

*Divisors and Sandpiles:
An Introduction to Chip-Firing*

Scott Corry
David Perkinson

This is a preliminary version of the book *Divisors and Sandpiles* published by the American Mathematical Society (AMS). This preliminary version is made available with the permission of the AMS and may not be changed, edited, or reposted at any other website without explicit written permission from the author and the AMS.

Divisors and Sandpiles

Scott Corry

David Perkinson

DEPARTMENT OF MATHEMATICS, LAWRENCE UNIVERSITY, APPLETON, WISCONSIN 54911

E-mail address: `corrys@lawrence.edu`

DEPARTMENT OF MATHEMATICS, REED COLLEGE, PORTLAND, OR 97202

E-mail address: `davidp@reed.edu`

2010 *Mathematics Subject Classification.* Primary 05C25

ABSTRACT.

To Madera and Diane

Contents

Preface	1
Part 1. Divisors	
Chapter 1. The dollar game	7
1.1. An initial game	7
1.2. Formal definitions	9
1.3. The Picard and Jacobian groups	14
Notes	16
Problems for Chapter 1	17
Chapter 2. The Laplacian	19
2.1. The discrete Laplacian	19
2.2. Configurations and the reduced Laplacian	24
2.3. Complete linear systems and convex polytopes	28
2.4. Structure of the Picard group	31
Notes	39
Problems for Chapter 2	40
Chapter 3. Algorithms for winning	43
3.1. Greed	43
3.2. q -reduced divisors	46
3.3. Superstable configurations	50
3.4. Dhar's algorithm and efficient implementation	50
3.5. The Abel-Jacobi map	54
Notes	58
Problems for Chapter 3	59

Chapter 4. Acyclic orientations	61
4.1. Orientations and maximal unwinnables	61
4.2. Dhar's algorithm revisited	63
Notes	67
Problems for Chapter 4	68
Chapter 5. Riemann-Roch	69
5.1. The rank function	69
5.2. Riemann-Roch for graphs	71
5.3. The analogy with Riemann surfaces	74
5.4. Alive divisors and stability	82
Notes	84
Problems for Chapter 5	86
Part 2. Sandpiles	
Chapter 6. The sandpile group	91
6.1. A first example	91
6.2. Directed graphs	96
6.3. Sandpile graphs	97
6.4. The reduced Laplacian	101
6.5. Recurrent sandpiles	104
6.6. Images of sandpiles on grid graphs	108
Notes	113
Problems for Chapter 6	114
Chapter 7. Burning and duality	119
7.1. Burning sandpiles	120
7.2. Existence and uniqueness	121
7.3. Superstables and recurrents	123
7.4. Forbidden subconfigurations	126
7.5. Dhar's burning algorithm for recurrents.	127
Notes	128
Problems for Chapter 7	129
Chapter 8. Threshold density	131
8.1. Markov Chains	132
8.2. The fixed-energy sandpile	141
8.3. The threshold density theorem	151
Notes	159
Problems for Chapter 8	160

Part 3. Topics

Chapter 9. Trees	165
9.1. The matrix-tree theorem	166
9.2. Consequences of the matrix-tree theorem	172
9.3. Tree bijections	175
Notes	190
Problems for Chapter 9	191
Chapter 10. Harmonic morphisms	195
10.1. Morphisms between graphs	195
10.2. Branched coverings of Riemann surfaces	204
10.3. Household-solutions to the dollar game	206
Notes	211
Problems for Chapter 10	212
Chapter 11. Divisors on complete graphs	215
11.1. Parking functions	215
11.2. Computing ranks on complete graphs	218
Problems for Chapter 11	227
Chapter 12. More about sandpiles	229
12.1. Changing the sink	229
12.2. Minimal number of generators for $\mathcal{S}(G)$	232
12.3. M -matrices	238
12.4. Self-organized criticality	240
Problems for Chapter 12	243
Chapter 13. Cycles and cuts	245
13.1. Cycles, cuts, and the sandpile group	245
13.2. Planar duality	250
Problems for Chapter 13	255
Chapter 14. Matroids and the Tutte polynomial	257
14.1. Matroids	257
14.2. The Tutte polynomial	259
14.3. 2-isomorphisms	262
14.4. Merino's Theorem	263
14.5. The Tutte polynomials of complete graphs	266
14.6. The h -vector conjecture	269
Notes	273
Problems for Chapter 14	274

Chapter 15. Higher dimensions	279
15.1. Simplicial homology	279
15.2. Higher-dimensional critical groups	287
15.3. Simplicial spanning trees	288
15.4. Firing rules for faces	292
Notes	296
Problems for Chapter 15	297
Appendix	
Appendix A.	303
A.1. Undirected multigraphs	303
A.2. Directed multigraphs	307
Appendix B.	311
B.1. Monoids, groups, rings, and fields	311
B.2. Modules	312
Glossary of Symbols	317
Bibliography	321

Preface

This book discusses the combinatorial theory of chip-firing on finite graphs, a subject that has its main sources in algebraic and arithmetic geometry on the one hand and statistical physics and probabilistic models for dispersion on the other. We have structured the text to reflect these two different motivations, with Part 1 devoted to the divisor theory of graphs (a discrete version of the algebraic geometry of Riemann surfaces) and Part 2 devoted to the abelian sandpile model (a toy model of a slowly-driven dissipative physical system). The fact that these seemingly different stories are in some sense two sides of the same coin is one of the beautiful features of the subject.

To provide maximal coherence, each of the first two parts focuses on a central result: Part 1 presents a quick, elegant, and self-contained route to M. Baker and S. Norine's Riemann-Roch theorem for graphs, while Part 2 does the same for L. Levine's threshold density theorem concerning the fixed-energy sandpile Markov chain. In the exposition of these theorems there are many tempting tangents, and we have collected our favorite tangential topics in Part 3. For instructors, the ability to include these topics should provide flexibility in designing a course, and the reader should feel free to pursue them at any time. As an example, a reader wanting an introduction to simplicial homology or matroids could turn to those chapters immediately. Of course, the choice of topics included in Part 3 certainly reflects the biases and expertise of the authors, and we have not attempted to be encyclopedic in our coverage. In particular, we have omitted a thorough discussion of self-organized criticality and statistical physics, the relationship to arithmetic geometry, pattern formation in the abelian sandpile model, and connections to tropical curves.

The audience we had in mind while writing this book was advanced undergraduate mathematics majors. Indeed, both authors teach at undergraduate liberal arts colleges in the US and have used this material multiple times for courses. In addition, the second author has used this subject matter in courses at the African Institute for Mathematical Sciences (AIMS) in South Africa, Ghana, and Cameroon.

The only prerequisites for reading the text are first courses in linear and abstract algebra, although Chapter 8 assumes rudimentary knowledge of discrete probability theory, as can be easily obtained online or through a standard text such as [82]. In fact, one of the charms of this subject is that much of it can be meaningfully and entertainingly presented even to middle school students! On the other hand, this text is also suitable for graduate students and researchers wanting an introduction to sandpiles or the divisor theory of graphs. We encourage all readers to supplement their reading with computer experimentation. A good option is the free open-source mathematics software system SageMath ([33], [77]), which has extensive built-in support for divisors and sandpiles.

In addition to presenting the combinatorial theory of chip-firing, we have the ulterior motive of introducing some generally useful mathematics that sits just at the edge of the standard undergraduate curriculum. For instance, most undergraduate mathematics majors encounter the structure theorem for finitely generated abelian groups in their abstract algebra course, but it is rare for them to work with the Smith normal form, which arises naturally in the computation of sandpile groups. In a similar way, the following topics all have connections to chip-firing, and the reader can find an introduction to them in the text: the matrix-tree theorem; Markov chains; simplicial homology; M -matrices; cycle and cut spaces for graphs; matroid theory; the Tutte polynomial. Further, for students intending to study algebraic geometry, learning the divisor theory of graphs is a fantastic stepping stone toward the theory of algebraic curves. In fact, some of the proofs of the corresponding results are the same in both the combinatorial and the algebro-geometric setting (e.g., Clifford's theorem).

We now provide a more detailed description of the text. A reader wanting to start with sandpile theory instead of divisors should be able to begin with Part 2, looking back to Part 1 when needed for vocabulary.

Part 1:

- Chapter 1 introduces the dollar game in which vertices of a finite graph trade dollars across the edges in an effort to eliminate debt. This simple game provides a concrete and tactile setting for the subject of divisors on graphs, and the chapter ends with a list of motivating questions.
- Chapter 2 introduces the Laplacian operator, which is really our central object of study. We reinterpret the dollar game in terms of the Laplacian, and introduce Smith normal form as a computational tool.
- Chapter 3 discusses algorithms for winning the dollar game or certifying unwinnability. The first section presents a greedy algorithm for winnability, while later sections describe the more sophisticated Dhar's algorithm along with the attendant concepts of q -reduced divisors and superstable configurations.
- Chapter 4 continues the study of Dhar's algorithm, using it to establish a crucial bijection between acyclic orientations of the graph and maximal unwinnable divisors.
- Chapter 5 draws on all of the previous chapters to establish the Riemann-Roch theorem for graphs. Section 5.3 lays out the striking analogy of the divisor theory of graphs with the corresponding theory for Riemann surfaces.

Part 2:

- In Chapter 6 we begin anew by imagining grains of sand stacked on the vertices of a graph, with new grains arriving slowly at random locations. When too many grains accumulate at a vertex, the vertex becomes unstable and topples, sending grains along edges to neighboring vertices. An avalanche may ensue as new vertices become unstable due to the toppling of other vertices. One vertex is designated as the sink, having the capacity to absorb an unlimited amount of sand—this ensures that every avalanche eventually comes to an end. As in Part 1, the Laplacian operator plays the central role in this story, but now our attention turns to the recurrent sandpiles: those configurations of sand on non-sink vertices that occur over and over again as we drop sand on the graph.
- Chapter 7 presents the burning algorithm for determining whether a sandpile is recurrent. As a consequence, we establish a duality between recurrent sandpiles and superstable configurations, thereby revealing one of many connections between the sandpile theory of Part 2 and the divisor theory of Part 1.
- Chapter 8 provides a brief introduction to Markov chains and then presents the threshold density theorem for the fixed-energy sandpile. In this model there is no sink vertex, so an avalanche might continue forever. Starting with a highly stable state, we imagine dropping additional grains of sand on the graph just as before, allowing time for the avalanches to occur in between successive grains. How much sand will be on the graph when it first passes the critical threshold where the avalanche never ends? The threshold density theorem provides a precise answer to this question in the limit where the starting state becomes “infinitely stable.”

Part 3:

- Chapter 9 contains two proofs of the matrix-tree theorem, which computes the number of spanning trees of a graph as the determinant of the reduced Laplacian matrix. This tree-number is also the number of recurrent sandpiles on the graph. The second section contains several corollaries, while the final two sections discuss tree-bijections and the remarkable rotor-router algorithm, providing an action of the sandpile group on the set of spanning trees.
- Chapter 10 returns to the setting of Part 1 and studies harmonic morphisms between graphs. These are discrete analogues of holomorphic mappings between Riemann surfaces, and we prove a graph-theoretic version of the classical Riemann-Hurwitz formula at the end of the first section, explaining the original result for surfaces in Section 10.2. In the final section, we interpret harmonic morphisms as generalized solutions to the dollar game in which individuals pool their money by forming households.
- Chapter 11 presents two related topics concerning the divisor theory of complete graphs. In the first section, the superstable configurations on a complete graph are shown to be essentially the same as parking functions—certain integer sequences arising frequently in combinatorics, introduced here via a simple story about parking cars. In the second section, we present the Cori-Le Borgne algorithm for computing ranks of divisors on complete graphs, exploiting the connection with parking functions in an appealing way.

- Chapter 12 addresses several additional topics related to sandpiles: the dependence of the sandpile group $\mathcal{S}(G, s)$ on the choice of sink vertex s ; the minimal number of generators for $\mathcal{S}(G, s)$; and a generalization of the sandpile dynamics in which the reduced Laplacian is replaced by a non-singular M -matrix. The chapter concludes with a brief discussion of the concept of self-organized criticality.
- Chapter 13 introduces the algebraic theory of cycles and cuts, and discusses the connection to the sandpile group. As an application, we prove that the sandpile group of a plane graph is isomorphic to the sandpile group of its planar dual.
- Chapter 14 provides a brief introduction to matroids and their Tutte polynomials. As applications, we show that the sandpile group of a graph G depends only on the cycle matroid of G ; prove Merino's Theorem identifying the number of superstable of each degree with the coefficients of the Tutte polynomial; state Stanley's conjecture concerning h -vectors of matroid complexes; and finally present Merino's proof of Stanley's conjecture in the case of cographic matroids.
- Chapter 15 provides a brief introduction to higher-dimensional versions of many of our topics. We begin by introducing the relevant setting of simplicial complexes and their associated homology theory. We then define higher-dimensional critical groups, generalizing the Jacobian of a graph. We define simplicial spanning trees and present a generalized matrix-tree theorem due to A. Duval, C. Klivans, and J. Martin. We conclude with some brief comments about higher-dimensional versions of the dollar game and the sandpile model.

Acknowledgments. We would like to thank: Matt Baker, Lionel Levine, Spencer Backman, Maddie Brant, Melody Chan, Art Duval, Laura Florescu, Munyo Frey-Edwards, Luis David García Puente, Rebecca Garcia, Darij Grinberg, Sam Hopkins, Tessa Hoppenfeld, Caroline Klivans, Jeremy Martin, Ina Metta, Dani Morar, Wesley Pegden, Jacob Perlman, Jamie Pommersheim, Jim Propp, Vic Reiner, Tom Roby, Farbod Shokrieh, Avi Steiner, Irena Swanson, Riley Thornton, John Wilmes, Qiao-yu Yang, Kuai Yu, and the many students with whom we have shared this subject at Lawrence University, Reed College, and the African Institute for Mathematics in South Africa, Ghana, and Cameroon.

Part 1

Divisors

The dollar game

1.1. An initial game

Consider the graph pictured in Figure 1. As suggested by the labels, we think of the vertices V as individuals, the edges E as relationships between individuals, and the entire graph $G = (V, E)$ as a community. The number of edges between two vertices indicates the strength of the relationship between the corresponding individuals, with zero edges indicating no relationship at all. For instance, Willa knows Yoshi twice as well as Xenia; she and Zelda are strangers.

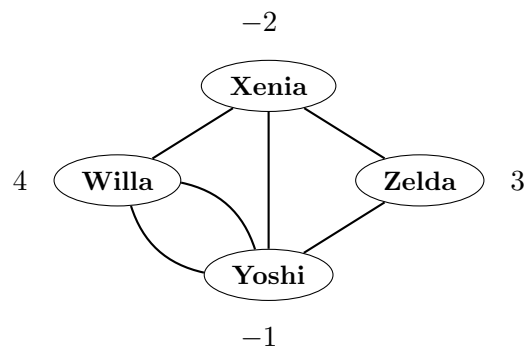


Figure 1. A distribution of wealth in a community.

As in most communities, the individuals represented in G are not equally wealthy. To record their varying degrees of prosperity, we place an integer representing a number of dollars next to each vertex, interpreting negative values as debt¹. We represent such a distribution of dollars as a formal sum of vertices. For example, the distribution of wealth D in Figure 1 is given by $D = 4W - 2X - Y + 3Z$,

¹Debt is owed to an outside entity not represented in the graph—perhaps a credit card company.

using first initials to label vertices. So Willa has \$4, Xenia owes \$2, etc., and the net wealth is $4 - 2 - 1 + 3 = 4$ dollars.

To redistribute wealth, each individual v may choose to lend a dollar along each edge incident to v (possibly going into debt), or else borrow a dollar along each edge (possibly forcing someone else into debt). Note that v never lends along some edges while borrowing along others: she is either in a philanthropic mood (in which case she lends to *all* vertices in proportion to the strength of their bond) or she is feeling needy (in which case she similarly borrows from all vertices). Figure 2 shows a lending move by Willa.




Figure 2. Willa lends to her friends.

The goal of the community is simple: find a sequence of lending/borrowing moves so that everyone is debt-free. This is called the *dollar game on G starting from D* , and if such a sequence exists, the game is said to be *winnable*. After Willa makes the lending move shown in Figure 2, only Xenia is in debt. If Xenia then borrows, the dollar game is won (Figure 3). Note there may be multiple paths to victory in the dollar game and multiple winning states. Can you see this in our running example?

Important: before proceeding, play the dollar games in Exercise 1.1. In each case, you are given an initial distribution of wealth. If possible, win by finding a sequence of lending and borrowing moves resulting in a state in which no vertex has debt.





Figure 3. Xenia borrows, and the dollar game is won.

Exercise 1.1. Can you win the dollar game?



1.2. Formal definitions

We now formalize the ideas presented in the previous section. See Appendix A for a brief review of definitions and terminology from graph theory.

Definition 1.2. A *multigraph* $G = (V, E)$ is a pair consisting of a set of *vertices* V and a multiset of *edges* E comprised of unordered² pairs $\{v, w\}$ of vertices. The prefix “multi-” means that a pair $\{v, w\}$ may occur multiple times in E . For ease of notation, we generally write vw for the edge $\{v, w\}$. A multigraph G is *finite* if both V and E are finite, and it is *connected* if any two vertices are joined by a sequence of edges.

In Part 1 of this book, we will simply use the term *graph* to mean a finite, connected, undirected multigraph without loop edges.

Definition 1.3. A *divisor* on a graph G is an element of the free abelian group on the vertices:

$$\text{Div}(G) = \mathbb{Z}V = \left\{ \sum_{v \in V} D(v)v : D(v) \in \mathbb{Z} \right\}.$$

²Here we are defining an *undirected* multigraph, which provides the setting for Part 1 of this book; in Part 2 we will also consider *directed* multigraphs, where the edge multiset consists of ordered pairs of vertices.

Divisors represent distributions of wealth on G : if $D = \sum_{v \in V} D(v)v \in \text{Div}(G)$, then each vertex/person v has $D(v)$ dollars, interpreting a negative value for $D(v)$ as debt. The net amount of money determined by D is called the *degree* of the divisor.

Definition 1.4. The *degree* of $D = \sum_{v \in V} D(v)v \in \text{Div}(G)$ is the integer

$$\deg(D) = \sum_{v \in V} D(v).$$

The collection of divisors of degree k is denoted $\text{Div}^k(G)$, and the collection of divisors of nonnegative degree is denoted $\text{Div}_+(G)$.

For $v \in V$, the word “degree” has two meanings. If we think of v as a divisor, then its degree is 1. However, as a vertex in the graph, its degree is its number of incident edges. To avoid ambiguity, we use $\deg_G(v)$ to denote the latter.

Definition 1.5. Let $D, D' \in \text{Div}(G)$ and $v \in V$. Then D' is obtained from $D \in \text{Div}(G)$ by a *lending move at v* , denoted $D \xrightarrow{v} D'$, if

$$D' = D - \sum_{vw \in E} (v - w) = D - \deg_G(v)v + \sum_{vw \in E} w.$$

Similarly, D' is obtained from D by a *borrowing move at v* , denoted $D \xrightarrow{-v} D'$, if

$$D' = D + \sum_{vw \in E} (v - w) = D + \deg_G(v)v - \sum_{vw \in E} w.$$

(Note that in the above sums, an edge vw will occur more than once if there are multiple edges between v and w .)

1.2.1. The abelian property and set-lendings. As is easily verified from Definition 1.5, the order of lending and borrowing moves does not matter: if Willa lends and then Xenia lends, the result is the same as if Xenia lent first, then Willa. This is known as the *abelian property* of the dollar game. To state the property formally, let $D \in \text{Div}(G)$ and $v, w \in V$. Then we have a commutative diagram:

$$\begin{array}{ccc} & D & \\ v \swarrow & & \searrow w \\ D' & & D'' \\ w \swarrow & & \nwarrow v \\ & D''' & \end{array}$$

By “commutative” we mean that following the lending moves down either the left or the right side of this diagram leads to the same final divisor, D''' .

In fact, not only may lendings and borrowings be made in any order without affecting the outcome, they may be made *simultaneously*. Again, this follows easily from Definition 1.5.

Definition 1.6. Let $D, D' \in \text{Div}(G)$ and suppose that D' is obtained from D by lending from all of the vertices in $W \subseteq V$. Then D' is the result of a *set-lending* or *set-firing* by W , denoted $D \xrightarrow{W} D'$.

In Figure 4, we lend from the circled set of vertices. Vertices within the circle trade dollars; only vertices incident to an edge that crosses the circle feel a net effect.




Figure 4. Set-lending on the diamond graph.

It is usually best to think in terms of set-lendings when playing the dollar game. Armed with this new set-lending technology, the reader is encouraged to replay the games in Exercise 1.1.

Exercise 1.7. Referring to Figure 1, verify that if Xenia, Yoshi, and Zelda all perform lending moves, the result is the same as if Willa had performed a single borrowing move. In general,


Borrowing from $v \in V$ is the same as set-lending by $V \setminus \{v\}$, and a set-lending by V has no net effect.

Thus, **from now on** we will feel free to use the phrase *lending moves* instead of *lending and borrowing moves*.

1.2.2. Linear equivalence. After a lending move, the distribution of wealth may change, but in some sense, the game has not. We develop language to make this notion precise.

Definition 1.8. Let $D, D' \in \text{Div}(G)$. Then D is *linearly equivalent* to D' if D' may be obtained from D by a sequence of lending moves. In that case, we write $D \sim D'$.

Exercise 1.9. Verify the following equivalence of divisors on the diamond graph:



Exercise 1.10. Check that \sim is an equivalence relation: that is, for all divisors D, D', D'' on G , the following properties hold:

- (1) reflexivity: $D \sim D$,
- (2) symmetry: if $D \sim D'$, then $D' \sim D$,
- (3) transitivity: if $D \sim D'$ and $D' \sim D''$, then $D \sim D''$.


Definition 1.11. The *divisor class* determined by $D \in \text{Div}(G)$ is

$$[D] = \{D' \in \text{Div}(G) : D' \sim D\}.$$

A divisor class may be thought of as a (closed) economy. We start with an initial distribution of wealth D . Through lending and borrowing we arrive at another distribution D' . The collection of all distributions obtainable from D is the divisor class $[D]$. Of course, since lending and borrowing are reversible, we could have just as well started at D' . In other words,

$$D \sim D' \iff [D] = [D'].$$

Example 1.12. Consider the divisor $D = 3v_1 - v_2$ on the segment graph S below:



By lending from v_1 twice, we see that $D \sim v_1 + v_2$. The whole divisor class for D is

$$[D] = \{av_1 + bv_2 : a, b \in \mathbb{Z}, a + b = 2\} = \{D' \in \text{Div}(S) : \deg(D') = 2\}.$$

Any two divisors of degree 2 on S are linearly equivalent.

To win the dollar game, we seek a distribution in which every vertex is debt-free. To state this in the language of divisors, first define

$$D \geq D'$$

for $D, D' \in \text{Div}(G)$, if $D(v) \geq D'(v)$ for all $v \in V$. In particular, we write $D \geq 0$ if $D(v) \geq 0$ for all $v \in V$.

Definition 1.13. A divisor D on G is *effective* if $D \geq 0$. The set of effective divisors on G is denoted $\text{Div}_+(G)$; it is not a subgroup of $\text{Div}(G)$ because it lacks inverses. Such an algebraic structure is called a *commutative monoid*.

We may restate our game as follows:

The dollar game: Is a given divisor linearly equivalent to an effective divisor?

Accordingly, we make the following definition.

Definition 1.14. The divisor $D \in \text{Div}(G)$ is *winnable* if D is linearly equivalent to an effective divisor; otherwise, D is *unwinnable*.

Earlier, you probably noticed that the game in part (c) of Example 1.1 is unwinnable for the following reason.


Proposition 1.15. Let $D, D' \in \text{Div}(G)$. If $D \sim D'$, then $\deg(D) = \deg(D')$.

Proof. Each lending move preserves the total number of dollars on the graph. \square

Corollary 1.16. If $\deg(D) < 0$, then D is unwinnable.

Exercise 1.17. Give an example showing that the converse to Corollary 1.16 is not generally true.

Example 1.18. Consider the divisor $D = 2v_1 + v_2 - 2v_3$ on the triangle graph C_3 :



Borrowing by v_3 reveals that $D \sim v_1$. So D is winnable and $v_1 \in [D]$. One might suspect that the divisor class $[D]$ consists of all divisors of degree 1, but that is not the case. For instance, $v_1 \not\sim v_2$. To see this, note that by Exercise 1.7, any divisor linearly equivalent to v_1 may be obtained through a sequence of moves by the vertices v_1 and v_2 only. Suppose that $D' \sim v_1$ is obtained through such a sequence, and that $D'(v_3) = 0$. Then in our sequence of moves, every lending move by v_1 must be balanced by a borrowing move by v_2 , and vice-versa. In particular, D' results from either (i) lending k times from v_1 and borrowing k times from v_2 , or (ii) borrowing k times from v_1 and lending k times from v_2 . But

$$v_1 \xrightarrow{v_1} -v_1 + v_2 + v_3 \xrightarrow{-v_2} -2v_1 + 3v_2 = v_1 - (3v_1 - 3v_2)$$

and

$$v_1 \xrightarrow{-v_1} 3v_1 - v_2 - v_3 \xrightarrow{v_2} 4v_1 - 3v_2 = v_1 + (3v_1 - 3v_2).$$

Iterating either of these procedures k times shows that

$$D' = v_1 \pm (3kv_1 - 3kv_2) \neq v_2 \quad \text{for any } k.$$

Hence, $v_1 \not\sim v_2$ as claimed. Similarly, $v_1 \not\sim v_3$. Moreover, we have determined part of the divisor class $[D]$:

$$\{D' \in [D] : D'(v_3) = 0\} = \{(1 + 3k)v_1 - 3kv_2 : k \in \mathbb{Z}\}.$$

More generally, a similar argument (cf. Problem 1.1) starting with the assumption that $D'(v_3) = m$ shows that

$$[D] = \{(1 + 3k + m)v_1 - (3k + 2m)v_2 + mv_3 : k, m \in \mathbb{Z}\}.$$

The argument in the previous example was somewhat cumbersome, despite the extreme simplicity of the triangle graph. It would be helpful to develop algebraic tools to clarify the computations and allow for work with more complicated graphs. In the next section, we begin this task by defining an abelian group that is closely related to the dollar game.

1.3. The Picard and Jacobian groups

The sum of $D, F \in \text{Div}(G)$ is defined vertex-wise:

$$D + F = \sum_{v \in V} (D(v) + F(v)) v.$$

Moreover, this sum respects linear equivalence. That is, if $D \sim D'$ and $F \sim F'$, then

$$D + F \sim D' + F'.$$

To see this, combine a lending sequence leading from D to D' with one leading from F to F' . This allows us to turn the set of divisor classes into a group.

Definition 1.19. The *Picard group* of G is the set of linear equivalence classes of divisors

$$\text{Pic}(G) = \text{Div}(G)/\sim$$

with addition

$$[D] + [F] = [D + F].$$

The *Jacobian group* is the subgroup of $\text{Pic}(G)$ consisting of divisor classes of degree 0:

$$\text{Jac}(G) = \text{Div}^0(G)/\sim.$$

Note that the Jacobian group is well-defined since linearly equivalent divisors have the same degree, and in general,

$$\deg(D + F) = \deg(D) + \deg(F).$$

In particular, if two divisors have degree 0, so does their sum.

The following proposition is the first step in analyzing the structure of $\text{Pic}(G)$.

Proposition 1.20. Fix $q \in V$. There is an isomorphism of groups,

$$\begin{aligned} \text{Pic}(G) &\rightarrow \mathbb{Z} \times \text{Jac}(G) \\ [D] &\mapsto (\deg(D), [D - \deg(D)q]). \end{aligned}$$

Proof. Problem 1.4. □

Example 1.21. Consider the triangle graph C_3 from Example 1.18. Set $q = v_3$ to get the isomorphism $\text{Pic}(C_3) \rightarrow \mathbb{Z} \times \text{Jac}(C_3)$ from Proposition 1.20. We have seen that v_1, v_2 , and v_3 are linearly inequivalent divisors on C_3 , which implies that $v_1 - v_3$ and $v_2 - v_3$ are also linearly inequivalent. Hence, $[v_1 - v_3]$ and $[v_2 - v_3]$ are distinct, non-identity elements of $\text{Jac}(G)$. On the other hand, an arbitrary divisor of degree zero on C_3 looks like $av_1 + bv_2 - (a+b)v_3 = a(v_1 - v_3) + b(v_2 - v_3)$, which shows that $[v_1 - v_3]$ and $[v_2 - v_3]$ together generate $\text{Jac}(G)$. Finally,

$$2(v_1 - v_3) \xrightarrow{v_1} v_2 - v_3,$$

and

$$3(v_1 - v_3) \xrightarrow{v_1} v_1 + v_2 - 2v_3 \xrightarrow{-v_3} 0$$

so $2[v_1 - v_3] = [v_2 - v_3]$ and $3[v_1 - v_3] = 0$ in $\text{Jac}(G)$. It follows that $\text{Jac}(G) \simeq \mathbb{Z}_3$ via the map $[v_1 - v_3] \mapsto 1 \pmod{3}$. The inverse of the isomorphism $\text{Pic}(G) \rightarrow \mathbb{Z} \times \mathbb{Z}_3$ is given by

$$(d, a \pmod{3}) \mapsto [a(v_1 - v_3) + dv_3] = [av_1 + (d - a)v_3].$$

(Problem 1.2 asks for a generalization of this argument to the cycle graph C_n .)

So far, we have been asking if a game is winnable. A more refined question is: How many different winning debt-free redistributions of wealth are there?

Definition 1.22. The *complete linear system* of $D \in \text{Div}(G)$ is

$$|D| = \{E \in \text{Div}(G) : E \sim D \text{ and } E \geq 0\}.$$

Thus, $|D|$ consists of all the winning states associated with the dollar game for D . In particular, D is winnable if and only if $|D|$ is nonempty, so we have yet another restatement of our game:

The dollar game: Is the complete linear system of a given divisor nonempty?

There are three divisors in the complete linear system for the divisor on the segment graph S in Example 1.12:

$$\{2v_1, v_1 + v_2, 2v_2\},$$

while the whole class for the divisor is infinite. For the triangle graph C_3 in Example 1.18, the complete linear system of the divisor $D = 2v_1 + v_2 - 2v_3$ is the singleton $|D| = \{v_1\}$. Indeed, v_1, v_2 , and v_3 are the only effective divisors of degree 1, and we have seen that $v_1 \sim D$ is not linearly equivalent to either v_2 or v_3 .

1.3.1. Some initial questions. As a preview of the upcoming chapters, we list here some initial questions about the dollar game:

- (1) Are there efficient algorithms for deciding whether a divisor $D \in \text{Div}(G)$ is winnable, and if so, for finding a winning sequence of lending moves? (Chapter 3)
- (2) Our intuition suggests that if there is enough money in the distribution D , then it should be possible to share the wealth and win the dollar game. But how much is enough? In particular, does there exist a number $W(G)$ (depending in

- a simple way on the graph G), such that every divisor D with $\deg(D) \geq W(G)$ is winnable? (Chapter 4)
- (3) If D is winnable, how many different debt-free winning distributions are there? More precisely, how can we compute the complete linear system $|D|$? Are there better ways of measuring the “size” of this linear system than simply determining its cardinality? (Chapters 2, 5)
- (4) What exactly is the relationship between the dollar game on G and the abelian groups $\text{Pic}(G)$ and $\text{Jac}(G)$? (Chapters 2, 3, 4, and 5)
-

Notes

The dollar game described in this chapter is one of many different “chip-firing games” that have been explored by various authors in the literature. In particular, N. Biggs studied a somewhat different dollar game in [14], while our game is that described by M. Baker and S. Norine in their seminal paper [6].

The algebro-geometric language of divisors and the Jacobian group appears in the context of finite graphs in [2], although our Jacobian group is there referred to as the Picard group. We have adopted the language of divisors, linear systems, and Jacobian/Picard groups as set out systematically in [6].

Problems for Chapter 1

1.1. Generalize the argument in Example 1.18 to show that if $D = 2v_1 + v_2 - 2v_3$ on the triangle graph C_3 , then

$$[D] = \{(1 + 3k + m)v_1 - (3k + 2m)v_2 + mv_3 : k, m \in \mathbb{Z}\}.$$

1.2. Let C_n denote the cycle graph with vertices v_0, \dots, v_{n-1} arranged consecutively. Define divisors $D_0 = 0$ and $D_i = v_i - v_0$ for $i = 1, \dots, n - 1$.

- Directly from set-lendings, show by induction that $k[D_1] = [D_k]$ for $k = 0, 1, \dots, n - 1$ and that $n[D_1] = 0$. Thus, if we take indices modulo n , we have $k[D_1] = [D_k]$ for all $k \in \mathbb{Z}$.
- Consider the homomorphism $\psi: \mathbb{Z}_n \rightarrow \text{Jac}(C_n)$ given by $k \mapsto k[D_1]$ (which is well-defined by part (a)). First show that ψ is surjective. Then show that $\text{Jac}(C_n)$ is isomorphic to \mathbb{Z}_n by showing that ψ is also injective: if $D_k \sim 0$, then $k \equiv 0 \pmod n$.
- Find, without proof, the complete linear system of the divisor $D = -2v_0 + v_1 + 2v_2 + v_3$ on C_4 .

1.3. Let P_n be the *path graph* on n vertices consisting of $n - 1$ edges connected to form a line as in Figure 5.



Figure 5. The path graph of length $n - 1$.

- Given a divisor $D = \sum_{i=1}^n a_i v_i \in \text{Div}(P_n)$, describe a sequence of set-lendings (and borrowings) which show $D \sim \deg(D)v_n$.
- Prove $\text{Pic}(P_n) \simeq \mathbb{Z}$.
- For more of challenge, show that the Picard group of every tree is isomorphic to \mathbb{Z} . (A *tree* is a connected graph without any cycles.)

1.4. Prove Proposition 1.20.

The Laplacian

A central tenet of the dollar game is that vertices share equitably: a lending move sends one dollar along each incident edge. There is an operator called the *Laplacian*, ubiquitous in the natural sciences, which measures the equitability or evenness of a diffusive process. It appears in equations modeling heat and fluid flow, chemical diffusion, properties of electrostatic and gravitational potentials, and quantum processes.

In this chapter, we associate to each graph a discrete version of the Laplacian and use it to reinterpret everything we have done so far. From this new perspective, the Laplacian is our central object of study, while the dollar game and later, in Part 2, the abelian sandpile model are convenient devices for studying its properties. The Laplacian connects our subject with the larger mathematical landscape.

The final section of the chapter discusses the structure theorem for finitely generated abelian groups and concludes with a practical method of computing Picard and Jacobian groups using the Smith normal form of the Laplacian.

2.1. The discrete Laplacian

Consider the space $\mathbb{Z}^V := \{f: V \rightarrow \mathbb{Z}\}$ of \mathbb{Z} -valued functions on the vertices of G .

Definition 2.1. The (*discrete*) *Laplacian operator* on G is the linear mapping $L: \mathbb{Z}^V \rightarrow \mathbb{Z}^V$ defined by

$$L(f)(v) := \sum_{vw \in E} (f(v) - f(w)).$$

We now spend some time showing the relevance of the Laplacian to the dollar game.

In the previous chapter, we introduced the operation of *set-lending*, in which every vertex in a subset $W \subseteq V$ makes a lending move. Of course, we could also allow some vertices in W to lend multiple times, or to borrow instead of lend. In all of these scenarios, the essential information may be encoded in a function $\sigma: V \rightarrow \mathbb{Z}$, where the value of σ at $v \in V$ is the number of times that v lends,

with negative values interpreted as borrowing. This idea motivates the following definition. In it, we also replace the word *lending* with *firing*. The justification for this shift of language is that in Part 2 we will use these same concepts in the context of *sandpiles*, and we would like to have a general terminology to account for either interpretation: lending dollars or toppling sand.

Definition 2.2. A *firing script* is a function $\sigma: V \rightarrow \mathbb{Z}$, and the collection of all firing scripts forms an abelian group, denoted $\mathcal{M}(G)$.

Thus, $\mathcal{M}(G)$ is just another notation for \mathbb{Z}^V , but writing $\sigma \in \mathcal{M}(G)$ emphasizes the specific interpretation of σ as a firing script.

