular, $\#G = \#Y$.

Theorem 2.5. (Group Actions Arising From Monoid Actions) *Let M be a finite commutative monoid and $\mu : M \times X \rightarrow X$ an irreducible monoid action. Then the restriction of μ to $eM \times eX$ is a transitive group action*

$$e\mu : eM \times eX \rightarrow eX.$$

If μ is also faithful, then $e\mu$ is free.

Proof. To show transitivity, let $R = \bigcap_{y \in X} (My)$. Then for any $x \in eX$ we have

$$R \subset Mx = M(ex) = (Me)x = (eM)x = e(Mx) \subset eX.$$

But $R = eX$ by the equivalence of (1) and (4) in Lemma 2.3, so the above inclusions are equalities. In particular, $(eM)x = eX$ for all $x \in eX$, which shows that eM acts transitively on eX .

To show freeness, suppose that $g(ex) = ex$ for some $g \in eM$ and some $x \in X$. Now fix $y \in X$. By transitivity, $ey = h(ex)$ for some $h \in eM$, hence

$$gy = gey = ghex = hgex = hex = ey.$$

If μ is faithful, it follows that $g = e$. Hence eM acts freely on eX . \square

2.2. Dickson's Lemma. The following lemma can be proved by induction on k using the infinite pigeonhole principle. We remark that it is also a case of the Hilbert basis theorem applied to the monomial ideal $(\mathbf{t}^{\mathbf{x}_1}, \mathbf{t}^{\mathbf{x}_2}, \dots)$ in the polynomial ring $\mathbb{Q}[\mathbf{t}] = \mathbb{Q}[t_1, \dots, t_k]$ (In fact the proof of the basis theorem now found in many textbooks, due to Gordan in 1900, uses this special case as a stepping stone; see [Eis95, Ex. 15.15] and [DdJ98] for some history.)

Lemma 2.6. (Dickson's Lemma, [Dic13]) *Let $k \geq 1$ be an integer. For any sequence $\mathbf{x}_1, \mathbf{x}_2, \dots \in \mathbb{N}^k$, there exist indices $m < n$ such that $\mathbf{x}_m \leq \mathbf{x}_n$ in the coordinatewise partial ordering.*

See [FFSS11] for an effective version, giving bounds on n in terms of k and the largest jump in the sequence \mathbf{x}_i . Another bound is obtained in [BS15] from a formalization of an existence proof.

2.3. Toppling matrices. We will use the following form of the Perron-Frobenius theorem (see [HJ90, §8] for part (i) and [Ash87] for part (ii)).

Lemma 2.7. (Perron-Frobenius) *Let P be a square matrix with nonnegative real entries.*

- (i) *P has a nonnegative real eigenvector \mathbf{x} with nonnegative real eigenvalue λ , such that the absolute values of all other eigenvalues of P are $\leq \lambda$.*
- (ii) *If P has rational entries and λ is rational, then \mathbf{x} can be taken to have integer entries.*

Following [PS04], we call L a *toppling matrix* if it satisfies the equivalent conditions of the following lemma. See [FK62, Theorem 4.3] for a proof of the equivalence. Such matrices are also known as “ M -matrices.”

For a vector \mathbf{x} , we write $\mathbf{x} > \mathbf{0}$ to mean that all coordinates of \mathbf{x} are nonnegative and $\mathbf{x} \neq \mathbf{0}$.

Lemma 2.8. (Toppling Matrices) *Let L be a square matrix with real entries such that $L_{ij} \leq 0$ for all $i \neq j$. The following are equivalent.*

- (1) *All principal minors of L are positive.*
- (2) *All eigenvalues of L have positive real part.*
- (3) *There exists a vector $\mathbf{x} > \mathbf{0}$ such that $L\mathbf{x} > \mathbf{0}$.*
- (4) *There exists a vector $\mathbf{y} > \mathbf{0}$ such that $L^T\mathbf{y} > \mathbf{0}$.*
- (5) *L is invertible, and all entries of L^{-1} are nonnegative.*

3. ABELIAN NETWORKS; HALTING DICHOTOMY

We now recall the definition of an abelian network, referring the reader to [BL13] for details. In an abelian network on a directed graph $G = (V, E)$, each vertex $v \in V$ has a *processor* \mathcal{P}_v which is an automaton with input alphabet A_v and state space Q_v . For each letter $a \in A_v$ there is a state transition map $t_a : Q_v \rightarrow Q_v$, and these maps are required to commute: $t_a t_b = t_b t_a$ for all $a, b \in A_v$. For each edge $(v, u) \in E$ there is a message-passing function $A_v \times Q_v \rightarrow A_u^*$ specifying the word (possibly empty) sent to processor \mathcal{P}_u in the event that \mathcal{P}_v in state $q \in Q_v$ processes letter $a \in A_v$. Each message-passing function is required to satisfy a commutativity condition: namely, if two input words to \mathcal{P}_v are permutations of one another, then for each outgoing edge (v, u) the resulting messages passed to \mathcal{P}_u must be permutations of one another.

To give a concrete example, let L be an integer $V \times V$ matrix with positive diagonal entries nonpositive off-diagonal entries. The (locally recurrent) *toppling network* $\text{Topp}(L)$ has $A_v = \{v\}$ and $Q_v = \mathbb{Z}/L_{vv}\mathbb{Z}$. The state transition is $t_v(q) = q + 1 \pmod{L_{vv}}$. Whenever processor \mathcal{P}_v transitions from state $L_{vv} - 1$ to state 0 it sends $-L_{uv}$ letters u to each processor \mathcal{P}_u for $u \in V$. The *sandpile network* $\text{Sand}(G)$ of a directed graph G without self-loops is the special case where L_{vv} is the outdegree of vertex v and $-L_{uv}$ is the number of edges from v to u . The commutativity conditions in the previous paragraph hold vacuously for a toppling network because it is *unary*: each alphabet A_v has just one letter. See [BL13] for many other examples of abelian networks, including non-unary examples.

The total state of an abelian network $\mathcal{N} = (\mathcal{P}_v)_{v \in V}$ is described by an element $\mathbf{q} \in Q := \prod_{v \in V} Q_v$ giving the internal states of the processors, together with a vector $\mathbf{x} \in \mathbb{Z}^A$ where $A = \sqcup_{v \in V} A_v$ indicating how many letters of each type are waiting to be processed. We use the notation \mathbf{x}, \mathbf{q} for this pair. Note that A is a disjoint union, so \mathbf{x} also specifies the locations of the letters (the \mathbf{x}_a letters a are located at the unique vertex v such that $a \in A_v$).

An *execution* is a word $w = a_1 \cdots a_r \in A^*$. It prescribes an order in which letters are to be processed. We write $\pi_w(\mathbf{x}, \mathbf{q})$ for the result of executing w starting from \mathbf{x}, \mathbf{q} . By definition π_w is the composition $\pi_{a_1} \circ \cdots \circ \pi_{a_r}$. Each π_a has three effects: change the internal state \mathbf{q}_v to $t_a \mathbf{q}_v$, where v is the unique vertex such that $a \in A_v$; decrement \mathbf{x}_a by one; and increment each coordinate \mathbf{x}_b by the number of letters b passed (as specified by the message passing function $T_{(v,u)}$ where $b \in A_u$). Note that decrementing \mathbf{x}_a may cause it to become negative, which is the reason for taking \mathbf{x} in \mathbb{Z}^A rather than \mathbb{N}^A . Writing $\pi_{a_1 \cdots a_i}(\mathbf{x}, \mathbf{q}) = \mathbf{x}^i, \mathbf{q}^i$, we say that w is *legal* for \mathbf{x}, \mathbf{q} if $\mathbf{x}_{a_i}^{i-1} \geq 1$ for all $i = 1, \dots, r$. We say that w is *complete* for \mathbf{x}, \mathbf{q} if $\mathbf{x}^r \leq \mathbf{0}$. The least action principle [BL13, Lemma 4.3] says that if w is any legal execution for \mathbf{x}, \mathbf{q} and w' is any complete execution for \mathbf{x}, \mathbf{q} , then $|w|_a \leq |w'|_a$ for all $a \in A$, where $|w|_a$ is the number of letters a in the word w . A consequence of particular importance in this paper is the following.

Lemma 3.1. (Halting Dichotomy, [BL13, Lemma 4.4]) *For a given initial state \mathbf{q} and input \mathbf{x} to an abelian network \mathcal{N} , either*

- (1) *There does not exist a finite complete execution for \mathbf{x}, \mathbf{q} ; or*

- (2) *Every legal execution for \mathbf{x}, \mathbf{q} is finite, and any two complete legal executions w, w' for \mathbf{x}, \mathbf{q} satisfy $|w| = |w'|$.*

It follows readily from the axioms of an abelian network (see [BL13, Lemma 4.2]) that $\pi_a \circ \pi_b = \pi_b \circ \pi_a$ for all $a, b \in A$ and hence π_w depends only on $|w|$. For $\mathbf{y} \in \mathbb{N}^A$, write $\pi_{\mathbf{y}}$ to mean π_w for any word w such that $|w| = \mathbf{y}$. We record here two basic properties of $\pi_{\mathbf{y}}$.