The set-lending (or set-firing) by a subset $W \subseteq V$ corresponds to the characteristic function $\chi_W \in \mathcal{M}(G)$ defined by

$$\chi_W(v) = \begin{cases} 1 & v \in W \\ 0 & v \notin W. \end{cases}$$

(In the sequel, we write χ_v for $\chi_{\{v\}}$.) Given any firing script σ , the result of applying the corresponding collection of lending moves to a divisor D will be the divisor D' given by

$$\begin{aligned} D' &= D - \sum_{v \in V} \sigma(v) \left(\deg_G(v) v - \sum_{vw \in E} w \right) \\ &= D - \sum_{v \in V} \sigma(v) \sum_{vw \in E} (v - w) \\ &= D - \sum_{v \in V} \sum_{vw \in E} (\sigma(v) - \sigma(w)) v \\ &= D - \sum_{v \in V} \left(\deg_G(v) \sigma(v) - \sum_{vw \in E} \sigma(w) \right) v. \end{aligned}$$

Definition 2.3. If $\sigma: V \rightarrow \mathbb{Z}$ is a firing script, then the *divisor of σ* is

$$\operatorname{div}(\sigma) := \sum_{v \in V} \left(\deg_G(v) \sigma(v) - \sum_{vw \in E} \sigma(w) \right) v.$$

In terms of this definition, we see that the effect of implementing a firing script σ is to replace a divisor D by the linearly equivalent divisor $D' = D - \operatorname{div}(\sigma)$. We denote this process by $D \xrightarrow{\sigma} D'$, and refer to it as a *script-firing*.

Exercise 2.4. Show that $\operatorname{div}: \mathcal{M}(G) \rightarrow \operatorname{Div}(G)$ is a group homomorphism. Moreover, show that $\deg(\operatorname{div}(\sigma)) = 0$ for all firing scripts σ .

Divisors of the form $\operatorname{div}(\sigma)$ are called *principal*. By the previous exercise, the set of principal divisors forms a subgroup $\operatorname{Prin}(G) < \operatorname{Div}^0(G)$. Moreover, the divisors obtainable from D by a sequence of borrowing and lending moves are exactly those divisors D' such that $D - D' \in \operatorname{Prin}(G)$. This means that the linear equivalence class of a divisor D is the coset of $\operatorname{Prin}(G)$ represented by D :

$$[D] = D + \operatorname{Prin}(G).$$

Hence, we can express the Picard and Jacobian groups as quotients:

$$\text{Pic}(G) = \text{Div}(G)/\text{Prin}(G), \quad \text{Jac}(G) = \text{Div}^0(G)/\text{Prin}(G).$$

We now explain the sense in which $\text{div}: \mathcal{M}(G) \rightarrow \text{Div}(G)$ and $L: \mathbb{Z}^V \rightarrow \mathbb{Z}^V$ are essentially the same. First note that the mapping $v \mapsto \chi_v$ determines an isomorphism of $\text{Div}(G)$ with \mathbb{Z}^V which fits into the following commutative diagram:

$$\begin{array}{ccc} \mathcal{M}(G) & \xrightarrow{\text{div}} & \text{Div}(G) \\ \parallel & & \downarrow \simeq \\ \mathbb{Z}^V & \xrightarrow{L} & \mathbb{Z}^V. \end{array}$$

Exercise 2.5. Prove commutativity by showing that $\text{div}(\chi_u)$ corresponds with $L(\chi_u)$ for each $u \in V$ under our isomorphism of $\text{Div}(G)$ with \mathbb{Z}^V .

Thus, we will feel free to blur the distinction between the two mappings, sometimes calling div “the Laplacian” of G . Nevertheless, these mappings are not really the same: their codomains, while isomorphic, are not equal. The divisor homomorphism div arises via the connection we are developing between the theory of Riemann surfaces and the dollar game (cf. Section 5.3) while the Laplacian operator L plays its traditional role as an operator taking functions to functions.

To realize these mappings as a single matrix, fix an ordering v_1, \dots, v_n of the vertices of G , providing an ordered basis for the free abelian group $\text{Div}(G)$ and a corresponding isomorphism with \mathbb{Z}^n . Writing $\chi_j := \chi_{v_j}$, it follows that $\{\chi_1, \dots, \chi_n\}$ is the *dual basis* of the free abelian group $\mathcal{M}(G) = \mathbb{Z}^V$ given by

$$\chi_j(v_i) = \begin{cases} 1 & i = j \\ 0 & i \neq j. \end{cases}$$

The firing script χ_j corresponds to vertex v_j making a single lending move.

Definition 2.6. The *Laplacian matrix*, which we also denote L , is the matrix representation of the Laplacian operator $L: \mathbb{Z}^V \rightarrow \mathbb{Z}^V$ with respect to the basis $\{\chi_j\}$. It is also the matrix representation of $\text{div}: \mathcal{M}(G) \rightarrow \text{Div}(G)$ with respect to the dual bases $\{\chi_j\}$ and $\{v_i\}$.

Explicitly, L is the $n \times n$ integer matrix with ij -entry

$$L_{ij} = L(\chi_j)(v_i) = \begin{cases} \deg_G(v_i) & i = j \\ -(\# \text{ of edges between } v_j \text{ and } v_i) & i \neq j. \end{cases}$$

Defining $\text{Deg}(G)$ to be the diagonal matrix listing the vertex-degrees of G , we see that¹

$$L = \text{Deg}(G) - A^t$$

where A is the *adjacency matrix* of the graph G , defined as

$$A_{ij} = \# \text{ of edges between } v_i \text{ and } v_j.$$

The matrix L encodes all of the lending moves for G , because a lending move by vertex v_j corresponds to subtracting the j th column of L from a divisor. See Figure 1 for an example.

¹For undirected multigraphs, the adjacency matrix is symmetric, so the transpose is unnecessary. But for directed multigraphs (which we will study in Part 2), the transpose is essential.




Figure 1. The Laplacian encodes firing rules.

The following commutative diagram has exact rows, meaning that the kernel of each horizontal arrow is the image of the previous arrow. (See Appendix B for more about exact sequences.) The vertical arrows are the isomorphisms resulting from the choice of bases $\{\chi_j\}$ and $\{v_i\}$.

$$\begin{array}{ccccccc}
 \mathcal{M}(G) & \xrightarrow{\text{div}} & \text{Div}(G) & \longrightarrow & \text{Pic}(G) & \longrightarrow & 0 \\
 \{\chi_j\} \downarrow & & \{v_i\} \downarrow & & \downarrow & & \\
 \mathbb{Z}^n & \xrightarrow{L} & \mathbb{Z}^n & \longrightarrow & \mathbb{Z}^n / \text{im}(L) & \longrightarrow & 0
 \end{array}$$

From the diagram, we see that the Picard group may be computed as the cokernel of the Laplacian matrix:

$$\text{Pic}(G) \simeq \text{cok}(L) := \mathbb{Z}^n / \text{im}(L).$$

Exercise 2.7. Explicitly, the isomorphism above works as follows: starting with an equivalence class $[D] \in \text{Pic}(G)$, choose a representative divisor $D \in \text{Div}(G)$, which may be written as $D = \sum_{i=1}^n D(v_i) v_i$. Use the coefficients to form a vector $(D(v_1), \dots, D(v_n)) \in \mathbb{Z}^n$, and then compute the coset $(D(v_1), \dots, D(v_n)) + \text{im}(L)$ in the quotient group $\mathbb{Z}^n / \text{im}(L)$. Show that this recipe does not depend on the choice of representative divisor D and defines an isomorphism of $\text{Pic}(G)$ with $\mathbb{Z}^n / \text{im}(L)$.

In concrete terms, this means that the Picard group is given by the free abelian group \mathbb{Z}^n modulo the subgroup generated by the columns of the Laplacian matrix. This description yields another restatement of the dollar game:

The dollar game: Given a divisor D (identified with an integer vector) does there exist a \mathbb{Z} -linear combination l of columns of the Laplacian matrix such that $D + l \geq 0$?

Recall from Example 1.7 that a set-lending by V has no net effect. In terms of the divisor homomorphism, this is the same as saying that

$$\text{div}(\chi_1 + \chi_2 + \dots + \chi_n) = 0.$$

In terms of the Laplacian matrix, this is just the statement that the all 1's vector $\vec{1}$ is in the kernel of L . In fact, since G is connected, this element generates the kernel, as shown in the next proposition. In Theorem 9.14, we will generalize this result to the case of certain directed multigraphs called *sandpile graphs*.

Proposition 2.8. *If $G = (V, E)$ is an undirected, connected multigraph, then the kernel of its Laplacian matrix L is generated by the all 1's vector $\vec{1} \in \mathbb{Z}^n$. Equivalently, the kernel of the divisor homomorphism div (equal to the kernel of the Laplacian operator) consists of the constant functions $c: V \rightarrow \mathbb{Z}$.*

Proof. We will prove that the kernel of div consists of the constant functions. So suppose that $\sigma: V \rightarrow \mathbb{Z}$ satisfies $\text{div}(\sigma) = 0$. Choose a vertex $v \in V$ where σ attains its maximum, $k \in \mathbb{Z}$. Since the divisor of σ is zero, we must have

$$\deg_G(v)k = \sum_{vw \in E} \sigma(w).$$

But since $\sigma(w) \leq k$ for all vertices w , the equality above will only obtain if $\sigma(w) = k$ for all w adjacent to v . Since G is connected, it follows that σ must take the value k at every vertex of G , so that $\sigma = k$ is a constant function. \square

The argument in the previous proof shows more generally that the kernel of div (for a possibly disconnected multigraph) consists of functions that are constant on connected components. This result is a first indication of an interplay between graph-theoretic properties of G (like connectivity) and algebraic structures related to the abelian group $\text{Pic}(G)$ (such as the kernel of the Laplacian matrix). As we will see in later chapters, this interplay extends to the *edge-connectivity* of the graph, defined as the minimal number of edges that must be removed in order to disconnect the graph. For example, trees are 1-edge connected, cycles are 2-edge connected, and the complete graph K_4 is 3-edge connected.

2.1.1. Relation to the continuous Laplacian. Recall the differential operator $\nabla = (\partial/\partial x_1, \dots, \partial/\partial x_n)$. If $\phi: \mathbb{R}^n \rightarrow \mathbb{R}$ is a differentiable function, then the *gradient* of ϕ is the vector field

$$\nabla\phi = \left(\frac{\partial\phi}{\partial x_1}, \dots, \frac{\partial\phi}{\partial x_n} \right).$$

At each point p , the vector $\nabla\phi(p)$ points in the direction of quickest increase of ϕ , and its magnitude is the rate of increase in that direction. The function ϕ (or sometimes its negative) is called a *potential* for the vector field $\nabla\phi$. Going the other direction, for a given vector field on \mathbb{R}^n , the potential (if it exists) is determined up to an additive constant.

If $F = (F_1, \dots, F_n): \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a vector field, then the *divergence* of F is

$$\nabla \cdot F = \frac{\partial F_1}{\partial x_1} + \dots + \frac{\partial F_n}{\partial x_n}.$$

Imagine a sphere of radius r centered at a point p . The *flux* of F through this sphere is the surface integral of the component of F in the direction of the outward-pointing normal to the sphere. Thinking of F as defining a flow, the flux measures the rate at which the flowing substance leaves the sphere. By Gauss' theorem, $\nabla \cdot F$ is the limit of the ratio of this flux to the volume enclosed by the sphere as $r \rightarrow 0$; it measures *flux density* and hence the “divergence” of F at the point p .

The traditional Laplacian is the differential operator

$$\Delta := \nabla^2 = \nabla \cdot \nabla := \frac{\partial^2}{\partial x_1^2} + \cdots + \frac{\partial^2}{\partial x_n^2}.$$

Thus, the Laplacian of $\phi: \mathbb{R}^n \rightarrow \mathbb{R}$ is

$$\Delta\phi = \nabla^2\phi = \nabla \cdot (\nabla\phi) = \frac{\partial^2\phi}{\partial x_1^2} + \cdots + \frac{\partial^2\phi}{\partial x_n^2}.$$

So $\Delta\phi(p)$ is a scalar representing the tendency of a vector field with potential ϕ to diverge from the point p . Roughly, it will be zero if the value $\phi(p)$ is the average of the values of ϕ at points on a tiny sphere surrounding p .

Now let $f: V \rightarrow \mathbb{Z}$, and apply the discrete Laplacian at a vertex v :

$$L(f)(v) := \sum_{vw \in E} (f(v) - f(w)) = \deg_G(v)f(v) - \sum_{vw \in E} f(w).$$

Dividing by the degree gives

$$\frac{1}{\deg_G(v)}L(f)(v) = f(v) - \frac{1}{\deg_G(v)} \sum_{vw \in E} f(w),$$

showing that the discrete Laplacian measures how much the value of f at v deviates from its average among the neighbors of v . In this way, L is a discrete analogue of the continuous Laplacian. Note in particular that $L(f)(v) = 0$ if and only if the value of f at v is the average of its values at the neighbors of v —in this case f is said to be *harmonic* at v . We will return to the subject of harmonic functions in Section 3.5 and Chapter 10.

Even if you were only interested in the continuous Laplacian, the discrete Laplacian would arise naturally from the application of numerical methods. To do calculations with the Laplacian on a computer, one first discretizes by replacing \mathbb{R}^n with a grid, which up to scaling, we can think of as \mathbb{Z}^n . For instance, let's first consider the case $n = 1$. So $f: \mathbb{R} \rightarrow \mathbb{R}$, and $\Delta f = d^2f/dx^2$. The discrete derivative of f at $i \in \mathbb{Z}$ is $f(i) - f(i - 1)$, and hence, the second discrete derivative is

$$(f(i + 1) - f(i)) - (f(i) - f(i - 1)) = -2f(i) + f(i + 1) + f(i - 1).$$

Thus, $\Delta f(i)$ is approximated by $-L(f)(i)$ where L is the discrete Laplacian of the infinite path graph (which would of course be finite in practice) and f is restricted to this graph. In two dimensions, we replace \mathbb{R}^2 with \mathbb{Z}^2 , and to approximate $\Delta f = \partial^2 f / \partial x_1^2 + \partial^2 f / \partial x_2^2$ at the grid point (i, j) , we similarly compute

$$\partial^2 f / \partial x_1^2 \approx -2f(i, j) + f(i + 1, j) + f(i - 1, j)$$

$$\partial^2 f / \partial x_2^2 \approx -2f(i, j) + f(i, j + 1) + f(i, j - 1).$$

Adding these gives $-L(f)(i, j)$ where L is the discrete Laplacian of a large 2-dimensional grid.

2.2. Configurations and the reduced Laplacian

We have seen that the Picard group can be computed as the cokernel of the Laplacian matrix. But what about the Jacobian group? To establish the relationship

between the Jacobian and the Laplacian, we first introduce the notion of a *configuration* on G .

Definition 2.9. Fix a vertex $q \in V$, and set $\tilde{V} = V \setminus \{q\}$. Then a *configuration* on G (with respect to q) is an element of the subgroup

$$\text{Config}(G, q) := \text{Config}(G) := \mathbb{Z}\tilde{V} \subset \mathbb{Z}V = \text{Div}(G).$$

As with divisors, we write $c \geq c'$ for $c, c' \in \text{Config}(G)$ if $c(v) \geq c'(v)$ for all $v \in \tilde{V}$, and write $c \geq 0$ if $c(v) \geq 0$ for all $v \in \tilde{V}$. In the latter case, c is said to be *nonnegative*. The *degree* of c is $\deg(c) := \sum_{v \in \tilde{V}} c(v)$.

Given a configuration $c \in \text{Config}(G, q)$, we may perform lending and borrowing moves at any vertex $v \in \tilde{V}$ as if c were a divisor, except we do not record how much money is on q .

Example 2.10. The picture below shows a lending move by the vertex v_2 , starting with the configuration $c = v_1 + 2v_2 + v_3$.




Figure 2. Effect of a lending move by v_2 on a configuration.

Just as with divisors, we will say two configurations are equivalent if they may be obtained from each other via lending and borrowing moves. For example, Figure 2 shows that $c = v_1 + 2v_2 + v_3$ is linearly equivalent to $c' = 2v_1 - v_2 + 2v_3$, which we indicate by writing $c \sim c'$. We make these notions precise in the following definition.

Definition 2.11. If $D \in \text{Div}(G)$, denote by $D|_{\tilde{V}}$ the configuration obtained by ignoring the vertex q . Let $\tilde{\mathcal{M}}(G) \subset \mathcal{M}(G)$ denote the subgroup of firing scripts with support contained in \tilde{V} :

$$\tilde{\mathcal{M}}(G) := \{\sigma : V \rightarrow \mathbb{Z} \mid \sigma(q) = 0\},$$

and define the homomorphism

$$\begin{aligned} \widetilde{\text{div}} : \tilde{\mathcal{M}}(G) &\rightarrow \text{Config}(G) \\ \sigma &\mapsto \text{div}(\sigma)|_{\tilde{V}}. \end{aligned}$$

Say two configurations $c, c' \in \text{Config}(G)$ are *equivalent*, denoted $c \sim c'$ if $c = c' + \widetilde{\text{div}}(\sigma)$ for some firing script σ , i.e., if $c = c'$ modulo the image of $\widetilde{\text{div}}$, in which case we write $c \xrightarrow{\sigma} c'$. Using characteristic functions, we get the special cases of *vertex-firings*, where $\sigma = \chi_v$ for $v \in \tilde{V}$, and *set-firings*, where $\sigma = \chi_W$ for $W \subseteq \tilde{V}$. In these cases, we sometime abuse notation and write $c \xrightarrow{v} c'$ and $c \xrightarrow{W} c'$, respectively.

We now wish to show that $\text{Jac}(G)$ is isomorphic to the group of configurations modulo equivalence. To see this, consider the group homomorphism

$$\begin{aligned} \varphi: \text{Config}(G) &\rightarrow \text{Jac}(G) \\ c &\mapsto [c - \deg(c)q]. \end{aligned}$$

This map φ is surjective, since if D is an arbitrary divisor of degree zero, then $D = \tilde{D} - \deg(\tilde{D})q$ for the configuration $\tilde{D} := D|_{\tilde{\mathcal{V}}}$. To emphasize this point: a divisor D of degree 0 is determined by its corresponding configuration, $D|_{\tilde{\mathcal{V}}}$.

To determine the kernel of φ , first note that by Proposition 2.8 the addition of a constant function to any firing script does not affect the corresponding principal divisor, so every principal divisor may be obtained from a script in $\tilde{\mathcal{M}}(G)$. With this in mind, we see that

$$\begin{aligned} \varphi(c) = 0 &\iff c - \deg(c)q \text{ is principal} \\ &\iff c - \deg(c)q = \text{div}(\sigma) \text{ for some } \sigma \in \tilde{\mathcal{M}}(G) \\ &\iff c = \tilde{\text{div}}(\sigma). \end{aligned}$$

Thus, we see that $\ker(\varphi) = \text{im}(\tilde{\text{div}})$. Therefore, φ induces the following isomorphism, where we use the notation $[c]$ for the equivalence class of a configuration.

Proposition 2.12. *There is an isomorphism of groups*

$$\begin{aligned} \text{Config}(G)/\text{im}(\tilde{\text{div}}) &\simeq \text{Jac}(G) \\ [c] &\mapsto [c - \deg(c)q] \\ [D|_{\tilde{\mathcal{V}}}] &\mapsto [D]. \end{aligned}$$

2.2.1. The reduced Laplacian. Fixing an ordering v_1, v_2, \dots, v_n for the vertices of G identifies configurations with integer vectors, just as for divisors. Of course, we must remove the vertex q and its corresponding firing script χ_q , which has the effect of removing the q -row and χ_q -column from the Laplacian matrix, L .

Definition 2.13. The *reduced Laplacian matrix*, \tilde{L} , is the matrix representation of $\tilde{\text{div}}: \tilde{\mathcal{M}}(G) \rightarrow \text{Config}(G)$ with respect to the dual bases $\{\chi_j\} \setminus \{\chi_q\}$ and $\{v_i\} \setminus \{q\}$. Explicitly, \tilde{L} is the $(n-1) \times (n-1)$ integer matrix obtained by removing the row corresponding to q and the column corresponding to χ_q from the Laplacian matrix L .

The following commutative diagram summarizes the setup:

$$\begin{array}{ccccccc} \tilde{\mathcal{M}}(G) & \xrightarrow{\tilde{\text{div}}} & \text{Config}(G) & \xrightarrow{\varphi} & \text{Jac}(G) & \longrightarrow & 0 \\ \downarrow & & \downarrow & & \downarrow & & \\ \mathbb{Z}^{n-1} & \xrightarrow{\tilde{L}} & \mathbb{Z}^{n-1} & \longrightarrow & \mathbb{Z}^{n-1}/\text{im}(\tilde{L}) & \longrightarrow & 0. \end{array}$$

As before, the rows are exact, while the vertical arrows are the isomorphisms defined by the choice of vertex-ordering. Just as the Picard group is the cokernel of the full Laplacian matrix, the Jacobian group may be computed as the cokernel of the reduced Laplacian:

$$(2.1) \quad \text{Jac}(G) \simeq \text{cok}(\tilde{L}) := \mathbb{Z}^{n-1}/\text{im}(\tilde{L}).$$

Remark 2.14. Note that $\text{Jac}(G)$ does not depend on a choice of vertex q . Therefore, the isomorphism in (2.1) shows that even though the reduced Laplacian depends on a choice of vertex, its cokernel $\text{cok}(\tilde{L}) := \mathbb{Z}^{n-1} / \text{im}(\tilde{L})$ does not (cf. Problem 2.1). In Chapter 12 we will extend the notions of the Jacobian group and the reduced Laplacian to directed graphs. In that case, the isomorphism of (2.1) and the independence of $\text{cok}(\tilde{L})$ from the choice of excluded vertex no longer hold in general. See Section 12.1 for details.


We have the following corollary to Proposition 2.8, which says that the reduced Laplacian is always invertible over the rational numbers.

Corollary 2.15. *If $G = (V, E)$ is an undirected, connected multigraph, then the kernel of its reduced Laplacian matrix \tilde{L} is zero. Consequently, \tilde{L} is invertible as a linear operator on \mathbb{Q}^{n-1} .*

Proof. By Proposition 2.8, the kernel of the full Laplacian L is generated by the all ones vector $\vec{1} \in \mathbb{Z}^n$. This implies that the final column of L is the negative of the sum of the first $n - 1$ columns. Since L is symmetric, this means that the final row of L is the negative of the sum of the first $n - 1$ rows. In particular, any vector orthogonal to the first $n - 1$ rows of L is automatically orthogonal to the last row of L . Now suppose that $\tilde{a} \in \mathbb{Z}^{n-1}$ is in the kernel of the reduced Laplacian \tilde{L} , obtained from L by removing the last row and column. Then the vector $(\tilde{a}, 0) \in \mathbb{Z}^n$ has the property that its dot product with the first $n - 1$ rows of L is zero. By the previous comments, this implies that the dot product of $(\tilde{a}, 0)$ with the last row of L is also zero, so $(\tilde{a}, 0) \in \ker(L)$. By Proposition 2.8, it follows that $\tilde{a} = 0$, so that the kernel of \tilde{L} is zero. Viewed as a linear operator on \mathbb{Q}^{n-1} , the injectivity of \tilde{L} implies its invertibility, so \tilde{L} is invertible over the rational numbers. \square

Remark 2.16. In Chapter 8, we will discover an interesting interpretation of the rational entries of the matrix \tilde{L}^{-1} (Theorem 8.29).

Example 2.17. Let's use equation 2.1 to give another computation of the Jacobian of the triangle graph C_3 (compare Example 1.21).



The Laplacian matrix of C_3 is

$$L = \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix},$$

so the reduced Laplacian is (using $q = v_3$):

$$\tilde{L} = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}.$$

Hence, $\text{im}(\tilde{L}) = \text{span}\{(2, -1), (-1, 2)\} = \text{span}\{(3, 0), (1, 1)\} \subset \mathbb{Z}^2$. Now consider the homomorphism $\mathbb{Z}^2 \rightarrow \mathbb{Z}_3$ defined by $(1, 0) \mapsto 1 \pmod 3$ and $(1, 1) \mapsto 0$. Then for

an arbitrary integer vector (a, b) , we have

$$(a, b) = (b, b) + (a - b, 0) \mapsto 0 + (a - b) \bmod 3,$$

so (a, b) is in the kernel if and only if $a - b = 3k$ for some $k \in \mathbb{Z}$. It follows that the kernel is spanned by $(1, 1)$ and $(3, 0)$. From equation 2.1

$$\text{Jac}(C_3) \simeq \mathbb{Z}^2 / \text{im}(\tilde{L}) \simeq \mathbb{Z}_3.$$

Explicitly, the inverse of this isomorphism is defined by

$$a \bmod 3 \mapsto [av_1 - av_3] \in \text{Jac}(C_3).$$

Section 2.4 introduces an algorithm for determining the structure of the Picard and Jacobian groups, in general.

2.3. Complete linear systems and convex polytopes

The problem of finding the complete linear system of a divisor D (i.e., finding all winning distributions obtainable from D) turns out to be equivalent to finding lattice points in a polytope determined by D and the Laplacian of the underlying graph. Before making these ideas precise, we give a first glimpse through a couple of simple examples.

Example 2.18. Consider the divisor $D = 3v_1 - v_2$ on the segment graph S from Example 1.12:

$$3 \textcircled{v_1} \text{---} \textcircled{v_2} - 1.$$

The Laplacian matrix is

$$L = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

We identify the divisor D with the integer vector $(3, -1) \in \mathbb{Z}^2$. In order to compute the complete linear system of D , we are looking for $\vec{x} = (x_1, x_2) \in \mathbb{Z}^2$ such that $D - L\vec{x} \geq 0$. This yields the pair of inequalities

$$\begin{aligned} x_1 - x_2 &\leq 3 \\ -x_1 + x_2 &\leq -1, \end{aligned}$$

and we see that our problem is to find the points with integer coordinates contained in the intersection of two half-spaces in \mathbb{R}^2 ; see Figure 3. (Such an intersection of finitely many half-spaces in \mathbb{R}^n is called a *convex polyhedron*).

Since $\vec{1} = (1, 1)$ is in the kernel of L , any solution \vec{x} to this system yields infinitely many other solutions of the form $\vec{x} + k\vec{1}$ for $k \in \mathbb{Z}$, all of which yield the same effective divisor $D - L\vec{x}$ in the complete linear system of D . Hence, we may as well look for solutions with $x_2 = 0$. Imposing this extra constraint yields the pair of inequalities $1 \leq x_1 \leq 3$. In terms of the dollar game, this says that the complete linear system of D may be obtained by lending from vertex v_1 once, twice, and three times. As before, we find that

$$|D| = \{2v_1, v_1 + v_2, 2v_2\}.$$





Figure 3. The convex polyhedron of winning firing-scripts for $D = 3v_1 - v_2$ on the segment graph S .

To compute the Picard group of S , note that the image of L is generated by the vector $(1, -1) \in \mathbb{Z}^2$. This vector also generates the kernel of the surjective degree homomorphism $\text{deg}: \mathbb{Z}^2 \rightarrow \mathbb{Z}$ defined by $\text{deg}(a, b) = a + b$, so it follows that

$$\text{Pic}(S) \simeq \mathbb{Z}^2 / \text{im}(L) \simeq \mathbb{Z}.$$

The inverse of this isomorphism is given by $d \mapsto [dv_1]$. This generalizes the fact obtained for degree 2 divisors in Example 1.12: two divisors are linearly equivalent on S if and only if they have the same degree.

Example 2.19. Consider the divisor $D = 2v_1 + v_2 - 2v_3$ on the triangle graph C_3 from Example 1.18:



The Laplacian matrix of C_3 is

$$L = \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix},$$

and we are looking for $\vec{x} = (x_1, x_2, x_3) \in \mathbb{Z}^3$ such that $D - L\vec{x} \geq 0$, where we have identified D with $(2, 1, -2) \in \mathbb{Z}^3$. Explicitly, we have the system of inequalities

$$\begin{aligned} 2x_1 - x_2 - x_3 &\leq 2 \\ -x_1 + 2x_2 - x_3 &\leq 1 \\ -x_1 - x_2 + 2x_3 &\leq -2. \end{aligned}$$

Once again, we are looking for the integer lattice points inside a convex polyhedron defined by the intersection of three half-spaces in \mathbb{R}^3 . By the same argument as

given in Example 2.18, we may utilize the fact that $\vec{1} \in \ker(L)$ to restrict attention to solutions with $x_3 = 0$. With this constraint, the system becomes

$$\begin{aligned} 2x_1 - x_2 &\leq 2 \\ -x_1 + 2x_2 &\leq 1 \\ -x_1 - x_2 &\leq -2, \end{aligned}$$

which defines the bounded triangle in \mathbb{R}^2 , pictured in Figure 4. By inspection, the only integer lattice point in this triangle is $(1, 1)$. This says that the complete linear system of D may be obtained by set-lending $\{v_1, v_2\}$, yielding $|D| = \{v_1\}$, as before.




Figure 4. Triangle for Example 2.19.

We are now ready to see how these ideas work out in general. Given a divisor D , we want to determine the complete linear system of D ,

$$|D| = \{E \in \text{Div}(G) : E \sim D \text{ and } E \geq 0\}.$$

As usual, choosing a vertex-ordering identifies divisors and firing scripts with vectors in \mathbb{Z}^n , so finding the complete linear system $|D|$ requires finding the firing scripts $\vec{x} \in \mathbb{Z}^n$ such that $D - L\vec{x} \geq 0$. If $D = a_1v_1 + \cdots + a_nv_n$, then we must solve the system of linear inequalities:

$$\begin{aligned} L_{11}x_1 + \cdots + L_{1n}x_n &\leq a_1 \\ L_{21}x_1 + \cdots + L_{2n}x_n &\leq a_2 \\ &\vdots \\ L_{n1}x_1 + \cdots + L_{nn}x_n &\leq a_n. \end{aligned}$$

Each solution $\vec{x} \in \mathbb{Z}^n$ yields an effective divisor $E = D - L\vec{x} \in |D|$, and all such effective divisors are obtained this way.

Geometrically, each inequality above defines a half-space in \mathbb{R}^n , and the intersection of the n half-spaces yields a convex polyhedron, $P \subset \mathbb{R}^n$. To say that P

is convex just means that if p_1 and p_2 are any two points in P , then the line segment between p_1 and p_2 is entirely contained within P as well. Since half-spaces are clearly convex, so is the intersection of half-spaces, P . If, in addition, P is a bounded set, then we say that P is a convex *polytope*.

The firing scripts we are looking for correspond to $\mathbb{Z}^n \cap P$, the intersection of the integer lattice with the polyhedron. If this intersection is nonempty, then it is infinite, since $\mathbb{Z}^n \cap P$ is invariant under translation by $\vec{1} \in \ker L$. In particular, since the firing scripts \vec{x} and $\vec{x} + k\vec{1}$ determine the same divisor linearly equivalent to D , it suffices to determine the intersection $\mathbb{Z}^n \cap \tilde{P}$ where \tilde{P} is the intersection of P with the hyperplane defined by setting the final coordinate equal to zero:

$$\tilde{P} := P \cap \{\vec{x} \in \mathbb{R}^n : x_n = 0\}.$$

Proposition 2.20. *The set \tilde{P} is a convex polytope, so $\mathbb{Z}^n \cap \tilde{P}$ is a finite set.*

Proof. The convexity of \tilde{P} follows immediately from the convexity of P , so we just need to show boundedness. Note that we may think of \tilde{P} as contained in \mathbb{R}^{n-1} , where it is defined by the following system of inequalities

$$\begin{aligned} \tilde{L}\tilde{x} &\leq \tilde{a} \\ L_n\tilde{x} &\leq a_n, \end{aligned}$$

where \tilde{L} is the reduced Laplacian consisting of the first $n - 1$ columns and rows of L , the vectors $\tilde{x}, \tilde{a} \in \mathbb{Z}^{n-1}$ are obtained by omitting the n th components of \vec{x} and \vec{a} , and $L_n \in \mathbb{Z}^{n-1}$ denotes the n th row of the Laplacian, with its final entry removed.

Now change coordinates by defining $\tilde{y} := -\tilde{L}\tilde{x} + \tilde{a}$, for which the matrix inequality becomes simply $\tilde{y} \geq 0$. Using $\tilde{x} = -\tilde{L}^{-1}(\tilde{y} - \tilde{a})$, the final inequality becomes

$$-\left(L_n\tilde{L}^{-1}\right)\tilde{y} \leq a_n - \left(L_n\tilde{L}^{-1}\right)\tilde{a} =: d.$$

Note that since the sum of the rows of L is zero, we have $-L_n = \tilde{L}_1 + \dots + \tilde{L}_{n-1}$, the sum of the rows of \tilde{L} . It follows that

$$-L_n\tilde{L}^{-1} = (\tilde{L}_1 + \dots + \tilde{L}_{n-1})\tilde{L}^{-1} = (\tilde{1}\tilde{L})\tilde{L}^{-1} = \tilde{1},$$

where $\tilde{1} \in \mathbb{Z}^{n-1}$ denotes the all ones row vector. Thus, in the new coordinates, our system of inequalities becomes

$$\begin{aligned} \tilde{y} &\geq 0 \\ \tilde{1}\tilde{y} &\leq d, \end{aligned}$$

where $d = a_n + \tilde{1}\tilde{a} = \deg(D)$ is the degree of the original divisor, D . This is clearly a bounded simplex in \mathbb{R}^{n-1} . But the original coordinates \tilde{x} are obtained from \tilde{y} by an affine-linear transformation, so \tilde{P} is also bounded. \square

2.4. Structure of the Picard group

For motivation, the reader is encouraged to quickly peruse Example 2.35 near the end of this section.

2.4.1. Finitely generated abelian groups. A *finitely generated* abelian group is an abelian group for which there exists a finite set $\{a_1, \dots, a_m\}$ of elements (*generators*) of A such that each $a \in A$ may be written as

$$a = k_1 a_1 + \dots + k_m a_m$$

for some $k_i \in \mathbb{Z}$. We write $A := \langle a_1, \dots, a_m \rangle$. Some examples follow.

- Every finite abelian group is generated by the finite set consisting of all of its elements. For instance, $\mathbb{Z}_4 = \langle 0, 1, 2, 3 \rangle$. Of course, we also have $\mathbb{Z}_4 = \langle 1 \rangle = \langle 3 \rangle$. We do not require the set of generators to be minimal (with respect to inclusion).
- A *cyclic group* is by definition generated by a single element, and every cyclic group is abelian. Every cyclic group is isomorphic to \mathbb{Z}_n for some nonnegative integer n . The case $n = 0$ yields the infinite cyclic group $\mathbb{Z} = \mathbb{Z}_0$.
- A finite product of cyclic groups is a finitely generated abelian group. A typical instance is $\mathbb{Z}_4 \times \mathbb{Z}_6 \times \mathbb{Z}^2$. Recall that addition is defined component-wise, e.g., $(2, 5, 8, -2) + (3, 3, 7, 2) = (1, 2, 15, 0)$ in this group.

In fact—as will be stated precisely in Theorem 2.23—the last example exhausts all possibilities: every finitely generated abelian group is a product of a finite number of cyclic groups, i.e., has the form $\prod_{i=1}^n \mathbb{Z}_{n_i}$ for suitable n_i . The *structure* is then represented by the list of the n_i . Different choices for the n_i may produce isomorphic groups, but it turns out that the ambiguity is accounted for by the following well-known result.

Theorem 2.21 (Chinese remainder theorem). *Let $m, n \in \mathbb{Z}$. Then*

$$\mathbb{Z}_{mn} \simeq \mathbb{Z}_m \times \mathbb{Z}_n$$

if and only if m and n are relatively prime. If $\gcd(m, n) = 1$, then an isomorphism is provided by $a \mapsto (a \bmod m, a \bmod n)$.

Thus, $\mathbb{Z}_{24} \simeq \mathbb{Z}_8 \times \mathbb{Z}_3$, but $\mathbb{Z}_4 \not\simeq \mathbb{Z}_2 \times \mathbb{Z}_2$.

Exercise 2.22. Find the simplest direct (without using the Chinese remainder theorem) proof that $\mathbb{Z}_4 \not\simeq \mathbb{Z}_2 \times \mathbb{Z}_2$.

A proof of the following is presented after our discussion of the Smith normal form, below.

Theorem 2.23 (Structure theorem for f.g. abelian groups). *A group is a finitely generated abelian group if and only if it is isomorphic to*

$$\mathbb{Z}_{n_1} \times \dots \times \mathbb{Z}_{n_k} \times \mathbb{Z}^r$$

for some list (possibly empty) of integers n_1, \dots, n_k with $n_i > 1$ for all i and some integer $r \geq 0$. These integers may be required to satisfy either of the following two conditions, and in either case they are uniquely determined by the isomorphism class of the group.

condition 1: $n_i | n_{i+1}$ (n_i evenly divides n_{i+1}) for all i . In this case, the n_i are the invariant factors of the group.

condition 2: There exist primes $p_1 \leq \dots \leq p_k$ and positive integers m_i such that $n_i = p_i^{m_i}$ for all i . In this case, the n_i are the elementary divisors and the \mathbb{Z}_{n_i} are the primary factors of the group.

The number r is the rank of the group.

Exercise 2.24. Find the rank, the invariant factors, and the primary factors of $\mathbb{Z}_4 \times \mathbb{Z}_6 \times \mathbb{Z}^2$.

Let A be a finitely generated abelian group. An element $a \in A$ is a *torsion* element if $n \cdot a = 0$ for some nonzero integer n .

Exercise 2.25. Show that the collection of torsion elements forms a subgroup of A .

The subgroup of torsion elements is called the *torsion part* of A and denoted A_{tor} . If $A \simeq \left(\prod_{i=1}^k \mathbb{Z}_{n_i}\right) \times \mathbb{Z}^r$, then $A_{\text{tor}} \simeq \prod_{i=1}^k \mathbb{Z}_{n_i}$ and $A/A_{\text{tor}} \simeq \mathbb{Z}^r$. Hence, the structure theorem guarantees

$$A \simeq A_{\text{tor}} \times A/A_{\text{tor}}.$$

2.4.2. Smith normal form. How does one go about computing the rank and invariant factors of a finitely generated abelian group A ? Let $\{a_1, \dots, a_m\}$ be generators, and define the group homomorphism determined by

$$\begin{aligned} \mathbb{Z}^m &\xrightarrow{\pi} A \\ e_i &\mapsto a_i \end{aligned}$$

where e_i is the i -th standard basis vector. Saying that the a_i generate A is the same as saying that π is surjective. Next, by a standard theorem from algebra², every subgroup of \mathbb{Z}^m is finitely generated. In particular, there exists a finite set of generators $\{b_1, \dots, b_n\}$ for the kernel of π . Define

$$\mathbb{Z}^n \xrightarrow{M} \mathbb{Z}^m$$

where M is the $m \times n$ integer matrix with i -th column b_i . Combining these two mappings yields a *presentation* of A :

$$\mathbb{Z}^n \xrightarrow{M} \mathbb{Z}^m \longrightarrow A$$

Hence, π induces an isomorphism

$$\begin{aligned} \text{cok}(M) &:= \mathbb{Z}^m / \text{im}(M) \simeq A \\ e_i &\mapsto a_i, \end{aligned}$$

where $\text{cok}(M)$ denotes the cokernel of $M: \mathbb{Z}^n \rightarrow \mathbb{Z}^m$. In this way, A is determined by the single matrix M .

We have just seen that each finitely generated abelian group is the cokernel of an integer matrix. Conversely, each integer matrix determines a finitely generated abelian group. However, the correspondence is not bijective. The construction of M , above, depended on arbitrary choices for generators of A and of the kernel of π . Making different choices creates a different matrix representing A . This is especially obvious if we choose a different number of generators for A or $\ker \pi$. However, there can be a difference even if the number of generators is kept constant.

²The key point is that abelian groups are modules over \mathbb{Z} , and \mathbb{Z} is a Noetherian ring. The reader wishing to fill in this detail is encouraged to consult Appendix B and/or consult a standard introductory algebra text.

In that case, changing the choice of generators corresponds to integer changes of coordinates for the codomain and domain of M , or equivalently, to performing integer row and column operations on M .

Definition 2.26. The *integer row (resp., column) operations* on an integer matrix consist of the following:

- (1) swapping two rows (resp., columns);
- (2) negating a row (resp., column);
- (3) adding one row (resp., column) to a different row (resp., column).

Write $M \sim N$ for integer matrices M and N if one may be obtained from the other through a sequence of integer row and column operations. Since the operations are reversible, \sim is an equivalence relation.

Suppose M is an $m \times n$ integer matrix and $M \sim N$. Start with identity matrices $P = I_m$ and $Q = I_n$, and consider the sequence of integer row and column operations transforming M into N . Whenever a row operation is performed in this sequence, apply the same row operation to P . Similarly, whenever a column operation is made, apply the same column operation to Q .

Exercise 2.27. Explain why the resulting matrices P and Q are invertible over the integers and why $PMQ = N$. The converse of this statement is also true: given any matrices P and Q , invertible over the integers and such that $PMQ = N$, it follows that $M \sim N$. However, the proof of this converse requires the existence of the Smith normal form (Theorem 2.33).

The relation $PMQ = N$ can be expressed in terms of a commutative diagram with exact rows:

$$(2.2) \quad \begin{array}{ccccccc} \mathbb{Z}^n & \xrightarrow{M} & \mathbb{Z}^m & \longrightarrow & \text{cok } M & \longrightarrow & 0 \\ & & \downarrow Q^{-1} \cong & & \downarrow \cong & & \\ & & \mathbb{Z}^n & \xrightarrow{N} & \mathbb{Z}^m & \longrightarrow & \text{cok } N \longrightarrow 0. \end{array}$$

The mapping $\text{cok}(M) \rightarrow \text{cok}(N)$ is induced by P .

Proposition 2.28. Let M and N be $m \times n$ integer matrices. Then if $M \sim N$, it follows that $\text{cok}(M) \simeq \text{cok}(N)$.

Proof. Since P and Q in the commutative diagram are isomorphisms the mapping of cokernels induced by P is an isomorphism (Problem 2.3). \square

Exercise 2.29. Suppose $M \sim N$ where $N = \text{diag}(m_1, \dots, m_\ell)$, a diagonal integer matrix with nonnegative entries. Show that

$$\text{cok}(M) \simeq \prod_{i=1}^{\ell} \mathbb{Z}_{m_i}.$$

(See Example 2.31.)

The previous exercise shows that to determine the structure of $\text{cok}(M)$, we should seek to transform M through integer row and column operations into a diagonal matrix of a particularly nice form.

Definition 2.30. An $m \times n$ integer matrix M is in *Smith normal form* if $M = \text{diag}(s_1, \dots, s_k, 0, \dots, 0)$, a diagonal matrix, where s_1, \dots, s_k are positive integers such that $s_i | s_{i+1}$ for all i . The s_i are called the *invariant factors* of M .

Example 2.31. The matrix

$$M := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 12 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

is in Smith normal form with invariant factors $s_1 = 1$, $s_2 = 2$, and $s_3 = 12$.

We have

$$\text{cok}(M) := \mathbb{Z}^5 / \text{im}(M) \simeq \mathbb{Z}_1 \times \mathbb{Z}_2 \times \mathbb{Z}_{12} \times \mathbb{Z}^2 \simeq \mathbb{Z}_2 \times \mathbb{Z}_{12} \times \mathbb{Z}^2.$$

So $\text{cok}(M)$ has rank $r = 2$ and its invariant factors are 2 and 12.

Note that 1 is an invariant factor of M but not of $\text{cok}(M)$. By definition, the invariant factors of a finitely generated abelian group are greater than 1; the invariant factors of M equal to 1 do not affect the isomorphism class of $\text{cok}(M)$ since \mathbb{Z}_1 is the trivial group.

Definition 2.32. Let M be an $m \times n$ integer matrix. For each $i = 0, \dots, \min\{m, n\}$, the i -th *determinantal divisor* of M , denoted $d_i := d_i(M)$, is the greatest common divisor of the $i \times i$ minors of M . By definition $d_0 := 1$.

Theorem 2.33. *Each equivalence class under \sim has a unique representative in Smith normal form: if M is an $m \times n$ integer matrix, there exists a unique matrix N in Smith normal form such that $M \sim N$. For $i = 1, \dots, \text{rk}(M)$, the i -th determinantal divisor is*

$$d_i = \prod_{\ell=1}^i s_\ell$$

where s_ℓ is the ℓ -th invariant factor of M .

Proof. We prove existence by providing an algorithm that transforms M into Smith normal form through integer row and column operations. Let $M = (m_{ij})$ be an $m \times n$ integer matrix. If $M = 0$, we are done. Otherwise, proceed as follows.

Step 1. By permuting rows and columns we may assume that m_{11} is the smallest nonzero entry in absolute value. By adding integer multiples of the first row to other rows or the first column to other columns, attempt to make all entries in the first row and first column except the $(1,1)$ -entry equal to 0. If during the process any nonzero entry in the matrix appears with absolute value less than m_{11} , permute rows and columns to bring that entry into the $(1,1)$ -position. In this way, m_{11} remains the smallest nonzero entry. Since the succession of values for m_{11} are nonzero and decreasing in magnitude, the process eventually terminates with a matrix of the form

$$\left(\begin{array}{c|cccc} m_{11} & 0 & 0 & \cdots & 0 \\ \hline 0 & & & & \\ 0 & & & & \\ \vdots & & & & \\ 0 & & & & \end{array} M' \right)$$

where M' is an $(m-1) \times (n-1)$ integer matrix. Negating the first row, if necessary, we take $m_{11} > 0$.

Step 2. If there is an entry of M' that is not divisible by m_{11} , say m_{ij} , then add column j to column 1 and go back to Step 1. Again, since the $(1,1)$ -entry is nonzero and decreases in magnitude, this new process terminates. Therefore, we may assume that m_{11} divides every entry of M' .

Step 3. Apply Steps 1 and 2 to M' , and thus, by recursion, we produce an equivalent matrix in Smith normal form.

Thus, each integer matrix is equivalent to a matrix in Smith normal form. To prove the statement about determinantal divisors, first note that for each i , elementary row and column operations do not affect the set of $i \times i$ minors, except possibly for signs. Hence, they do not affect the i -th determinantal divisor. Therefore, we may assume the matrix is in Smith normal form, in which case the result is clear.

Uniqueness of the Smith normal form follows from the expression of the invariant factors in terms of determinantal divisors. \square

Proof. (*Structure theorem for finitely generated abelian groups.*) The above discussion provides a proof of the structure theorem, Theorem 2.23. In sum: we have seen that every finitely generated abelian group A is isomorphic to $\text{cok}(M)$ for some $m \times n$ integer matrix M . By making integer row and column operations, we may assume that M is in Smith normal form. It follows that

$$A \simeq \mathbb{Z}_{s_1} \times \cdots \times \mathbb{Z}_{s_k} \times \mathbb{Z}^r$$

for some list of positive integers s_1, \dots, s_k with $s_i | s_{i+1}$ for all i and some integer $r \geq 0$. Since $\mathbb{Z}_1 = 0$, we may assume that $s_i > 1$ for all i . This takes care of the existence part of the structure theorem.

For uniqueness, start by letting $B := A/A_{\text{tor}} \simeq \mathbb{Z}^r$. Then $B/2B \simeq \mathbb{Z}_2^r$ as vector spaces over the field \mathbb{Z}_2 . Since isomorphic vector spaces have the same dimension, this defines r intrinsically in terms of A .

Next, replacing A by A_{tor} , we may assume $r = 0$. Suppose there are also integers $s'_1, \dots, s'_{k'}$, greater than 1 and with $s'_i | s'_{i+1}$ for all i such that

$$A \simeq \mathbb{Z}_{s_1} \times \cdots \times \mathbb{Z}_{s_k} \simeq \mathbb{Z}_{s'_1} \times \cdots \times \mathbb{Z}_{s'_{k'}}.$$

By induction on k , we show that $k = k'$ and $s_i = s'_i$ for all i . We have $A = \{0\}$ if and only if $k = k' = 0$, so the result holds in this case. Suppose $k > 0$ and, hence, $k' > 0$. Let s be the smallest positive integer such that $sa = 0$ for all $a \in A$. By considering the element $a \in A$ corresponding to $(0, \dots, 0, 1) \in \prod_{i=1}^k \mathbb{Z}_{s_i}$ under the isomorphism, we see $s = s_k$. Similarly, $s = s'_{k'}$. Replace A by A/aA , and the result follows by induction.

The equivalence of the uniqueness statements in cases 1 and 2 of Theorem 2.23 follows from the Chinese remainder theorem (cf. Problem 2.4). \square

Remark 2.34. Suppose that A is an abelian group of rank r with k invariant factors. Then the corresponding expression for A as a product of cyclic groups makes it clear A can be generated by $k + r$ elements. In fact, this is the minimal number of elements required to generate A . To see this, suppose A is generated by ℓ elements. From the discussion in Section 2.4.2, we see that A is isomorphic to the cokernel of a matrix, M , with ℓ rows. From the Smith normal form of this matrix and the uniqueness statement in the structure theorem, we see that $A \simeq \left(\prod_{i=1}^k \mathbb{Z}_{s_i}\right) \times \mathbb{Z}^r$ where the s_i are the invariant factors of M that are not equal to 1. If t is the total number of invariant factors of M , then $t \geq k$ and $\ell = t + r$. Hence, $\ell \geq k + r$.

2.4.3. Computing $\text{Jac}(G)$ and $\text{Pic}(G)$. The following example illustrates the relevance of the structure theorem and the computation of Smith normal forms to our subject.

Example 2.35. Consider the complete graph K_4 consisting of 4 vertices with each pair of distinct vertices connected by an edge (Figure 5).




Figure 5. The complete graph K_4 .

The reduced Laplacian for K_4 with respect to any vertex is

$$\tilde{L} = \begin{pmatrix} 3 & -1 & -1 \\ -1 & 3 & -1 \\ -1 & -1 & 3 \end{pmatrix}.$$

Figure 6 shows the reduction of \tilde{L} to Smith normal form. Row and column operations are recorded in matrices P and Q , invertible over the integers. One may check that $P\tilde{L}Q = D$ where D is the Smith normal form for \tilde{L} .

Since $\text{Jac}(K_4) \simeq \text{cok}(\tilde{L})$, we conclude that the invariant factors of $\text{Jac}(G)$ are $s_1 = 4$ and $s_2 = 4$. From Exercise 2.27 and the commutative diagram in (2.2), we get the following isomorphism induced by P :

$$\begin{aligned} \text{Jac}(K_4) \simeq \text{cok}(\tilde{L}) &\xrightarrow{P} \mathbb{Z}_1 \times \mathbb{Z}_4 \times \mathbb{Z}_4 &\rightarrow \mathbb{Z}_4 \times \mathbb{Z}_4 \\ (x, y, z) &\mapsto \begin{pmatrix} 1 & 0 & 0 \\ 3 & 1 & 0 \\ 2 & 1 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} &\mapsto (3x + y, 2x + y + z). \end{aligned}$$

The final mapping is formed by dropping the first component.

Exercise 2.36. Directly compute the determinantal divisors for the reduced Laplacian of K_4 and use them to confirm that $\text{Jac}(K_4) \simeq \mathbb{Z}_4 \times \mathbb{Z}_4$.

$$\begin{array}{l}
\text{row ops : } \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \longrightarrow \begin{pmatrix} 1 & 0 & 0 \\ 3 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix} \\
\\
\tilde{L} : \quad \begin{pmatrix} 3 & -1 & -1 \\ -1 & 3 & -1 \\ -1 & -1 & 3 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 3 & -1 \\ -3 & -1 & -1 \\ 1 & -1 & 3 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ -3 & 8 & -4 \\ 1 & -4 & 4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 8 & -4 \\ 0 & -4 & 4 \end{pmatrix} \\
\\
\text{column ops : } \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \longrightarrow \begin{pmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \longrightarrow \begin{pmatrix} 0 & 1 & 0 \\ -1 & 3 & -1 \\ 0 & 0 & 1 \end{pmatrix} \longrightarrow \dots \\
\\
\text{-----} \\
\text{continued:} \\
\text{row ops : } \quad \begin{pmatrix} 1 & 0 & 0 \\ 3 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix} \longrightarrow \begin{pmatrix} 1 & 0 & 0 \\ 3 & 1 & 0 \\ 2 & 1 & 1 \end{pmatrix} =: P \\
\\
\tilde{L} : \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 8 & -4 \\ 0 & -4 & 4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 4 & 8 \\ 0 & -4 & -4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & -4 & 4 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 4 \end{pmatrix} =: D \\
\\
\text{column ops : } \quad \begin{pmatrix} 0 & 1 & 0 \\ -1 & 3 & -1 \\ 0 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 1 \\ -1 & 1 & 3 \\ 0 & -1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 0 & 1 \\ -1 & 1 & 1 \\ 0 & -1 & 2 \end{pmatrix} =: Q
\end{array}$$

Figure 6. Reduction of the reduced Laplacian for K_4 to Smith normal form.

To summarize what the previous example helps to illustrate, let G be a graph and let \tilde{L} be its reduced Laplacian with respect to any vertex. In light of Exercise 2.29 and the fact that $\text{Jac}(G) \simeq \text{cok}(\tilde{L})$, it follows that the invariant factors of $\text{Jac}(G)$ are determined by the invariant factors of \tilde{L} . Since \tilde{L} is invertible, none of its invariant factors are 0. So the rank of $\text{Jac}(G)$ is 0. The invariant factors of \tilde{L} equal to 1 have no effect on the isomorphism class of $\text{cok}(\tilde{L})$. Suppose m_1, \dots, m_k are the invariant factors of \tilde{L} that are greater than 1. The m_i are the invariant factors of $\text{Jac}(G)$, and hence,

$$(2.3) \quad \text{Jac}(G) \simeq \mathbb{Z}_{m_1} \times \cdots \times \mathbb{Z}_{m_k}.$$

The structure of $\text{Pic}(G)$ is then determined since $\text{Pic}(G) \simeq \mathbb{Z} \times \text{Jac}(G)$: the rank of $\text{Pic}(G)$ is 1, and

$$\text{Pic}(G) \simeq \mathbb{Z}_{m_1} \times \cdots \times \mathbb{Z}_{m_k} \times \mathbb{Z}.$$

As a final note, we have the following important formula for calculating the size of $\text{Jac}(G)$.

Proposition 2.37. *The order of the Jacobian group of G is*

$$|\text{Jac}(G)| = |\det(\tilde{L})|.$$

Proof. Since $\text{Jac}(G) \simeq \text{cok}(\tilde{L})$, Exercise 2.29 shows that the number of elements of $\text{Jac}(G)$ is the product of the invariant factors of \tilde{L} . Integer row and column

operations do not affect the determinant of a matrix except, possibly, for the sign. Thus, $\det(\tilde{L})$ is also the product of its invariant factors, up to sign (but see the following remark). \square

Remark 2.38. A *spanning tree* of G is a connected subgraph containing all of the vertices of G and no cycles. In Chapter 9 we will prove the matrix-tree theorem which shows that $\det(\tilde{L})$ counts something—the number of spanning trees of G ! Hence, $\det(\tilde{L})$ is necessarily positive, and

$$|\text{Jac}(G)| = \det(\tilde{L}).$$

Notes

The notation div for the Laplacian viewed as the divisor homomorphism first appears in the paper [7] by Baker and Norine. For more about lattice points in convex polyhedra and polytopes, see [10]. In a series of papers ([69], [70], [71], [72]), D. Lorenzini analyzed the Smith normal form of the Laplacian to relate the structure of a graph to the structure of its Jacobian group. We discuss Lorenzini's work further in Section 12.2, where we extend some results from [70] concerning the minimal number of generators of $\text{Jac}(G)$.

Problems for Chapter 2

2.1. Let $\tilde{L} := \tilde{L}_q$ and $\tilde{L}' := \tilde{L}_{q'}$ be the reduced Laplacians of a graph G with respect to vertices q and q' , respectively. Exhibit an explicit isomorphism $\text{cok}(\tilde{L}) \simeq \text{cok}(\tilde{L}')$ stemming from Proposition 2.12.

2.2. Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be graphs, and pick vertices $v_1 \in V_1$ and $v_2 \in V_2$. Let G be the graph obtained from G_1 and G_2 by identifying vertices v_1 and v_2 , thus gluing the two graphs together at a vertex.

- Prove that $\text{Jac}(G) \simeq \text{Jac}(G_1) \times \text{Jac}(G_2)$.
- Prove that every finite abelian group is the Jacobian of some graph.
- Does there exist a simple graph G such that $\text{Jac}(G) \simeq \mathbb{Z}_2$? Here, *simple* means that there is at most one edge between every pair of vertices.
- Use part (a) to show that the Jacobian group of a tree is trivial.

2.3. Provide details for the proof of Proposition 2.28 by showing the dashed mapping in commutative diagram 2.2 is well-defined and bijective.

2.4. The uniqueness statement of Theorem 2.23 comes in two versions: case 1 and case 2. Use the Chinese remainder theorem to show the cases are equivalent.

2.5. The *house graph* H is displayed in Figure 7. Determine the structure of $\text{Jac}(H)$




Figure 7. The house graph.

by computing the Smith normal form for the reduced Laplacian of H .

2.6. Compute the Smith normal form for the reduced Laplacian of the complete graph, K_n . (Hint: start the computation by adding columns 2 through $n - 1$ to column 1.) Conclude that

$$\text{Jac}(K_n) \simeq \mathbb{Z}_n^{n-2}.$$

2.7. Let C_n denote the cycle graph with n vertices (n vertices placed on a circle). Use the computation of the Smith normal form of the reduced Laplacian for C_n to give an explicit isomorphism between $\text{Jac}(C_n)$ and \mathbb{Z}_n .

2.8. For each graph G described below, find the structure of $\text{Jac}(G)$ by computing the Smith normal form of the reduced Laplacian of G .

- Let K_n^- be the graph obtained by deleting one edge from the complete graph on n vertices. (See (b) for a generalization.)
- Let $G(n, k)$ be the graph obtained by deleting k pairwise-disjoint edges from the complete graph on n vertices. (Thus, $2k \leq n$.)

- (c) Let $H(n, k)$ be the graph obtained by deleting $k < n - 1$ edges adjacent to the same vertex in the complete graph on n vertices.
- (d) Consider the graph G_n formed by n triangles sharing an edge. Figure 8 illustrates G_4 .




Figure 8. G_4 for Problem 2.8 (d).

- (e) Let P_n be the path graph with vertices v_1, \dots, v_n , as described in Problem 1.3. Let T_n be the graph formed from P_n by adding a new vertex v_0 , and edges $\{v_0, v_i\}$ for $i = 1, \dots, n$. Figure 9 illustrates T_5 .




Figure 9. T_5 for Problem 2.8 (e).

2.9. Let L be the Laplacian matrix of $G = (V, E)$ with vertices $V = \{1, \dots, n\}$. Prove that for all column vectors $x \in \mathbb{Z}^n$

$$x^t L x = \sum_{ij \in E} (x_i - x_j)^2.$$

Thus, if we think of x as a function $V \rightarrow \mathbb{Z}$, then $x^t L x$ is a measure of how far that function is from being constant.

2.10. Let B_k denote the *banana graph* consisting of two vertices v, w connected by k edges. Fix $d \geq 0$ and consider the divisor $D = dv$. As discussed in Section 2.3, the set of firing scripts σ such that $dv - \text{div}(\sigma) \geq 0$ form the lattice points in a convex polyhedron $P \subset \mathbb{R}^2$. Moreover, the set of such firing scripts with $\sigma(w) = 0$ form the lattice points in a convex polytope $\tilde{P} \subset \mathbb{R}$ (in this 1-dimensional case, \tilde{P} is a closed interval). Describe P and \tilde{P} explicitly.

2.11. Let C_n denote the cycle graph with vertices v_0, \dots, v_{n-1} arranged counterclockwise in a circle. Identify each divisor D on C_n with a vector $(D(v_0), \dots, D(v_{n-1}))$ in \mathbb{Z}^n , as usual. Let $D_i = v_i - v_0$ for $i \in \{0, 1, \dots, n-1\}$ be representatives for the elements of $\text{Jac}(G)$ (Problem 1.2).

- (a) Show that two divisors D and D' on C_n of the same degree are linearly equivalent if and only if there is an equality of dot products

$$D \cdot (0, 1, \dots, n-1) = D' \cdot (0, 1, \dots, n-1) \pmod{n}.$$

- (b) Fix a nonnegative integer d . For each i , show that $D \geq 0$ is in the complete linear system $|D_i + dv_0|$ if and only if

$$D \cdot (1, 1, \dots, 1) = d \quad \text{and} \quad D \cdot (0, 1, \dots, n-1) = i \pmod{n}.$$

Algorithms for winning

3.1. Greed

One way to play the dollar game is for each in-debt vertex to attempt to borrow its way out of debt. Figure 1 reveals a problem with this strategy, though: borrowing once from each vertex is the same as not borrowing at all (see Exercise 1.7).




Figure 1. A greedy cycle.

Although it isn't obvious, the good news is that this is the *only* problem with the strategy. This gives an algorithm for the dollar game: repeatedly choose an in-debt vertex and make a borrowing move at that vertex until either (i) the game is won, or (ii) it is impossible to go on without reaching a state in which all of the vertices have made borrowing moves. For example, in Figure 1, after borrowing at v_1 and v_2 , only v_3 is in debt, and borrowing there would mean that we have borrowed at every vertex. Hence, this game is unwinnable.

The procedure is formalized as Algorithm 1.

Proof of validity of the algorithm. First suppose that $D \in \text{Div}(G)$ is winnable. Take any effective divisor D' such that $D \sim D'$, and choose a firing script $\sigma \in \mathcal{M}(G)$

Algorithm 1 Greedy algorithm for the dollar game.

```

1: INPUT:  $D \in \text{Div}(G)$ .
2: OUTPUT: TRUE if  $D$  is winnable; FALSE if not.
3: initialization:  $M = \emptyset \subset V$ , the set of marked vertices.
4: while  $D$  not effective do
5:   if  $M \neq V$  then
6:     choose any vertex in debt:  $v \in V$  such that  $D(v) < 0$ 
7:     modify  $D$  by performing a borrowing move at  $v$ 
8:     if  $v$  is not in  $M$  then
9:       add  $v$  to  $M$ 
10:   else /* required to borrow from all vertices*/
11:     return FALSE /* unwinnable */
12: return TRUE /* winnable */

```

such that

$$D \xrightarrow{\sigma} D'.$$

By adding a constant function, we may assume that $\sigma \leq 0$ and that σ does not have full support, i.e., the following set is nonempty:

$$Z := \{v \in V : \sigma(v) = 0\}.$$

Thus, σ transforms D into an effective divisor by only making borrowing moves and never borrowing at any vertex in Z .

Given this knowledge about D , we now apply the greedy algorithm. If D is effective, we are done. Otherwise, choose a vertex u such that $D(u) < 0$. In order to get u out of debt it will be necessary to make a borrowing move there. Thus, it is necessarily the case that $\sigma(u) < 0$. The algorithm modifies D by borrowing at u . We keep track of this by replacing σ by $\sigma + u$ (in other words, replacing $\sigma(u)$ by $\sigma(u) + 1$). If $\sigma(u) = 0$, then add u to Z . Repeat. Since $\sigma \leq 0$ at each step and the sum of the values of σ increases by one at each step, the algorithm must eventually terminate with D transformed into a linearly equivalent effective divisor.

Now suppose that $D \in \text{Div}(G)$ is unwinnable. Let D_1, D_2, D_3, \dots be any sequence of divisors obtained by starting at $D = D_1$ and then repeatedly borrowing at in-debt vertices. We must show that, eventually, each vertex must borrow. First note that for all $v \in V$ and for all i ,

$$D_i(v) \leq \max\{D(v), \deg_G(v) - 1\} =: B_v.$$

This is because a vertex v will only borrow if it is in debt, in which case it will gain $\deg_G(v)$ dollars. Letting $B := \max\{B_v : v \in V\}$, we see that $D_i(v) \leq B$ for all $v \in V$ and for all i . Let $n = |V|$. Then, since $\deg(D_i) = \deg(D)$ for all i , we see that for each vertex v we have

$$\deg(D) = D_i(v) + \sum_{w \neq v} D_i(w) \leq D_i(v) + (n-1)B.$$

It follows that for all i and for all v , we have

$$\deg(D) - (n-1)B \leq D_i(v) \leq B.$$

Hence, there are only finitely many possibilities for the divisors D_i , so there must exist j and k with $j < k$ such that $D_j = D_k$. Since the kernel of the Laplacian

is generated by $\vec{1}$ (Proposition 2.8), the sequence of vertices at which borrowing moves were made in going from D_j to D_k contains every vertex. \square

3.1.1. Uniqueness of greedy algorithm script. The greedy algorithm can be modified to produce a firing script in the case its input is winnable. Initialize by setting $\sigma = 0$, and then each time step 6 is invoked, replace σ by $\sigma - v$. It turns out that the resulting script is independent of the order in which vertices are added.

Proposition 3.1. *Suppose D is winnable, and let σ_1 and σ_2 be firing scripts produced by applying the greedy algorithm to D , so that firing these scripts from D produces effective divisors E_1 and E_2 respectively. Then $\sigma_1 = \sigma_2$, so that $E_1 = E_2$ as well.*

Proof. Suppose on the contrary that $\sigma_1 \neq \sigma_2$. Then without loss of generality, we may assume that there is a vertex that borrows more times according to σ_2 than according to σ_1 . Let $\{w_1, \dots, w_m\}$ be a greedy-sequence of borrowings corresponding to σ_2 . By our assumption, as we perform this borrowing sequence, there will be a first step, k , when vertex w_k borrows more than $|\sigma_1(w_k)|$ times. Let $\tilde{\sigma}_2$ be the firing script associated with $\{w_1, \dots, w_{k-1}\}$, the first $k-1$ steps of the σ_2 -borrowing sequence. Then (remembering that $\sigma_1, \sigma_2 \leq 0$), we have $\tilde{\sigma}_2 \geq \sigma_1$ and $\tilde{\sigma}_2(w_k) = \sigma_1(w_k)$. Hence, after performing the first $k-1$ steps of the σ_2 -borrowing sequence, the amount of money at vertex w_k is:

$$\begin{aligned} (D - \operatorname{div}(\tilde{\sigma}_2))(w_k) &= D(w_k) - \deg_G(w_k) \tilde{\sigma}_2(w_k) + \sum_{w_k u \in E} \tilde{\sigma}_2(u) \\ &\geq D(w_k) - \deg_G(w_k) \sigma_1(w_k) + \sum_{w_k u \in E} \sigma_1(u) \\ &= (D - \operatorname{div}(\sigma_1))(w_k) \\ &= E_1(w_k) \\ &\geq 0. \end{aligned}$$

So w_k has no debt after the first $k-1$ steps of the σ_2 -borrowing sequence, contradicting the assumption that w_k borrows at the next (k th) step. \square

Given a divisor D , suppose that two players separately employ the greedy algorithm in an attempt to win the dollar game. Then our work in this section has shown that they will either both discover that D is unwinnable (if they are forced to borrow at every vertex), or they will both win the game. Moreover, despite the fact that they may choose different *sequences* of borrowing moves, their final firing scripts will be the same, and they will both end up with the same effective divisor $E \sim D$. At this point, it is natural to ask how E is special compared to the other effective divisors in the complete linear system $|D|$. It turns out that this is best answered in terms of sandpiles: in Chapter 6 we will introduce a notion of *duality*, and show that E is dual to the *stabilization* of the divisor that is dual to the original divisor D .

3.2. q -reduced divisors


In the previous section, we saw how greed can lead to a resolution of the dollar game. One might think that benevolence, in the form of debt-free vertices lending to their in-debt neighbors, might also work, but it does not in general (cf. Problem 3.1). Nonetheless, in this section we present a particular version of benevolence that does solve the dollar game and which will have theoretical significance for us later on.

Start with $D \in \text{Div}(G)$. We would like to find $E \in |D|$, i.e., an effective divisor linearly equivalent to D . Proceed as follows.

- (1) Pick some “benevolent vertex” $q \in V$. Call q the *source* vertex, and let $\tilde{V} := V \setminus \{q\}$ be the set of *non-source* vertices.
- (2) Let q lend so much money that the non-source vertices, sharing among themselves, are out of debt.
- (3) At this stage, only q is in debt, and it makes no further lending or borrowing moves. It is now the job of the non-source vertices to try to relieve q of its debt. Look for $S \subseteq \tilde{V}$ with the property that if all members of S simultaneously lend, then none go into debt. Having found such an S , make the corresponding set-lending move. Repeat until no such S remains. The resulting divisor is said to be *q -reduced*.

Result: In the end, if q is no longer in debt, we win. Otherwise, $|D| = \emptyset$, i.e., D is unwinnable.

Exercise 3.2. Find linearly equivalent q -reduced divisors for the following two divisors and thus determine whether each is winnable:



Several questions immediately come to mind with regard to the above strategy:

- Is it always possible to complete step 2?
- Is step 3 guaranteed to terminate?
- If the strategy does not win the game, does this mean the game is unwinnable? (After all, the moves in step 3 are constrained.)
- Is the resulting q -reduced divisor unique?
- Can the strategy be efficiently implemented?

The main goal of this chapter is to show that the answer to all of these questions is “yes.”

3.2.1. Existence and uniqueness.

Definition 3.3. Let $D \in \text{Div}(G)$, and let $S \subseteq V$. Suppose D' is obtained from D by firing each of the vertices in S once. Then $D \xrightarrow{S} D'$ is a *legal set-firing* if $D'(v) \geq 0$ for all $v \in S$, i.e., after firing S , none of the vertices in S are in debt. In this case, we say it is *legal* to fire S . [Note: if it is legal to fire S , then the vertices in S must also be out of debt *before* firing.]

Definition 3.4. Let $D \in \text{Div}(G)$ and $q \in V$. Then D is *q -reduced* if

- (1) $D(v) \geq 0$ for all $v \in \tilde{V} := V \setminus \{q\}$, and
- (2) for all nonempty $S \subseteq \tilde{V}$, it is not legal to fire S . In other words, for each such S , there exists $v \in S$ such that

$$D(v) < \text{outdeg}_S(v),$$

where $\text{outdeg}_S(v)$ denotes the number of edges of the form $\{v, w\}$ with $w \notin S$.

Recall that a *spanning tree* of G is a connected subgraph T of G containing all the vertices and no cycles. Fixing $q \in V$, we refer to the pair (T, q) as the spanning tree T *rooted at* q .

Let T be a spanning tree of G rooted at the source vertex q . An ordering of the vertices v_1, v_2, \dots, v_n is called a *tree ordering* compatible with T if $i < j$ whenever v_i lies on the unique path from v_j to q in T . In particular, $v_1 = q$ for any tree ordering.

Example 3.5. Consider the spanning tree T of the house graph pictured in Figure 2. Let $q = v_1$ be the root vertex. Two of the many tree orderings of the vertices compatible with T are indicated by the vertex labels v_i and w_i .




Figure 2. A spanning tree of the house graph with two compatible tree orderings.

We can extend a tree ordering of the vertices to a total ordering of all divisors as follows: given $D, D' \in \text{Div}(G)$, set $D' \prec D$ if (i) $\deg(D') < \deg(D)$ or (ii) $\deg(D') = \deg(D)$ and $D'(v_k) > D(v_k)$ for the smallest index k such that $D'(v_k) \neq D(v_k)$. Roughly, $D' \prec D$ if it has “more cash closer to q .” We say that \prec is a *tree ordering rooted at* q .

The essential property of this ordering is that when a set of non-source vertices fires, some dollars move towards q , producing a smaller divisor with respect to the tree ordering. This idea is made precise in the following proposition.

Proposition 3.6. *Let $D \in \text{Div}(G)$, and let \prec be a tree ordering rooted at $q \in V$. Let D' be the divisor obtained from D by firing a nonempty set $S \subseteq \tilde{V}$. Then, $D' \prec D$.*

Proof. Let $v_1 \prec \cdots \prec v_n$ be the tree ordering, and let k be the smallest index such that v_k is adjacent to a vertex of S . Since $q \notin S$ and every element of S has a path to q , it is necessarily the case that $v_k \notin S$. Then $\deg(D') = \deg(D)$ and $D'(v_i) = D(v_i)$ for $i = 1, \dots, k-1$, but $D'(v_k) > D(v_k)$. \square

Theorem 3.7 (Existence and uniqueness for q -reduced divisors). *Let $D \in \text{Div}(G)$, and fix $q \in V$. Then there exists a unique q -reduced divisor linearly equivalent to D .*

Proof. Fix a tree ordering $v_1 \prec \cdots \prec v_n$ compatible with a spanning tree T rooted at q . We first show that by a sequence of lending moves, all vertices except q can be brought out of debt. Each vertex $v \neq q$ has a unique neighbor w in T such that $w \prec v$. Call this neighbor $\epsilon(v)$. Now bring v_n, v_{n-1}, \dots, v_2 out of debt successively, in the order listed, by having $\epsilon(v_i)$ lend to v_i .

We may now assume that $D(v) \geq 0$ for all vertices $v \neq q$. If D is not q -reduced, perform a legal set-firing of non-source vertices to produce an equivalent divisor, still nonnegative away from q , and strictly smaller than D with respect to the tree ordering by Proposition 3.6. Repeat. Say the resulting sequence of divisors is D_1, D_2, D_3, \dots with $D = D_1$. Write $D_i = c_i + k_i q$ where $c_i \in \text{Config}(G, q)$ for all i . We have $\deg(c_1) \geq \deg(c_i) \geq 0$ for all i . Hence, there can be only finitely many distinct divisors D_i in the sequence. But $D_i \neq D_j$ if $i \neq j$ since $D_{i+1} \prec D_i$. Hence, the process stops, producing a q -reduced divisor.