Lemma 3.2. *For all $\mathbf{y}, \mathbf{z} \in \mathbb{N}^A$ we have*

- (i) $\pi_{\mathbf{y}+\mathbf{z}} = \pi_{\mathbf{y}} \circ \pi_{\mathbf{z}}$.
- (ii) *For all $\mathbf{x}, \mathbf{w} \in \mathbb{Z}^A$ and all $\mathbf{q} \in Q$, if $\pi_{\mathbf{y}}(\mathbf{x}, \mathbf{q}) = \mathbf{x}', \mathbf{q}'$, then*

$$\pi_{\mathbf{y}}((\mathbf{x} + \mathbf{w}), \mathbf{q}) = (\mathbf{x}' + \mathbf{w}), \mathbf{q}'.$$

Proof. For part (i), if $|w| = \mathbf{y}$ and $|w'| = \mathbf{z}$ then $\pi_{\mathbf{y}+\mathbf{z}} = \pi_{ww'} = \pi_w \circ \pi_{w'} = \pi_{\mathbf{y}} \circ \pi_{\mathbf{z}}$. Part (ii) is immediate from [BL13, eq. (4)]; it says that any additional letters \mathbf{w} present during an execution affect neither the messages passed $\mathbf{x}' - \mathbf{x}$ nor the final state \mathbf{q}' . \square

4. TOTAL KERNEL AND PRODUCTION MATRIX

In this section we continue the development of the foundations of abelian networks begun in [BL13]. We associate two algebraic objects to an abelian network, the *total kernel* K and *production matrix* P . Only P figures in our criterion for halting on all inputs, but we will see that K is the natural domain of P considered as a \mathbb{Z} -linear map. In the sequel [BL14], both K and P play an essential role in analyzing the critical group of an abelian network that halts on all inputs.

4.1. The local action. Given $\mathbf{x} \in \mathbb{N}^A$ and $\mathbf{q} \in Q$, define

$$\mathbf{x} \triangleright \mathbf{q} := \pi_{\mathbf{x}}(\mathbf{x}, \mathbf{q}).$$

In words, $\mathbf{x} \triangleright \mathbf{q} \in \mathbb{N}^A \times Q$ is the result of performing the following operation starting from state \mathbf{q} : “for each $a \in A$ add \mathbf{x}_a letters of type a and process each letter once.” If $\mathbf{x} \triangleright \mathbf{q} = \mathbf{y}, \mathbf{r}$ then message passing produced a total of \mathbf{y}_a letters of type a for each $a \in A$, and the resulting state of the processor at vertex v was \mathbf{r}_v for each $v \in V$.

For any $\mathbf{z} \in \mathbb{Z}^A$ we also write

$$\mathbf{x} \triangleright (\mathbf{z}, \mathbf{q}) := \pi_{\mathbf{x}}((\mathbf{x} + \mathbf{z}), \mathbf{q}).$$

The next lemma shows that \triangleright defines a monoid action of \mathbb{N}^A on $\mathbb{Z}^A \times Q$. We call this the *local action* because each processor \mathcal{P}_v processes only the letters \mathbf{x}_v that were added at v .

Lemma 4.1. *For any $\mathbf{x}, \mathbf{y} \in \mathbb{N}^A$ and $\mathbf{z} \in \mathbb{Z}^A$ and any $\mathbf{q} \in Q$,*

$$\mathbf{x} \triangleright (\mathbf{y} \triangleright (\mathbf{z}, \mathbf{q})) = (\mathbf{x} + \mathbf{y}) \triangleright (\mathbf{z}, \mathbf{q}).$$

Proof. Write $\mathbf{y} \triangleright \mathbf{q} = \mathbf{y}' \cdot \mathbf{q}'$. By Lemma 3.2(i), since $\pi_{\mathbf{x}+\mathbf{y}} = \pi_{\mathbf{x}} \circ \pi_{\mathbf{y}}$, we have

$$\begin{aligned} (\mathbf{x} + \mathbf{y}) \triangleright (\mathbf{z} \cdot \mathbf{q}) &= \pi_{\mathbf{x}+\mathbf{y}}((\mathbf{x} + \mathbf{y} + \mathbf{z}) \cdot \mathbf{q}) \\ &= \pi_{\mathbf{x}}(\pi_{\mathbf{y}}((\mathbf{x} + \mathbf{y} + \mathbf{z}) \cdot \mathbf{q})) \\ &= \pi_{\mathbf{x}}((\mathbf{x} + \mathbf{y}' + \mathbf{z}) \cdot \mathbf{q}') \\ &= \mathbf{x} \triangleright ((\mathbf{y}' + \mathbf{z}) \cdot \mathbf{q}') \\ &= \mathbf{x} \triangleright (\mathbf{y} \triangleright (\mathbf{z} \cdot \mathbf{q})) \end{aligned}$$

where in the third and last equalities we have used Lemma 3.2(ii). \square

The *local monoid* M_v of a vertex v is the set of maps $Q_v \rightarrow Q_v$ generated by the maps t_a for $a \in A_v$ under composition. Write $t_v : \mathbb{N}^{A_v} \rightarrow M_v$ for the monoid homomorphism sending $\mathbf{1}_a \mapsto t_a$ for $a \in A_v$. Denote by

$$t : \mathbb{N}^A \rightarrow \prod_{v \in V} M_v$$

the Cartesian product of the maps t_v . Each t_v is surjective by the definition of M_v , so t is surjective. Note that if $\mathbf{y} \triangleright \mathbf{q} = \mathbf{z} \cdot \mathbf{r}$, then $\mathbf{r} = t(\mathbf{y})\mathbf{q}$. Given $\mathbf{m} \in \prod_{v \in V} M_v$, write $\mathbf{m}\mathbf{q} = (\mathbf{m}_v \mathbf{q}_v)_{v \in V}$. In general, knowing $\mathbf{y} \triangleright \mathbf{q} = \mathbf{z} \cdot \mathbf{r}$ does not determine $\mathbf{y} \triangleright (\mathbf{m}\mathbf{q})$, but the next lemma shows that it does in the case $\mathbf{r} = \mathbf{q}$.

Lemma 4.2. *If $\mathbf{y} \triangleright \mathbf{q} = \mathbf{z} \cdot \mathbf{q}$, then $\mathbf{y} \triangleright (\mathbf{m}\mathbf{q}) = \mathbf{z} \cdot (\mathbf{m}\mathbf{q})$ for all $\mathbf{m} \in \prod_{v \in V} M_v$.*

Proof. Since t is surjective there exists $\mathbf{u} \in \mathbb{N}^A$ such that $t(\mathbf{u}) = \mathbf{m}$. Then $\mathbf{u} \triangleright \mathbf{q} = \mathbf{w} \cdot \mathbf{m}\mathbf{q}$ for some $\mathbf{w} \in \mathbb{N}^A$. By Lemma 4.1,

$$\mathbf{y} \triangleright (\mathbf{w} \cdot \mathbf{m}\mathbf{q}) = \mathbf{y} \triangleright (\mathbf{u} \triangleright \mathbf{q}) = \mathbf{u} \triangleright (\mathbf{y} \triangleright \mathbf{q}) = \mathbf{u} \triangleright (\mathbf{z} \cdot \mathbf{q}) = (\mathbf{w} + \mathbf{z}) \cdot \mathbf{m}\mathbf{q}.$$

From Lemma 3.2(ii) it follows that $\mathbf{y} \triangleright \mathbf{m}\mathbf{q} = \mathbf{z} \cdot \mathbf{m}\mathbf{q}$. \square

4.2. Production matrix. From here on we assume that \mathcal{N} is a *finite* abelian network: that is, the underlying graph $G = (V, E)$ is finite, and the alphabet and state space of each vertex are finite sets. The main ingredients that rely on finiteness are the results from §2.1.

Let e_v be the minimal idempotent of the local monoid M_v . The recurrent elements (Lemma 2.3) of the monoid action $M_v \times Q_v \rightarrow Q_v$ play a special role in this section.

Definition 4.3. A state $\mathbf{q} \in Q$ is *locally recurrent* if $\mathbf{q}_v \in e_v Q_v$ for all $v \in V$. (Equivalently, $\mathbf{q}_v = e_v \mathbf{q}_v$ for all $v \in V$.)

By Lemma 2.4, every $m \in M_v$ acts invertibly on $e_v Q_v$. Thus for each $a \in A_v$ the map $q \mapsto T_v(a, q)$ is a permutation of $e_v Q_v$, so we have a group action

$$\mathbb{Z}^{A_v} \times e_v Q_v \rightarrow e_v Q_v.$$

Let K_v be the set of vectors in \mathbb{Z}^{A_v} that act as the identity on $e_v Q_v$.