For uniqueness, suppose D and D' are linearly equivalent q -reduced divisors. Then $D = D' - \text{div}(\sigma)$ for some firing script σ . Let $m = \max_{v \in V} \{\sigma(v)\}$, and

$$S = \{v \in V : \sigma(v) = m\}.$$

If $S = V$, then σ is a constant function and $D = D'$. So we may assume $S \neq V$. We may also assume $q \notin S$ (otherwise, swap D and D' , which exchanges σ and $-\sigma$).

If $v \in S$ and $u \notin S$, then $\sigma(v) - \sigma(u) \geq 1$. Therefore, for each $v \in S$, we have

$$0 \leq D(v) = D'(v) - \sum_{vu \in E} (\sigma(v) - \sigma(u)) \leq D'(v) - \text{outdeg}_S(v).$$

Hence, $S \subseteq \tilde{V}$ is a legal set-firing, contradicting the fact that D' is q -reduced. \square

Corollary 3.8. *Let $D \in \text{Div}(G)$, and let D' be the q -reduced divisor linearly equivalent to D . Then $|D| \neq \emptyset$ if and only if $D' \geq 0$. In other words, D is winnable if and only if $D'(q) \geq 0$.*

Proof. Suppose D is winnable, and let $E \in |D|$. From E , perform all legal set-firings of non-source vertices, arriving at the q -reduced divisor $E' \geq 0$. By uniqueness, $D' = E'$, and hence $D' \geq 0$. The converse is immediate. \square

Example 3.9. Let C_n be the cycle graph on n vertices, and let $D = v - q$ for any distinct vertices v and q of C_n . Then D is q -reduced, hence, unwinnable.

It seems reasonable to expect that the complete linear system of a divisor will be small if the degree of the divisor is small or if the graph has many interconnections among its vertices—both of these qualities make it harder to fire collections of vertices without going into debt. We now provide a sound basis for this intuition.

The graph G is d -edge connected if G remains connected whenever $d - 1$ of its edges are removed, and the *edge-connectivity* of G is the largest d such that G is d -edge connected. Equivalently, the edge-connectivity is the minimal number of edge-removals needed to disconnect the graph.

Proposition 3.10. *Suppose that the graph G is d -edge connected. Then every effective divisor on G of degree less than d is q -reduced.*

Proof. Suppose that E is an effective divisor of degree $k < d$, and let $S \subset \tilde{V}$ be an arbitrary non-empty subset of non-source vertices. Note that the number of edges connecting S to $V \setminus S \neq \emptyset$ is given by $\sum_{v \in S} \text{outdeg}_S(v) \geq d$, since G is d -edge connected. But then

$$\sum_{v \in S} E(v) \leq k < d \leq \sum_{v \in S} \text{outdeg}_S(v).$$

It follows that there exists $v \in S$ such that $E(v) < \text{outdeg}_S(v)$, so that E is q -reduced. \square

Proposition 3.11. *Suppose that G has edge-connectivity d . Let $D \in \text{Div}(G)$ be a winnable divisor of degree less than d . Then the complete linear system of D consists of a single effective divisor. For each degree $k \geq d$, there exists a complete linear system of degree k containing more than one effective divisor.*

Proof. Since D is winnable, the linear system of D is nonempty. Suppose that $E_1 \sim E_2 \sim D$ for effective divisors E_1, E_2 . Since

$$\deg(E_1) = \deg(E_2) = \deg(D) < d,$$

the previous proposition implies that E_1 and E_2 are each q -reduced. By the uniqueness of q -reduced divisors (Theorem 3.2.1), it follows that $E_1 = E_2$. Thus, $|D|$ consists of a single effective divisor as claimed.

Since G is not $(d + 1)$ -edge connected, there exists a subset $U \subseteq E(G)$ of size d whose removal disconnects the graph. Let W be a connected component of the disconnected graph $(V(G), E(G) \setminus U)$, and consider the set-firing script $\chi_{V(W)} \in \mathcal{M}(G)$. Then $\text{div}(\chi_{V(W)}) = F_1 - F_2$, where F_1 and F_2 are each effective of degree d , F_1 has support in W , and F_2 has support outside W . In particular, F_1 and F_2 are distinct linearly equivalent effective divisors of degree d . Thus, the complete linear system of F_1 contains at least two effective divisors. But then for any $k > d$, we have that $F_1 + (k - d)q \sim F_2 + (k - d)q$ are distinct linearly equivalent divisors, so that the complete linear system of the degree- k divisor $F_1 + (k - d)q$ contains at least two elements. \square




Figure 3. The superstabilization of a configuration on K_5 .

3.3. Superstable configurations

It will be convenient to translate the notion of a q -reduced divisor in terms of the configurations introduced in Section 2.2.

Definition 3.12. Let $c \in \text{Config}(G)$, and let $S \subseteq \tilde{V}$. Suppose c' is the configuration obtained from c by firing the vertices in S . Then $c \xrightarrow{S} c'$ is a *legal set-firing* if $c'(v) \geq 0$ for all $v \in S$.

Definition 3.13. The configuration $c \in \text{Config}(G)$ is *superstable* if $c \geq 0$ and has no legal nonempty set-firings, i.e., for all nonempty $S \subseteq \tilde{V}$, there exists $v \in S$ such that

$$c(v) < \text{outdeg}_S(v).$$

Example 3.14. In Figure 3, the configuration on the left is not superstable, since $S = \{v_1, v_2, v_3, v_4\}$ is a legal firing set. Firing S yields the superstable configuration on the right.

Remark 3.15. Every divisor D may be written as $c + kq$ where $c \in \text{Config}(G, q)$ and $k \in \mathbb{Z}$. In this form, D is q -reduced if and only if c is superstable, and if so, by Corollary 3.8, D is winnable if and only if $k \geq 0$.

We have the isomorphism of Proposition 2.12:

$$\begin{aligned} \text{Config}(G)/\text{im}(\widetilde{\text{div}}) &\rightarrow \text{Jac}(G) \\ [c] &\mapsto [c - \text{deg}(c)q]. \end{aligned}$$

So $c \sim c'$ as configurations exactly when $c - \text{deg}(c)q \sim c' - \text{deg}(c')q$ as divisors. It follows from existence and uniqueness of q -reduced divisors that each $c \in \text{Config}(G)$ is equivalent to a unique superstable configuration, which we call the *superstabilization* of c .

The collection of superstables of G forms an abelian group where addition is the usual addition of configurations followed by superstabilization. The isomorphism of Proposition 2.12 then identifies this group with $\text{Jac}(G)$.

3.4. Dhar's algorithm and efficient implementation

At the beginning of Section 3.2, we outlined a three-step procedure for determining whether a given divisor is winnable. The proof of the existence and uniqueness

theorem for q -reduced divisors and Corollary 3.8 guarantee the procedure works in principle. However, how quickly can it be made to work in practice?

Of particular concern is step 3. When it is reached, the original divisor has been replaced by a linearly equivalent divisor of the form $c + kq$ where c is a nonnegative configuration. The immediate task is then to find a legal nonempty set-firing for c . But how should we search through the $2^{|\tilde{V}|} - 1$ plausible subsets that might be fired? Amazingly, it can be done very efficiently!

3.4.1. Dhar's algorithm. Let $c \in \text{Config}(G, q)$ be nonnegative. To find a legal set-firing for c , imagine the edges of our graph are made of wood. Vertex q is ignited, and fire spreads along its incident edges. Think of the configuration c as a placement of $c(v)$ firefighters at each $v \in \tilde{V}$. Each firefighter can control the fire coming from a single edge, so a vertex is protected as long as it has at most $c(v)$ burning incident edges. If it ever happens that this number is greater than $c(v)$, the firefighters are overwhelmed and the vertex they are protecting ignites.¹ When a vertex ignites, fire spreads along its incident edges. In the end, the unburnt vertices constitute a set that may be legally fired from c . If this set is empty, then c is superstable.

Example 3.16. Figure 4 illustrates Dhar's algorithm on the house graph with a configuration c . The vertex q is ignited, and fire spreads along the incident edges. A vertex v ignites when the number of red (burning) edges incident to v is greater than $c(v)$. At the end, the two unburnt vertices form a legal firing set for the configuration c .





Figure 4. Dhar's algorithm.

Exercise 3.17. Run Dhar's algorithm on the configuration below and show that all vertices are burnt in the end. Hence, the configuration is superstable.



We formalize the preceding discussion as Algorithm 2. Recall that for $S \subset V$ and $v \in V$, the number of edges of G connecting v to a vertex outside of S is denoted $\text{outdeg}_S(v)$.

¹Don't worry—the firefighters are airlifted out by helicopter.

Algorithm 2 Dhar's algorithm.

```

1: INPUT: a nonnegative configuration  $c$ .
2: OUTPUT: a legal firing set  $S \subset \tilde{V}$ , empty iff  $c$  is superstable.
3: initialization:  $S = \tilde{V}$ 
4: while  $S \neq \emptyset$  do
5:   if  $c(v) < \text{outdeg}_S(v)$  for some  $v \in S$  then
6:      $S = S \setminus \{v\}$ 
7:   else
8:     return  $S$  /*  $c$  is not superstable */
9: return  $S$ 

```

Proof of validity. When S is returned, $c(v) \geq \text{outdeg}_S(v)$ for all $v \in S$. Hence, S is a legal firing set. If c is superstable, then it has no nonempty legal firing set, and hence S is empty in that case. Conversely, suppose that S is empty, and let U be a nonempty subset of \tilde{V} . We must show that U is not a legal firing set for c . At initialization, $U \subseteq S = \tilde{V}$, and during the while-loop, vertices are removed from S one at a time. Since S is empty upon termination, a vertex from U is eventually removed from S . Let u be the first such vertex, and consider the step of the algorithm just before u is removed from S . At that point, U is still a subset of S and $c(u) < \text{outdeg}_S(u)$. The key observation is that since $U \subseteq S$, we have $\text{outdeg}_S(u) \leq \text{outdeg}_U(u)$, and the result follows. \square

Returning to Example 3.16, label the vertices as in Figure 5. Then Figure 4 illustrates Dhar's algorithm applied to the configuration $c = v_4 + 2v_5$. At initialization, $S = \{v_2, v_3, v_4, v_5\}$, corresponding to the unburnt vertices in the first picture in the figure. During the first run of the while loop, vertex v_2 satisfies the condition $c(v_2) = 0 < \text{outdeg}_S(v_2) = 1$, so v_2 is removed from S to yield the new set $S = \{v_3, v_4, v_5\}$ corresponding to the second picture in the figure. In the second run of the while loop, vertex v_3 satisfies the condition $c(v_3) = 0 < \text{outdeg}_S(v_3) = 1$, so it is removed to yield $S = \{v_4, v_5\}$ corresponding to the unburnt vertices in the third picture. At this point, we have $c(v_4) = 1 = \text{outdeg}_S(v_4)$ and $c(v_5) = 2 = \text{outdeg}_S(v_5)$, so the legal firing set S is returned.




Figure 5. Vertex labeling of the house graph.

Remark 3.18. Instead of removing just a single vertex in step 6 of Dhar's algorithm, one could remove any number of vertices that meet the criterion of step 5. For instance, steps 5 and 6 of Dhar's algorithm could be replaced by

$$U = \{v \in \tilde{V} : c(v) < \text{outdeg}_S(v)\}$$

if $U \neq \emptyset$ then
 $S = S \setminus U$.

3.4.2. Greed and the debt-reduction trick. We would now like to further analyze and refine step 2 of our three-step procedure from the beginning of Section 3.2. Given $D \in \text{Div}(G)$, the proof of the existence and uniqueness theorem for q -reduced divisors uses a tree ordering on the vertices of the graph to show that all of the non-source vertices may be brought out of debt through lending moves. Problem 3.5 shows that a simpler greedy algorithm also works: if a non-source vertex is in debt, it borrows. Repeat until every non-source vertex is out of debt.

Example 3.19. Consider the divisor $D = -2v_3 + 3v_4 - v_5 + q$ on the house graph in Figure 5. Figure 6 illustrates the process of borrowing greedily at non-source, in-debt vertices.




Figure 6. Greedy borrowing from non-source vertices.

Before applying this greedy algorithm, you may want to quickly replace D with an equivalent divisor whose non-source vertices are not in too much debt by using the following *debt-reduction trick*. Write $D = c + kq$ with $c \in \text{Config}(G)$. Fixing an ordering of the vertices, v_1, \dots, v_n with $q = v_1$, identify c with a vector $c \in \mathbb{Z}^{n-1}$, as usual. We seek a firing script that will replace c by an equivalent configuration whose components are small. By Corollary 2.15, the reduced Laplacian \tilde{L} is invertible over the rationals. If we could use $\tilde{L}^{-1}c$ as a firing script, we would get the configuration $c - \tilde{L}(\tilde{L}^{-1}c) = 0$, whose components are as small as could be! However, $\tilde{L}^{-1}c$ will not be a valid firing script in general, since it is probably not an integer vector. Seeking the next best thing, we replace $\tilde{L}^{-1}c$ with the close-by smaller integer vector $\lfloor \tilde{L}^{-1}c \rfloor$ obtained by taking the floor of each component.

Firing

$$\sigma := \lfloor \tilde{L}^{-1}c \rfloor \in \mathbb{Z}^{n-1}$$

yields the equivalent configuration $c' := c - \tilde{L}\sigma$. Problem 3.6 shows $|c'(v)| < \deg_G(v)$ for all $v \in \tilde{V}$, which means that for c' any particular non-source vertex v

can be brought out of debt with a single borrowing move. (Of course, borrowing might then increase the debt of v 's neighbors).

Finally, we replace D by the linearly equivalent divisor,

$$D' := c' + (\deg(D) - \deg(c'))q.$$

Example 3.20. Let's apply the debt-reduction trick described above to the divisor $D = -10v_2 - 2v_3 + 11v_4 - 7v_5 + 10q$ on the house graph in Figure 5. Note that $\deg(D) = 2$ and $D = c + 10q$ for the configuration $c = -10v_2 - 2v_3 + 11v_4 - 7v_5$. The reduced Laplacian matrix and its inverse are given by

$$\tilde{L} = \begin{pmatrix} 2 & -1 & 0 & 0 \\ -1 & 3 & -1 & -1 \\ 0 & -1 & 2 & -1 \\ 0 & -1 & -1 & 3 \end{pmatrix}, \quad \tilde{L}^{-1} = \frac{1}{11} \begin{pmatrix} 8 & 5 & 4 & 3 \\ 5 & 10 & 8 & 6 \\ 4 & 8 & 13 & 7 \\ 3 & 6 & 7 & 8 \end{pmatrix}.$$

Applying \tilde{L}^{-1} to the vector $(-10, -2, 11, -7)$ corresponding to c yields $\tilde{L}^{-1}c = (-\frac{67}{11}, -\frac{24}{11}, \frac{38}{11}, -\frac{21}{11})$. Taking the floor of each component defines the firing script $\sigma = (-7, -3, 3, -2)$. Implementing this firing script yields the linearly equivalent configuration $c' = v_2 + v_3 - v_5$:

$$c' = c - \tilde{L}\sigma = (1, 1, 0, -1).$$

Hence, we may replace the original divisor D by the linearly equivalent divisor D' with smaller coefficients:

$$D' = c' + (\deg(D) - \deg(c'))q = v_2 + v_3 - v_5 + q.$$

3.4.3. Summary. We reformulate the procedure from the beginning of Section 3.2.

Algorithm 3 Find the linearly equivalent q -reduced divisor.

- 1: INPUT: $D \in \text{Div}(G)$ and $q \in V$.
 - 2: OUTPUT: the unique q -reduced divisor linearly equivalent to D .
 - 3: *optional*: Apply the debt-reduction trick to D .
 - 4: Use a greedy algorithm to bring each vertex $v \neq q$ out of debt, so we may assume $D(v) \geq 0$ for all $v \neq q$.
 - 5: Repeatedly apply Dhar's algorithm until D is q -reduced.
-

3.5. The Abel-Jacobi map

As seen in Proposition 3.11, the q -reduced concept has immediate implications for the size of complete linear systems on d -edge connected graphs. To discuss these implications in a larger context, we begin by formulating the dollar game with explicit reference to the Picard group:

The dollar game: Given a divisor $D \in \text{Div}(G)$, does there exist an effective divisor, $E \in \text{Div}_+(G)$, such that $[E] = [D]$ in $\text{Pic}(G)$?

In order to phrase this in terms of the Jacobian group, we need a way of shifting questions about degree- d divisors to questions about degree-0 divisors. Of course, $\text{Div}^0(G)$ is a subgroup of $\text{Div}(G)$, while $\text{Div}^d(G)$ is only a *subset* of $\text{Div}(G)$ for $d \neq 0$. But each of these subsets has a natural action by the group $\text{Div}^0(G)$:

$$\begin{aligned} \text{Div}^0(G) \times \text{Div}^d(G) &\rightarrow \text{Div}^d(G) \\ (Z, D) &\mapsto Z + D. \end{aligned}$$

This action is simply-transitive: Given any two degree- d divisors D and D' , there is a unique divisor Z of degree zero such that $D' = Z + D$, namely $Z := D' - D$. To describe this situation (a natural simply-transitive group action), we say that $\text{Div}^d(G)$ is a *torsor* for the group $\text{Div}^0(G)$.

Moreover, for each $d \in \mathbb{Z}$, the subgroup of principal divisors acts on $\text{Div}^d(G)$ by restriction of the natural $\text{Div}^0(G)$ -action:

$$\begin{aligned} \text{Prin}(G) \times \text{Div}^d(G) &\rightarrow \text{Div}^d(G) \\ (\text{div}(f), D) &\mapsto \text{div}(f) + D. \end{aligned}$$

Define $\text{Pic}^d(G) := \text{Div}^d(G)/\text{Prin}(G)$ to be the quotient set, i.e., the set of linear equivalence classes of degree- d divisors. For $d = 0$, this is the Jacobian group, $\text{Jac}(G) = \text{Div}^0(G)/\text{Prin}(G)$.

Exercise 3.21. Show that each $\text{Pic}^d(G)$ is a torsor for $\text{Jac}(G)$.

You should think of a torsor as a group that has forgotten its identity element. Indeed: choosing any linear equivalence class $\mathcal{Q} \in \text{Pic}^d(G)$ yields a bijection

$$\begin{aligned} \text{Jac}(G) &\rightarrow \text{Pic}^d(G) \\ \mathcal{Z} &\mapsto \mathcal{Z} + \mathcal{Q}, \end{aligned}$$

by which the group law on $\text{Jac}(G)$ may be transferred to $\text{Pic}^d(G)$. In this way $\text{Pic}^d(G)$ is identified with the group $\text{Jac}(G)$, with the chosen divisor class \mathcal{Q} serving as the identity.

The dollar game on G starting with $D \in \text{Div}^d(G)$ asks for an effective divisor $E \in \text{Div}_+^d(G)$ such that $[E] = [D]$ in $\text{Pic}^d(G)$. Subtracting the chosen divisor-class $\mathcal{Q} \in \text{Pic}^d(G)$ yields the following reformulation in terms of the group $\text{Jac}(G)$:

The dollar game: Given a divisor $D \in \text{Div}^d(G)$, does there exist an effective divisor, $E \in \text{Div}_+^d(G)$, such that $[E] - \mathcal{Q} = [D] - \mathcal{Q}$ in $\text{Jac}(G)$?

Thus, at the price of choosing an arbitrary divisor class $\mathcal{Q} \in \text{Pic}^d(G)$, we have succeeded in casting the dollar game as a question about the Jacobian group $\text{Jac}(G)$. At this point, it is reasonable to ask whether some choices for \mathcal{Q} are better than others. The following proposition suggests that the answer is yes.

Proposition 3.22. *The following sequence is exact:*

$$0 \rightarrow \text{Jac}(G) \rightarrow \text{Pic}(G) \xrightarrow{\text{deg}} \mathbb{Z} \rightarrow 0.$$

The set of splittings $\{s: \mathbb{Z} \rightarrow \text{Pic}(G) : \deg \circ s = \text{id}\}$ is in bijection with $\text{Pic}^1(G)$, the divisor-classes of degree 1. In particular, the choice of any vertex $q \in V$ determines a splitting s_q defined by $s_q(1) := [q]$.

Exercise 3.23. Prove Proposition 3.22 and show that each choice of splitting determines an isomorphism $\text{Pic}(G) \simeq \text{Jac}(G) \times \mathbb{Z}$.

The previous proposition and exercise suggest that in choosing the divisor class \mathcal{Q} , we should begin by fixing a vertex $q \in V$, and then define $\mathcal{Q} := [dq]$. This procedure simultaneously selects a divisor in each degree, thereby identifying each torsor $\text{Pic}^d(G)$ with $\text{Jac}(G)$ in a way compatible with the isomorphism induced by the splitting s_q :

$$\begin{aligned} \text{Pic}(G) &\xrightarrow{\sim} \text{Jac}(G) \times \mathbb{Z} \\ \mathcal{C} &\mapsto (\mathcal{C} - [\deg(\mathcal{C})q], \deg(\mathcal{C})). \end{aligned}$$

To summarize: by choosing a vertex $q \in V$, the dollar game on G is transformed into the following question about the Jacobian group $\text{Jac}(G)$. (Note that the choice of a vertex q is the same as the specification of a “benevolent vertex” in the strategy that led us to the notion of q -reduced divisors.)

The dollar game: Given a divisor $D \in \text{Div}^d(G)$, does there exist an effective divisor, $E \in \text{Div}_+^d(G)$, such that $[E - dq] = [D - dq]$ in $\text{Jac}(G)$?

We are thus led to consider the following family of maps from the graded pieces of the monoid $\text{Div}_+(G)$ to the Jacobian $\text{Jac}(G)$:

$$\begin{aligned} S_q^{(d)}: \text{Div}_+^d(G) &\rightarrow \text{Jac}(G) \\ E &\mapsto [E - dq]. \end{aligned}$$

These maps are completely determined by additive extension from $S_q := S_q^{(1)}$, which we think of as being defined on the vertices of G . This map is called the *Abel-Jacobi* map, $S_q: V \rightarrow \text{Jac}(G)$, and is defined as $S_q(v) := [v - q]$. As shown below, it enjoys a universal property related to *harmonic functions* on the graph G .

Definition 3.24. Let A be an abelian group, and $h: V \rightarrow A$ a function defined on the vertices of G . Then h is *harmonic* if for all $v \in V$,

$$\deg_G(v)h(v) = \sum_{vw \in E} h(w).$$

In words: the value of h at v is the average of the values of h at the neighbors of v in G .

Note that for $A = \mathbb{Z}$, the harmonic functions are exactly given by the kernel of the Laplacian $L: \mathcal{M}(G) \rightarrow \text{Div}(G)$. In Chapter 2, we saw that this kernel consists of functions that are constant on the connected components of G . More generally, for an arbitrary abelian group A , we can think of the integral Laplacian matrix as defining a mapping $L_A: \mathcal{M}(G, A) \rightarrow \text{Div}(G, A)$ from A -valued functions on G to divisors with A -coefficients. The A -valued harmonic functions are given by the kernel of L_A .

Exercise 3.25. Show that any constant function $a: V \rightarrow A$ is harmonic. Consequently, every harmonic function $h: V \rightarrow A$ may be written uniquely as $h = \tilde{h} + h(q)$, where \tilde{h} is harmonic with $\tilde{h}(q) = 0$.

Exercise 3.26. Show that the Abel-Jacobi map $S_q: V \rightarrow \text{Jac}(G)$ is harmonic. Moreover, given any group homomorphism $\phi: \text{Jac}(G) \rightarrow A$, show that the composition $\phi \circ S_q: V \rightarrow A$ is a harmonic function sending q to $0 \in A$.

By the previous exercise, we can find A -valued harmonic functions on G by instead finding homomorphisms from $\text{Jac}(G)$ to A . In fact, the next proposition shows that S_q is universal for harmonic functions on G . That is, every harmonic function (sending q to 0) arises from a group homomorphism. In effect, the Abel-Jacobi map reduces the harmonic function theory of the graph G to group theory.

Proposition 3.27. *Suppose that $h: V \rightarrow A$ is harmonic on G with $h(q) = 0$. Then there exists a unique homomorphism $\hat{h}: \text{Jac}(G) \rightarrow A$ such that $h = \hat{h} \circ S_q$:*

$$\begin{array}{ccc} V & \xrightarrow{S_q} & \text{Jac}(G) \\ h \downarrow & \swarrow \exists! \hat{h} & \\ A & & \end{array} .$$

Proof. Given the harmonic function h , consider the unique extension to the free abelian group $\text{Div}(G) = \mathbb{Z}V$, and denote by $h_0: \text{Div}^0(G) \rightarrow A$ the restriction to the subgroup of degree-0 divisors. For any vertex $v \in V$, we have

$$\begin{aligned} h_0(\text{div}(\chi_v)) &= h_0(\text{deg}_G(v)v - \sum_{vw \in E} w) \\ &= \text{deg}_G(v)h(v) - \sum_{vw \in E} h(w) = 0. \end{aligned}$$

Since the principal divisors of the form $\text{div}(\chi_v)$ generate the subgroup $\text{Prin}(G)$, this shows that h_0 induces a homomorphism $\hat{h}: \text{Jac}(G) = \text{Div}^0(G)/\text{Prin}(G) \rightarrow A$. For any $v \in V$,

$$(\hat{h} \circ S_q)(v) = \hat{h}([v - q]) = h(v) - h(q) = h(v),$$

where in the last step we use the assumption that $h(q) = 0$. The uniqueness of \hat{h} follows from the uniqueness of h_0 . \square

Proposition 3.11 may be reformulated as a simple statement about the Abel-Jacobi map.

Proposition 3.28. *The map $S_q^{(d)}: \text{Div}_+^d(G) \rightarrow \text{Jac}(G)$ is injective if and only if G is $(d + 1)$ -edge connected.*

Proof. First suppose that G is $(d + 1)$ -edge connected and that $E_1, E_2 \in \text{Div}_+^d(G)$ are effective of degree d with $S_q^{(d)}(E_1) = S_q^{(d)}(E_2)$. Explicitly, this means that $E_1 - dq \sim E_2 - dq$, which implies that $E_1 \sim E_2$, so that $E_2 \in |E_1|$. But by Proposition 3.11, we have $|E_1| = \{E_1\}$, so that $E_2 = E_1$ and $S_q^{(d)}$ is injective.

Now suppose that $S_q^{(d)}$ is injective, and that G is k -edge connected but not $(k + 1)$ -edge connected. We wish to show that $d + 1 \leq k$. By the injectivity of $S_q^{(d)}$, if $F_1 \neq F_2$ are distinct effective divisors of degree d , then F_2 is not linearly equivalent to F_1 . This means that every nonempty complete linear system of degree d consists

of exactly one effective divisor. But by Proposition 3.11, there exists a complete linear system of each degree greater than or equal to k containing more than one element. It follows that $d < k$, so that $d + 1 \leq k$ as required. \square

Having settled the question of the injectivity of the maps $S_q^{(d)}$, it is natural to wonder about their surjectivity. We will be able to provide a simple answer to this question in the next chapter (cf. Proposition 4.10).

Notes

Most of the results appearing in this chapter have their source in the paper [6] by Baker and Norine or the paper [8] by Baker and Shokrieh. In particular, see [8] for a thorough analysis of the complexity of the algorithms presented here. We learned of the greedy algorithm for the dollar game through conversations with Spencer Backman who further pointed us to the work by Björner et al. on the relation between greedoids and chip-firing games ([18]). The original version of Dhar's algorithm appears in [34].

Problems for Chapter 3

3.1. Consider the following proposal for a “benevolent” algorithm to solve the dollar game, mirroring the greedy algorithm presented in this chapter:

Let D be a divisor. Pick any pair of adjacent vertices v and w such that v is debt-free and w is in debt, and have v perform a lending move. Repeat until each vertex is out of debt or until a state is reached at which continuing would force each vertex to have made a lending move, at which point the game is declared unwinnable.

Give an example of a winnable divisor D which is declared unwinnable by the above procedure.

3.2. Let C_4 be the cycle graph with vertices v_1, v_2, v_3, v_4 listed in order around the cycle. Let $D = -2v_1 - 2v_2 + 2v_3 + 3v_4 \in \text{Div}(C_4)$.

- (a) Find the firing script σ determined by the greedy algorithm for the dollar game.
- (b) Let L be the Laplacian for C_4 . Verify that $D - L\sigma$ is effective.

3.3. Let C_n be the cycle graph with vertices v_1, \dots, v_n , listed in order around the cycle. Suppose $n = 2m$ with $m > 1$, and let $D = -2v_1 + 2v_{m+1} \in \text{Div}(C_n)$. What is the firing script produced by the greedy algorithm for the dollar game applied to D ?

3.4. Let G be the house graph pictured in Figure 7 in the Problems for Chapter 2. Let $q = v_1$, and compute linearly equivalent q -reduced divisors for the following divisors on G :

- (a) $D_1 = (-3, 2, 4, -2, 1)$
- (b) $D_2 = (2, 1, -5, 2, 2)$
- (c) $D_3 = (0, -2, -2, 0, 0)$.

3.5. Let $D \in \text{Div}(G)$ and fix a source vertex $q \in V$. The proof of the existence and uniqueness theorem for q -reduced divisors guarantees the existence of a firing script σ such that $D \xrightarrow{\sigma} D'$ where $D'(v) \geq 0$ for all $v \neq q$. Using this fact, mimic the proof of the validity of Algorithm 1 in Chapter 3 to verify that the greedy algorithm in Algorithm 4 brings the non-source vertices of D out of debt.

Algorithm 4 Greedy algorithm for Problem 3.5.

```

1: INPUT:  $D \in \text{Div}(G)$ .
2: OUTPUT:  $D' \sim D$  such that  $D'(v) \geq 0$  for all  $v \in \tilde{V}$ .
3: initialization:  $D' = D$ .
4: while  $D'|_{\tilde{V}} \not\geq 0$  do
5:   choose  $v \in \tilde{V}$  such that  $D'(v) < 0$ 
6:   modify  $D'$  by performing a borrowing move at  $v$ 
7: return  $D'$ 

```

3.6. Let $c \in \text{Config}(G, q)$. Fixing an ordering of the vertices, let \tilde{L} be the reduced Laplacian of G and identify c with an integer vector in \mathbb{Z}^{n-1} , as usual. Let

$$\sigma := \lfloor \tilde{L}^{-1} c \rfloor \in \mathbb{Z}^{n-1}$$

be the integer vector obtained from $\tilde{L}^{-1}c$ by taking the floor of each of its components. Define $c' = c - \tilde{L}\sigma$. Prove that $|c'(v)| < \deg_G(v)$ for all $v \in \tilde{V}$.

Acyclic orientations

In Section 3.4 we introduced Dhar's algorithm which provides an efficient method for computing the q -reduced divisor linearly equivalent to a given divisor D . In this chapter we investigate Dhar's algorithm further and show that it may be used to establish a bijection between two seemingly unrelated structures on a graph G : the set of *maximal unwinnable* q -reduced divisors and the set of *acyclic orientations with unique source* q . This bijection will play a crucial role in the proof of the Riemann-Roch theorem for graphs (Theorem 5.9) in Chapter 5.


4.1. Orientations and maximal unwinnables

If we are interested in determining which divisors on a graph are winnable and which are unwinnable, then a good place to start is with the *maximal unwinnable divisors*: those unwinnable divisors D such that given any unwinnable divisor D' such that $D \leq D'$, it follows that $D = D'$. Equivalently, D is a maximal unwinnable if D is unwinnable but $D + v$ is winnable for each $v \in V$. Similarly, a *maximal superstable* configuration is a superstable c such that given any superstable c' with $c \leq c'$, we have $c = c'$.

Exercise 4.1. Let D and D' be linearly equivalent divisors. Explain why D is maximal unwinnable if and only if D' is maximal unwinnable.

As a consequence of Remark 3.15, it is clear that if the divisor D is a maximal unwinnable, then its q -reduced form is $c - q$ where c is a maximal superstable. The converse is not immediately obvious but will emerge as a consequence of our work with acyclic orientations in this chapter (cf. Corollary 4.9). It then follows that every unwinnable divisor is dominated by some (not necessarily unique) maximal unwinnable divisor. Thus, we may find all of the unwinnable divisors on a graph by simply removing dollars from the maximal unwinnables.

One of the key ideas in the divisor theory of graphs is that Dhar's algorithm provides a connection between the maximal superstables and the *acyclic orientations*, defined below. If $\{u, v\}$ is an edge of G , we think of the ordered pair $e = (u, v)$ as an assignment of a direction or orientation for the edge. We write $e^- = u$ and $e^+ = v$ and think of e as an arrow from u to v :



Each edge has exactly two possible orientations. An *orientation* \mathcal{O} for G consists of a choice of direction for each edge. See Figure 1 for an example of an orientation for a graph with a multiset of edges.




Figure 1. An acyclic orientation with unique source v_1 .

A vertex u is a *source* for the orientation \mathcal{O} if every edge incident to u is directed away from u , i.e, appears as $e \in \mathcal{O}$ with $e^- = u$. Dually, a vertex v is a *sink* for \mathcal{O} if every edge incident to v is directed towards v , i.e, appears as $e \in \mathcal{O}$ with $e^+ = v$. An orientation is *acyclic* if, as in Figure 1, it contains no cycle of directed edges. Note that if there are multiple edges between two vertices, they must be oriented in the same direction in any acyclic orientation.

Exercise 4.2. Explain why an acyclic orientation must have at least one sink and one source. Is it possible to have more than one of either?

If $u \in V$ and \mathcal{O} is an orientation, define

$$\begin{aligned} \text{indeg}_{\mathcal{O}}(u) &= |\{e \in \mathcal{O} : e^+ = u\}| \\ \text{outdeg}_{\mathcal{O}}(u) &= |\{e \in \mathcal{O} : e^- = u\}|. \end{aligned}$$

The function $v \mapsto \text{indeg}_{\mathcal{O}}(v)$ is the *indegree sequence* for \mathcal{O} .

Lemma 4.3. *An acyclic orientation is determined by its indegree sequence: if \mathcal{O} and \mathcal{O}' are acyclic orientations of G and $\text{indeg}_{\mathcal{O}}(v) = \text{indeg}_{\mathcal{O}'}(v)$ for all $v \in V$, then $\mathcal{O} = \mathcal{O}'$.*

Proof. Given an acyclic orientation \mathcal{O} on G , let $V_1 \subset V$ be its (nonempty) set of source vertices. These are exactly the vertices v such that $\text{indeg}_{\mathcal{O}}(v) = 0$. Remove these vertices and their incident edges from G and from \mathcal{O} to get an acyclic orientation \mathcal{O}_1 on a subgraph G_1 of G . Repeat, letting V_2 be the sources of \mathcal{O}_1 , etc. In this way, we get a sequence (V_1, V_2, \dots) partitioning V . The result follows by noting that the sequence determines \mathcal{O} and is determined by the indegree sequence of \mathcal{O} . \square

Exercise 4.4. What is the acyclic orientation of the graph in Figure 1 with indegree sequence $(v_1, \dots, v_5) \mapsto (2, 0, 3, 2, 0)$?

4.2. Dhar's algorithm revisited

Given a superstable c , place $c(v)$ firefighters on each vertex v , ignite the source vertex, and let the fire spread according to Dhar's algorithm. This time, though, whenever fire spreads along an edge from vertex u to vertex v , record this fact by directing that edge from u to v . (Direct the edge even if there are enough firefighters on v at this point to prevent v from igniting.) In the end, since c is superstable, all the vertices are ignited, but what can be said about the oriented edges?

Example 4.5. Consider the superstable shown on the graph in Figure 2. The figure shows the process of running Dhar's algorithm as described above. Observe that every edge burns, and the resulting orientation is acyclic with unique source at q .




Figure 2. Running Dhar's algorithm on a superstable.

Exercise 4.6. Find a maximal superstable c on the house graph such that running Dhar's algorithm on c produces the acyclic orientation shown in Figure 3 below. Note that this orientation has a unique source q .




Figure 3. An acyclic orientation with unique source.