Definition 4.4. The *total kernel* of \mathcal{N} is the subgroup of \mathbb{Z}^A given by

$$K = \prod_{v \in V} K_v.$$

Lemma 4.5. *If \mathcal{N} is finite, then K is a subgroup of finite index in \mathbb{Z}^A . In particular, K is generated as a group by $K \cap \mathbb{N}^A$.*

Proof. Since each $e_v Q_v$ is a finite set, for any $\mathbf{x} \in \mathbb{Z}^{A_v}$ we have $n\mathbf{x} \in K_v$ for some $n \geq 1$. Thus K_v has finite index in \mathbb{Z}^{A_v} . Since V is a finite set, $K = \prod_{v \in V} K_v$ has finite index in \mathbb{Z}^A . In particular, K contains a vector with all coordinates strictly positive, which implies that K is generated as a group by $K \cap \mathbb{N}^A$. \square

Note that

$$K \cap \mathbb{N}^A = \{\mathbf{x} \in \mathbb{N}^A \mid t(\mathbf{x})\mathbf{q} = \mathbf{q} \text{ for all locally recurrent } \mathbf{q} \in Q\}. \quad (2)$$

Fix a locally recurrent state $\mathbf{q} \in Q$. For any $\mathbf{k} \in K \cap \mathbb{N}^A$ we have

$$\mathbf{k} \triangleright \mathbf{q} = P_{\mathbf{q}}(\mathbf{k}) \cdot \mathbf{q} \quad (3)$$

for some vector $P_{\mathbf{q}}(\mathbf{k}) \in \mathbb{N}^A$. Next we show that

$$P_{\mathbf{q}} : K \cap \mathbb{N}^A \rightarrow \mathbb{N}^A$$

extends to a group homomorphism.

Lemma 4.6. *Let \mathcal{N} be a finite abelian network. Then $P_{\mathbf{q}}$ extends to a group homomorphism $K \rightarrow \mathbb{Z}^A$.*

Proof. Write $P = P_{\mathbf{q}}$. Let $\mathbf{k}_1, \mathbf{k}_2 \in K \cap \mathbb{N}^A$. By Lemma 4.1,

$$(\mathbf{k}_1 + \mathbf{k}_2) \triangleright \mathbf{q} = \mathbf{k}_1 \triangleright (P(\mathbf{k}_2) \cdot \mathbf{q}) = (P(\mathbf{k}_1) + P(\mathbf{k}_2)) \cdot \mathbf{q}$$

hence

$$P(\mathbf{k}_1 + \mathbf{k}_2) = P(\mathbf{k}_1) + P(\mathbf{k}_2). \quad (4)$$

By Lemma 4.5, every $\mathbf{k} \in K$ can be written as $\mathbf{k}_1 - \mathbf{k}_2$ for $\mathbf{k}_1, \mathbf{k}_2 \in K \cap \mathbb{N}^A$. Define $P(\mathbf{k}) = P(\mathbf{k}_1) - P(\mathbf{k}_2)$. Equation (4) now implies that this extension is well defined and a group homomorphism. \square

By tensoring $P_{\mathbf{q}} : K \rightarrow \mathbb{Z}^A$ with \mathbb{Q} , we obtain a linear map $P_{\mathbf{q}} : \mathbb{Q}^A \rightarrow \mathbb{Q}^A$. To be more explicit, for any $\mathbf{x} \in \mathbb{Q}^A$, by Lemma 4.5 there is an integer $n \geq 1$ such that $n\mathbf{x} \in K$, and we define

$$P_{\mathbf{q}}(\mathbf{x}) := \frac{1}{n} P_{\mathbf{q}}(n\mathbf{x}).$$

So far we have defined $P_{\mathbf{q}}$ only for locally recurrent \mathbf{q} . We extend the definition to all states $\mathbf{q} = (q_v)_{v \in V}$ by setting $P_{\mathbf{q}} := P_{\hat{\mathbf{q}}}$, where $\hat{\mathbf{q}} = (e_v q_v)_{v \in V}$.

Definition 4.7. The *production matrix* of a finite abelian network \mathcal{N} with initial state \mathbf{q} is the matrix of the linear map $P_{\mathbf{q}} : \mathbb{Q}^A \rightarrow \mathbb{Q}^A$.

The (a, b) entry p_{ab} of the production matrix says ‘‘on average’’ how many letters a are created by processing the letter b : specifically, if $n1_b \in K$, then $p_{ab} = \frac{1}{n} p_{ab}(n)$, where $p_{ab}(n) = \mathbf{N}(b^n, q)_a$ is the number of a ’s created by executing the word b^n . We have chosen the term ‘‘production matrix’’ to evoke [DFR05]. Indeed the succession rules studied in that paper can be modeled by an abelian network whose underlying graph is a single vertex with a loop.

We remark that the production matrix can be also defined for some networks that are not finite by setting

$$p_{ab} := \lim_{n \rightarrow \infty} \frac{1}{n} p_{ab}(n)$$

if this limit exists.

4.3. Local components. An abelian finite automaton \mathcal{P} with state space Q and alphabet A is specified by transition maps $t_a : Q \rightarrow Q$ for $a \in A$, such that $t_a \circ t_b = t_b \circ t_a$ for all $a, b \in A$. The *transition monoid* of \mathcal{P} is the submonoid $M \subset \text{End}(Q)$ generated by $\{t_a\}_{a \in A}$, where $\text{End}(Q)$ denotes the monoid of all set maps $Q \rightarrow Q$ under composition. (For example, the local monoid M_v of §4.1 is the transition monoid of \mathcal{P}_v .)

By construction, the transition monoid of \mathcal{P} has a faithful action $M \times Q \rightarrow Q$. We say that \mathcal{P} is *irreducible* if this monoid action is irreducible (§2.1). The *irreducible components* \mathcal{P}_α of \mathcal{P} are the automata with alphabet A and state space Q_α , where $Q = \sqcup Q_\alpha$ is the partition of Q into equivalence classes under the relation $q \sim q'$ if there exist $m, m' \in M$ such that $mq = m'q'$ (see Lemma 2.1).

We say that an abelian network $\mathcal{N} = (\mathcal{P}_v)_{v \in V}$ is *locally irreducible* if each processor \mathcal{P}_v is irreducible. The *local components* of \mathcal{N} are the abelian networks $\mathcal{N}_\alpha = (\mathcal{P}_v^{\alpha v})_{v \in V}$ where each $\mathcal{P}_v^{\alpha v}$ is an irreducible component of \mathcal{P}_v . Note that the local components have the same underlying graph G and alphabet A , with a possibly smaller state space at each vertex.

The next lemma should be compared with (2).

Lemma 4.8. *If \mathcal{N} is finite and locally irreducible, then for any fixed locally recurrent $\mathbf{q} \in Q$ we have*

$$K \cap \mathbb{N}^A = \{\mathbf{k} \in \mathbb{N}^A \mid t(\mathbf{k})\mathbf{q} = \mathbf{q}\}.$$

Proof. By Theorem 2.5 the action of $e_v M_v$ on $e_v Q_v$ is free. If $\mathbf{k} \in \mathbb{N}^A$ and $t(\mathbf{k})\mathbf{q} = \mathbf{q}$ for one locally recurrent \mathbf{q} , then for all $v \in V$

$$t_v(\mathbf{k}_v)e_v \mathbf{q}_v = t_v(\mathbf{k}_v)\mathbf{q}_v = \mathbf{q}_v.$$

By freeness it follows that $t_v(\mathbf{k}_v)e_v = e_v$, so $t(\mathbf{k})\mathbf{r} = \mathbf{r}$ for all locally recurrent \mathbf{r} . Hence $\mathbf{k} \in K$. \square

For $\mathbf{q}, \mathbf{r} \in Q$ we write $\mathbf{q} \sim \mathbf{r}$ if \mathbf{q} and \mathbf{r} belong to the same local component (that is, $\mathbf{q}_v \sim \mathbf{r}_v$ for all $v \in V$). We denote by $\mathcal{N}_{\mathbf{q}}$ the local component of \mathcal{N} containing state \mathbf{q} . The next lemma shows that within a local component all states have the same production matrix.

Lemma 4.9.

- (i) *In a locally irreducible finite abelian network, the production matrix $P_{\mathbf{q}}$ does not depend on the initial state \mathbf{q} .*
- (ii) *In any finite abelian network, if $\mathbf{q} \sim \mathbf{r}$ then $P_{\mathbf{q}} = P_{\mathbf{r}}$.*

Proof. (i) Let $\mathbf{q}_1, \mathbf{q}_2 \in Q$ be locally recurrent states. By Theorem 2.5 each group action $e_v M_v \times e_v Q_v \rightarrow e_v Q_v$ is transitive, so there exists $\mathbf{x} \in \mathbb{N}^A$ such that $t(\mathbf{x})\mathbf{q}_1 = \mathbf{q}_2$. Let \mathbf{z} be such that

$$\mathbf{x} \triangleright \mathbf{q}_1 = \mathbf{z} \cdot \mathbf{q}_2.$$

Fix $\mathbf{k} \in K \cap \mathbb{N}^A$, and let $\mathbf{y}_i = P_{\mathbf{q}_i}(\mathbf{k})$ for $i = 1, 2$. Then

$$\mathbf{x} \triangleright (\mathbf{k} \triangleright \mathbf{q}_1) = \mathbf{x} \triangleright (\mathbf{y}_1 \cdot \mathbf{q}_1) = (\mathbf{y}_1 + \mathbf{z}) \cdot \mathbf{q}_2$$

while

$$\mathbf{k} \triangleright (\mathbf{x} \triangleright \mathbf{q}_1) = \mathbf{k} \triangleright (\mathbf{z} \cdot \mathbf{q}_2) = (\mathbf{y}_2 + \mathbf{z}) \cdot \mathbf{q}_2.$$

By Lemma 4.1,

$$\mathbf{k} \triangleright (\mathbf{x} \triangleright \mathbf{q}) = \mathbf{x} \triangleright (\mathbf{k} \triangleright \mathbf{q})$$

hence $\mathbf{y}_1 + \mathbf{z} = \mathbf{y}_2 + \mathbf{z}$, and hence $\mathbf{y}_1 = \mathbf{y}_2$. Since $\mathbf{k} \in K \cap \mathbb{N}^A$ was arbitrary we conclude that $P_{\mathbf{q}_1} = P_{\mathbf{q}_2}$.

Part (ii) follows by applying part (i) to the local component $\mathcal{N}_{\mathbf{q}}$. \square

4.4. Strong components. Let \mathcal{N} be a locally irreducible finite abelian network with production matrix $P = (p_{ab})_{a,b \in A}$. In this section we describe another way to break \mathcal{N} into smaller pieces by reducing its alphabet A . (The results of this section are not used anywhere else in this paper, so the reader who wishes to skip to the halting problem in §5 can safely do so.)

Definition 4.10. The *production graph* $\Gamma = \Gamma(\mathcal{N})$ of \mathcal{N} is the directed graph with vertex set A and edge set $\{(a, b) \mid p_{ba} > 0\}$.

Write $a \dashrightarrow b$ if there is a directed path in Γ from a to b . The *strong components* of Γ are the equivalence classes of the relation $\{(a, b) \mid a \dashrightarrow b \text{ and } b \dashrightarrow a\}$.

Definition 4.11. The *strong components* of \mathcal{N} are the subnetworks $\mathcal{N}^1, \dots, \mathcal{N}^s$ with alphabets A^1, \dots, A^s , where A^1, \dots, A^s are the strong components of Γ .

Note that the local components of §4.3 are defined by restricting the state space, whereas the strong components are defined by restricting the alphabet. The strong components of a locally irreducible network need not be locally irreducible (Figure 1). However, Lemma 4.13 below shows that the local components of a strong component do not decompose any further. Moreover, all local components of a given strong component are homotopic.

Definition 4.12. (Homotopy) Locally irreducible abelian networks \mathcal{N} and \mathcal{N}' on the same graph with the same total alphabet are called *homotopic*, written $\mathcal{N} \approx \mathcal{N}'$, if they have the same total kernel K and the same production matrix P .

The reason for calling this homotopy comes from a method of diagramming states described in [BL13]. The locally recurrent states of each processor correspond to points on the discrete torus \mathbb{Z}^{A_v}/K_v , with state transitions given by adding a basis vector $\mathbf{1}_a$ for $a \in A_v$. Message passing can be visualized by surfaces cutting between the states, with each surface labeled by the letter to be passed. If \mathcal{N} and \mathcal{N}' have the same total kernel K , then for each vertex v the state diagrams of processors \mathcal{P}_v and \mathcal{P}'_v live on the same discrete torus. Then \mathcal{N} and \mathcal{N}' have

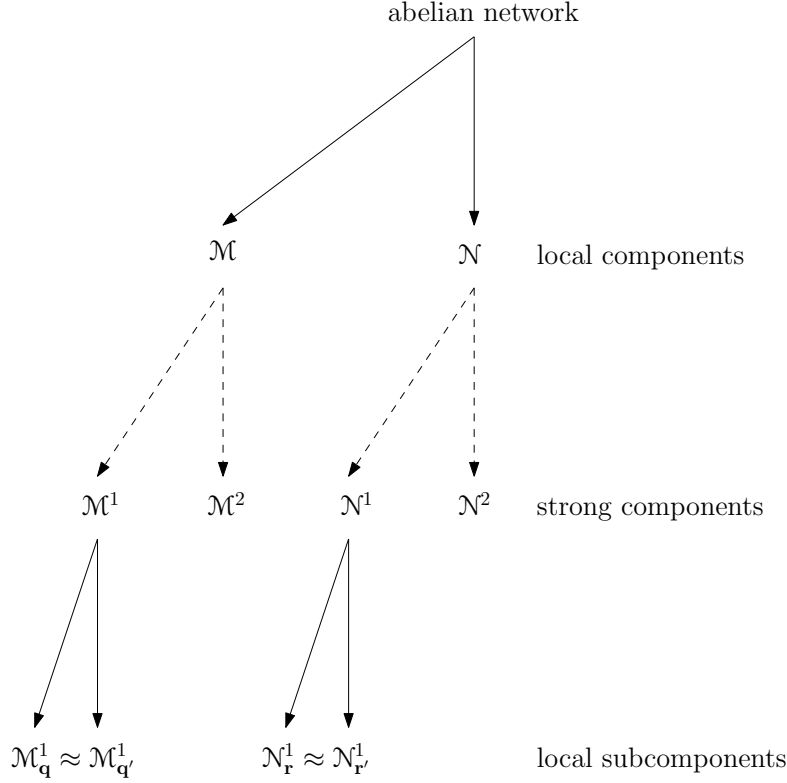


FIGURE 1. The local components of an abelian network may decompose into strong components, which may further decompose further into local subcomponents. Solid arrows represent restrictions of the state space, and dashed arrows represent restrictions of the alphabet.

the same production matrix if and only if for each v the state diagram of \mathcal{P}'_v can be obtained from the state diagram of \mathcal{P}_v by altering message surfaces without changing the homotopy type of any surface.

Lemma 4.13. *Let \mathcal{N} be a locally irreducible finite abelian network with strong components \mathcal{N}^i , total kernel K and production matrix P . Let $\mathcal{N}_{\mathbf{q}}^i$ be the local component of \mathcal{N}^i containing state \mathbf{q} . The following hold for all $\mathbf{q}, \mathbf{q}' \in Q$.*

- (i) *The total kernel of $\mathcal{N}_{\mathbf{q}}^i$ is the group generated by $K \cap \mathbb{N}^{A^i}$.*
- (ii) *The production matrix of $\mathcal{N}_{\mathbf{q}}^i$ is the $A^i \times A^i$ submatrix P_{ii} of P .*
- (iii) *$\mathcal{N}_{\mathbf{q}}^i$ has only one strong component.*
- (iv) *$\mathcal{N}_{\mathbf{q}}^i \approx \mathcal{N}_{\mathbf{q}'}^i$.*

Proof. Since $\mathcal{N}_{\mathbf{q}}^i = \mathcal{N}_{\widehat{\mathbf{q}}}^i$ where $\widehat{\mathbf{q}} = (e_v^i \mathbf{q}_v)_{v \in V}$, we may assume that \mathbf{q} is a locally recurrent state of \mathcal{N}^i .

(i) By Lemma 4.5, the total kernel $K_{\mathbf{q}}^i$ of $\mathcal{N}_{\mathbf{q}}^i$ is generated as a group by $K_{\mathbf{q}}^i \cap \mathbb{N}^{A^i}$. Since both $\mathcal{N}_{\mathbf{q}}^i$ and \mathcal{N} are locally irreducible, by Lemma 4.8 we have for $\mathbf{k} \in \mathbb{N}^{A^i}$

$$\mathbf{k} \in K_{\mathbf{q}}^i \iff t(\mathbf{k})\mathbf{q} = \mathbf{q} \iff \mathbf{k} \in K.$$

Hence

$$K_{\mathbf{q}}^i \cap \mathbb{N}^{A^i} = K \cap \mathbb{N}^{A^i}.$$

(ii) Write \triangleright for the local action of \mathcal{N} and \triangleright^i for the local action of \mathcal{N}^i . For $\mathbf{k} \in \mathbb{N}^{A^i}$ the only difference between these actions is that $\mathbf{k} \triangleright \mathbf{q}$ may produce some letters in $A - A^i$ in addition to the letters in A^i produced by $\mathbf{k} \triangleright^i \mathbf{q}$. Hence, writing $P_{\mathbf{q}}^i$ for the production matrix of $\mathcal{N}_{\mathbf{q}}^i$, if $\mathbf{k} \in K \cap \mathbb{N}^{A^i}$ then $\mathbf{k} \triangleright \mathbf{q} = \mathbf{y} \cdot \mathbf{q}$, where $\mathbf{y} = P_{\mathbf{q}}^i \mathbf{k} + \mathbf{z}$ for some $\mathbf{z} \in \mathbb{N}^{A - A^i}$. Letting $\mathbf{r} = (e_v \mathbf{q}_v)_{v \in V}$, we have $\mathbf{k} \triangleright \mathbf{r} = \mathbf{y} \cdot \mathbf{r}$ by Lemma 4.2. Since \mathbf{r} is locally recurrent for \mathcal{N} , it follows that $\mathbf{y} = P\mathbf{k}$. Writing ρ_i for the projection $\mathbb{N}^A \rightarrow \mathbb{N}^{A^i}$, we conclude that

$$P_{\mathbf{q}}^i \mathbf{k} = \rho_i(P\mathbf{k}) = P_{ii}\mathbf{k}$$

and hence $P_{\mathbf{q}}^i = P_{ii}$.