Definition 4.7. The divisor corresponding to an orientation \mathcal{O} of G is

$$\mathbf{D}(\mathcal{O}) = \sum_{v \in V} (\text{indeg}_{\mathcal{O}}(v) - 1) v.$$

The corresponding configuration (with respect to $q \in V$) is

$$\mathbf{c}(\mathcal{O}) = \sum_{v \in \tilde{V}} (\text{indeg}_{\mathcal{O}}(v) - 1) v.$$

For example, if \mathcal{O} is the orientation of Figure 1, then $\mathbf{D}(\mathcal{O}) = -v_1 + v_2 + 2v_5$ and $\mathbf{c}(\mathcal{O}) = v_2 + 2v_5$, taking $q = v_1$.

Theorem 4.8. *Fix $q \in V$. The correspondence $\mathcal{O} \mapsto \mathbf{c}(\mathcal{O})$ gives a bijection between acyclic orientations of G with unique source q and maximal superstables of G (with respect to q).*

Proof. Let \mathcal{O} be an acyclic orientation with unique source q , and let $c = \mathbf{c}(\mathcal{O})$. We first show that c is superstable using Dhar's algorithm as stated on page 52. To initialize the algorithm, start with S equal to an *ordered* list of the elements of \tilde{V} such that that if $u, v \in \tilde{V}$ and $(u, v) \in \mathcal{O}$, then u precedes v in S . Since \mathcal{O} is acyclic, this is always possible, and in fact, there may be several choices.

Let u be the first element of S . The first time through the while-loop of Dhar's algorithm, $\text{outdeg}_S(u)$ is the number of edges connecting u to q (which may be greater than 1 since we are allowing G to be a multigraph). However, since u is first in the list S , we know that $c(u) = \text{indeg}_{\mathcal{O}}(u) - 1 = \text{outdeg}_S(u) - 1$. Hence, we may remove u from S at step 6 and continue. Suppose, by induction, that the algorithm has proceeded until exactly all of the vertices preceding some vertex $v \in S$ have been removed from S in the order listed. If $\{w, v\}$ is an edge with $w \in \tilde{V} \setminus S$, then since w has already been removed from S , it precedes v in our list. Since every edge is oriented, it must be that $(w, v) \in \mathcal{O}$. Moreover, since all vertices preceding v have been removed from S , only the vertices outside of S contribute to the indegree of v . Hence, $\text{outdeg}_S(v) = \text{indeg}_{\mathcal{O}}(v) = c(v) + 1$. So v may be removed from S . In this way, the algorithm continues removing vertices in order until S is empty. It follows that c is superstable.

Algorithm 5, below, associates an acyclic orientation $\mathbf{a}(c)$ with unique sink q to each superstable configuration c . The idea is to run Dhar's algorithm on c , and just before a vertex v is removed from S , take each edge connecting v to a vertex outside of S and direct it into v . Each time Step 5 is invoked, a choice is made, but we will see below that in the case of a maximal superstable, the output is independent of these choices.

Algorithm 5 Acyclic orientation algorithm.

- 1: INPUT: a superstable configuration $c \in \text{Config}(G)$.
 - 2: OUTPUT: an acyclic orientation $\mathbf{a}(c)$ of G with unique sink q .
 - 3: **initialization:** $\mathbf{a}(c) = \emptyset$ and $S = \tilde{V}$
 - 4: **while** $S \neq \emptyset$ **do**
 - 5: choose $v \in S$ such that $c(v) < \text{outdeg}_S(v)$
 - 6: **for all** edges $\{u, v\} \in E$ such that $u \in V \setminus S$ **do**
 - 7: $\mathbf{a}(c) = \mathbf{a}(c) \cup \{(u, v)\}$
 - 8: $S = S \setminus \{v\}$
 - 9: return $\mathbf{a}(c)$
-

Considering steps 5–7 in the algorithm, just before a vertex v is removed from the set S we have

$$c(v) < \text{outdeg}_S(v) = \text{indeg}_{\mathbf{a}(c)}(v).$$

It follows that $c \leq \mathbf{c}(\mathbf{a}(c))$ for all superstables c . Hence, if c is maximal, then $c = \mathbf{c}(\mathbf{a}(c))$. So the composition $\mathbf{c} \circ \mathbf{a}$ is the identity mapping on the set of maximal superstables.

On the other hand, suppose that \mathcal{O} is any acyclic orientation with unique source q , and set $c = \mathbf{c}(\mathcal{O})$. Since $c \leq \mathbf{c}(\mathbf{a}(c))$,

$$(4.1) \quad \text{indeg}_{\mathcal{O}}(v) \leq \text{indeg}_{\mathbf{a}(c)}(v)$$

for all $v \in V$. However, since \mathcal{O} and $\mathbf{a}(c)$ are both acyclic orientations,

$$\sum_{v \in V} \text{indeg}_{\mathcal{O}}(v) = \sum_{v \in V} \text{indeg}_{\mathbf{a}(c)}(v) = |E|.$$

Therefore, we must have equality in (4.1), and thus, by Lemma 4.3 $\mathcal{O} = \mathbf{a}(c) = \mathbf{a}(\mathbf{c}(\mathcal{O}))$. We have shown \mathbf{c} and \mathbf{a} are inverse mappings between maximal superstable and acyclic orientations with unique source q . Consequently, \mathbf{c} is a bijection as claimed. \square

In the following corollary, we introduce the term *genus* for the quantity $|E| - |V| + 1$. Among graph theorists, this quantity is usually called the *cyclomatic number* or *cycle rank*, since it counts the number of independent cycles in the graph (see Chapter 13). Further, a graph theorist is more likely to use the word “genus” to refer to the smallest number γ such that the graph can be embedded without edge-crossings on a compact orientable surface of genus γ (i.e., a donut with γ holes). Thus, for example, a planar graph would have genus 0 since it can be embedded on a sphere with no crossings. Our use of the word “genus” is dictated by the main goal for this part of the book—the Riemann-Roch formula for graphs (Theorem 5.9)—in which it plays the role of the usual genus for Riemann surfaces.

Corollary 4.9. *Let $g = |E| - |V| + 1$, the genus of G .*

- (1) *A superstable c is maximal if and only if $\deg(c) = g$.*
- (2) *A divisor D is maximal unwinnable if and only if its q -reduced form is $c - q$ where c is a maximal superstable.*
- (3) *The correspondence $\mathcal{O} \mapsto \mathbf{D}(\mathcal{O})$ gives a bijection between acyclic orientations of G with unique source q and maximal unwinnable q -reduced divisors of G .*
- (4) *If D is a maximal unwinnable divisor, then $\deg(D) = g - 1$. Hence, if $\deg(D) \geq g$, then D is winnable.*

Proof. Let c be superstable. Applying Algorithm 5 gives an acyclic orientation $\mathcal{O} = \mathbf{a}(c)$ with unique sink q such that $c(v) \leq \mathbf{c}(\mathcal{O})(v)$ for all $v \in \tilde{V}$. We have seen that $\mathbf{c}(\mathcal{O})$ is a maximal superstable. Then

$$\begin{aligned} \deg(c) &\leq \deg(\mathbf{c}(\mathcal{O})) \\ &= \sum_{v \in \tilde{V}} (\text{indeg}_{\mathcal{O}}(v) - 1) \\ &= \sum_{v \in \tilde{V}} \text{indeg}_{\mathcal{O}}(v) - \sum_{v \in \tilde{V}} 1 \\ &= |E| - (|V| - 1), \end{aligned}$$

and the inequality on the first line is an equality if and only if $c(v) = \mathbf{c}(\mathcal{O})(v)$ for all $v \in \tilde{V}$. Part (1) follows.

As stated earlier, a consequence of Remark 3.15 is that if the divisor D is a maximal unwinnable, then its q -reduced form is $c - q$ where c a maximal superstable. For the converse, we may assume that $D = c - q$ with c a maximal superstable. Then D is unwinnable since $D(q) < 0$. To show D is a maximal unwinnable, let $v \in V$. If $v = q$, then clearly, $D + v$ is winnable. If $v \neq q$, then $D = (c + v) - q$ with $c + v \in \text{Config}(G)$. To compute the q -reduced form of $D + v$, we superstabilize $c + v$. By part (1), the degree of $c + v$ is $g + 1$ and the degree of its superstabilization is at most g . Hence, at least one dollar is sent to q , showing that $D + v$ is winnable. This completes the proof of part (2).

Parts (3) and (4) now follow immediately. \square

Note that the converse of Corollary 4.9 (4) is clearly false: any effective divisor of degree $g - 1$ on a graph of genus g is winnable. On the other hand, Corollary 4.9 provides an answer to the second question from the end of Chapter 1: if $\deg(D) \geq g = |E| - |V| + 1$, then D is winnable, and g is the minimal degree with this property. This degree-condition on winnability also allows us to answer our question from the end of Chapter 3 about the surjectivity of the Abel-Jacobi maps.

Proposition 4.10. *The map $S_q^{(d)}: \text{Div}_+^d(G) \rightarrow \text{Jac}(G)$ is surjective if and only if $d \geq g$, the genus of the graph G .*

Proof. Suppose that $[D] \in \text{Jac}(G)$ is an arbitrary divisor class of degree zero. For $d \geq g$, consider the divisor $D + dq$. By Corollary 4.9, there exists an effective divisor $E \in \text{Div}_+^d(G)$ such that $E \sim D + dq$. It follows that $S_q^{(d)}(E) = [E - dq] = [D]$. Since $[D] \in \text{Jac}(G)$ was arbitrary, this shows the surjectivity of $S_q^{(d)}$ for $d \geq g$.

On the other hand, by Corollary 4.9, there exists a maximal unwinnable divisor F of degree $g - 1$. Consider the degree-0 divisor class $[F - (g - 1)q] \in \text{Jac}(G)$. If this element were in the image of $S_q^{(d)}$ for some $d \leq g - 1$, then we would have $[E - dq] = [F - (g - 1)q]$ for some effective divisor E of degree d . But then F would be linearly equivalent to the effective divisor $E + (g - 1 - d)q$, contradicting the unwinnability of F . Hence, the maps $S_q^{(d)}$ are not surjective for $d \leq g - 1$. \square

The following proposition concerns acyclic orientations that do not necessarily have a unique source.

Proposition 4.11. *Let \mathcal{O} be any acyclic orientation. Then $\mathbf{D}(\mathcal{O})$ is a maximal unwinnable divisor.*

Proof. Let $D = \mathbf{D}(\mathcal{O})$. Let v be any source of \mathcal{O} , and let \mathcal{O}' be the acyclic orientation obtained by reversing the direction of each edge incident on v . We call this operation a *vertex reversal*. Then $D' = \mathbf{D}(\mathcal{O}')$ is the divisor obtained from D by firing vertex v . Hence, $D' \sim D$.

We now show that through a series of vertex reversals, \mathcal{O} can be transformed into an acyclic orientation with a unique source. The result then follows from Theorem 4.8. Start by fixing a source v for \mathcal{O} . Let S be the set consisting of v and its adjacent vertices. Clearly, v is the only source vertex in S . Next, choose any vertex $v' \in V \setminus S$ that is adjacent to a vertex in S . If v' is a source, reverse the directions of its incident edges. In any case, add v' to S . It is still the case that v is

the unique source vertex in S . Repeat until $S = V$ to obtain an acyclic orientation with unique source v . \square

Notes

The correspondence between maximal superstable configurations and acyclic orientations has been noted many times, in different forms ([11], [15], [49], [50]). See also [3] and [57].

Problems for Chapter 4

- 4.1. If D is an unwinnable divisor, why must there exist a maximal unwinnable divisor D' such that $D \leq D'$?
- 4.2. Let G be the house graph from Problem 2.5, and take $q = v_1$. Find all maximal superstables on G and their corresponding acyclic orientations.
- 4.3. Let G be a graph of genus g . Use Corollary 4.9 (4) to show that $g = 0$ if and only if $v \sim w$ for all vertices v, w in G . Thus, G is a tree if and only if $\text{Jac}(G) = 0$.

Riemann-Roch

We now reach the culmination of Part 1 of this book, revealing an elegant relationship between the *rank* of a divisor D and of a certain *canonically dual* divisor $K - D$. This formula was first established in the 2007 paper [6] by M. Baker and S. Norine, from which much of the material in this book is derived. The result immediately captured the attention of mathematicians from a variety of fields, due to the fact that it provides a discrete, graph-theoretic version of a central result from classical algebraic geometry: the Riemann-Roch theorem for Riemann surfaces, which we describe in Section 5.3. For this reason, Theorem 5.9 due to Baker and Norine is known as the Riemann-Roch theorem for graphs.

5.1. The rank function

In one sense, the “degree of winnability” of the dollar game is measured by the size of complete linear systems: D is “more winnable” than D' if $\#|D| > \#|D'|$. As shown in Section 2.3, determining the size of $|D|$ requires counting the number of lattice points in a certain convex polytope. However, there is another, less obvious, measure of winnability that turns out to be central to our main goal—the Riemann-Roch theorem for graphs. Instead of demanding an exact determination of $\#|D|$, we will introduce a quantity $r(D) \in \mathbb{Z}$ that measures “robustness of winnability.”

To begin, define $r(D) = -1$ if D is an unwinnable divisor. That is:

$$r(D) = -1 \iff |D| = \emptyset.$$

Next, define $r(D) = 0$ if D is *barely winnable* in the sense that D is winnable, but there exists a vertex $v \in V$ such that $D - v$ is unwinnable. That is, $r(D) = 0$ if and only if the winnability of D can be destroyed by a single vertex losing a dollar. In general, for $k \geq 0$ we define

$$r(D) \geq k \iff |D - E| \neq \emptyset \quad \text{for all effective } E \text{ of degree } k.$$

In words: $r(D) \geq k$ if and only if the dollar game on G is winnable starting from all divisors obtained from D by removing k dollars from the graph. It follows that

$r(D) = k$ if and only if $r(D) \geq k$ and there exists an effective divisor E of degree $k + 1$ such that $D - E$ is not winnable.

Exercise 5.1. Show that $r(D) \leq \max\{-1, \deg(D)\}$ for all divisors D .

Exercise 5.2. Show that if D is a divisor of degree 0, then $r(D) = 0$ if and only if D is principal.

Exercise 5.3. Show that $r(D) \leq r(D + v) \leq r(D) + 1$ for all $D \in \text{Div}(G)$ and all $v \in V$. That is, adding one dollar to a divisor can increase its rank by at most one.

Exercise 5.4. Show that if $r(D)$ and $r(D')$ are each nonnegative, then $r(D + D') \geq r(D) + r(D')$.

Proposition 5.5. *Suppose that G is a d -edge connected graph and $E \in \text{Div}_+(G)$ is an effective divisor of degree less than d . Then $r(E) = \min\{E(v) : v \in V(G)\}$.*

Proof. Choose a vertex v with minimal coefficient $E(v)$. Set $m := E(v) + 1$, and suppose that $|E - mv|$ is nonempty. Then there exists an effective divisor E' such that $E - mv \sim E'$, or equivalently $E \sim E' + mv$. But by Proposition 3.11, this linear equivalence must be an equality: $E = E' + mv = E' + (E(v) + 1)v$. This is a contradiction since v appears on the right hand side at least once more than on the left. It follows that $|E - mv| = \emptyset$, so $r(E) \leq E(v)$. On the other hand, the minimality of $E(v)$ implies that if F is any effective divisor of degree less than or equal to $E(v)$, then $E - F$ is effective. It follows that $r(E) = E(v)$ as claimed. \square

It might seem that the rank function, $D \mapsto r(D)$, is even more difficult to compute than the size of the complete linear system $|D|$. Indeed, to make a straightforward computation of $r(D)$, one would need to answer the entire sequence of questions:

Is $|D|$ nonempty?

If so, then are all $|D - v_i|$ nonempty?

If so, then are all $|D - v_i - v_j|$ nonempty?

If so, then what about all $|D - v_i - v_j - v_k|$?

etc.,

each of which involves the investigation of lattice points in a convex polytope. In fact, the problem of computing the rank of a general divisor on a general graph is **NP**-hard ([62]), which means it is likely that the time it takes for any algorithm to compute the rank of a divisor will grow exponentially with the size of the graph.

Of course, the fact that the rank computation is **NP**-hard in general doesn't mean that it is hard to compute the rank of *every* divisor on a graph G . Indeed, the following proposition provides a lower bound for the rank and shows that the bound is attained for divisors that dominate a maximal unwinnable.

Proposition 5.6. *For all divisors $D \in \text{Div}(G)$, we have the lower bound $r(D) \geq \deg(D) - g$. Moreover, if $D \geq N$ for some maximal unwinnable divisor N , then $r(D) = \deg(D) - g$.*

Proof. First note that, by Corollary 4.9, D is winnable if $\deg(D) \geq g$. Thus, if D is unwinnable, then $r(D) = -1 \geq \deg(D) - g$. Similarly, if D is winnable and $\deg(D) \leq g$, then $r(D) \geq 0 \geq \deg(D) - g$. So it only remains to consider winnable divisors of degree greater than g . Given such a divisor D , let E be any effective divisor of degree at most $\deg(D) - g > 0$. Then $\deg(D - E) \geq g$, so $D - E$ is winnable. By the definition of the rank function, we see that $r(D) \geq \deg(D) - g$ as claimed.

Finally, suppose that $D \geq N$ for some maximal unwinnable divisor N . Consider the effective divisor $E := D - N$ with $\deg(E) = \deg(D) - \deg(N) = \deg(D) - g + 1$. Then $|D - E| = |N| = \emptyset$, which shows that $r(D) \leq \deg(D) - g$. It follows that $r(D) = \deg(D) - g$. \square

In the next section we will see that the Riemann-Roch theorem provides an easy way to determine whether a given divisor D dominates a maximal unwinnable (cf. Corollary 5.12). Those divisors that do not dominate a maximal unwinnable are the ones for which the rank computation is difficult. Nevertheless, we will at least be able to improve upon Exercise 5.1 to provide a nontrivial upper bound for the ranks of such divisors in terms of their degrees (Clifford's theorem, Corollary 5.13).

5.2. Riemann-Roch for graphs

In the previous section we explained that it is difficult, in general, to compute the rank of a divisor. But in Proposition 5.6 we provided a linear lower bound for the rank in terms of the degree of the divisor and the genus of the graph:

$$r(D) \geq \deg(D) - g.$$

In this section we identify the “correction term” that turns this inequality into an equality. In order to do so, we need to introduce the *canonical divisor* of a graph G .

Definition 5.7. For any orientation \mathcal{O} on a graph G , denote by \mathcal{O}^{rev} the *reversed orientation* obtained by reversing the direction of each edge in \mathcal{O} . Define the *canonical divisor* of G to be the divisor

$$K := \mathbf{D}(\mathcal{O}) + \mathbf{D}(\mathcal{O}^{\text{rev}}).$$

Note that for every $v \in V(G)$,

$$K(v) = (\text{indeg}_{\mathcal{O}}(v) - 1) + (\text{outdeg}_{\mathcal{O}}(v) - 1) = \deg_G(v) - 2,$$

so that the canonical divisor depends only on the graph G and not on the orientation \mathcal{O} .

Exercise 5.8. Show that $\deg(K) = 2g - 2$, where $g = |E| - |V| + 1$ is the genus of the graph G .

Theorem 5.9 (Riemann-Roch). *Let D be a divisor on a (loopless, undirected) graph G of genus $g = |E| - |V| + 1$ with canonical divisor K . Then*

$$r(D) - r(K - D) = \deg(D) + 1 - g.$$

Proof. By definition of $r(D)$, there exists an effective divisor F of degree $r(D) + 1$ such that $|D - F| = \emptyset$, i.e., $D - F$ is unwinnable. Find the q -reduced divisor linearly equivalent to $D - F$:

$$D - F \sim c + kq,$$

where c is superstable and $k \in \mathbb{Z}$. Since $D - F$ is unwinnable, $k < 0$. Now pick a maximal superstable $\tilde{c} \geq c$, and consider the maximal unwinnable divisor $\tilde{c} - q$. Let \mathcal{O} be the corresponding acyclic orientation of G . So

$$\mathbf{D}(\mathcal{O}) = \tilde{c} - q \geq c + kq \sim D - F.$$

Define the effective divisor

$$H := (\tilde{c} - c) - (k + 1)q \sim \mathbf{D}(\mathcal{O}) - (D - F).$$

Add $\mathbf{D}(\mathcal{O}^{\text{rev}})$ to both sides of the above relation:

$$\mathbf{D}(\mathcal{O}^{\text{rev}}) + H \sim \underbrace{\mathbf{D}(\mathcal{O}^{\text{rev}}) + \mathbf{D}(\mathcal{O})}_K - (D - F),$$

and rearrange to get

$$K - D - H \sim \mathbf{D}(\mathcal{O}^{\text{rev}}) - F.$$

Since $\mathbf{D}(\mathcal{O}^{\text{rev}})$ is unwinnable and $F \geq 0$, it follows a fortiori that $\mathbf{D}(\mathcal{O}^{\text{rev}}) - F$, and hence, $K - D - H$, is unwinnable. Therefore, $r(K - D) < \deg(H)$. The degree of H is

$$\begin{aligned} \deg(H) &= \deg(\mathbf{D}(\mathcal{O}) - (D - F)) \\ &= \deg(\mathbf{D}(\mathcal{O})) - \deg(D) + \deg(F) \\ &= g - 1 - \deg(D) + r(D) + 1. \end{aligned}$$

Hence,

$$(\star) \quad r(K - D) < g - \deg(D) + r(D),$$

for all divisors D .

Since (\star) holds for all divisors D , we may substitute $K - D$ for D yielding

$$r(D) < g - \deg(K - D) + r(K - D).$$

Since $\deg(K) = 2g - 2$,

$$(\star\star) \quad r(D) < 2 - g + \deg(D) + r(K - D).$$

Combining (\star) and $(\star\star)$,

$$\deg(D) - g \stackrel{\star}{<} r(D) - r(K - D) \stackrel{\star\star}{<} \deg(D) - g + 2.$$

Since $r(D) - r(K - D)$ is an integer, we must have

$$r(D) - r(K - D) = \deg(D) + 1 - g.$$

□

Exercise 5.10. Show that $r(K) = g - 1$.

As first consequences of the Riemann-Roch formula, we prove the following results about maximal unwinnable divisors.

Corollary 5.11. *A divisor N is maximal unwinnable if and only if $K - N$ is maximal unwinnable.*

Proof. If N is maximal unwinnable, then $\deg(N) = g - 1$ and $r(N) = -1$. Riemann-Roch says

$$-1 - r(K - N) = \deg(N) + 1 - g = 0,$$

which implies that $r(K - N) = -1$ so $K - N$ is unwinnable. But $\deg(K - N) = \deg(K) - \deg(N) = 2g - 2 - g + 1 = g - 1$, so $K - N$ is maximal unwinnable. Replacing N by $K - N$ yields the other direction. \square

Corollary 5.12. *Suppose that $D \in \text{Div}(G)$ is a divisor. Then D dominates a maximal unwinnable divisor if and only if $r(K - D) = -1$.*

Proof. First suppose that $D \geq N$ for a maximal unwinnable divisor N . Then $K - D \leq K - N$, and $K - N$ is maximal unwinnable by the previous corollary. It follows that $K - D$ is unwinnable, so $r(K - D) = -1$.

Now suppose that $r(K - D) = -1$. Then $K - D$ is unwinnable, and we may choose a maximal unwinnable that dominates it: $K - D \leq N$. Then $K - N \leq D$, so that D dominates the divisor $K - N$, which is maximal unwinnable by the previous corollary. \square

To what extent does the Riemann-Roch theorem help us determine the rank of a divisor D ? First, we've seen that it is easy to check whether D is unwinnable, i.e., $r(D) = -1$: simply compute the q -reduced divisor linearly equivalent to D and check whether it is effective. Similarly, if D is winnable, it is easy to check if $r(K - D) = -1$. The Riemann-Roch formula then implies $r(D) = \deg(D) - g$. The difficult case is when D is winnable and $r(D - K) \geq 0$. Here, the Riemann-Roch formula only gives us a lower bound on the rank: $r(D) > \deg(D) - g$. The next result provides an upper bound in this case.

Corollary 5.13 (Clifford's Theorem). *Suppose that $D \in \text{Div}(G)$ is a divisor with $r(D) \geq 0$ and $r(K - D) \geq 0$. Then $r(D) \leq \frac{1}{2} \deg(D)$.*

Proof. By Exercises 5.4 and 5.10,

$$r(D) + r(K - D) \leq r(D + K - D) = r(K) = g - 1.$$

Adding this to the Riemann-Roch formula gives $2r(D) \leq \deg(D)$, and dividing by 2 yields the stated upper bound. \square

Corollary 5.14. *Let $D \in \text{Div}(G)$.*

- (1) *If $\deg(D) < 0$, then $r(D) = -1$.*
- (2) *If $0 \leq \deg(D) \leq 2g - 2$, then $r(D) \leq \frac{1}{2} \deg(D)$.*
- (3) *If $\deg(D) > 2g - 2$, then $r(D) = \deg(D) - g$.*

Proof. Part 1 is clear and part 3 follows from the Riemann-Roch formula. For part 2, suppose that $0 \leq \deg(D) \leq 2g - 2$. The result certainly holds if $r(D) = -1$,

so assume $r(D) \geq 0$. If $r(K - D) \geq 0$ as well, then the result holds by Clifford's theorem. Finally, if $r(K - D) = -1$, then by the Riemann-Roch formula,

$$r(D) = \deg(D) - g = \frac{1}{2} \deg(D) + \frac{1}{2}(\deg(D) - 2g) < \frac{1}{2} \deg(D).$$

□

5.3. The analogy with Riemann surfaces

The Riemann-Roch theorem for graphs derives its name from a theorem about surfaces established in the mid-nineteenth century by Bernhard Riemann and his student, Gustav Roch. In fact, the entire theory of divisors on graphs is modeled after the theory of divisors on Riemann surfaces, and in this section we give a brief and informal outline of the classical story. We begin by providing a quick summary of the divisor theory of graphs as presented in the previous chapters.

Our basic object of study has been a finite, undirected, connected multigraph, $G = (V, E)$, as pictured below: The *genus* of a graph is $g := |E| - |V| + 1$, which




Figure 1. A graph of genus $g = 4$.

counts the number of independent cycles or “holes” in the graph G . Every graph has a *divisor group*

$$\text{Div}(G) := \left\{ \sum_{v \in V} a_v v : a_v \in \mathbb{Z} \right\},$$

and the *degree homomorphism*, $\deg: \text{Div}(G) \rightarrow \mathbb{Z}$, sends a divisor to the sum of its coefficients:

$$\deg \left(\sum_{v \in V} a_v v \right) = \sum_{v \in V} a_v.$$

Every graph has a *canonical divisor*, K , defined by

$$K = \sum_{v \in V} (\deg_G(v) - 2) v.$$

The canonical divisor for the genus-four graph shown in Figure 1 is $K = 2v_1 + v_2 + 2v_3 + v_5$.

Every graph G also has a group of *firing scripts*, which are simply integer-valued functions on the vertices:

$$\mathcal{M}(G) := \{\sigma: V \rightarrow \mathbb{Z}\}.$$

The *divisor homomorphism*, $\text{div}: \mathcal{M}(G) \rightarrow \text{Div}(G)$, sends a firing script to the resulting degree-zero divisor:

$$\text{div}(\sigma) := \sum_{v \in V} \left(\deg_G(v) \sigma(v) - \sum_{vw \in E} \sigma(w) \right) v.$$

The kernel of the divisor homomorphism is the subgroup of constant functions, \mathbb{Z} . These homomorphisms fit into a sequence in which the image of each map is contained in the kernel of the next map:

$$0 \rightarrow \mathbb{Z} \rightarrow \mathcal{M}(G) \xrightarrow{\text{div}} \text{Div}(G) \xrightarrow{\text{deg}} \mathbb{Z} \rightarrow 0.$$

This sequence is exact except at $\text{Div}(G)$, where (in general) the kernel of deg is strictly larger than the image of div , which we call the subgroup of *principal divisors*, $\text{Prin}(G)$. Taking the quotient, we obtain the *Picard group* and its subgroup the *Jacobian*:

$$\text{Pic}(G) = \text{Div}(G)/\text{Prin}(G) \quad \text{and} \quad \text{Jac}(G) = \text{Div}^0(G)/\text{Prin}(G).$$

The Picard and Jacobian groups are related by the split-exact sequence

$$0 \rightarrow \text{Jac}(G) \rightarrow \text{Pic}(G) \xrightarrow{\text{deg}} \mathbb{Z} \rightarrow 0.$$

We call two divisors $D, D' \in \text{Div}(G)$ *linearly equivalent* and write $D \sim D'$ if they define the same element of the Picard group. That is, $D \sim D'$ when $D = D' - \text{div}(\sigma)$ for some firing script σ . A divisor with all coefficients nonnegative is called *effective*. For any divisor D , its *complete linear system* $|D|$ is the set of linearly equivalent effective divisors:

$$|D| := \{E \in \text{Div}(G) : E = D + \text{div}(\sigma) \geq 0 \text{ for some } \sigma \in \mathcal{M}(G)\}.$$

The *rank* $r(D)$ of a divisor D is defined to be one less than the minimal degree of an effective divisor E such that $|D - E| = \emptyset$. Equivalently:

$$r(D) = \max \{r \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all effective } E \text{ of degree } r\}.$$

The Riemann-Roch theorem for graphs expresses a relation between the rank of a divisor, D , and the rank of the canonically dual divisor $K - D$:

$$r(D) - r(K - D) = \text{deg}(D) + 1 - g.$$

Our goal now is to sketch an entirely analogous story in the context of Riemann surfaces. There, the basic objects of study are smooth, compact, connected Riemann surfaces as pictured in figure 2.




Figure 2. A surface of genus $g = 2$.

Roughly speaking, a Riemann surface is a topological space such that every point has an open neighborhood that has been identified with an open subset of




Figure 3. Stereographic projection.

the complex plane, \mathbb{C} . Specifying these identifications in a compatible way is called putting a *complex structure* on the underlying topological surface. The *genus* of a surface is its number of “handles.” Viewed only as topological spaces, any two surfaces of genus g are the same (i.e., homeomorphic). But for $g \geq 1$, there are infinitely many different complex structures, so that there are infinitely many different Riemann surfaces of genus g . To illustrate the ideas in this section, we will provide a running example with $g = 0$; we will see below (as a consequence of the Riemann-Roch theorem!) that this is actually the only example of a Riemann surface of genus zero.

Example 5.15. The topological surface of genus zero is the sphere, S^2 . In order to put a complex structure on the sphere, we use the technique of *stereographic projection* as shown in Figure 3. Thinking of S^2 as the subset of \mathbb{R}^3 defined by the equation $x^2 + y^2 + z^2 = 1$, we identify the complex plane \mathbb{C} with the equatorial xy -plane by writing $x + iy = (x, y, 0)$. Consider the line through the north pole $(0, 0, 1)$ and any other point (x, y, z) on the sphere; this line intersects the complex plane at a unique point $u + iv$ given by the formula

$$u + iv = \frac{x + iy}{1 - z}.$$

The mapping $(x, y, z) \mapsto u + iv$ defines a homeomorphism from the punctured sphere $S^2 - \{(0, 0, 1)\}$ to the complex plane. Similarly, we may instead consider lines originating at the south pole $(0, 0, -1)$, which yields a homeomorphism from $S^2 - \{(0, 0, -1)\}$ to \mathbb{C} defined by

$$u' + iv' = \frac{x - iy}{1 + z}.$$

Thus, we have succeeded in identifying an open neighborhood of each point in S^2 with the complex plane. But in order to qualify as a complex structure, the identifications must be compatible on overlapping neighborhoods. To check this compatibility in the present case, note that if $(x, y, z) \neq (0, 0, \pm 1)$, so that both mappings are defined, we see that

$$(u + iv)(u' + iv') = \frac{x^2 + y^2}{1 - z^2} = \frac{1 - z^2}{1 - z^2} = 1,$$

so that the *transition function* $u + iv \mapsto u' + iv'$ is simply inversion of complex numbers, $\zeta \mapsto \zeta^{-1}$, which is a holomorphic, i.e., complex-differentiable, function. In general, the compatibility we require for a complex structure is that all transition functions are holomorphic. As mentioned above, we will later show that the Riemann-Roch theorem implies that every complex structure on the sphere is equivalent to the one we have just described via stereographic projection.

Every Riemann surface determines a *divisor group*:

$$\text{Div}(S) := \left\{ \sum_{p \in S} a_p p : a_p \in \mathbb{Z}, a_p = 0 \text{ for all but finitely many } p \right\}.$$

Hence, a divisor on S is a finite, formal, integer-linear combination of points of the surface. A divisor is *effective* if all of its coefficients are nonnegative. The *degree homomorphism*, $\text{deg}: \text{Div}(S) \rightarrow \mathbb{Z}$, sends a divisor to the (finite) sum of its coefficients:

$$\text{deg} \left(\sum_{p \in S} a_p p \right) = \sum_{p \in S} a_p.$$

If S is a Riemann surface, then we may consider complex-valued functions $h: S \rightarrow \mathbb{C}$. Since S has a complex structure, near each point $p \in S$ we may think of h as a function of a complex variable w , where $w = 0$ corresponds to the point p . We may now check whether $h(w)$ is complex-differentiable near $w = 0$, in which case we say that h is holomorphic at p (the fact that all transition functions are holomorphic guarantees that this notion is well-defined). If h is holomorphic at p , then we may express $h(w)$ as a power-series:

$$h(w) = \sum_{k \geq 0} c_k w^k.$$

Of course, not all functions on S are holomorphic. In fact, a theorem from complex analysis asserts that (since S is compact), only the constant functions are holomorphic at every point of S . In order to obtain a richer collection of functions on S , we say that $f: S \rightarrow \mathbb{C} \cup \{\infty\}$ is *meromorphic* if near each point $p \in S$, the function $f(w)$ may be expressed as a quotient of holomorphic functions $g(w)/h(w)$ with $h(w)$ not identically zero. Equivalently, f is meromorphic at p if $f(w)$ can be represented by a Laurent series:

$$f(w) = \sum_{k \geq M} c_k w^k.$$

In terms of this Laurent series, we define the *order* of f at p to be $\text{ord}_p(f) := \min\{k \mid c_k \neq 0\}$. Note that

$$f(p) = \begin{cases} 0 & \text{if } \text{ord}_p(f) > 0 & \text{(zero of } f) \\ \infty & \text{if } \text{ord}_p(f) < 0 & \text{(pole of } f) \\ c_0 \neq 0 & \text{if } \text{ord}_p(f) = 0 & \text{.} \end{cases}$$

We denote the field of meromorphic functions on S by $\mathcal{M}(S)$:

$$\mathcal{M}(S) := \{f: S \rightarrow \mathbb{C} \cup \{\infty\} : f \text{ is meromorphic}\}.$$

The associated multiplicative group (obtained by removing the zero function) is denoted $\mathcal{M}(S)^\times$. From the compactness of the surface S , it follows that every nonzero meromorphic function has only finitely many zeros and poles. Hence, we may define a *divisor homomorphism* by sending each nonzero meromorphic function to the weighted sum of its zeros and poles:

$$\operatorname{div}(f) = \sum_{p \in S} \operatorname{ord}_p(f) p.$$

Example 5.16. Consider the Riemann sphere from Example 5.15, which we identify with $\mathbb{C} \cup \{\infty\}$ via stereographic projection from the north pole, associating ∞ with the north pole $(0, 0, 1)$ itself. The following rational function is an example of a meromorphic function on the sphere:

$$f(w) := \frac{w(w-1)}{(w+1)^2}.$$

The value of this function at $w = \infty$ is obtained by writing it in terms of $1/w$ and evaluating at $1/w = 0$:

$$f(1/w) = \frac{(1-1/w)}{(1+1/w)^2} \implies f(\infty) = 1.$$

We see that f has zeros of order 1 at $w = 0$ and $w = 1$ as well as a pole of order 2 at $w = -1$. It follows that the divisor of f is given by

$$\operatorname{div}(f) = p_0 + p_1 - 2p_{-1},$$

where (for clarity) we have denoted the point of the Riemann surface with coordinate w by the symbol p_w . This example is entirely typical, because every meromorphic function on the Riemann sphere is actually a rational function: $\mathcal{M}(\mathbb{C} \cup \{\infty\}) = \mathbb{C}(w)$.

Returning to an arbitrary Riemann surface S , the image of the divisor homomorphism is the subgroup of *principal divisors*, $\operatorname{Prin}(S)$. It follows from the compactness of S that all principal divisors have degree zero, and that the kernel of div is the nonzero constant functions, \mathbb{C}^\times . As in the case of graphs, these homomorphisms fit into a sequence

$$0 \rightarrow \mathbb{C}^\times \rightarrow \mathcal{M}(S)^\times \xrightarrow{\operatorname{div}} \operatorname{Div}(S) \xrightarrow{\operatorname{deg}} \mathbb{Z} \rightarrow 0.$$

Again, this sequence is exact except at the term $\operatorname{Div}(S)$, from which we obtain the Picard and Jacobian groups:

$$\operatorname{Pic}(S) = \operatorname{Div}(S)/\operatorname{Prin}(S) \quad \text{and} \quad \operatorname{Jac}(S) = \operatorname{Div}^0(S)/\operatorname{Prin}(S).$$

Once again, the Picard and Jacobian groups are related by the split-exact sequence

$$0 \rightarrow \operatorname{Jac}(S) \rightarrow \operatorname{Pic}(S) \xrightarrow{\operatorname{deg}} \mathbb{Z} \rightarrow 0.$$

Linear equivalence of divisors is defined in the familiar way:

$$D \sim D' \quad \text{if and only if} \quad D = D' - \operatorname{div}(f) \quad \text{for some meromorphic } f.$$

(For convenience, we define the divisor of the zero function to be zero.) The complete linear system of a divisor D is the collection of linearly equivalent effective divisors:

$$|D| = \{E \mid E = D + \operatorname{div}(f) \geq 0 \text{ for some } f \in \mathcal{M}(S)\}.$$

Note that, just as for graphs, if $\deg(D) < 0$, then the complete linear system $|D|$ is empty since linear equivalence preserves the degree of divisors.

For a given divisor D , the requirement that $D + \operatorname{div}(f)$ is effective places restrictions on the zeros and poles of the meromorphic function f :

$$\operatorname{ord}_p(f) \geq -D(p) \quad \text{for all } p \in S.$$

Explicitly: if $D(p) < 0$, then f must vanish to order *at least* $-D(p)$ at p , and if $D(p) > 0$, then f can have a pole of order *at most* $D(p)$ at p . (If $D(p) = 0$, then f cannot have a pole at p , although it may have a zero of any order or take a finite, nonzero value). We say that such meromorphic functions have *poles at worst* D . Moreover, the collection of meromorphic functions with poles at worst D forms a complex vector subspace of the field of all meromorphic functions: $\mathcal{L}(D) \subset \mathcal{M}(S)$. A fundamental result of algebraic geometry states that the vector spaces $\mathcal{L}(D)$ are finite-dimensional. Moreover, two functions $f, f' \in \mathcal{L}(D)$ yield the same effective divisor $E = D + \operatorname{div}(f) = D + \operatorname{div}(f')$ if and only if one is a scalar multiple of the other: $f' = cf$. Thus, the complete linear system $|D|$ is the *projectivization* of the vector space $\mathcal{L}(D)$, defined as the space of lines through the origin in $\mathcal{L}(D)$. Hence, as a measure of the size of $|D|$, we define the *rank* of D to be one less than the dimension of the vector space $\mathcal{L}(D)$:

$$r(D) := \dim_{\mathbb{C}} \mathcal{L}(D) - 1.$$

In particular, $|D|$ is empty if and only if $r(D) = -1$.

Every Riemann surface has a *canonical divisor class*, $[K] \in \operatorname{Pic}(S)$. A representative divisor, K , may be obtained by choosing any *meromorphic differential 1-form* on S and taking its divisor. Roughly speaking, a differential 1-form is an object, ω , that may be integrated over paths in S . Due to the complex structure on S , near each point p we may express ω as $f dw$, where f is a meromorphic function of the local coordinate w . Then K , the divisor of ω , is given near p by $\operatorname{div}(f)$, and we obtain the full divisor K by working in local patches that cover S . Choosing a different 1-form ω' would lead to a linearly equivalent divisor K' .

Example 5.17. Consider the differential form defined by dw on the Riemann sphere $\mathbb{C} \cup \{\infty\}$, where w is the local coordinate on the punctured sphere obtained by stereographic projection from the north pole. This differential form has no zeros or poles in the patch obtained by removing the north pole. To investigate its behavior at the north pole, we must instead use the local coordinate w' obtained by stereographic projection from the south pole. Recall that these two local coordinates are related by inversion: $w = 1/w'$. Hence we have

$$dw = d(1/w') = -\frac{1}{w'^2} dw'.$$

It follows that dw has a double pole at ∞ (where $w' = 0$), and no other zeros or poles. Hence, the canonical divisor class of the Riemann sphere is represented by the divisor $K = -2p_{\infty}$.

We are now able to state the Riemann-Roch theorem for Riemann surfaces:

Theorem 5.18 (Riemann-Roch, 1865). *Let $D \in \operatorname{Div}(S)$. Then*

$$r(D) - r(K - D) = \deg(D) + 1 - g.$$

The reader will notice that the Riemann-Roch theorem for surfaces has exactly the same form as the Riemann-Roch theorem for graphs. Moreover, as the next proposition shows, the rank function for Riemann surfaces, although defined as a vector space dimension, can be computed in the same way as for graphs:

Proposition 5.19. *Let D be a divisor on a Riemann surface S . Then the rank function $r(D) = \dim_{\mathbb{C}} \mathcal{L}(D) - 1$ may be computed in the following way:*

$$r(D) = \max\{r \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all effective } E \text{ of degree } r\}.$$

Proof. First note that the right hand side is -1 if and only if $|D|$ is empty which is equivalent to the statement that $r(D) = -1$. So we may suppose $r(D) \geq 0$. In this case, there exists a nonzero meromorphic function $f \in \mathcal{L}(D)$. Then choose a point $p \in S$ that is not a zero of f or in the support of D . Consider the linear evaluation map $\alpha_p: \mathcal{L}(D) \rightarrow \mathbb{C}$ defined by $g \mapsto g(p)$. The kernel of α_p is exactly $\mathcal{L}(D - p)$, i.e., those functions with poles at worst D that also vanish at the additional point p . But since $f(p) \neq 0$, the map α_p is surjective, so that by rank-nullity, the dimension of $\mathcal{L}(D - p)$ is one less than the dimension of $\mathcal{L}(D)$. Hence, $r(D - p) = r(D) - 1$. We may continue to subtract points from D in this way, lowering the rank by one each time, until we reach a divisor of rank -1 . It follows that there exists an effective divisor E of degree $r(D) + 1$ such that $|D - E| = \emptyset$, so that

$$r(D) \leq \max\{r \in \mathbb{Z} : |D - E| \neq \emptyset \text{ for all effective } E \text{ of degree } r\}.$$

To establish the other inequality, we need to consider the relationship between $r(D)$ and $r(D - p)$ for arbitrary points p , not just the ones chosen in the previous paragraph. So far we have shown that $r(D - p) = r(D) - 1$ for all points p not contained in the support of D and for which there exists a nonzero meromorphic function f satisfying $f(p) \neq 0$. So now suppose that p is not in the support of D , but every function $f \in \mathcal{L}(D)$ vanishes at p . In this case $\mathcal{L}(D - p) = \mathcal{L}(D)$, so the rank does not change: $r(D - p) = r(D)$. On the other hand, if p is in the support of D , then we claim that $r(D) \leq r(D - p) + 1$. Indeed, suppose that $r(D) > r(D - p)$, so that $\mathcal{L}(D - p)$ is a proper subspace of $\mathcal{L}(D)$. Then choose two functions $f, g \in \mathcal{L}(D) - \mathcal{L}(D - p)$; we wish to show that $f - \lambda g \in \mathcal{L}(D - p)$ for some $\lambda \in \mathbb{C}$. Now, the fact that f and g each have poles at worst D but not poles at worst $D - p$ implies that they have the same order at p , namely $-D(p)$. But then (f/g) is a meromorphic function with neither a zero nor a pole at p ; define $\lambda := (f/g)(p) \in \mathbb{C}^\times$. Then $f - \lambda g$ is a meromorphic function with poles at worst D that also vanishes at p , so $f - \lambda g \in \mathcal{L}(D - p)$ as required. This shows that the quotient vector space $\mathcal{L}(D)/\mathcal{L}(D - p)$ has dimension 1, and hence $r(D) - r(D - p) = 1$ as claimed. Putting this all together, we see that the rank of a divisor drops by at most 1 every time we remove a point. It follows that if E is an effective divisor of degree at most $r(D)$, then $\dim_{\mathbb{C}} \mathcal{L}(D - E) \geq 1$, so that $|D - E| \neq \emptyset$. It follows that

$$r(D) \geq \max\{r \in \mathbb{Z} \mid |D - E| \neq \emptyset \text{ for all effective } E \text{ of degree } r\}.$$

□

We have seen in this section that the divisor theory of graphs is exactly analogous to the classical theory of divisors on Riemann surfaces. However, there is a shift of emphasis between the two contexts. In the case of graphs, the preceding chapters

have utilized firing scripts in a supporting role: from the point of view of the dollar game, we are fundamentally interested in the divisors themselves, and firing scripts provide a useful tool for talking about lending and borrowing moves. But in the context of Riemann surfaces, the situation is reversed: a divisor D is a discrete object that imposes conditions on zeros and poles, cutting down the infinite-dimensional (over \mathbb{C}) field of all meromorphic functions to a finite-dimensional subspace $\mathcal{L}(D)$. From this point of view, divisors are a useful tool for talking about the function theory of Riemann surfaces.

For a given Riemann surface S , one of the most important uses of the function spaces $\mathcal{L}(D)$ is to provide mappings $\varphi: S \rightarrow \mathbb{P}^r$, where \mathbb{P}^r is the projective space of dimension $r = r(D)$ introduced below. Under suitable conditions on the divisor D , the mapping φ will be an embedding. Such embeddings are interesting and useful, because they allow us to study the abstract surface S using techniques and results that are special to subvarieties of projective spaces.

The projective space \mathbb{P}^r is the projectivization of the vector space \mathbb{C}^{r+1} , i.e., the space of lines through the origin. Explicitly, we may describe \mathbb{P}^r as the result of removing the origin from \mathbb{C}^{r+1} and then identifying nonzero vectors that are complex-scalar multiples of each other:

$$\mathbb{P}^r = (\mathbb{C}^{r+1} - \{\mathbf{0}\}) / \{\mathbf{v} \sim \lambda \mathbf{v}\}.$$

We use the symbol $[v_1 : v_2 : \dots : v_{r+1}]$ to denote the equivalence class in \mathbb{P}^r of a nonzero vector $(v_1, v_2, \dots, v_{r+1}) \in \mathbb{C}^{r+1}$. Consider the case $r = 1$. If $(v_1, v_2) \in \mathbb{C}^2 - \{(0, 0)\}$ and $v_2 \neq 0$, then $[v_1 : v_2] = [v_1/v_2 : 1]$. On the other hand, if $v_2 = 0$, then $v_1 \neq 0$, and $[v_1 : v_2] = [1 : 0]$. Hence, we may describe \mathbb{P}^1 as

$$\mathbb{P}^1 = \{[v : 1] \mid v \in \mathbb{C}\} \cup \{[1 : 0]\} = \mathbb{C} \cup \{\infty\},$$

and we see that the projective line \mathbb{P}^1 is simply the Riemann sphere. A similar argument shows that for $r > 1$ we have

$$\mathbb{P}^r = \mathbb{C}^r \cup \mathbb{P}^{r-1}.$$

For instance, the projective plane \mathbb{P}^2 looks like \mathbb{C}^2 together with a projective line \mathbb{P}^1 “at infinity”.

We now explain how a divisor D with $r = r(D) > 0$ on a Riemann surface S leads to a mapping from S to the projective space \mathbb{P}^r . Given such a divisor D , choose a basis f_1, f_2, \dots, f_{r+1} for the vector space $\mathcal{L}(D)$. Then consider the mapping $\varphi: S \rightarrow \mathbb{P}^r$ defined by

$$\varphi(p) := [f_1(p) : f_2(p) : \dots : f_{r+1}(p)].$$

This mapping is defined at p provided at least one of the functions f_i does not vanish at p ; points where all the f_i vanish are called *basepoints* of the linear system $|D|$. So, if $|D|$ has no basepoints, then $\varphi: S \rightarrow \mathbb{P}^r$ is a (holomorphic) mapping of the Riemann surface S to projective space. As mentioned above, under suitable conditions on the divisor D , this mapping will be an embedding, so that the image $\varphi(S) \subset \mathbb{P}^r$ is an isomorphic copy of the Riemann surface S .

At the other extreme from being an embedding, if $r(D) = 1$, then $\varphi: S \rightarrow \mathbb{P}^1$ is a holomorphic many-to-one mapping between Riemann surfaces—if $\deg(D) = d$, then a generic point of \mathbb{P}^1 has exactly d preimages. The finitely many points with

fewer than d preimages are called the *branch points* of φ . Such mappings are called *branched coverings* of \mathbb{P}^1 , and they provide a powerful way of studying Riemann surfaces. In the special case where $d = r = 1$, we obtain an isomorphism between S and \mathbb{P}^1 . In the next example, we use this construction to show that every Riemann surface of genus zero is isomorphic to the Riemann sphere.

Example 5.20. Suppose that S is a Riemann surface of genus zero—that is, S is the result of specifying a complex structure on the topological sphere S^2 . Pick any point $p \in S$, and consider the degree-one divisor $D := p$ on S . Then by Riemann-Roch we have

$$r(p) = r(K - p) + \deg(p) + 1 - 0 \geq -1 + 2 = 1,$$

so $\mathcal{L}(p)$ has dimension at least 2. Hence, there must be a nonconstant meromorphic function f on S having at most a single pole of order 1 at p . On the other hand, since only the constants are holomorphic, f must have at least one pole, which means that it has exactly one pole (of order 1) at p . If we think of $f: S \rightarrow \mathbb{C} \cup \{\infty\}$ as a function from S to the Riemann sphere, we see that p is the unique point mapping to ∞ . Now consider any complex number $\zeta \in \mathbb{C}$. Then the nonconstant meromorphic function $f - \zeta: S \rightarrow \mathbb{C} \cup \{\infty\}$ also has a single pole of order 1 at p , and (since its divisor has degree 0) must therefore have a unique zero at some point $q \in S$. But this means that $f(q) = \zeta$, and q is the unique point of S mapping to ζ . It follows that (viewed as a map of Riemann surfaces rather than a meromorphic function) $f: S \rightarrow \mathbb{C} \cup \{\infty\}$ is a holomorphic bijection from S to the Riemann sphere. It follows that the inverse function is also holomorphic, so S is isomorphic to the Riemann sphere.

In Chapter 10, we introduce *harmonic morphisms* between graphs, which provide an analogue of branched coverings of Riemann surfaces—this analogy is described in Section 10.2. Moreover, in Theorem 10.24 we show that divisors of rank 1 furnish harmonic branched covers of trees, which (as graphs of genus zero) play the role of the Riemann sphere in the graph-theoretic analogy.

Riemann surfaces are the one-dimensional (over \mathbb{C}) case of a more general type of geometric object: compact complex manifolds. Moreover, there is a higher-dimensional version of the Riemann-Roch theorem called the Hirzebruch-Riemann-Roch theorem, proved by Friedrich Hirzebruch in the 1950s. Shortly thereafter, Alexandre Grothendieck established a generalization of this result (appropriately called the Grothendieck-Riemann-Roch theorem) which applies to mappings between smooth schemes. A major open question in the graph-theoretic context is whether there is a higher-dimensional version of the Riemann-Roch theorem for simplicial complexes (of which graphs are the one-dimensional case). Chapter 15 starts to address this topic by providing a brief introduction to higher-dimensional chip-firing.

5.4. Alive divisors and stability

The Riemann-Roch formula expresses a precise relation between the ranks of divisors D and $K - D$, where $K = \sum_{v \in V} (\deg_G(v) - 2)v$ is the canonical divisor of the graph G . This suggests that we should pay attention to the map from

$\text{Div}(G)$ to itself defined by $D \mapsto K - D$. In particular, we have seen that this map preserves the set of maximal unwinnables (Corollary 5.11) and reveals a duality between unwinnable divisors and divisors that dominate a maximal unwinnable (Corollary 5.12). In this section, we will define another closely related map from $\text{Div}(G)$ to itself that expresses a further duality for divisors.

Definition 5.21. Let $D \in \text{Div}(G)$ be a divisor. A vertex $v \in V$ is *stable for D* if $D(v) < \deg_G(v)$; otherwise v is *unstable for D* . Note that a vertex is unstable for D exactly when it can lend without going into debt. If every vertex is stable for D , then we say that the divisor D is *stable*; otherwise D is *unstable*. The divisor D is *alive* or *unstabilizable* if every member of the linear equivalence class $[D]$ is unstable.

Exercise 5.22. Prove that alive divisors are winnable.

If we add 1 to each vertex in the canonical divisor, then we obtain the *maximal stable divisor*

$$D_{\max} := K + \vec{1} = \sum_{v \in V} (\deg_G(v) - 1) v.$$

This is the unique largest divisor in which no vertex can lend without going into debt. The next proposition shows that the mapping $D \mapsto D^* := D_{\max} - D$ yields a duality between unwinnable and alive divisors; for this reason we call D^* the *dual divisor* to the original divisor D .

Proposition 5.23. *A divisor D is (maximal) unwinnable if and only if D^* is (minimal) alive.*

Proof. Suppose D is unwinnable. Then for all $D' \sim D$, there exists $v \in V$ such that $D'(v) < 0$, so $(D_{\max} - D')(v) = \deg_G(v) - 1 - D'(v) \geq \deg_G(v)$. Hence, v is unstable for $D_{\max} - D'$. Since $[D^*] = \{D_{\max} - D' \mid D' \sim D\}$, it follows that D^* is alive.

Now suppose D^* is alive, and let $D' \in [D]$. Then $D_{\max} - D' \sim D^*$, so there exists $v \in V$ such that $\deg_G(v) \leq (D_{\max} - D')(v) = \deg_G(v) - 1 - D'(v)$. This is only possible if $D'(v) < 0$, so D' is not effective. Thus, D is unwinnable.

Now suppose that D is maximal with respect to being unwinnable. This means that $D + v$ is winnable for all vertices v . But we have just shown that this is equivalent to the statement that for all v , the divisor $(D + v)^* = D^* - v$ is not alive. This means that D^* is minimal with respect to being alive. \square

Corollary 5.24. *If $D \in \text{Div}(G)$ is an alive divisor, then D is minimal alive if and only if $\deg(D) = |E(G)|$.*

Proof. A divisor D is minimal alive if and only if D^* is maximal unwinnable, and hence if and only if D^* is unwinnable of degree $g - 1$ (cf. Corollary 4.9). But

$$\begin{aligned} \deg(D^*) &= \sum_{v \in V} (\deg_G(v) - 1) - \deg(D) \\ &= 2|E(G)| - |V(G)| - \deg(D) \\ &= g - 1 + |E(G)| - \deg(D). \end{aligned}$$

It follows that D^* has degree $g - 1$ if and only if $\deg(D) = |E(G)|$. \square

Proposition 5.25. *Let $D \in \text{Div}(G)$ be a divisor. Then D is dominated by a minimal alive divisor if and only if $D - \vec{1}$ is unwinnable.*

Proof. Let $A \in \text{Div}(G)$ denote a minimal alive divisor, so that A^* is maximal unwinnable. Then

$$\begin{aligned} D \leq A &\iff A^* \leq D^* \\ &\iff r(K - D^*) = -1 \quad (\text{Corollary 5.12}) \\ &\iff r(K - D_{\max} + D) = -1 \\ &\iff r(D - \vec{1}) = -1 \quad (\text{since } D_{\max} = K + \vec{1}). \end{aligned}$$

□

Exercise 5.26. Prove that D is dominated by a minimal alive divisor if and only if no element of the complete linear system $|D|$ has full support.

The duality $D \leftrightarrow D^*$ allows us to see the greedy algorithm from Section 3.1 in a different light. In the case where the dollar game on G starting with a divisor D is winnable, the algorithm produces a unique firing script, $\sigma \in \mathcal{M}(G)$, that wins the game (Proposition 3.1.1). But note that $D(v) < 0$ iff $D^*(v) \geq \deg_G(v)$, so that the debt-vertices of D correspond to the unstable vertices of D^* . Moreover, $(D + Lv)^* = D^* - Lv$, so a borrowing move for D corresponds to a lending move for D^* . Thus, applying the greedy algorithm to win the dollar game starting with D is the same as attempting to stabilize D^* by sequentially lending from unstable vertices. Moreover, this greedy lending algorithm applied to D^* will either return a stable divisor, or else certify that no stabilization exists (if every vertex eventually becomes unstable). Proposition 3.1.1 immediately yields a dual uniqueness result for the stabilization of divisors:

Proposition 5.27. *Suppose that D^* is not alive, so a stabilization of D^* exists. Then both the stabilization and the corresponding firing script are unique.*

In Part 2 we will further pursue the notions of stability and stabilization by reinterpreting the dollar game as a toy physical system involving the toppling of grains of sand. Moreover, the duality $D \leftrightarrow D^*$ will reappear in the form of Theorem 7.12, expressing a duality between two special classes of sandpiles: the recurrents and the superstables.

Notes

The Riemann-Roch theorem for graphs was first established by Baker and Norine in [6] via a different proof from the one presented here. In particular, the original proof proceeds by showing ([6, Theorem 2.2]) that the Riemann-Roch formula is equivalent to the conjunction of two properties (RR1) and (RR2), corresponding in our text to Problem 5.2 and Exercise 5.10 respectively.

The proof of Riemann-Roch presented here, using acyclic orientations, is due to Cori and Le Borgne ([27]). Independently, Backman ([3]) provides a similar proof

using acyclic orientations and provides a thorough investigation of the relationship between the dollar game and orientations of a graph.

Problems for Chapter 5

5.1. Suppose that $D \in \text{Div}(G)$ has the property that for every vertex v , there exists an effective divisor $E \in |D|$ such that $E(v) \geq 1$. Show that $r(D) \geq 1$.

5.2. For each divisor $D \in \text{Div}(G)$, show that there exists a maximal unwinnable divisor N with the property that D is winnable if and only if $N - D$ is unwinnable.

5.3. Let \mathcal{N} denote the set of maximal unwinnable divisors on a graph G . For a divisor $D \in \text{Div}(G)$, define $\deg^+(D) := \sum_{v \in V} \max\{D(v), 0\}$, the sum of the non-negative coefficients of D . This problem establishes the following formula for the rank function, which appears as Lemma 2.7 in [6].

$$r(D) = \left(\min_{D' \in [D], N \in \mathcal{N}} \deg^+(D' - N) \right) - 1 =: R(D).$$

(a) First show that $r(D) \geq R(D)$ by filling in the details of the following proof by contradiction. If $r(D) < R(D)$, then there must exist an effective divisor E of degree $R(D)$ such that $D - E$ is unwinnable. Now apply Problem 5.2 to the divisor $D - E$ and obtain a contradiction of the definition of $R(D)$.

(b) To show that $r(D) \leq R(D)$, start by choosing $D' \in [D]$ and $N \in \mathcal{N}$ so that $\deg^+(D' - N) = R(D) + 1$. Then write $D' - N = E - E'$ for effective divisors E, E' with $\deg(E) = R(D) + 1$ and show that $D - E$ is unwinnable.

5.4. If D is a divisor on a tree and $\deg(D) \geq -1$, show that $r(D) = \deg(D)$ in two ways: (i) directly from the definition of rank, and (ii) from Riemann-Roch. (Note that a graph is a tree if and only if its genus is 0.)

5.5. Let v be a vertex on a graph G of genus g . Show that

$$r(v) = \begin{cases} 1 & \text{if } g = 0, \text{ i.e., } G \text{ is a tree,} \\ 0 & \text{if } g > 0. \end{cases}$$

5.6. Use the Riemann-Roch theorem to determine the rank of an arbitrary divisor on the cycle graph with n vertices, C_n .

5.7. Let \mathcal{N} denote the set of maximal unwinnable divisors on a graph G , and let K be the canonical divisor of G . By Corollary 5.11, there is an involution $\iota: \mathcal{N} \rightarrow \mathcal{N}$ given by $\iota(N) = K - N$. Illustrate how this involution permutes the maximal unwinnables for the diamond graph, displayed in Figure 4. List all maximal unwinnables in q -reduced form.




Figure 4. The diamond graph.

5.8. Find all effective divisors of degree 2 and rank 1 on the house graph pictured in Figure 7 in the problems for Chapter 2. Explain how you know your list is complete.

5.9. The *Weierstrass sequence* for $D \in \text{Div}(G)$ at a vertex $v \in V$ is the sequence of ranks

$$(r(D - kv) : k \in \mathbb{N}) = r(D), r(D - v), r(D - 2v), \dots$$

The integer $k \geq 1$ is a *Weierstrass gap* for D at v if $r(D - (k - 1)v) \neq r(D - kv)$. By Exercise 5.3, we know that for all $k \in \mathbb{Z}$,

$$r(D - (k - 1)v) - r(D - kv) \in \{0, 1\}.$$

So the sequence starts at $r(D)$, decreases by at most 1 at each step, and is eventually constant at -1 , resulting in $r(D) + 1$ gaps in total. Let k_i be the i -th gap. The *Weierstrass weight* for v in D is then

$$\text{wt}(v, D) := \sum_{i=1}^{r(D)+1} (k_i - 1),$$

where $\text{wt}(v, D) := 0$ if $r(D) = -1$. Thus, the Weierstrass weight is a measure of how far the Weierstrass sequence differs from

$$r(D), r(D) - 1, r(D) - 2, \dots, 0, -1, -1, \dots$$

The vertex v is a Weierstrass point with respect to D or a *D -Weierstrass point* if $\text{wt}(v, D) > 0$. Since the rank of a divisor depends only on its linear equivalence class, it follows that if $D \sim D'$ and $v \sim w$ (in particular if $v = w$), then $\text{wt}(v, D) = \text{wt}(w, D')$.

- Show that if G is a tree, then no divisor on G has Weierstrass points.
- Let G be any graph. Show there are g Weierstrass gaps for the canonical divisor K and that v is a Weierstrass point for K if and only if $r(gv) \geq 1$.
- Let $n \geq 4$. Show that every vertex of the complete graph K_n is a Weierstrass point for the canonical divisor on K_n .
- Now consider the case of the cycle graph C_n with vertices v_0, \dots, v_{n-1} arranged consecutively. We saw in Problem 1.2 that we have an isomorphism

$$\begin{aligned} \mathbb{Z}/n\mathbb{Z} &\rightarrow \text{Jac}(C_n) \\ i &\mapsto i[v_1 - v_0] \end{aligned}$$

and that $i[v_1 - v_0] = [v_i - v_0]$ for $i \in \{0, 1, \dots, n - 1\}$. For each $i \in \{0, 1, \dots, n - 1\}$ and $d \in \mathbb{Z}$, define

$$D_i(d) := (v_i - v_0) + dv_0 = v_i + (d - 1)v_0.$$

From Proposition 1.20 it follows that $[D_0(d)], \dots, [D_{n-1}(d)]$ is a complete list of distinct divisor classes of degree d , i.e., the elements of $\text{Pic}^d(C_n)$.

- Let $D \in \text{Div}(C_n)$ and $v \in V(C_n)$. Use the Riemann-Roch theorem to show that v is a Weierstrass point for D if and only if

$$\deg(D) \geq 1 \quad \text{and} \quad r(D - \deg(D)v) = 0.$$

- Let $d \geq 1$. Show the vertex v_i is a Weierstrass point for $D_j(d)$ if and only if $j = di \pmod{n}$.

- (iii) Show that if $d \geq 1$, then each vertex of C_n is a Weierstrass point for a unique divisor class of degree d .
- (iv) Show that every vertex is a Weierstrass point for $D_0(n)$ and that $D_i(n)$ has no Weierstrass points if $i \neq 0$.

The definitions given in this problem come from the theory of Riemann surfaces via the analogy outlined in Section 5.3. As explained in that section, the complete linear system of a divisor on a Riemann surface may be used to construct a mapping of the surface into projective space. Weierstrass points on the surface are those at which the mapping exhibits “inflectionary” behavior. For more information on Weierstrass points for graphs, the reader is referred to [7] and [5].

Part 2

Sandpiles

The sandpile group

In Part 1, we interpreted divisors as distributions of wealth in a community represented by a graph G . This interpretation suggested the dollar game, and ultimately led to the Riemann-Roch theorem for ranks of divisors on graphs. Along the way, we uncovered q -reduced divisors and Dhar's algorithm, and established a tight connection with acyclic orientations on G . Behind the scenes stood the discrete Laplacian and its reduced form, together with the associated Picard and Jacobian groups.

In Part 2, we will change our interpretation of divisors from distributions of dollars to configurations of sand. Instead of regarding G as a community that lends and borrows, we will think of it as a hillside where sand topples. This new interpretation will suggest new questions and allow us to view old friends (such as q -reduced divisors, Dhar's algorithm, and the Jacobian group) in a new light.

In the next section, we introduce the sandpile interpretation via a simple example, before giving the general definitions in Sections 6.2 and 6.3.

6.1. A first example

Let G be the diamond graph pictured in Figure 1. Imagine that grains of sand can




Figure 1. The diamond graph.

be piled onto the vertices of G to form a *sandpile*. Figure 2 shows G with 4 grains of

sand on v_1 , one grain on each of v_2 and v_3 , and no sand on s . Note that a negative




Figure 2. A sandpile on G .

amount of sand doesn't make much sense, so unlike in the dollar game (where we interpret negative numbers as debt), we allow only nonnegative amounts of sand at each vertex. If a vertex gets too much sand, it becomes unstable, at which point it may *fire* or *topple*, sending grains of sand to its neighboring vertices. The other vertices may also be unstable or may become unstable due to the toppling of its neighbors. Thus, dropping a single grain of sand onto an existing sandpile can lead to an avalanche of firings.

How much sand does it take to make a vertex unstable? There is some choice here, but we will always use the following rule: a vertex is unstable if it has at least as many grains of sand as its degree. Recall that by the *degree* of a vertex, we mean the number of edges emanating from that vertex. Thus, a vertex v is unstable if it has enough sand to send a grain along each edge to its neighbors. (This corresponds to a lending move by v in the dollar game, except that v is now forbidden to lend if it will be in debt afterwards.)

If sand were conserved, then configurations with too many grains of sand would never stabilize: no matter how vertices were fired, there would always be at least one unstable vertex. To prevent this, we designate s to be the *sink* vertex, stipulating that any grains of sand landing on s immediately disappear. Figure 3 shows what happens to the configuration in Figure 2 when vertex v_1 fires. A grain of sand is delivered to each neighbor, the grain sent to s being lost.




Figure 3. Firing a vertex.

Starting with an initial unstable sandpile, what happens if we repeatedly choose and fire unstable vertices? Whenever a vertex adjacent to the sink fires, a grain of sand is lost, which tends to make the sandpile more stable. But is it possible to have a sequence of firings that can be repeated in a never-ending cycle? The answer is given as the following exercise.

Exercise 6.1. Explain why every initial sandpile on G will, through repeated firings, eventually stabilize.

Figure 4 shows the stabilization of a sandpile on G . At each step in the sta-




Figure 4. Stabilization.

bilization process observed in Figure 4 there is only one unstable vertex to fire. Starting with a different sandpile, though, it could easily happen that multiple vertices would be simultaneously unstable. A choice would then be required as to which vertex to fire first. The question naturally arises: would these choices make a difference in the end? The answer is the first pleasant coincidence in the subject, accounting for the word “abelian” in the abelian sandpile model. We leave it, at this point, as the following (perhaps premature) exercise.

Exercise 6.2. Explain why the choice of ordering in which to fire unstable vertices has no effect on the eventual stabilization of a sandpile.

In light of Exercises 6.1 and 6.2, by firing unstable vertices, every sandpile c on G eventually reaches a unique stable state. We denote the stabilization of c by c° and write

$$c \rightsquigarrow c^\circ.$$

Here and in the remainder of the book, we use the jagged arrow \rightsquigarrow to distinguish *legal* sequences of unstable vertex firings from more general processes that might involve firing stable vertices; these more general processes will generically be denoted by $c \rightarrow c'$.

6.1.1. Additive structure. Identify each sandpile on G with a 3-tuple $c = (c_1, c_2, c_3)$ of integers by letting c_i be the number of grains of sand on vertex v_i . Using this notation, the stabilization in Figure 4 could be written

$$(4, 1, 1) \xrightarrow{v_1} (1, 2, 2) \xrightarrow{v_3} \dots \xrightarrow{v_3} (1, 2, 0)$$

or, for short,

$$(4, 1, 1) \rightsquigarrow (1, 2, 0).$$

Let $\text{Stab}(G)$ denote the collection of the 18 (why 18?) stable sandpiles on G . Define the sum, $a \otimes b$, of two sandpiles $a, b \in \text{Stab}(G)$ by first adding vertex-wise, i.e., adding them as integer vectors, then stabilizing:

$$a \otimes b := (a + b)^\circ.$$

For example, $(2, 1, 1) \otimes (2, 0, 0) = (1, 2, 0)$ as illustrated in Figure 5.




Figure 5. Stable addition.

Exercise 6.3. Argue that the operation \otimes is commutative and associative and that the zero sandpile, $(0, 0, 0)$, serves as an identity element.

By Exercise 6.3, $(\text{Stab}(G), \otimes)$ is almost an abelian group—it only lacks inverses. Such a structure is called a *commutative monoid*.

Exercise 6.4. Show that only the zero sandpile has an inverse in $(\text{Stab}(G), \otimes)$. Would this still be true if G were replaced by a different graph?

It turns out that there is a special subset of $\text{Stab}(G)$ that does form a group, which we now describe. By way of motivation, note that in the dollar game on G starting with a divisor D , there is an explicit objective (to find an effective divisor $E \sim D$), and the vertices $v \in V$ are free to perform lending and borrowing moves to achieve that goal. Indeed, the combination of this freedom and an explicit goal is what lends that scenario the flavor of a *game*. For sandpiles, on the other hand, we constrain the ability of vertices to topple (requiring them to be unstable), while eliminating the effectivity goal by working from the outset with nonnegative amounts of sand. Moreover, our decision to include a sink ensures that every sandpile will stabilize, so there is no longer a game to win or lose. Instead, we have a model for a simple physical system: beginning with the zero sandpile on G , drop grains of sand randomly onto the non-sink vertices, producing a sequence of sandpiles. Whenever we reach an unstable sandpile c , it topples (producing an *avalanche*) to yield a stable sandpile c° . This model is *dissipative* because sand is lost into the sink s , and it is *slowly-driven* because of the ongoing addition of sand at a rate that allows for stabilization between each additional grain. Such systems

often exhibit the complex behavior of *self-organized criticality*, a topic that we will discuss in Section 12.4.

These ideas are implemented in the following experiment: pick a vertex of G at random; drop a grain of sand on that vertex and stabilize the resulting sandpile; repeat until 100 grains of sand have been dropped. The results from ten trials of this experiment are shown in Table 1. Interestingly, there are eight sandpiles that

Sandpile	Trials									
(0, 0, 0)	0	0	0	0	0	0	0	0	0	0
(0, 0, 1)	0	0	0	1	0	1	0	1	1	0
(0, 1, 0)	0	0	1	0	0	0	1	0	0	1
(0, 1, 1)	0	0	0	0	0	0	0	0	0	1
(0, 2, 0)	0	0	1	0	0	0	1	0	0	0
(0, 2, 1)	11	8	11	11	16	14	13	12	9	16
(1, 0, 0)	1	1	0	0	1	0	0	0	0	0
(1, 0, 1)	0	1	0	0	1	0	0	1	0	0
(1, 1, 0)	0	0	0	1	0	1	0	0	1	0
(1, 1, 1)	0	0	0	0	1	0	0	0	0	0
(1, 2, 0)	12	14	13	15	9	11	10	12	18	15
(1, 2, 1)	16	14	16	12	13	7	13	12	12	13
(2, 0, 0)	1	0	0	0	0	0	0	0	0	0
(2, 0, 1)	15	11	9	16	8	17	7	10	10	15
(2, 1, 0)	7	12	15	13	11	16	16	17	9	11
(2, 1, 1)	17	15	13	10	7	15	14	15	6	8
(2, 2, 0)	6	11	9	12	16	12	12	10	21	10
(2, 2, 1)	14	13	12	9	17	6	13	10	13	10

Table 1. Number of sandpiles formed by randomly dropping one grain of sand at a time and stabilizing.

appear more than once in each trial (cf. Figure 6). These are the *recurrent* sandpiles for G , denoted $\mathcal{S}(G)$. As we will see in general later, they form a group under the

$$(0, 2, 1), (1, 2, 0), (1, 2, 1), (2, 0, 1), (2, 1, 0), (2, 1, 1), (2, 2, 0), (2, 2, 1).$$

Figure 6. Recurrent sandpiles for G .

operation inherited from $\text{Stab}(G)$. This group, $\mathcal{S}(G)$, is our central object of study. You might be asking yourself: how could $\mathcal{S}(G)$ be a group if it does not contain the zero sandpile, $(0, 0, 0)$? The answer is that $\mathcal{S}(G)$ has an identity, but it is not the identity of $\text{Stab}(G)$. In other words, $\mathcal{S}(G)$ is not a submonoid of $\text{Stab}(G)$ even though they share the same operation.