(iii) By part (ii), the production graph of $\mathcal{N}_{\mathbf{q}}^i$ is the strong component A^i of Γ , so $\mathcal{N}_{\mathbf{q}}^i$ has only one strong component.

(iv) By parts (i) and (ii), $K_{\mathbf{q}}^i$ and $P_{\mathbf{q}}^i$ do not depend on \mathbf{q} , so all local components of \mathcal{N}^i are homotopic. \square

The strong components are partially ordered by the accessibility relation \dashrightarrow . If we label them so that $A^i \dashrightarrow A^j$ implies $i \geq j$, then the production matrix is block triangular

$$P = \begin{pmatrix} P_{11} & P_{12} & \cdots & P_{1s} \\ 0 & P_{22} & \cdots & P_{2s} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & P_{ss} \end{pmatrix}$$

and the diagonal block P_{ii} is the production matrix of \mathcal{N}^i .

5. HALTING CRITERION

Let \mathcal{N} be an abelian network with total state space Q , and fix a state $\mathbf{q} \in Q$. If case (2) of Lemma 3.1 holds for all inputs $\mathbf{x} \in \mathbb{N}^A$, then we say that \mathcal{N} *halts on all inputs to initial state \mathbf{q}* . If this is the case for all $\mathbf{q} \in Q$, then we say that \mathcal{N} *halts on all inputs*. The main result of this section is Theorem 5.6, which gives an efficient way to decide whether a finite abelian network halts on all inputs to a given initial state.

If $\mathbf{x} \leq \mathbf{y}$, then any legal execution for $\mathbf{x} \cdot \mathbf{q}$ is a legal execution for $\mathbf{y} \cdot \mathbf{q}$. It follows that halting is a monotone property:

$$\text{If } \mathcal{N} \text{ halts on input } \mathbf{y} \cdot \mathbf{q}, \text{ then } \mathcal{N} \text{ halts on all inputs } \mathbf{x} \cdot \mathbf{q} \text{ for } \mathbf{x} \leq \mathbf{y}. \quad (5)$$

Recall the equivalence relation \sim on Q introduced in §4.3.

Lemma 5.1. *If $\mathbf{q}_1 \sim \mathbf{q}_2$, then \mathcal{N} halts on all inputs to initial state \mathbf{q}_1 if and only if \mathcal{N} halts on all inputs to initial state \mathbf{q}_2 .*

Proof. If $\mathbf{q}_1 \sim \mathbf{q}_2$ then $t(\mathbf{y}_1)\mathbf{q}_1 = t(\mathbf{y}_2)\mathbf{q}_2$ for some $\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{N}^A$. Thus it suffices to show for all $\mathbf{y} \in \mathbb{N}^A$ and all $\mathbf{q} \in Q$ that \mathcal{N} halts on all inputs to initial state \mathbf{q} if and only if \mathcal{N} halts on all inputs to initial state $t(\mathbf{y})\mathbf{q}$.

Let $\mathbf{y} \triangleright \mathbf{q} = \mathbf{z}.\mathbf{r}$. Then $\mathbf{r} = t(\mathbf{y})\mathbf{q}$. For any $\mathbf{x} \in \mathbb{N}^A$, we have

$$(\mathbf{x} + \mathbf{y}) \triangleright \mathbf{q} = \mathbf{x} \triangleright (\mathbf{y} \triangleright \mathbf{q}) = \mathbf{x} \triangleright (\mathbf{z}.\mathbf{r})$$

so both $(\mathbf{x} + \mathbf{y}).\mathbf{q}$ and $(\mathbf{x} + \mathbf{z}).\mathbf{r}$ have legal executions resulting in $(\mathbf{x} + \mathbf{y}) \triangleright \mathbf{q}$. Thus \mathcal{N} halts on input $(\mathbf{x} + \mathbf{y}).\mathbf{q}$ if and only if \mathcal{N} halts on input $(\mathbf{x} + \mathbf{z}).\mathbf{r}$. By monotonicity (5), it follows that \mathcal{N} halts on all inputs to initial state \mathbf{q} if and only if \mathcal{N} halts on all inputs to initial state \mathbf{r} . \square

Definition 5.2. A state $\mathbf{x}.\mathbf{q}$ is an *amplifier* if $\mathbf{x} \in \mathbb{N}^A$ and there exists a nonempty legal execution w from $\mathbf{x}.\mathbf{q}$ such that $\pi_w(\mathbf{x}.\mathbf{q}) = \mathbf{y}.\mathbf{q}$ for some $\mathbf{y} \geq \mathbf{x}$.

Definition 5.3. A state $\mathbf{x}.\mathbf{q}$ is a *strong amplifier* if $\mathbf{x} \in \mathbb{N}^A - \{\mathbf{0}\}$ and $\mathbf{x} \triangleright \mathbf{q} = \mathbf{y}.\mathbf{q}$ for some $\mathbf{y} \geq \mathbf{x}$.

In words, a strong amplifier is a pair $\mathbf{x}.\mathbf{q}$ with the property that after processing all letters once, the network has returned to the same state \mathbf{q} with at least as many letters of each type as before.

Example. In the sandpile network $\text{Sand}(G)$ of an undirected graph with no sink, let $\mathbf{x} = (d_v)_{v \in V}$ be the configuration where each vertex has the same number of letters (“chips”) as its degree. For any initial state \mathbf{q} , processing all letters once causes each vertex to topple once, so that each vertex v receives one letter from each of its d_v neighbors. Hence $\mathbf{x} \triangleright \mathbf{q} = \mathbf{x}.\mathbf{q}$, and $\mathbf{x}.\mathbf{q}$ is a strong amplifier.

Lemma 5.4. *The following are equivalent for a finite abelian network \mathcal{N} .*

- (1) \mathcal{N} has an amplifier.
- (2) \mathcal{N} has a strong amplifier.
- (3) \mathcal{N} fails to halt on some input.

Proof. If \mathcal{N} has an amplifier $\mathbf{x}.\mathbf{q}$, then there is a legal execution $\pi_{\mathbf{u}}(\mathbf{x}.\mathbf{q}) = \mathbf{y}.\mathbf{q}$ for some $\mathbf{y} \geq \mathbf{x}$ and $\mathbf{u} \in \mathbb{N}^A - \{\mathbf{0}\}$. Then $\mathbf{u} \triangleright \mathbf{q} = \pi_{\mathbf{u}}(\mathbf{u}.\mathbf{q}) = (\mathbf{u} + \mathbf{y} - \mathbf{x}).\mathbf{q}$, where the second equality follows from Lemma 3.2(ii). Therefore $\mathbf{u}.\mathbf{q}$ is a strong amplifier, which shows that (1) \Rightarrow (2).

Next if \mathcal{N} has a strong amplifier $\mathbf{u}.\mathbf{q}$, then there is a legal execution w with $|w| = \mathbf{u} \in \mathbb{N}^A - \{\mathbf{0}\}$ starting with $\mathbf{u}.\mathbf{q}$ and ending with $(\mathbf{u} + \mathbf{v}).\mathbf{q}$ for some $\mathbf{v} \geq \mathbf{0}$. Then for any $n \geq 0$ the same w is a legal execution starting with $(\mathbf{u} + n\mathbf{v}).\mathbf{q}$ and ending with $(\mathbf{u} + (n+1)\mathbf{v}).\mathbf{q}$. Hence w^n is a legal execution starting from $\mathbf{u}.\mathbf{q}$ for all $n \geq 0$. Since there exist arbitrarily long legal executions, we conclude from Lemma 3.1 that \mathcal{N} does not halt on input $\mathbf{u}.\mathbf{q}$, which shows that (2) \Rightarrow (3).