Exercise 6.5. Create the addition table for $\mathcal{S}(G)$. What is the identity element? (See Section 6.5.1 for comments about identity elements in sandpile groups.)

Exercise 6.6. Compute the Jacobian group, $\text{Jac}(G)$, and show that it is isomorphic to $\mathcal{S}(G)$.

6.2. Directed graphs

In Part 1, we restricted attention to undirected multigraphs. Ultimately, this was because our goal was the Riemann-Roch theorem for graphs, which doesn't hold in general for directed multigraphs. However, the abelian sandpile model works nicely in the directed context, and we introduce the relevant definitions here.

From now on, a graph will be a finite, connected, multidigraph. Thus, a *graph* is an ordered pair $G = (V, E)$ where V is a finite set of *vertices* and E is a finite multiset of *directed edges*. Each element of E is an ordered pair $e = (u, v) \in V \times V$ where u and v are the *tail* and *head* of e , respectively. We use the notation $e^- = u$ and $e^+ = v$ and say e *emanates* from u . We will sometimes write $uv \in E$ to indicate the directed edge (u, v) . If $u = v$, then e is a *loop*. These are allowed but, we will see, do not add much to the theory of sandpiles.

The “multi” part of “multidigraph” indicates that E is a multiset, so that an edge may occur multiple times. Thus, each edge $e = (u, v) \in E$ has an integer *multiplicity*, $\text{mult}(e) \geq 1$. Figure 7 illustrates our conventions for drawing graphs. Edges are labeled by their multiplicities. An unlabeled edge like (u, s) is assumed to have multiplicity 1. The undirected edge between v and s represents the pair of directed edges, (v, s) and (s, v) , each with multiplicity 2.




Figure 7. A graph.

There are several different notions of connectedness for a digraph. For us, *connected* will mean what is sometimes called *weakly connected*: we assume that the underlying undirected graph for G is connected. (This undirected graph has the same vertex set as G and an edge $\{u, v\}$ for each directed edge (u, v) or (v, u) of G .) Usually, in fact, we will impose an even stronger condition than connectedness on G . This condition—the existence of a globally accessible vertex—is described in the next section.

The *outdegree* of a vertex v of G is

$$\text{outdeg}(v) = |\{e \in E : e^- = v\}|,$$

where each edge is counted according to its multiplicity. Similarly, the *indegree* of v in G is

$$\text{indeg}(v) = |\{e \in E : e^+ = v\}|.$$

For the graph in Figure 7, we have $\text{outdeg}(v) = 5$ and $\text{indeg}(v) = 7$.

Let $u, v \in V$. A *path of length k* from u to v is an ordered list of edges, e_1, \dots, e_k such that (1) $e_1^- = u$, (2) $e_i^+ = e_{i+1}^-$ for $i = 1, \dots, k-1$, and (3) $e_k^+ = v$. The *distance* from u to v is the minimum of the lengths of all paths from u to v , denoted $d(u, v)$. If there is no path from u to v , take $d(u, v) = \infty$. If $d(u, v) < \infty$, we say v is *accessible* from u . For the graph in Figure 7, we have $d(s, u) = 2$ and $d(u, s) = 1$.

Undirected graphs (without loops) provided the context for the dollar game in Part 1. By replacing each undirected edge $\{v, w\}$ by a pair of directed edges vw and wv , we may think of an undirected graph as a special type of directed graph. In terms of sandpiles, we might think of an undirected graph as a flat field where sand topples in all directions. We could then interpret directed graphs as hillsides, where the edges emanating from a vertex point “downhill.” Note that according to the above procedure, an undirected loop serves to increase the stability of the corresponding vertex by adding 2 to its outdegree.

6.3. Sandpile graphs

A *sandpile graph* is a triple $G = (V, E, s)$ where (V, E) is a graph, also referred to as G , and $s \in V$ is the *sink*. Normally, calling a vertex in a directed graph a “sink” would mean that its outdegree is zero. But that is not what we mean here. Instead, for a sandpile graph we require the designated sink vertex to be *globally accessible*. That is, there is some directed path from each vertex to s . Edges emanating from s are allowed (though for much of what we do, there would be no harm in deleting these outgoing edges).

Notation. We let \tilde{V} denote the non-sink vertices:

$$\tilde{V} := V \setminus \{s\}.$$

Let G be a sandpile graph with sink s . A *configuration* of sand on G is an element of the free abelian group on its non-sink vertices:

$$\text{Config}(G, s) := \text{Config}(G) := \mathbb{Z}\tilde{V} := \{\sum_{v \in \tilde{V}} c(v)v : c(v) \in \mathbb{Z} \text{ for all } v\}.$$

If c is a configuration and $v \in \tilde{V}$, then $c(v)$ will always denote the coefficient of v in c . We think of $c(v)$ as the number of grains of sand sitting on vertex v , even though this number might be negative. That is, we use the term *configuration* in the same way as we did in the context of divisors in Part 1. The reason for introducing configurations here is that, despite our true interest in nonnegative configurations, for which we reserve the suggestive term *sandpile*, our analysis will often require the consideration of general configurations. We let $0 := \sum_{v \in \tilde{V}} 0 \cdot v$ and $1_{\tilde{V}} := \sum_{v \in \tilde{V}} v$, the zero and all ones configurations, respectively. Of course, these two configurations are sandpiles.

The *degree* of a configuration c is the net amount of sand in c :

$$\text{deg}(c) = \sum_{v \in \tilde{V}} c(v) \in \mathbb{Z}.$$

Define a partial order on the group of configurations: if $a, b \in \mathbb{Z}\tilde{V}$, then $a \leq b$ if $a(v) \leq b(v)$ for all $v \in \tilde{V}$, and $a \leq b$ if $a \leq b$ and $a \neq b$. In terms of this partial

order, a configuration c is a *sandpile* if $0 \leq c$. We also sometimes use the adjective *nonnegative* to indicate that a given configuration is a sandpile.

6.3.1. Vertex firing. Let c be a configuration on G , and let $v \in \tilde{V}$. *Firing* or *toppling* v from c produces a new configuration c' by the rule

$$c' = c - \text{outdeg}(v)v + \sum_{vw \in E: w \neq v} w,$$

and we write

$$c \xrightarrow{v} c'.$$

Note that loops $vv \in E$ have no effect.

A vertex firing is pictured in Figure 3. We think of firing v as sending one grain of sand along each of the edges emanating from v . If the head of e happens to be the sink vertex, then that sand disappears. The vertex $v \in \tilde{V}$ is *stable* in the configuration c if $c(v) < \text{outdeg}(v)$; otherwise, it is *unstable*. For instance, if $c(v)$ is negative, then it is stable. If each vertex is stable in c , then c is *stable*; otherwise, it is *unstable*.

If v is unstable in c , we say that firing v is *legal* for c ; this means that the amount of sand on v after firing v would be nonnegative. We use a jagged arrow to emphasize the legality of a vertex firing, writing $c \xrightarrow{\sim} c'$. Firing a sequence of non-sink vertices v_1, \dots, v_k from $c = c^{(1)}$ produces a sequence of configurations

$$c^{(1)} \xrightarrow{v_1} c^{(2)} \xrightarrow{v_2} \dots \xrightarrow{v_k} c^{(k)}.$$

We say v_1, \dots, v_k is a *legal firing sequence* for c if firing v_i from $c^{(i)}$ is legal for all i .


If c' is reached from c after a sequence $\sigma = v_1, \dots, v_k$ of vertex firings, we write

$$c \xrightarrow{\sigma} c',$$

sometimes omitting the label, σ . If σ is a legal firing sequence, we may write $c \xrightarrow{\sim} c'$. Further, if c' is stable, it is called a *stabilization* of c . As a consequence of the *least action principle* described below, we shall see that every configuration has a unique stabilization. So we can speak of *the* stabilization of a configuration c , denoted c° .

Vertex firing corresponds exactly with the notion of vertex lending for configurations in the context of the dollar game of Part 1. For sandpiles, what we called a borrowing move in the dollar game is now called a *reverse-firing*. Thus, if $c, c' \in \text{Config}(G)$, then c' is obtained from c by firing vertex v if and only if c is obtained from c' through a reverse-firing of v . More generally, a mix of vertex firings and reverse-firings constitutes a *firing script*, $\sigma: \tilde{V} \rightarrow \mathbb{Z}$. These are elements of the abelian group $\tilde{\mathcal{M}}(G)$ introduced in Chapter 2.

A sequence of vertex firings and reverse-firings determines a firing script, although a given firing script will generally arise from many different sequences. Just as in Part 1, the configuration resulting from a sequence of vertex firings is independent of the ordering of the sequence, since these firings just involve addition and subtraction in the abelian group $\text{Config}(G)$. So we have the abelian property for vertex firing just as we did for divisors in Part 1: if $c \in \text{Config}(G)$ and $v, w \in \tilde{V}$, then there is a commutative diagram:



Further, if v and w are both unstable in c , then both firing sequences,

$$c \xrightarrow{v} c' \xrightarrow{w} c''' \quad \text{and} \quad c \xrightarrow{w} c'' \xrightarrow{v} c''',$$

are *legal*. That’s because firing a vertex never removes sand from another vertex and hence can never stabilize another vertex. This stronger version of the abelian property is sometimes referred to as the *local confluence property* or *diamond property* of the sandpile model. Nevertheless, not all rearrangements of a legal firing sequence will be legal.

It turns out that legal firing sequences are distinguished by an efficiency property with respect to stabilization—this is the least action principle described in the next section.

6.3.2. Least action principle. Figure 8 depicts a sequence of vertex firings on

the triangle graph . The vertex v_1 is legal for c , and firing it produces




Figure 8. Vertex firings, not all legal.

the stable configuration c' . But note that the non-legal firing sequence v_1, v_1, v_2 transforms c into the zero configuration, $\tilde{c} = 0$, which is also stable. This example shows that, in general, there may be many stable configurations reachable from a given configuration through a sequence of vertex firings. However, the following theorem shows that, just as in this example, the shortest sequence leading to a stable configuration will be legal.

Theorem 6.7 (Least action principle). *Let $c \in \text{Config}(G)$, and suppose $\sigma, \tau \geq 0$ are firing scripts such that σ arises from a legal firing sequence for c and $c \xrightarrow{\tau} \tilde{c}$ with \tilde{c} stable. Then $\sigma \leq \tau$.*

Proof. Let v_1, \dots, v_k be a legal firing sequence corresponding to σ , so that $\sigma = \sum_i v_i$. The proof goes by induction on k , with the base case $k = 0$ being obvious. Suppose $k > 0$. Since v_1 is unstable in c , the only way that v_1 could be stable in \tilde{c} is if it fired at least once according to τ —it could not have become stable through the firing of other vertices alone. Thus $\tau(v_1) > 0$. Fire v_1 to get a configuration c' ,

and let $\tau' := \tau - v_1$. Then v_2, \dots, v_k is a legal firing sequence for c' , and $c' \xrightarrow{\tau'} \tilde{c}$. By induction, $\sigma - v_1 \leq \tau'$, and the result follows:

$$\sigma \leq \tau' + v_1 = \tau. \quad \square$$

Corollary 6.8 (Uniqueness of stabilization). *Let c be a configuration and $\sigma, \sigma' \geq 0$ firing scripts corresponding to legal firing sequences for c . Suppose that $c \xrightarrow{\sigma} \tilde{c}$ and $c \xrightarrow{\sigma'} \tilde{c}'$, with \tilde{c} and \tilde{c}' both stable. Then $\sigma = \sigma'$ and $\tilde{c} = \tilde{c}'$.*

6.3.3. Existence of stabilizations. The previous corollary shows that stabilizations, if they exist, are unique. In this section, we show existence: every configuration can be transformed into a stable configuration through a sequence of legal vertex firings.

To do this we define an order relation on $\mathbb{Z}\tilde{V}$ in which first, a configuration is judged to be large if it has lots of sand, and second, in the case of two configurations with the same amount of sand, the smaller is the one with *more* sand near the sink. (This mimics the tree orderings used to prove existence and uniqueness of q -reduced divisors in Chapter 3.)

Definition 6.9. Let u_1, \dots, u_n be an ordering of the non-sink vertices of G such that $i < j$ if $d(u_i, s) < d(u_j, s)$, i.e., if u_i is closer to the sink. The *sandpile ordering* of $\text{Config}(G) = \mathbb{Z}\tilde{V}$ is the total ordering \prec defined as follows. Given distinct configurations a, b , let $c := a - b = \sum_{i=1}^n c_i u_i$. Then $a \prec b$ if

- (1) $\deg(a) < \deg(b)$, or
- (2) $\deg(a) = \deg(b)$ and $c_k > 0$ for the smallest k such that $c_k \neq 0$.

Given a sandpile ordering \prec , we employ the usual conventions: writing $a \preceq b$ means $a = b$ or $a \prec b$; writing $a \succ b$ means $b \prec a$; and so on. In particular, $u_1 \prec \dots \prec u_n$. Be careful not to confuse \prec with the partial ordering of configurations, $<$, defined at the beginning of Section 6.3. One may easily check the following:

Exercise 6.10.

- (1) Find configurations $a \prec b$ and infinitely many c such that $a \prec c \prec b$.
- (2) Prove the following:

Property 1. If a, b, c are configurations and $a \prec b$, then $a + c \prec b + c$.

Property 2. If c is a configuration with $0 \leq c$, then $0 \preceq c$, and there are only finitely many configurations c' such that $0 \leq c'$ and $c' \preceq c$.

Lemma 6.11. *Let \prec be a sandpile ordering, and let $c, \tilde{c} \in \mathbb{Z}\tilde{V}$. If $c \rightarrow \tilde{c}$ via a sequence of vertex firings, then $\tilde{c} \prec c$.*

Proof. We may assume that $c \xrightarrow{v} \tilde{c}$ for some vertex v . If $(v, s) \in E$, then when v fires, some sand is lost to the sink. So in that case $\tilde{c} \prec c$ since $\deg(\tilde{c}) < \deg(c)$. Otherwise, since the sink is globally accessible, there exists $(v, u) \in E$ for some vertex u closer to the sink than v , i.e., $u \prec v$. When v fires, the vertex u receives sand and no vertex besides v loses sand. Hence, $\tilde{c} \prec c$ since \tilde{c} has more sand closer to the sink. \square

Theorem 6.12 (Existence). *Every configuration has a stabilization (unique by Corollary 6.8).*

Proof. Fix a sandpile ordering. If c is a sandpile, i.e., if $c \geq 0$, then the fact that c has a stabilization follows immediately from Lemma 6.11 and Property 2 of sandpile orderings, above.

Given an arbitrary configuration c , define the sandpile c^+ by

$$c^+(v) = \max\{0, c(v)\}$$

for each $v \in \tilde{V}$. We have just seen that c^+ is stabilizable. Further, by the least action principle, every legal firing sequence for c^+ is finite. Since every legal firing sequence for c is also a legal firing sequence for c^+ , it follows that c has no infinite legal firing sequence, and is therefore stabilizable. \square

Corollary 6.13. *Let c be a configuration, and suppose σ and τ are two legal firing sequences for c resulting in the same configuration \tilde{c} :*

$$c \xrightarrow{\sigma} \tilde{c} \quad \text{and} \quad c \xrightarrow{\tau} \tilde{c}.$$

Then σ and τ are rearrangements of each other.

Proof. By Theorem 6.12, there is a legal firing sequence μ stabilizing \tilde{c} . Then the concatenated sequences σ, μ and τ, μ are both legal firing sequences stabilizing c . By Corollary 6.8, these two concatenated sequences are rearrangements of each other; therefore, the same holds for σ and τ . \square

6.4. The reduced Laplacian

Just as for the divisor-theory in Part 1, the theory of sandpiles is really the study of the *reduced Laplacian* \tilde{L} , which we now recall and extend to the context of directed sandpile graphs. If $\sigma: \tilde{V} \rightarrow \mathbb{Z}$ is any firing script, then the net effect of implementing σ is to replace a configuration c by a new configuration c' given by

$$\begin{aligned} c' &= c - \sum_{v \in \tilde{V}} \sigma(v) \left(\text{outdeg}(v)v - \sum_{vw \in E: w \neq s} w \right) \\ &=: c - \tilde{L}(\sigma), \end{aligned}$$

where $\tilde{L}: \tilde{\mathcal{M}}(G) \rightarrow \mathbb{Z}\tilde{V}$ is defined implicitly via the sum in the first line. Ordering the non-sink vertices, $\tilde{V} = \{v_1, \dots, v_n\}$, identifies $\mathbb{Z}\tilde{V}$ with \mathbb{Z}^n :

$$\sum_{i=1}^n c_i v_i \in \mathbb{Z}\tilde{V} \quad \longleftrightarrow \quad (c_1, \dots, c_n) \in \mathbb{Z}^n.$$

Consider the basis $\{\chi_1, \dots, \chi_n\}$ for the group of firing scripts $\tilde{\mathcal{M}}(G)$ that is dual to the basis $\{v_1, \dots, v_n\}$ for $\mathbb{Z}\tilde{V}$:

$$\chi_j(v_i) := \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Then in terms of the bases $\{\chi_j\}$ and $\{v_i\}$, the reduced Laplacian \tilde{L} becomes an $n \times n$ integer matrix:

$$\begin{aligned} \tilde{L}_{ij} = \tilde{L}(\chi_j)_i &= \left(\sum_{v \in \tilde{V}} \chi_j(v) \left(\text{outdeg}(v) v - \sum_{vw \in E: w \neq s} w \right) \right)_i \\ &= \left(\text{outdeg}(v_j) v_j - \sum_{v_j w \in E: w \neq s} w \right)_i \\ &= \begin{cases} \text{outdeg}(v_i) - \#(\text{loops at } v_i) & \text{if } i = j \\ -(\# \text{ edges from } v_j \text{ to } v_i) & \text{otherwise.} \end{cases} \end{aligned}$$

Let $\widetilde{\text{out}}(G) = \text{diag}(\text{outdeg}(v_1), \dots, \text{outdeg}(v_n))$ be the diagonal matrix of non-sink vertex outdegrees, and let \tilde{A} be the (*reduced*) *adjacency matrix* for the sandpile graph G , where \tilde{A}_{ij} is the number of edges from v_i to v_j . Then


$$\tilde{L} = \widetilde{\text{out}}(G) - \tilde{A}^t.$$

Of course, the definition of the reduced Laplacian depends on the sink vertex, which is part of the structure of the sandpile graph $G = (V, E, s)$. But the graph (V, E) has a full Laplacian $L: \mathcal{M}(G) \rightarrow \mathbb{Z}V$ given in matrix form by

$$L = \text{out}(G) - A^t,$$

where $\text{out}(G)$ is the $(n + 1) \times (n + 1)$ diagonal matrix of *all* vertex outdegrees, and A is the full adjacency matrix of the graph (V, E) . As in Part 1, the reduced Laplacian \tilde{L} is obtained from the Laplacian L by removing the row and column corresponding to the sink vertex, s .

Example 6.14. The Laplacian for the graph G from Figure 7,



with vertex order u, v, s is

$$\begin{aligned} L = \text{out}(G) - A^t &= \begin{pmatrix} 6 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 3 \end{pmatrix} - \begin{pmatrix} 0 & 3 & 0 \\ 5 & 0 & 2 \\ 1 & 2 & 1 \end{pmatrix} \\ &= \begin{pmatrix} 6 & -3 & 0 \\ -5 & 5 & -2 \\ -1 & -2 & 2 \end{pmatrix}. \end{aligned}$$

Fixing s as the sink vertex, the reduced Laplacian for G is

$$\tilde{L} = \begin{pmatrix} 6 & -3 \\ -5 & 5 \end{pmatrix}.$$

Definition 6.15. The *Laplacian lattice* \mathcal{L} and the *reduced Laplacian lattice* $\tilde{\mathcal{L}}$ are the subgroups formed by the images of the Laplacian and reduced Laplacian, respectively:

$$\mathcal{L} = \text{im } L \subseteq \mathbb{Z}V \quad \text{and} \quad \tilde{\mathcal{L}} = \text{im } \tilde{L} \subseteq \mathbb{Z}\tilde{V}.$$

In the context of the Laplacian, we will usually abuse notation and write the vertex v when we really mean the firing script χ_v . Thus, if $v \in \tilde{V}$, then

$$\tilde{L}v = \text{outdeg}(v)v - \sum_{vw \in E: w \neq s} w.$$

As an immediate consequence of the above discussion, we have the following proposition and its corollary.

Proposition 6.16. *The configuration c' is obtained from the configuration c by firing $v \in \tilde{V}$ if and only if $c' = c - \tilde{L}v$. Equivalently, c is obtained from c' by reverse-firing v if and only if $c = c' + \tilde{L}v$.*

Corollary 6.17. *The configuration c' is obtained from c through a sequence of vertex firings and reverse-firings if and only if*

$$c = c' \text{ mod } \tilde{\mathcal{L}}.$$

Example 6.18. Continuing with the graph from Figure 7 and Example 6.14 with vertex order u, v, s , consider the configuration $c = 2u + 7v = (2, 7)$. Fire vertex v :

$$c = (2, 7) \xrightarrow{v} (5, 2).$$

The corresponding calculation via Proposition 6.16 is

$$\begin{pmatrix} 5 \\ 2 \end{pmatrix} = \begin{pmatrix} 2 \\ 7 \end{pmatrix} - \begin{pmatrix} 6 & -3 \\ -5 & 5 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

6.4.1. Uniqueness of the firing script. By the matrix-tree theorem (Theorem 9.3) or the fact that the reduced Laplacian is an M -matrix (cf. Section 12.3), we will see that the reduced Laplacian for a sandpile graph is invertible. As a consequence we have the following result.

Proposition 6.19. *Let $c, c' \in \text{Config}(G)$ with $c = c' \text{ mod } \tilde{\mathcal{L}}$. Then there exists a unique firing script σ such that $c \xrightarrow{\sigma} c'$.*

Proof. The firing script is $\tilde{L}^{-1}(c - c')$. □

Thus, if $c = c' \text{ mod } \tilde{\mathcal{L}}$, it makes sense to refer to *the* firing script for the ordered pair (c, c') . For each $v \in \tilde{V}$, the v -th component of this firing script records the net number of times v fires in any sequence of firings and reverse-firings leading from c to c' . Negative components correspond to reverse-firings.

6.5. Recurrent sandpiles

In Section 6.1 we repeatedly dropped grains of sand onto randomly chosen vertices of a graph, one at a time, allowing the sandpile to stabilize after each added grain. We observed that some stable configurations appeared many times and others appeared at most once (cf. Table 1). In this section, we explain this phenomenon.

Definition 6.20. A configuration c on G is *recurrent* if

- (1) c is a sandpile, i.e., $c \geq 0$,
- (2) c is stable,
- (3) for every configuration a , there exists a configuration $b \geq 0$ such that $c = (a + b)^\circ$, the stabilization of $a + b$.

The collection of all recurrent elements of G —the *recurrents* of G —is denoted $\mathcal{S}(G)$.

Condition 3 says that starting from any configuration a , we can get back to any given recurrent configuration by adding an appropriate amount of sand and stabilizing. Note that it would make no difference to the definition if we insisted that $a \geq 0$ (Problem 6.4). By the end of this section (cf. Definition 6.32), the set of recurrents will be endowed with a group structure, and $\mathcal{S}(G)$ will be used to denote this group.

Usually, it is not easy to explicitly describe the recurrent configurations of a graph. For instance, how would one have known, a priori, that the eight configurations in Figure 6 are recurrent? In general, there is only one recurrent configuration that is immediately recognizable: the maximal stable configuration.

Definition 6.21. The *maximal stable configuration* on G is

$$c_{\max} := \sum_{v \in \tilde{V}} (\text{outdeg}(v) - 1) v.$$

It is clear that c_{\max} is stable and that $c_{\max} \geq c$ for all stable configurations c . It is also not hard to see that c_{\max} is recurrent. Indeed, given any configuration a , let $b = c_{\max} - a^\circ$. Let σ be a legal firing sequence that stabilizes a . Then $b \geq 0$ and σ is a legal firing sequence for $a + b$. But $(a + b) \xrightarrow{\sigma} (a^\circ + b) = c_{\max}$, which shows that $(a + b)^\circ = c_{\max}$. In fact, the following result shows that c_{\max} is the key to finding *all* the recurrent configurations.

Proposition 6.22. *A configuration c is recurrent if and only if there exists a configuration $b \geq 0$ such that $c = (c_{\max} + b)^\circ$.*

Proof. Problem 6.5. □

Exercise 6.23. Consider the sequence of stable sandpiles produced by repeatedly dropping a grain of sand on a vertex chosen uniformly at random and stabilizing. Why do you expect to eventually reach a recurrent sandpile? Show that from that point on, every sandpile in the sequence is recurrent (thus accounting for what we observed in Table 1).

The zero configuration, $0 := \sum_{v \in \tilde{V}} 0 \cdot v$, is recurrent only under very special circumstances. The graph G is *acyclic* if it has no cycles, i.e., there does not exist

a sequence of edges e_1, \dots, e_k with $e_i^+ = e_{i+1}^-$ for $i = 1, \dots, k - 1$ and $e_k^+ = e_1^-$. We say G is *s-acyclic* if its only cycles are ones that pass through the sink, s .

Proposition 6.24. *The following are equivalent for a sandpile graph G .*

- (1) *The zero configuration is recurrent.*
- (2) *Every stable sandpile is recurrent.*
- (3) *The graph G is s-acyclic.*

Proof. Problem 6.6. □

Definition 6.25. The *stable addition* of sandpiles a and b , denoted $a \otimes b$, is defined as ordinary addition of group elements followed by stabilization:

$$a \otimes b := (a + b)^\circ.$$

Stable addition is clearly commutative with the zero configuration serving as an identity element.

Proposition 6.26. *Stable addition of sandpiles is associative.*

Proof. The result follows by uniqueness of stabilization, Corollary 6.8. Let a, b, c be sandpiles. Since $c \geq 0$, the firing sequence that stabilizes $a + b$ is still legal for $a + b + c$, so

$$a + b + c \rightsquigarrow (a + b)^\circ + c \rightsquigarrow ((a + b)^\circ + c)^\circ = (a \otimes b) \otimes c.$$

Similarly, the sequence that stabilizes $b + c$ is legal for $a + b + c$, and we have

$$a + b + c \rightsquigarrow a + (b + c)^\circ \rightsquigarrow (a + (b + c)^\circ)^\circ = a \otimes (b \otimes c).$$

By uniqueness of stabilization, it follows that

$$(a \otimes b) \otimes c = (a + b + c)^\circ = a \otimes (b \otimes c).$$

□

Definition 6.27. The *sandpile monoid* is the set of nonnegative, stable configurations $\text{Stab}(G)$, with the operation of stable addition.

We now come to a central result for the abelian sandpile model. Recall that $\tilde{\mathcal{L}}$ denotes the image of the reduced Laplacian (Definition 6.15).

Theorem 6.28. *The set of recurrents, $\mathcal{S}(G)$, is a group under stable addition, and*

$$(6.1) \quad \begin{aligned} \mathcal{S}(G) &\rightarrow \mathbb{Z}\tilde{V}/\tilde{\mathcal{L}} \\ c &\mapsto c + \tilde{\mathcal{L}} \end{aligned}$$

is an isomorphism of groups.

The proof of this theorem is approached through several lemmas. For that purpose, consider $1_{\tilde{V}} := \sum_{v \in \tilde{V}} v$, the all ones configuration, and define the two configurations

$$\begin{aligned} c_{\text{big}} &:= c_{\text{max}} + 1_{\tilde{V}} = \sum_{v \in \tilde{V}} \text{outdeg}(v) v \\ c_{\text{null}} &:= c_{\text{big}} - c_{\text{big}}^\circ. \end{aligned}$$

The configuration c_{null} has two useful properties: $c_{\text{null}} = 0 \pmod{\tilde{\mathcal{L}}}$ and $c_{\text{null}} \geq 1_{\tilde{V}}$.

Lemma 6.29. *Each element of $\mathbb{Z}\tilde{V}/\tilde{\mathcal{L}}$ is represented by a recurrent configuration.*

Proof. Given any configuration c , take $k \gg 0$ so that

$$c + kc_{\text{null}} \geq c_{\text{max}}.$$

Working modulo $\tilde{\mathcal{L}}$,

$$(c + kc_{\text{null}})^\circ \equiv c + kc_{\text{null}} \equiv c,$$

and $(c + kc_{\text{null}})^\circ$ is recurrent since it can be formed by adding (a nonnegative amount of) sand to c_{max} and stabilizing:

$$(c + kc_{\text{null}})^\circ = (c_{\text{max}} + (c + kc_{\text{null}} - c_{\text{max}}))^\circ.$$

□

Lemma 6.30. *If c is recurrent, then $(c_{\text{null}} + c)^\circ = c$.*

Proof. Let c be recurrent. From Proposition 6.22, there exists a sandpile $b \geq 0$ such that $(b + c_{\text{max}})^\circ = c$. Setting $a = b - 1_{\tilde{V}}$, it follows that $(a + c_{\text{big}})^\circ = c$. Then

$$\begin{aligned} a + c_{\text{big}} + c_{\text{null}} &\rightsquigarrow (a + c_{\text{big}})^\circ + c_{\text{null}} \\ &= c + c_{\text{null}} \rightsquigarrow (c + c_{\text{null}})^\circ, \end{aligned}$$

and

$$\begin{aligned} a + c_{\text{big}} + c_{\text{null}} &= a + c_{\text{big}} + c_{\text{big}} - c_{\text{big}}^\circ \\ &\rightsquigarrow a + c_{\text{big}} + c_{\text{big}}^\circ - c_{\text{big}}^\circ \\ &= a + c_{\text{big}} \\ &\rightsquigarrow (a + c_{\text{big}})^\circ = c. \end{aligned}$$

The result follows by uniqueness of stabilization. □

Lemma 6.31. *There is a unique recurrent configuration in each equivalence class of $\mathbb{Z}\tilde{V}$ modulo $\tilde{\mathcal{L}}$.*

Proof. By Lemma 6.29 there exists at least one recurrent configuration in each equivalence class of $\mathbb{Z}\tilde{V}$ modulo $\tilde{\mathcal{L}}$. Suppose c' and c'' are recurrents and that $c' = c'' \pmod{\tilde{\mathcal{L}}}$. Then

$$c' = c'' + \sum_{v \in \tilde{V}} n_v \tilde{L}v$$

for some $n_v \in \mathbb{Z}$. Let $J_- := \{v : n_v < 0\}$ and $J_+ := \{v : n_v > 0\}$, and define

$$c := c' + \sum_{v \in J_-} (-n_v) \tilde{L}v = c'' + \sum_{v \in J_+} n_v \tilde{L}v.$$

Take $k \gg 0$ so that for every $v \in \tilde{V}$,

$$(c + kc_{\text{null}})(v) \geq \max_{w \in \tilde{V}} \{|n_w| \text{ outdeg}(w)\}.$$

Thus, each vertex v of $c + kc_{\text{null}}$ can be legally fired $|n_v|$ times. Therefore,

$$c + kc_{\text{null}} = c' + \sum_{v \in J_-} (-n_v) \tilde{L}v + kc_{\text{null}} \rightsquigarrow c' + kc_{\text{null}} \rightsquigarrow c',$$

the last step following from repeat applications of Lemma 6.30. Similarly, we have $c + kc_{\text{null}} \rightsquigarrow c''$. By uniqueness of stabilization, $c' = c''$. \square

We now prove Theorem 6.28.

Proof. The mapping

$$\begin{aligned} \mathcal{S}(G) &\rightarrow \mathbb{Z}\tilde{V}/\tilde{\mathcal{L}} \\ c &\mapsto c + \tilde{\mathcal{L}} \end{aligned}$$

respects addition by Proposition 6.16. It is surjective by Lemma 6.29 and injective by Lemma 6.31. \square

Definition 6.32. The *sandpile group* of G , denoted $\mathcal{S}(G)$, is the set of recurrent configurations with stable addition.

As a consequence of Theorem 6.28, we can determine the structure of $\mathcal{S}(G)$ by computing the *Smith normal form* of \tilde{L} . Details appear in Chapter 2. The size of the sandpile group may be calculated as follows.

Proposition 6.33.


$$|\mathcal{S}(G)| = |\det(\tilde{L})|.$$

Proof. The proof given for Proposition 2.37 holds equally well here, for directed graphs. \square

Remark 6.34. As pointed out after Proposition 2.37, $\det(\tilde{L})$ is the number of spanning trees (in this case directed, rooted spanning trees into the sink), and hence positive. Thus we actually have

$$|\mathcal{S}(G)| = \det(\tilde{L}).$$

Exercise 6.35. Let G be a triangle with an edge of multiplicity 3 to the sink, s .



- (1) Find the number of elements in $\mathcal{S}(G)$ using Proposition 6.33.
- (2) Find all the recurrents on G .

6.5.1. The identity. We saw in Proposition 6.24 that the zero configuration is rarely recurrent. Hence, even though $(0 + c)^\circ = c$ for all configurations c , it is seldom the case that 0 is the identity of the sandpile group. And if the identity is not 0, then what is it? That's an unusual question to ask about a group! A glance at the images of identity elements on grid graphs in Section 6.6 indicates some of the complexity.

We at least have a way of calculating the identity:

Proposition 6.36. *The identity of $\mathcal{S}(G)$ is*

$$(2c_{\max} - (2c_{\max})^\circ)^\circ.$$

Proof. The displayed element is recurrent since

$$(2c_{\max} - (2c_{\max})^\circ)^\circ = (c_{\max} + \underbrace{(c_{\max} - (2c_{\max})^\circ)}_{\geq 0})^\circ,$$

and as an element of $\mathbb{Z}\tilde{V}/\tilde{\mathcal{L}}$ it is equal to (the equivalence class of) 0, the identity of $\mathbb{Z}\tilde{V}/\tilde{\mathcal{L}}$. This completes the proof since the isomorphism $\mathcal{S}(G) \rightarrow \mathbb{Z}\tilde{V}/\tilde{\mathcal{L}}$ of (6.1) preserves the identity. \square

Example 6.37. Let G be the diamond graph pictured in Figure 1 (reproduced below) with vertex order v_1, v_2, v_3, s .




Fig. 1

Then $c_{\max} = (2, 2, 1)$. One may check that

$$2c_{\max} = (4, 4, 2) \rightsquigarrow (2, 2, 0) = (2c_{\max})^\circ,$$

with firing script $(3, 3, 4)$, and

$$2c_{\max} - (2c_{\max})^\circ = (2, 2, 2) \rightsquigarrow (2, 2, 0) = (2, 2, 0)^\circ.$$

Hence, the sandpile identity for G is $(2, 2, 0)$. Note that

$$(2, 2, 0) = \tilde{L}(2v_1 + 2v_2 + 2v_3),$$

confirming that $(2, 2, 0) = 0 \pmod{\tilde{\mathcal{L}}}$.

6.6. Images of sandpiles on grid graphs

The $m \times n$ sandpile grid graph has non-sink vertices

$$\tilde{V} := \{(i, j) \in \mathbb{Z}^2 : 1 \leq i \leq m, 1 \leq j \leq n\}$$

with edges between horizontal and vertical neighbors: $(i, j), (i', j') \in \tilde{V}$ are adjacent if

$$|i - i'| + |j - j'| = 1.$$

Vertices on the boundary are connected to the sink vertex s . So if $i \in \{1, m\}$ or $j \in \{1, n\}$, then (i, j) is adjacent to the sink. Further, each of the four corner vertices, $(1, 1)$, $(1, n)$, $(m, 1)$, and (m, n) , is connected by an extra edge to the sink (two in total). Thus, every non-sink vertex has degree 4. Figure 9 shows the 4×4 sandpile grid graph.

Sandpiles on grids display conveniently on computer screens. Each pixel is taken to represent a vertex, and the color of a pixel represents the number of grains of sand on the corresponding vertex. See Figure 10 for images of the identity elements for various sandpile grid graphs.




Figure 9. The 4×4 sandpile grid graph. The dashed boundary edges connect to the sink vertex, which is not pictured.


One may consider more general sandpile grid graphs. We could change the boundary, for instance: Cut out a connected region of the integer lattice \mathbb{Z}^2 containing a finite number of points (i, j) . Points in the interior have edges to their horizontal and vertical neighbors, and points on the boundary have edges to the sink so that each non-sink vertex has degree 4. Figure 11 displays identity elements for grid graphs with diamond (rotated square) and circular boundaries, respectively. Another possible generalization is to consider non-rectangular grids. Figure 13 displays the identity element on a hexagonal grid for which each vertex now has 6 neighbors.