Lastly, suppose that \mathcal{N} fails to halt on some input $\mathbf{y}.\mathbf{q}$. Then there is an infinite word $w_1 w_2 \cdots$ such that $w_1 \cdots w_n$ is a legal execution for all $n \geq 1$. Let $\mathbf{y}_n.\mathbf{q}_n = \pi_{w_1 \cdots w_n}(\mathbf{y}.\mathbf{q})$. Since the total state space Q is finite, there exists $\mathbf{q} \in Q$ such that $\mathbf{q}_n = \mathbf{q}$ for infinitely many n . By Dickson’s Lemma 2.6, there exist indices $j < k$

such that $\mathbf{q}_j = \mathbf{q}_k = \mathbf{q}$ and $\mathbf{y}_j \leq \mathbf{y}_k$. This $\mathbf{y}_j \cdot \mathbf{q}_j$ is an amplifier, which shows that (3) \Rightarrow (1). \square

The next lemma is a variant of Lemma 5.4 with distinguished initial state. Recall from §4.3 the local component $\mathcal{N}_{\mathbf{q}}$ containing state \mathbf{q} .

Lemma 5.5. *The following are equivalent for a finite abelian network \mathcal{N} and state \mathbf{q} .*

- (1) $\mathcal{N}_{\mathbf{q}}$ has an amplifier.
- (2) $\mathcal{N}_{\mathbf{q}}$ has a strong amplifier.
- (3) \mathcal{N} fails to halt on some input to initial state \mathbf{q} .

Proof. By Lemma 5.4 it suffices to show that (3) is equivalent to the statement that $\mathcal{N}_{\mathbf{q}}$ fails to halt on some input. Any execution for $\mathbf{x} \cdot \mathbf{q}$ in \mathcal{N} is also an execution for $\mathbf{x} \cdot \mathbf{q}$ in $\mathcal{N}_{\mathbf{q}}$, so (3) is equivalent to the statement that $\mathcal{N}_{\mathbf{q}}$ fails to halt on some input to initial state \mathbf{q} . By Lemma 5.1, if $\mathcal{N}_{\mathbf{q}}$ fails to halt on some input, then $\mathcal{N}_{\mathbf{q}}$ fails to halt on some input to initial state \mathbf{q} , which completes the proof. \square

Theorem 5.6. (Halting Criterion 1) *A finite abelian network \mathcal{N} halts on all inputs to initial state \mathbf{q} if and only if every eigenvalue of the production matrix $P_{\mathbf{q}}$ has absolute value strictly less than 1.*

Proof. Let $P = P_{\mathbf{q}}$, and let \mathbf{x} and λ be the Perron-Frobenius eigenvector and eigenvalue of P ; if $\lambda = 1$ then we may take \mathbf{x} to have integer entries (Lemma 2.7). By Lemma 5.5 it suffices to show that (1) if $\lambda \geq 1$, then $\mathcal{N}_{\mathbf{q}}$ has an amplifier; and (2) if $\mathcal{N}_{\mathbf{q}}$ has a strong amplifier, then $\lambda \geq 1$.

(1) If $\lambda \geq 1$, then there is a vector $\mathbf{y} \in \mathbb{Q}^A$ such that $\mathbf{x} \leq \mathbf{y} \leq \lambda \mathbf{x}$. Then $P\mathbf{y} \geq P\mathbf{x} = \lambda \mathbf{x} \geq \mathbf{y}$. Choosing $n \geq 1$ such that $n\mathbf{y} \in K$, we have $n\mathbf{y} \triangleright \mathbf{q} = P(n\mathbf{y}) \cdot \mathbf{q}$, so $n\mathbf{y} \cdot \mathbf{q}$ is an amplifier.

(2) If $\mathcal{N}_{\mathbf{q}}$ has a strong amplifier $\mathbf{y} \cdot \mathbf{q}'$, then by Lemma 4.2, $\mathbf{y} \cdot \mathbf{m} \mathbf{q}'$ is also a strong amplifier for any $\mathbf{m} \in \prod_{v \in V} M_v$. In particular, taking $\mathbf{m} = (e_v)_{v \in V}$ we obtain a strong amplifier $\mathbf{y} \cdot \mathbf{r}$ with \mathbf{r} locally recurrent and $\mathbf{r} \sim \mathbf{q}' \sim \mathbf{q}$. By Lemma 4.9, $P_{\mathbf{q}}(\mathbf{y}) = P_{\mathbf{r}}(\mathbf{y}) \geq \mathbf{y}$, which shows that $\lambda \geq 1$. \square

Remark. In the case that \mathcal{N} is locally irreducible, the production matrix $P_{\mathbf{q}}$ does not depend on \mathbf{q} by Lemma 4.9. In this case Theorem 5.6 gives a criterion for \mathcal{N} to halt on all inputs regardless of initial state.

If $P_{\mathbf{q}}$ has Perron-Frobenius eigenvalue $\lambda \geq 1$, then \mathcal{N} runs forever on some inputs to initial state \mathbf{q} . In the next section we examine how to tell whether this is the case for a given input.

6. LAPLACIAN MATRIX; SANDPILIZATION

An abelian network \mathcal{N} is called *locally finite* if the state space and alphabet of each processor are finite. Let \mathcal{N} be a locally finite and locally irreducible abelian network with total alphabet $A = \sqcup A_v$ and total kernel $K = \prod K_v$. Then each K_v has finite index in \mathbb{Z}^{A_v} . For each letter $a \in A_v$, let r_a be the smallest positive integer such that $r_a 1_a \in K_v$.

Denote by D the $A \times A$ diagonal matrix with diagonal entries r_a , and by I the $A \times A$ identity matrix. Let P be the production matrix of \mathcal{N} , which is well defined by Lemma 4.9.

Definition 6.1. The *Laplacian* of \mathcal{N} is the $A \times A$ matrix

$$L = (I - P)D.$$

Lemma 6.2. L has integer entries.

Proof. If $\mathbf{x} \in \mathbb{N}^A$ then $D\mathbf{x} \in K$, so $(D\mathbf{x}) \triangleright \mathbf{q} = (PD\mathbf{x}) \cdot \mathbf{q}$ by the definition of the production matrix P . In particular, $PD\mathbf{x} \in \mathbb{N}^A$. It follows that $L\mathbf{x} \in \mathbb{N}^A$ for all $\mathbf{x} \in \mathbb{N}^A$, so L has integer entries. \square

The class of toppling matrices, defined in §2.3, is one of several extensions of the notion of “positive definite” to non-symmetric matrices. The next lemma encapsulates how toppling matrices arise in our setting.

Lemma 6.3. Let $L = (I - P)D$, where P is a nonnegative $n \times n$ matrix, D is an $n \times n$ positive diagonal matrix, and I is the $n \times n$ identity matrix. Then L is a toppling matrix if and only if the Perron-Frobenius eigenvalue of P is strictly less than 1.

Proof. The off-diagonal entries of L are nonpositive, so Lemma 2.8 applies. Let \mathbf{x} and λ be a Perron-Frobenius eigenvector and eigenvalue of P , and let $\mathbf{y} = D^{-1}\mathbf{x}$. Then $\mathbf{y} > \mathbf{0}$ and

$$L\mathbf{y} = (I - P)\mathbf{x} = (1 - \lambda)\mathbf{x}. \quad (6)$$

If $\lambda < 1$, then $L\mathbf{y} > \mathbf{0}$, so L is a toppling matrix. Conversely, if L is a toppling matrix, then L^{-1} has nonnegative entries. Applying L^{-1} to (6) yields $\mathbf{y} = (1 - \lambda)L^{-1}\mathbf{x}$. Since $L^{-1}\mathbf{x} > \mathbf{0}$ and $\mathbf{y} > \mathbf{0}$, we conclude that $\lambda < 1$. \square

If \mathcal{N} is any finite abelian network and $\mathbf{q} \in Q$, let $L_{\mathbf{q}} = (I - P_{\mathbf{q}})D_{\mathbf{q}}$ be the Laplacian of the local component $\mathcal{N}_{\mathbf{q}}$. Using Lemma 6.3, we can rephrase Theorem 5.6 in terms of the Laplacian as follows. The case of a toppling network is due to Gabrielov [Gab94].

Corollary 6.4. (Halting Criterion 2) *A finite abelian network \mathcal{N} halts on all inputs to initial state \mathbf{q} if and only if $L_{\mathbf{q}}$ is a toppling matrix.*

Example. Let \mathcal{N} be a toppling network with three vertices a, b, c and thresholds $r_a = 3$, $r_b = 4$, $r_c = 5$. For the messages passed we have:

$$\begin{aligned} a \text{ topples} &\rightarrow b \text{ receives 2 chips, } c \text{ receives 2 chips.} \\ b \text{ topples} &\rightarrow a \text{ receives 1 chip, } c \text{ receives 2 chips.} \\ c \text{ topples} &\rightarrow a \text{ receives 0 chips, } b \text{ receives 2 chips.} \end{aligned}$$

Note that no vertex is a sink, and vertex a is productive (it “creates” a chip each time it topples). The production matrix and Laplacian of this network are:

$$P = \begin{pmatrix} 0 & \frac{1}{4} & 0 \\ \frac{2}{3} & 0 & \frac{2}{5} \\ \frac{3}{3} & \frac{1}{2} & 0 \end{pmatrix}, \quad L = \begin{pmatrix} 3 & -1 & 0 \\ -2 & 4 & -2 \\ -2 & -2 & 5 \end{pmatrix}.$$

Recall the least action principle [BL13, Lemma 4.3], which says that if w is a legal execution for \mathbf{x}, \mathbf{r} and w' is a complete execution for \mathbf{x}, \mathbf{r} , then $|w| \leq |w'|$. If \mathcal{N} halts on input \mathbf{x}, \mathbf{r} , then the *odometer* of \mathbf{x}, \mathbf{r} is defined as $[\mathbf{x}, \mathbf{r}] = |w|$, where w is any complete legal execution for \mathbf{x}, \mathbf{r} .

Lemma 6.7. *Let \mathcal{N} be a locally irreducible finite abelian network. Suppose that α, \mathbf{q} is a strong amplifier for \mathcal{N} . If \mathbf{r} is locally recurrent and \mathcal{N} halts on input \mathbf{x}, \mathbf{r} , then any legal execution w for \mathbf{x}, \mathbf{r} satisfies $|w|_a < \alpha_a$ for some $a \in A$.*

Proof. Suppose for a contradiction that \mathcal{N} halts on input \mathbf{x}, \mathbf{r} but $[\mathbf{x}, \mathbf{r}] \geq \alpha$. Write $\mathbf{u} = [\mathbf{x}, \mathbf{r}] = \alpha + \mathbf{v}$ for some $\mathbf{v} \geq \mathbf{0}$. Write $\pi_{\mathbf{v}}(\mathbf{x}, \mathbf{r}) = \mathbf{y}, \mathbf{s}$ where $\mathbf{s} = t(\mathbf{v})\mathbf{r}$. We will show that $\mathbf{y} \leq \mathbf{0}$, so there is a complete execution w' for \mathbf{x}, \mathbf{r} with $|w'| = \mathbf{v}$. However, by the definition of the odometer there is also a legal execution w for \mathbf{x}, \mathbf{r} with $|w| = \mathbf{u}$. This contradicts the least action principle since $\mathbf{u} - \mathbf{v} = \alpha \in \mathbb{N}^A - \{\mathbf{0}\}$.

Since α, \mathbf{q} is a strong amplifier we have $\alpha \triangleright \mathbf{q} = \beta, \mathbf{q}$ for some $\beta \geq \alpha$. Now by Lemma 2.3(1) (which applies to each local action of M_v on Q_v , as \mathcal{N} is finite and locally irreducible) since \mathbf{r} is locally recurrent we have $\mathbf{r} = \mathbf{m}\mathbf{q}$ for some $\mathbf{m} \in \prod M_v$. So $\mathbf{s} = (t(\mathbf{v})\mathbf{m})\mathbf{q}$, and by Lemma 4.2 it follows that $\alpha \triangleright \mathbf{s} = \beta, \mathbf{s}$. By the definition of the odometer, no messages remain after executing w , so $\pi_{\mathbf{u}}(\mathbf{x}, \mathbf{r}) = \mathbf{0}, \mathbf{s}'$ for some state \mathbf{s}' . Using $\pi_{\mathbf{u}} = \pi_{\alpha} \circ \pi_{\mathbf{v}}$, we have

$$\mathbf{0}, \mathbf{s}' = \pi_{\alpha}(\mathbf{y}, \mathbf{s}) = (\mathbf{y} + \beta - \alpha), \mathbf{s}$$

where the last equality uses Lemma 3.2(ii). Hence $\mathbf{y} = \alpha - \beta \leq \mathbf{0}$ as desired. \square

In the special case of a locally recurrent toppling network $\text{Topp}(L)$, if $\mathbf{y} \in \mathbb{N}^V - \{\mathbf{0}\}$ and $L\mathbf{y} \leq \mathbf{0}$ then $(D\mathbf{y}), \mathbf{0}$ is a strong amplifier. In this case Lemma 6.7 implies that if for a particular input \mathbf{x}, \mathbf{r} there is a legal execution in which each vertex v processes at least $L_{vv}\mathbf{y}_v$ letters, then $\text{Topp}(L)$ does not halt on input \mathbf{x}, \mathbf{r} . In particular, we recover the criterion of Björner and Lovász [BL92, Prop. 4.4]: for the sandpile network $\text{Sand}(G)$ on a directed graph G , if $\mathbf{y} \in \mathbb{N}^V - \{\mathbf{0}\}$ and $L\mathbf{y} = \mathbf{0}$ and each vertex v topples at least \mathbf{y}_v times, then toppling persists forever. When G is undirected (or Eulerian directed) we can take $\mathbf{y} = \mathbf{1}$ and we recover the criterion of Tardos [Tar88, Lemma 4]: if each vertex topples at least once, then toppling persists forever.

7. CONCLUDING REMARKS

We indicate here a few directions for further research on abelian networks.

Halting problem for a given input. Theorem 5.6 gives a polynomial time algorithm to check whether a finite abelian network \mathcal{N} halts on all inputs. Under what conditions is there an efficient algorithm to check whether \mathcal{N} halts on a *given* input $\mathbf{x}_0, \mathbf{q}_0$?

An inefficient algorithm runs as follows. Let $\mathbf{x}_n, \mathbf{q}_n = \mathbf{x}_{n-1} \triangleright \mathbf{q}_{n-1}$ for $n \geq 1$. By Dickson's Lemma 2.6 there exist $m < n$ such that $\mathbf{q}_m = \mathbf{q}_n$ and $\mathbf{x}_m \leq \mathbf{x}_n$. Each time we generate a new state $\mathbf{x}_n, \mathbf{q}_n$, exhaustively check for such an $m < n$. When

we find one, if $\mathbf{x}_n = \mathbf{0}$ then \mathcal{N} has already halted, and if $\mathbf{x}_n \neq \mathbf{0}$ then \mathcal{N} will never halt.

Bounds obtained from Dickson’s Lemma grow very quickly [FFSS11]. Lemma 6.7 suggests a possibly more efficient approach. Given a strong amplifier $\alpha.\mathbf{q}$, what is a bound for the time it takes for \mathcal{N} either to halt or be certified by Lemma 6.7 to run forever? With such a bound in hand, there remains the question of which abelian networks have small (i.e. polynomial in the size of description of \mathcal{N}) strong amplifiers.

Infinite abelian networks. Questions about the recurrence or transience of rotor walk [LL09, AH11a, AH11b, FGLP14, FLP14] and the explosiveness of sandpiles [FLP10] are cases of the halting problem for spatially infinite abelian networks. In this setting, “halting” means that *each processor* processes only finitely many letters, even though the total number of letters processed may be infinite. Generalizing the least action principle to infinite executions, along the lines of [FMR09], may be useful in approaching these and related questions.

Among the many hard questions in this area, let us single out one. Suppose \mathcal{N} is a toppling network whose underlying graph is the square grid \mathbb{Z}^2 . Let \mathbf{q} be an initial state, and suppose that \mathcal{N} and \mathbf{q} are periodic in the sense that there is a full rank sublattice $\Lambda \subset \mathbb{Z}^2$ such that \mathcal{P}_v and \mathbf{q}_v depend only on $v + \Lambda$. Given the finite data of \mathcal{P}_v and \mathbf{q}_v on a fundamental domain, is it decidable whether there exists an input \mathbf{x} with finite support such that \mathcal{N} does not halt on $\mathbf{x}.\mathbf{q}$? Cairns [Cai15], using sandpile circuits designed by Moore and Nilsson [MN99], has shown that the analogous problem in \mathbb{Z}^3 is undecidable.

Homotopy via embedding. Does Lemma 4.13 have a converse? Given irreducible, strongly connected abelian networks \mathcal{N}_1 and \mathcal{N}_2 with $\mathcal{N}_1 \approx \mathcal{N}_2$ is there an irreducible \mathcal{N} such that \mathcal{N}_1 and \mathcal{N}_2 are local components of a strong component of \mathcal{N} ?

Abelian networks with coefficients. We can define an abelian network purely in terms of monoids and without any reference to automata. The free commutative monoid \mathbb{N}^A played an important role in our theory. In particular, we used heavily the fact that \mathbb{N}^A is cancellative (for instance in the proof of Lemma 4.9). What happens if we replace \mathbb{N}^A by a different monoid?

To make this question more precise, suppose M and M' are commutative monoids, written additively. Define an **action of M on Q metered by M'** as a monoid action $\nu : M \times Q \rightarrow Q$ together with a map

$$\mu : M \times Q \rightarrow M'$$

satisfying $\mu(\mathbf{0}, \mathbf{q}) = \mathbf{0}'$ and

$$\mu(\mathbf{x} + \mathbf{y}, \mathbf{q}) = \mu(\mathbf{x}, \nu(\mathbf{y}, \mathbf{q})) + \mu(\mathbf{y}, \mathbf{q}) \tag{7}$$

for all $\mathbf{x}, \mathbf{y} \in M$ and all $\mathbf{q} \in Q$. The interpretation is that $\mu(\mathbf{x}, \mathbf{q})$ measures the “cost” (or “byproduct”) of \mathbf{x} acting on \mathbf{q} , and that costs are additive.

An example of a metered action is the local action \triangleright of an abelian network (§4.1): we take $M = M' = \mathbb{N}^A$ with μ and ν defined by

$$\mathbf{x} \triangleright \mathbf{q} = \mu(\mathbf{x}, \mathbf{q}) \cdot \nu(\mathbf{x}, \mathbf{q}).$$

The byproduct of \mathbf{x} acting on \mathbf{q} is that some messages are passed, namely $\mu(\mathbf{x}, \mathbf{q})$.

Let us define an *abstract abelian network* on a directed graph $G = (V, E)$ as a collection of 4-tuples (M_v, Q_v, μ_v, ν_v) indexed by $v \in V$, such that M_v is a commutative monoid, Q_v is a set, and (μ_v, ν_v) is an action of M_v on Q_v metered by

$$\prod_{(v,u) \in E} M_u.$$

As a special case, fix a commutative monoid C and an alphabet $A = \sqcup_{v \in V} A_v$. An *abelian network with coefficients in C* is an abstract abelian network with $M_v = C^{A_v}$ for all $v \in V$.

It would be interesting to compare the computational power of such networks for different monoids C . For example, taking $C = \mathbb{Z}$ we obtain the class of locally recurrent abelian networks (i.e., those satisfying $Q_v = e_v Q_v$ for all $v \in V$). Taking $C = \mathbb{R}_+$ gives a class of networks with continuous input, which includes the abelian avalanche model of [Gab93] and the divisible sandpile of [LP09]. The latter computes the linear program relaxation of the integer programs that sandpiles compute [BL13, Remark 4.9]. Other \mathbb{R}_+ -networks (analogous to the oil-and-water model of [BL13]) should compute the linear programs of [Tse90]. What about $C = \mathbb{Z}/p\mathbb{Z}$?

Let us point out that the definition (7) of a metered action makes sense for arbitrary monoids M and M' , which allows us to define networks with coefficients in an arbitrary monoid C . Are there interesting examples with C noncommutative?

ACKNOWLEDGMENTS

This research was supported by an NSF postdoctoral fellowship and NSF grants DMS-1105960 and DMS-1243606, and by the UROP and SPUR programs at MIT.

REFERENCES

- [AH11a] Omer Angel and Alexander E. Holroyd, Rotor walks on general trees. *SIAM J. Discrete Math.* **25**(1):423–446, 2011. [arXiv:1009.4802](#)
- [AH11b] Omer Angel and Alexander E. Holroyd, Recurrent rotor-router configurations. 2011. [arXiv:1101.2484](#)
- [Ash87] Jonathan Ashley, On the Perron-Frobenius eigenvector for nonnegative integral matrices whose largest eigenvalue is integral, *Lin. Alg. Appl.* **94**:103–108, 1987.
- [BS15] Josef Berger and Helmut Schwichtenberg, A bound for Dickson’s lemma. [arXiv:1503.03325](#)
- [BLS91] Anders Björner, László Lovász and Peter Shor, Chip-firing games on graphs, *European J. Combin.* **12**(4):283–291, 1991.
- [BL92] Anders Björner and László Lovász, Chip-firing games on directed graphs, *J. Algebraic Combin.* **1**(4):305–328, 1992.
- [BL13] Benjamin Bond and Lionel Levine, Abelian networks I. Foundations and examples. [arXiv:1309.3445](#)

- [BL14] Benjamin Bond and Lionel Levine, Abelian networks III. The critical group. [arXiv:1409.0170](#)
- [Cai15] Hannah Cairns, Some halting problems for abelian sandpiles are undecidable in dimension three. [arXiv:1508.00161](#)
- [DdJ98] Wolfram Decker and Theo de Jong, “Gröbner bases and invariant theory.” In Buchberger, B., Winkler, F. eds, *Gröbner Bases and Applications*, LNS 251, pp. 61–89, 1998.
- [DFR05] Emeric Deutsch, Luca Ferrari and Simone Rinaldi, Production matrices, *Adv. Appl. Math.* **34**(1):101–122, 2005.
- [Dha99] Deepak Dhar, The abelian sandpile and related models, *Physica A* **263**:4–25, 1999. [arXiv:cond-mat/9808047](#)
- [Dic13] Leonard E. Dickson, Finiteness of the odd perfect and primitive abundant numbers with distinct prime factors, *Amer. J. Math.* **35**(4):413–422, 1913.
- [Eis95] David Eisenbud, *Commutative Algebra With A View Toward Algebraic Geometry*, Springer, 1995.
- [FL15] Matthew Farrell and Lionel Levine, CoEulerian graphs. [arXiv:1502.04690](#)
- [FLP10] Anne Fey, Lionel Levine and Yuval Peres, Growth rates and explosions in sandpiles, *J. Stat. Phys.* 138:143–159, 2010. [arXiv:0901.3805](#)
- [FMR09] Anne Fey, Ronald Meester, and Frank Redig, Stabilizability and percolation in the infinite volume sandpile model, *Ann. Probab.* **37**(2):654–675, 2009. [arXiv:0710.0939](#)
- [FK62] M. Fiedler and V. Ptak, On matrices with non-positive off-diagonal elements and positive principal minors, *Czechoslovak Math. J.* **12**:382–400, 1962.
- [FFSS11] Diego Figueira, Santiago Figueira, Sylvain Schmitz, and Philippe Schnoebelen, Ackermannian and primitive-recursive bounds with Dickson’s lemma, *26th Annual IEEE Symposium on Logic in Computer Science*, 2011. [arXiv:1007.2989](#)
- [FGLP14] Laura Florescu, Shirshendu Ganguly, Lionel Levine and Yuval Peres, Escape rates for rotor walks in \mathbb{Z}^d , *SIAM Journal on Discrete Mathematics* **28**(1):323–334, 2014. [arXiv:1301.3521](#).
- [FLP14] Laura Florescu, Lionel Levine and Yuval Peres, The range of a rotor walk. [arXiv:1408.5533](#)
- [Gab93] Andrei Gabrielov, Abelian avalanches and Tutte polynomials, *Physica A* **195**:253–274, 1993.
- [Gab94] Andrei Gabrielov, Asymmetric abelian avalanches and sandpiles. Preprint, 1994. <http://www.math.purdue.edu/~agabriel/asym.pdf>
- [Gre51] J. A. Green, On the structure of semigroups, *Ann. of Math.* 54:163–172, 1951.
- [Gri01] Pierre A. Grillet, *Commutative semigroups*, Kluwer Academic Publishers, 2001.
- [Gri07] Pierre A. Grillet, Commutative actions, *Acta Sci. Math. (Szeged)* **73**:91–112, 2007.
- [HJ90] Roger A. Horn and Charles R. Johnson, *Matrix Analysis*, Cambridge Univ. Press, 1990.
- [LL09] Itamar Landau and Lionel Levine, The rotor-router model on regular trees, *J. Combin. Theory A* **116**: 421–433, 2009. [arXiv:0705.1562](#)
- [LP09] Lionel Levine and Yuval Peres, Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile, *Potential Anal.* **30**:1–27, 2009. [arXiv:0704.0688](#)
- [MN99] Cristopher Moore and Martin Nilsson. The computational complexity of sandpiles. *J. Stat. Phys.* **96**:205–224, 1999.
- [PS04] Alexander Postnikov and Boris Shapiro, Trees, parking functions, syzygies, and deformations of monomial ideals. *Trans. Amer. Math. Soc.* **356**(8):3109–3142, 2004. [arXiv:math.CO/0301110](#)
- [Sch57] Marcel-Paul Schützenberger, \overline{D} représentation des demi-groupes, *C. R. Acad. Sci. Paris* 244:1994–96, 1957.
- [Ste10] Benjamin Steinberg, A theory of transformation monoids: combinatorics and representation theory, *Electr. J. Combin.* **17**:R164, 2010. [arXiv:1004.2982](#)
- [Tar88] Gábor Tardos, Polynomial bound for a chip firing game on graphs, *SIAM J. Disc. Math.* **1**(3):1988.

[Tse90] Paul Tseng, Distributed computation for linear programming problems satisfying a certain diagonal dominance condition, *Mathematics of Operations Research* **15**(1):33–48, 1990.

BENJAMIN BOND, DEPARTMENT OF MATHEMATICS, STANFORD UNIVERSITY, STANFORD, CALIFORNIA 94305. <http://stanford.edu/~benbond>

LIONEL LEVINE, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14853. <http://www.math.cornell.edu/~levine>