Figure 12 represents the stabilization of a sandpile initially consisting of 2^{24} grains of sand in the center of an enormous rectangular grid graph—large enough so that no grain falls into the sink during stabilization.


6.6.1. Patterns. People are just beginning to understand the patterns in these sandpiles. For instance, generalizing Figure 12, imagine placing n grains of sand on a single vertex of an enormous grid graph and stabilizing. As n grows, rescale (by a factor of \sqrt{n}) so that the image of the resulting sandpile does not change size. Pegden and Smart ([76]) have shown that these scaled sandpiles approach a limit as $n \rightarrow \infty$. Patterns in the limit persist and, in fact, are parametrized by a space related to Apollonian circle packings ([67], [68])!

In some of the images we’ve discussed above, there are thin filaments—patterns of thin line segments. For example, look inside the curved triangles in the identity element for the 4000×4000 grid in Figure 10. If you do a web search for images of “tropical curves”, you will see a similarity. Tropical curves are discrete versions of Riemann surfaces (a graph, as studied in Part 1 of this text, being an even “more discrete” version). A concrete connection between these patterns in sandpiles and tropical curves has been forged by Kalinin and Shkolnikov ([61]).

There are many open problems concerning these patterns. One is the following: characterize the large rectangle of vertices containing 2 grains of sand inside the identity element for an $m \times n$ sandpile grid graph, and prove that it exists (cf. Figure 10).



4000 × 4000



4 × 4

16 × 16

50 × 50

100 × 100

grains: 0 = ■, 1 = ■, 2 = ■, 3 = ■.

Figure 10. The identity elements for the sandpile group of several rectangular grid sandpiles. The existence of the interior rectangle of height 2 for all rectangular grid graphs is an open problem.




Figure 11. Identity elements on sandpile grid graphs with a diamond (a square rotated with respect to the underlying grid) or circular boundary. Color coding as in Figure 10.




Figure 12. The stabilization of 2^{24} grains of sand placed on a single vertex in a grid graph without boundary. Color coding as in Figure 10. (Image due to Wesley Pegden.)

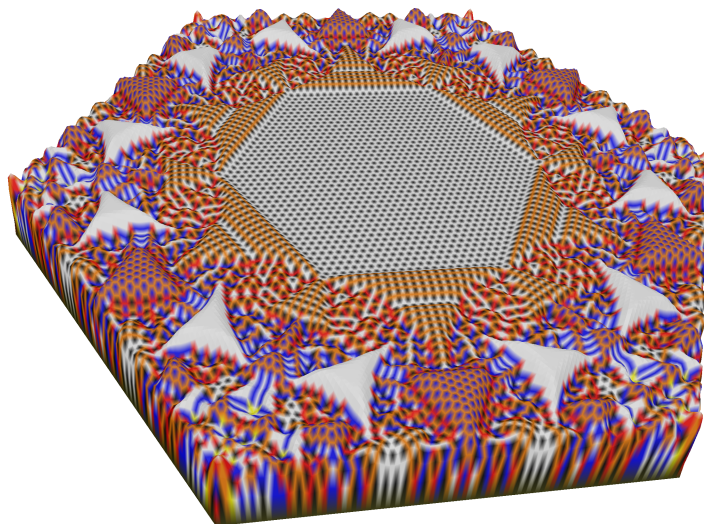


Figure 13. The identity element for a hexagonal grid graph shown in three dimensions with heights smoothed.

Notes

The idea underlying the abelian sandpile model, in which firing rules are determined by the Laplacian, is something of a folk science, having roots spreading in many directions. We will not attempt to summarize the history here. The point of view presented in this chapter comes from Dhar ([34]), who was inspired by Bak, Tang, and Wiesenfeld's work ([4]) on self-organized criticality (cf. Section 12.4). Some of the other earlier work that has informed our understanding of the fundamentals is by: Biggs ([14], [13]); Björner, Lovász, and Shor ([17], [18]); Cori, Rossin, and Salvy ([28], [29]); and Lorenzini ([69], [70]). We are also much in debt to the paper *Chip-firing and rotor-routing on directed graphs*, by Holroyd, Levine, Mészáros, Peres, Propp, and Wilson ([56]).

The images of sandpile identities were created using software available at [44] and [55].

Problems for Chapter 6

6.1. Let $G = (V, E, s)$ be an undirected sandpile graph. Use the greedy lending algorithm described at the end of Section 5.4 to give an alternate proof of Corollary 6.8 in the undirected case.

6.2. This problem gives an alternate proof of Theorem 6.12, the existence of a stabilization. Let c be a configuration, and let $\eta = v_1, v_2, \dots$ be any legal sequence of vertex firings for c . Given $v \in \tilde{V}$, choose a directed path e_1, \dots, e_k from v to the sink, s . Argue that the most number of times that v can appear in η is strictly less than

$$(\deg(c) + 1) \prod_{i=2}^k \text{outdeg}(e_i^-).$$

6.3. Let a and b be configurations.

- (a) Give an example showing that $(a + b)^\circ$ is not necessarily equal to $(a^\circ + b^\circ)^\circ$.
- (b) Show that if a and b are both sandpiles, then $(a + b)^\circ = (a^\circ + b^\circ)^\circ$.

6.4. Show that a configuration c is recurrent if and only if for each *nonnegative* configuration a , there exists a configuration $b \geq 0$ such that $(a + b)^\circ = c$.

6.5. Prove Proposition 6.22.

6.6. Prove Proposition 6.24.

6.7. Prove that a recurrent configuration must have at least one grain of sand on each directed cycle not containing the sink vertex.

6.8. Find the recurrent configurations for the graph in Figure 1 but with v_1 taken as the sink vertex rather than s .

6.9. Let e be the identity for the sandpile group for a sandpile graph G , and let c be a sandpile on G . Show that c is recurrent if and only if $(e + c)^\circ = c$.

6.10. If $c \in \mathcal{S}(G)$ and k is a positive integer, as usual, define


$$kc := \underbrace{c \otimes \cdots \otimes c}_{k \text{ times}} = \underbrace{(c + \cdots + c)^\circ}_{k \text{ times}}.$$

[If $k = 0$, then kc is defined to be the identity of $\mathcal{S}(G)$, and for $k < 0$, define $kc := -k(-c)$, where $-c$ is the inverse of c in $\mathcal{S}(G)$.] The *order* of $c \in \mathcal{S}(G)$ is the least positive integer k such that kc is the identity.

- (a) Prove that the order of c is the least common multiple of the denominators of the entries of $\tilde{L}^{-1}c$.
- (b) Illustrate this with the recurrent element $(1, 1, 1)$ on the cycle graph with 4 vertices, C_4 .

6.11. Describe an algorithm for computing the inverse of an element $c \in \mathcal{S}(G)$.

6.12. Let \tilde{L} be the reduced Laplacian of the following sandpile graph G with sink s :



- (a) By hand, find 3×3 integer matrices U and V , invertible over the integers, such that $U\tilde{L}V = D$ where D is the Smith normal form of \tilde{L} .
- (b) Use the computation you just made to give an explicit isomorphism

$$\mathbb{Z}^3 / \tilde{\mathcal{L}} \approx \mathbb{Z}/d_1\mathbb{Z} \times \mathbb{Z}/d_2\mathbb{Z} \times \mathbb{Z}/d_3\mathbb{Z}$$

for some integers d_i satisfying $d_1|d_2|d_3$.


6.13. Compute the invariant factors for the sandpile group of the complete graph K_n .

6.14. This problem presents a characterization of recurrent configurations due to Babai and Toumpakari ([1]). A *semigroup* is a set N with an associative binary operation. A *monoid* is a semigroup with an identity element. The product of subsets $A, B \subseteq N$ is defined as it is for groups: $AB := \{ab : a \in A, b \in B\}$. A nonempty subset $I \subseteq N$ is an *ideal* if $IN = NI = I$. An ideal is *minimal* if the only ideal it contains is itself. The intersection of all ideals of N is either empty or the unique minimal ideal of N .

Let \mathcal{M} be a finite commutative monoid. Prove the following:

- (a) \mathcal{M} has a unique minimal ideal \mathcal{S} .
- (b) $c \in \mathcal{S}$ if and only if for all $a \in \mathcal{M}$, there exists $b \in \mathcal{M}$ such that $c = ab$.
- (c) \mathcal{S} is an abelian group. (Hint: Take $a \in \mathcal{S}$ and consider the sequence a, a^2, \dots . Finiteness implies there exists $i < j$ such that $a^i = a^j$, i.e., $a^i = a^i a^{j-i}$. Define $e := a^{i(j-i)}$ and argue that e is an *idempotent* of \mathcal{S} , meaning $e^2 = e$. This e will be the identity for the group.)
- (d) Suppose \mathcal{M} is the sandpile monoid for a graph G . Then,
 - (i) $\mathcal{S} = \mathcal{S}(G)$, the sandpile group;
 - (ii) $\mathcal{S}(G)$ is a *principal ideal* in the sandpile monoid, i.e., $\mathcal{S}(G) = \mathcal{M}c$ for some sandpile c .

6.15. *Computing with sandpiles.* Consider the following sandpile on a long path graph:



There is a single unstable vertex, having 2 grains of sand. When that vertex topples, the resulting sandpile is essentially the same but shifted to the right one vertex. We regard the traveling 0-2 pair as a signal propagating along a wire.

Next, Figure 14 shows how to combine a “red” and a “blue” copy of this wire to create a new kind of wire capable of sending *information*. The underlying graph G is a pair of long path graphs. The figure shows two sandpiles on G , each having a pair of unstable vertices. Simultaneously firing these vertices produces essentially the same configurations but translated one step to the right, as before. One can imagine Alice starting with the all ones configuration on G . To encode a bit of information, she picks up a grain of sand on each of the wires and displaces these to the right. This information then travels rightward down the wire to Bob who can decode it based on the order in which the 0-2 pairs arrive.




Figure 14. Encoding bits of information in a sandpile. The underlying graph is a pair of long path graphs.

We would now like to apply logical operations to the information traveling in our wires. Figure 15 shows two types of logic gates. At the top is a NOT-gate. It is shown with the bit 0 encoded on the left. After passing through the gate, the bit is changed to a 1. Conversely, if the bit had started as a 1, it would become a 0 after traveling through the gate.

The bottom of Figure 15 shows an OR-gate. Its input on the left is a pair of (encoded) bits. In general, the gate transforms the pair of bits x, y to 0 if $x = y = 0$ and otherwise it transforms the pair to 1. The input of 0 and 1 shown in the figure will emerge from the gate as a 1.

- Verify that the OR-gate pictured in Figure 15 actually works. (You will notice that bits in the two incoming wires must be synchronized.)
- An AND-gate would take bits x and y and output a 1 if and only if $x = y = 1$. Explain how to construct an AND-gate from NOT-gates and an OR-gate.
- Create a more elegant AND-gate by slightly modifying the OR-gate of Figure 15.
- Create a “splitter”. One wire comes into the splitter from the left and two leave to the right. An incoming bit x is split into two (synchronized) copies of x leaving to the right.

After seeing these logic gates, readers familiar with the theory of computation might expect that these wires could be used to construct a Turing machine (an idealized computer), and indeed that it that case. See [48]—from which the ideas

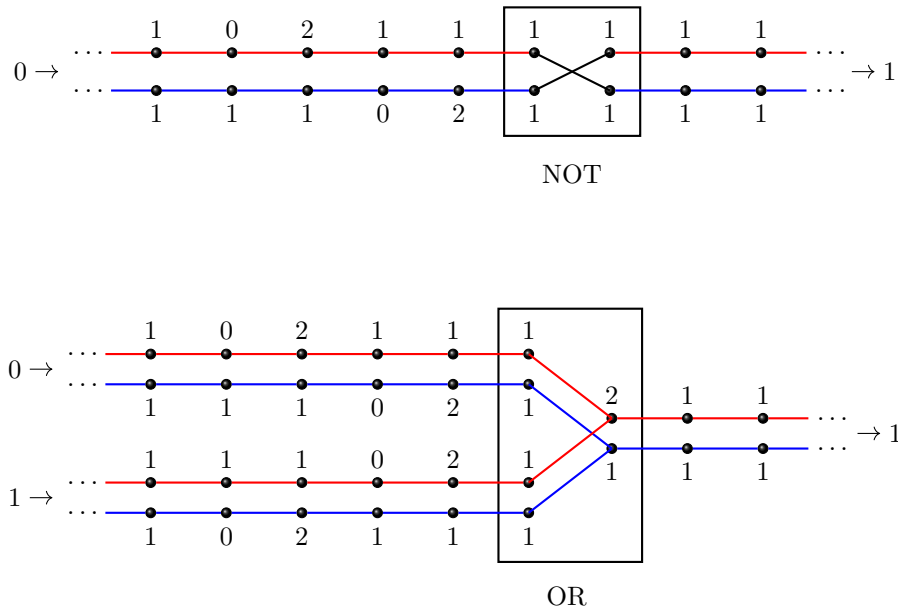


Figure 15. Sandpile logic gates.

for this problem were derived—for details. For an extension of the sandpile model to a model of computation in a network, see [21].

Burning and duality

How would you determine if a sandpile c on a sandpile graph is recurrent? One way would be to start generating all of the recurrents—say, by systematically dropping grains of sand onto the maximal stable sandpile and stabilizing—checking whether c appears. Another way is to check whether $(e + c)^\circ = c$ where e is the identity element of the sandpile group (cf. Problem 6.9). However, it may take a long time to compute e if the graph is large. The good news is that each graph has a *burning sandpile*, quickly computed, that can be substituted for e in the above test. That is, a sandpile c is recurrent if and only if adding the burning sandpile to c and stabilizing returns c .

As an example, consider the sandpile graph G of Figure 1 with sink s and vertex order v_1, v_2, v_3, s . (The edge from v_1 to v_3 has multiplicity 2.) The burning




Figure 1. Sandpile graph G .

sandpile for G is $\mathbf{b} = (1, 0, 1) = v_1 + v_3$. Thus, one may test if a sandpile on G is recurrent by adding a grain of sand to vertices v_1 and v_3 and stabilizing. For instance, if $c = (0, 1, 2)$, then $(\mathbf{b} + c)^\circ = (2, 2, 1) \neq c$. So c is not recurrent. On the other hand, $(\mathbf{b} + (2, 2, 1))^\circ = (2, 2, 1)$, and hence, $(2, 2, 1)$ is recurrent. Further, it turns out that here, as in general, $\mathbf{b} \in \tilde{\mathcal{L}}$, the image of the reduced Laplacian, and hence, $(2, 2, 1) = (c + 2\mathbf{b})^\circ$ is the unique recurrent obtainable from c via firings and reverse-firings (cf. Theorem 6.28).

Exercise 7.1. Let G be the graph in Figure 1.

- (1) Prove that $\mathbf{b} = (1, 0, 1) \in \tilde{\mathcal{L}}$.
- (2) Show that $(2, 2, 1)$ is recurrent by finding a sandpile a such that $(2, 2, 1) = (c_{\max} + a)^\circ$.

As hinted at by this example, not only does the burning sandpile provide a test for recurrence, it also allows us to construct the recurrent equivalent to a given configuration c modulo $\tilde{\mathcal{L}}$: starting at c construct a sequence of configurations by continually adding the burning sandpile and stabilizing. Eventually, the sequence becomes constant at a recurrent sandpile.

We will present an algorithm for constructing the burning sandpile, \mathbf{b} , but for an undirected graph, it is easy to describe: just fire the sink. So in that case, $\mathbf{b}(v)$ is the number of edges from v to the sink for each $v \in \tilde{V}$.

7.1. Burning sandpiles

Let $G = (V, E, s)$ be a sandpile graph with sink s . If $W \subseteq V$ and $v \in V$, we define the *indegree of v with respect to W* to be

$$\text{indeg}_W(v) = |\{w \in W : (w, v) \in E\}|.$$

Definition 7.2. A vertex $v \in \tilde{V} = V \setminus \{s\}$ is *selfish* if $\text{indeg}_{\tilde{V}}(v) > \text{outdeg}(v)$.

Graphs with no selfish vertices are particularly amenable to the methods we are about to introduce. These graphs include undirected or, more generally, Eulerian graphs (see Appendix A, Definition A.28 and Proposition A.29). In fact, if G has no selfish vertices, then it is essentially Eulerian in the following sense: by adjusting the number of edges emanating from the sink, we can form an Eulerian graph G' on the same vertex set as G with the property that the mapping $c \mapsto c$ gives an isomorphism of sandpile groups, $\mathcal{S}(G) \approx \mathcal{S}(G')$ (Problem 7.1).

Definition 7.3. The *support* of a configuration c on G is

$$\text{supp}(c) := \{v \in \tilde{V} : c(v) \neq 0\}.$$

The *closure of the support of c* , denoted $\overline{\text{supp}}(c)$, is the set of non-sink vertices accessible from $\text{supp}(c)$ via a directed path in G that avoids the sink.

Definition 7.4. A nonnegative configuration b on G is a *burning sandpile* if

- (1) $b \equiv 0 \pmod{\tilde{\mathcal{L}}}$, and
- (2) $\overline{\text{supp}}(b) = \tilde{V}$.

If b is a burning sandpile, we call $\sigma_b := \tilde{L}^{-1}b$ its (*burning*) *script*.

Theorem 7.5. Let b be a burning sandpile for G with burning script σ_b . Let e be the identity of $\mathcal{S}(G)$. Then

- (1) $(kb)^\circ = e$ for some $k \gg 0$.
- (2) A sandpile c is recurrent if and only if $(b + c)^\circ = c$.

- (3) A sandpile c is recurrent if and only if the firing script for the stabilization $(b+c) \rightsquigarrow (b+c)^\circ$ is σ_b .
- (4) $\sigma_b \geq \chi_{\tilde{V}}$ where $\chi_{\tilde{V}}$ is the firing script corresponding to firing all non-sink vertices.
- (5) Suppose c is a stable configuration and β is the firing script for $(b+c) \rightsquigarrow (b+c)^\circ$. Then $\beta \leq \sigma_b$.

Proof. (1) By choosing k large enough and selectively firing unstable vertices, property 2 of the definition of a burning sandpile says $kb \rightsquigarrow c + c_{\max}$ for some sandpile c . Thus, $(kb)^\circ$ is recurrent since it can be obtained from c_{\max} by adding sand and stabilizing. The unique recurrent configuration equal to 0 modulo $\tilde{\mathcal{L}}$ is the identity element. Hence, $(kb)^\circ = e$.

(2) (\Rightarrow) If c is recurrent, then so is $(b+c)^\circ$. However, since $c \equiv (b+c)^\circ \pmod{\tilde{\mathcal{L}}}$, we conclude $c = (b+c)^\circ$ by uniqueness of recurrent representatives, as just above.

(\Leftarrow) Suppose $c = (b+c)^\circ$. Using part (1), fix $k \gg 0$ so that $(kb)^\circ = e$. Then

$$c = (kb+c)^\circ = (e+c)^\circ.$$

Since e is recurrent, so is c .

(3) Let ϕ be the firing script for $b+c \rightsquigarrow (b+c)^\circ$. Then

$$\begin{aligned} c \text{ is recurrent} &\iff (b+c)^\circ = c \\ &\iff b+c - \tilde{\mathcal{L}}\phi = c \\ &\iff b = \tilde{\mathcal{L}}\phi \\ &\iff \phi = \tilde{\mathcal{L}}^{-1}b = \sigma_b. \end{aligned}$$

(4) Since c_{\max} is recurrent, the firing script for

$$(b+c_{\max}) \rightsquigarrow (b+c_{\max})^\circ = c_{\max}$$

is σ_b by part (3). Let $v \in \tilde{V}$. Since b is a burning sandpile, there exists $w \in \text{supp}(b)$ and a directed path v_1, v_2, \dots, v_m in G with $w = v_1$ and $v_m = v$. Then v_1, \dots, v_m is a legal firing sequence for $b+c_{\max}$. Recall that the firing script is independent of any particular firing sequence. Thus, each non-sink vertex fires at least once in the stabilization of $b+c_{\max}$. So $\sigma_b \geq \chi_{\tilde{V}}$.

(5) Suppose c is a stable configuration. Let β be the firing script for the stabilization of $b+c$, and let F be a sequence of legal vertex firings stabilizing $b+c$. Then F is also a legal sequence of firings for $b+c_{\max}$, yielding the configuration $a := b+c_{\max} - \tilde{\mathcal{L}}\beta$. Let γ be the firing vector for the stabilization of a . Since $b \equiv 0 \pmod{\tilde{\mathcal{L}}}$ and c_{\max} is recurrent, the stabilization of a is c_{\max} and the firing script for the stabilization of $b+c_{\max}$ is σ_b . By uniqueness of the firing script for a stabilization, $\sigma_b = \beta + \gamma$. Since $\gamma \geq 0$, it follows that $\beta \leq \sigma_b$. \square

7.2. Existence and uniqueness

Theorem 7.6. *There exists a unique burning sandpile \mathbf{b} for G with minimal script $\sigma_{\mathbf{b}} = \tilde{\mathcal{L}}^{-1}\mathbf{b}$: if $\sigma_{b'}$ is the script for another burning sandpile b' , then $\sigma_{b'} \geq \sigma_{\mathbf{b}}$. Henceforth, we call \mathbf{b} the burning sandpile and $\sigma_{\mathbf{b}}$ the burning script for G . For \mathbf{b} ,*

- (1) for all $v \in \tilde{V}$, we have $\mathbf{b}(v) < \text{outdeg}(v)$ unless $\text{indeg}_{\tilde{V}}(v) = 0$, (i.e., not counting edges from the sink, v is a source), in which case $\mathbf{b}(v) = \text{outdeg}(v)$;
- (2) $\sigma_{\mathbf{b}} \geq \chi_{\tilde{V}}$ with equality if and only if G has no selfish vertices.

Proof. We create $\sigma_{\mathbf{b}}$ and $\mathbf{b} = \tilde{L}\sigma_{\mathbf{b}}$ with a greedy algorithm. Start by defining $\sigma^0 = \chi_{\tilde{V}}$ and $b^0 = \tilde{L}\sigma^0$ (the superscript is just an index), then proceed to build a sequence of scripts σ^i and configurations b^i . For $i \geq 0$, if $b^i \geq 0$, stop, letting $\mathbf{b} = b^i$ and $\sigma_{\mathbf{b}} = \sigma^i$. Otherwise, choose a vertex w such that $b^i(w) < 0$ and define

$$\sigma^{i+1} := \sigma^i + w \quad \text{and} \quad b^{i+1} := b^i + \tilde{L}w = \tilde{L}\sigma^{i+1}.$$

In other words, b^{i+1} is obtained from b^i by reverse-firing w .


Since $b^0(v) = \text{outdeg}(v) - \text{indeg}_{\tilde{V}}(v)$ for all v , it follows for all i by induction that $b^i(v) \leq \text{outdeg}(v)$ with equality if and only if $\text{indeg}_{\tilde{V}}(v) = 0$. Thus, part 1 follows once we prove that the process halts. Part 2 is then immediate by construction.

To show that the process halts, define a sequence of configurations $c^i := c_{\max} - b^i$, and note that b^i is nonnegative if and only if c^i is stable. Moreover, the greedy algorithm that constructs the sequence b^i corresponds to the greedy toppling algorithm that stabilizes $c^0 = c_{\max} - b^0$. Since stabilizations exist and are unique, we see that the process does halt at some uniquely determined sandpile \mathbf{b} as claimed. Since $\mathbf{b} = \tilde{L}\sigma_{\mathbf{b}} \in \tilde{\mathcal{L}}$, we just need to show that $\overline{\text{supp}}(\mathbf{b}) = \tilde{V}$, and that $\sigma_{b'} \geq \sigma_{\mathbf{b}}$ for any other burning sandpile b' .

We show that the closure of the support of \mathbf{b} is all of \tilde{V} by contradiction. So suppose that $W := \tilde{V} \setminus \overline{\text{supp}}(\mathbf{b}) \neq \emptyset$, and let \tilde{L}_W denote the square matrix obtained from \tilde{L} by retaining only the rows and columns labeled by vertices in W . If we order the vertices \tilde{V} so that those in W appear first, then the reduced Laplacian \tilde{L} has the block form:

$$\tilde{L} = \left[\begin{array}{c|c} \tilde{L}_W & 0 \\ \hline \star & \star \end{array} \right].$$

The upper right block is 0 because, by assumption, there are no directed edges pointing from a vertex of $\overline{\text{supp}}(\mathbf{b})$ into W :



Moreover, \tilde{L}_W is itself the reduced Laplacian matrix of the sandpile graph obtained from G by identifying all vertices in $\overline{\text{supp}}(\mathbf{b})$ with the sink of G . Therefore, \tilde{L}_W is invertible.

Write \mathbf{b}_W and $\sigma_{\mathbf{b},W}$ for the subvectors of \mathbf{b} and $\sigma_{\mathbf{b}}$, respectively, with components indexed by W . Since $W \cap \text{supp}(\mathbf{b}) = \emptyset$, we see that $\mathbf{b}_W = 0$. Looking at the W -rows of the equation $\mathbf{b} = \tilde{L}\sigma_{\mathbf{b}}$, we see that $0 = \mathbf{b}_W = \tilde{L}_W\sigma_{\mathbf{b},W}$, which implies that $\sigma_{\mathbf{b},W} = 0$ by the invertibility of \tilde{L}_W . But this contradicts the fact that

$\sigma_{\mathbf{b}} \geq \chi_{\tilde{V}}$. Thus we see that $W = \emptyset$ and $\overline{\text{supp}}(\mathbf{b}) = \tilde{V}$, so \mathbf{b} is a burning sandpile as claimed.

Finally, let b' be another burning sandpile for G . We have $\sigma_{b'} \geq \chi_{\tilde{V}}$ by Theorem 7.5 (4), so $\tau' := \sigma_{b'} - \chi_{\tilde{V}} \geq 0$. Now note that

$$c^0 - \tilde{L}\tau' = c_{\max} - \tilde{L}\chi_{\tilde{V}} - \tilde{L}(\sigma_{b'} - \chi_{\tilde{V}}) = c_{\max} - b'$$

which is stable. But $\tau := \sigma_{\mathbf{b}} - \chi_{\tilde{V}}$ comes from a legal firing sequence for c^0 , so by the least action principle (Theorem 6.7) we have $\tau \leq \tau'$, which implies that $\sigma_{\mathbf{b}} \leq \sigma_{b'}$ as required. □

Remark 7.7 (Burning script algorithm). The proof of Theorem 7.6 provides an algorithm for computing the burning sandpile and its script. Briefly, the procedure is as follows: Let b be the sum of the columns of the reduced Laplacian, \tilde{L} . If $b(v) < 0$ for some vertex v , update b by adding in the v -th column of \tilde{L} . Continue until b first becomes nonnegative. Keeping track of which columns of \tilde{L} were added to get the burning sandpile gives the burning script.

Example 7.8. The reduced Laplacian of the graph in Figure 1 with respect to the vertex order v_1, v_2, v_3 is

$$\tilde{L} = \begin{pmatrix} 4 & -1 & -1 \\ -1 & 3 & -1 \\ -2 & -1 & 2 \end{pmatrix}.$$

Adding its columns gives $b = (2, 1, -1)$. Since $b(v_3) < 0$, add the third column of \tilde{L} to b to get $b = (1, 0, 1)$. Since b is now nonnegative, it is the burning sandpile. To form b we first added the columns of \tilde{L} , then added in the third column; so the burning script is $(1, 1, 2)$.

Example 7.9 (Undirected graphs). If G is an undirected sandpile graph with sink s , then the algorithm for constructing the burning sandpile halts immediately with burning script $\sigma_{\mathbf{b}} = \chi_{\tilde{V}}$. The burning sandpile is

$$\mathbf{b} = \sum_{v \in \tilde{V}: vs \in E} v.$$

It is the configuration that would be produced on G by firing the sink vertex. Thus, if all edges in G have multiplicity 1, to test if a configuration c is recurrent, add one grain to c at each vertex connected to the sink and stabilize. Then c is recurrent if and only if the stabilization is c or, equivalently, if and only if each non-sink vertex fires exactly once during the stabilization.

7.3. Superstables and recurrents

As defined in Section 3.3, a configuration c on an *undirected* graph is superstable if $c \geq 0$ (i.e., c is a sandpile) and c has no legal set-firings. We now make the appropriate generalization for directed graphs (cf. Remark 7.13 (3)).

Let $c \in \text{Config}(G)$. Then *script-firing* by the script $\sigma \in \mathbb{Z}\tilde{V}$ results in the configuration $c' = c - \tilde{L}\sigma$. If c is a sandpile, we say the script-firing is *legal* for c if $c' \geq 0$; in other words, after the script is fired, no vertices have a negative amount of sand. Note that this is weaker than our earlier notion of a firing script arising from a legal sequence of vertex firings: a firing script is legal for c provided that the final configuration is a sandpile—there may be no legal sequence of firings that implements the script. We say a sandpile c has *no legal script-firings* if there is no firing script $\sigma \geq 0$ that is legal for c .

Definition 7.10. A sandpile c is *superstable* if it has no legal script-firings.

Exercise 7.11. Give an example of a stable sandpile that is not superstable.

Theorem 7.12. *The following are equivalent for a sandpile c :*

- (1) c is recurrent.
- (2) $c_{\max} - c$ is superstable.
- (3) $c + \tilde{L}\sigma$ is unstable for all $\sigma \geq 0$.
- (4) $c + \tilde{L}\sigma$ is unstable for all $0 \leq \sigma \leq \sigma_{\mathbf{b}}$.

Remark 7.13.

- (1) Theorem 7.12 implies a duality between recurrents and superstable:

c is recurrent if and only if $c_{\max} - c$ is superstable.

Note that it easily follows that c is superstable if and only if $c_{\max} - c$ is recurrent, as well.

- (2) From part 3, we see that a sandpile is recurrent if and only if performing a sequence of reverse firings always results in a configuration with an unstable vertex.
- (3) In the case of an undirected graph, $\sigma_{\mathbf{b}} = \chi_{\tilde{V}}$ (Example 7.9). So from part 4:

$$\begin{aligned} c \text{ is superstable} &\iff c_{\max} - c \text{ is recurrent} \\ &\iff c_{\max} - c + \tilde{L}\chi_W \text{ is unstable for all } \emptyset \neq W \subseteq \tilde{V} \\ &\iff c - \tilde{L}\chi_W \not\geq 0 \text{ for all } \emptyset \neq W \subseteq \tilde{V} \\ &\iff c \text{ has no legal set firings.} \end{aligned}$$

Thus, our definition of superstable generalizes that given in Part 1 for undirected graphs.

Lemma 7.14. *If $\tilde{L}\sigma \geq 0$, then $\sigma \geq 0$.*

Proof. Suppose $\tilde{L}\sigma \geq 0$. Let $N := \{v \in \tilde{V} : \sigma(v) < 0\}$ and assume that $N \neq \emptyset$. Since \tilde{L} is nonpositive off the diagonal and $\sigma(w) \geq 0$ for $w \notin N$,

$$\begin{aligned} 0 \leq \sum_{v \in N} (\tilde{L}\sigma)(v) &= \sum_{v \in N} \sum_{w \in \tilde{V}} \tilde{L}_{vw} \sigma(w) \\ &\leq \sum_{v \in N} \sum_{w \in N} \tilde{L}_{vw} \sigma(w) \end{aligned}$$

$$= \sum_{w \in N} \left(\sum_{v \in N} \tilde{L}_{vw} \right) \sigma(w).$$

However, since the sum of the elements in any column of \tilde{L} is nonnegative and positive elements occur only along the diagonal, if $w \in N$, then $\sum_{v \in N} \tilde{L}_{vw} \geq 0$. So the above calculation implies that $\sum_{v \in N} \tilde{L}_{vw} = 0$ for all $w \in N$. It follows that no vertex in N is connected to a vertex not in N , including the sink vertex. Since s is globally accessible in G , this is a contradiction. So we must have $N = \emptyset$. \square

Remark 7.15. To appreciate the significance of the previous lemma in a larger context, see Section 12.3, which describes a class of M -matrices possessing nice properties for chip-firing, thereby generalizing the class of reduced Laplacians.

Proof of Theorem 7.12. [(1) \Rightarrow (2)] Suppose that c is recurrent, and for the sake of contradiction, suppose that $c_{\max} - c$ is not superstable. Then there exists $\sigma \succeq 0$ such that $c_{\max} - c - \tilde{L}\sigma \geq 0$. By the definition of recurrent, there exists a sandpile $m \geq c_{\max}$ such that $m \rightsquigarrow (m)^\circ = c$. Letting τ be the corresponding firing script, we have

$$0 \leq c_{\max} - c - \tilde{L}\sigma = c_{\max} - (m - \tilde{L}\tau) - \tilde{L}\sigma = c_{\max} - m + \tilde{L}(\tau - \sigma).$$

It follows that $\tilde{L}(\tau - \sigma) \geq m - c_{\max} \geq 0$, and hence, by Lemma 7.14, $\tau - \sigma \geq 0$. Moreover, we see that $c_{\max} \geq m - \tilde{L}(\tau - \sigma)$, so $m - \tilde{L}(\tau - \sigma)$ is stable. By the least action principle (Theorem 6.7), we have $\tau - \sigma \geq \tau$, yielding $\sigma \leq 0$, a contradiction.

[(2) \Rightarrow (3)] If $c_{\max} - c$ is superstable and $\sigma \succeq 0$, then firing σ must produce a negative vertex: $(c_{\max} - c - \tilde{L}\sigma)(v) < 0$ for some $v \in \tilde{V}$. It follows that $c_{\max}(v) < (c + \tilde{L}\sigma)(v)$; so v is unstable in $c + \tilde{L}\sigma$.

[(3) \Rightarrow (4)] Obvious.

[(4) \Rightarrow (1)] Suppose c is a nonrecurrent sandpile. Let \mathbf{b} be the burning sandpile for G with burning script $\sigma_{\mathbf{b}}$. Let v_1, \dots, v_k be a legal firing sequence that stabilizes $c + \mathbf{b}$:

$$c + \mathbf{b} \xrightarrow{v_1, \dots, v_k} (c + \mathbf{b})^\circ.$$

Then $\sigma := \sum_{i=1}^k v_i \neq \sigma_{\mathbf{b}}$ since c is not recurrent. Define $\tau := \sigma_{\mathbf{b}} - \sigma$, and note that $0 \preceq \tau \leq \sigma_{\mathbf{b}}$ by Theorem 7.5 (5). However,

$$c + \tilde{L}\tau = c + \tilde{L}\sigma_{\mathbf{b}} - \tilde{L}\sigma = (c + \mathbf{b})^\circ,$$

which has no unstable vertices. \square




Figure 2. A forbidden subconfiguration.

7.4. Forbidden subconfigurations

It is sometimes possible to determine that a configuration is not recurrent by looking locally. For instance, Figure 2 displays a portion of a configuration c on an undirected graph. We will see shortly that simply by looking at this portion of c , one can tell it is not recurrent.

Let $G = (V, E, s)$ be a sandpile graph with sink s .

Definition 7.16. Let c be a configuration on G and let W be a nonempty subset of \tilde{V} . The *subconfiguration* of c corresponding to W is

$$c|_W := \sum_{w \in W} c(w) w.$$

The subconfiguration $c|_W$ is a *forbidden subconfiguration (FSC)* if

$$c(v) < \text{indeg}_W(v) := |\{w \in W : (w, v) \in E\}|$$

for all $v \in W$.

Proposition 7.17. *If c is recurrent, then it has no FSC.*

Proof. Suppose that c is recurrent, and let W be a nonempty subset of \tilde{V} . Let \mathbf{b} be the burning sandpile for G , and pick a firing sequence $F = u_1, u_2, \dots$ for the stabilization $\mathbf{b} + c \rightsquigarrow c$. By Theorem 7.5 (4), each non-sink vertex appears in F , possibly multiple times. For each $u \in \tilde{V}$, let $\ell(u)$ be the last time u appears in F , i.e., $\ell(u)$ is the largest index i such that $u_i = u$. Let v be the element of W with smallest ℓ -value, i.e., $\ell(v) < \ell(w)$ for all $w \in W \setminus \{v\}$. In F , after v fires for the last time, $c + \mathbf{b}$ will have partially stabilized. The amount of sand on v at that point is nonnegative and will get no smaller as the firing sequence proceeds. However, all the other vertices of W will fire, adding $\text{indeg}_W(v)$ grains of sand to v . Thus, after F has fired, $c + \mathbf{b}$ has stabilized to c , and $c(v) \geq \text{indeg}_W(v)$. So $c|_W$ is not an FSC. \square

Example 7.18. The set $W = \{u, v, w\}$ is an FSC for the configuration pictured in Figure 2. Hence, the configuration cannot be recurrent.

Example 7.19. There is no converse to Proposition 7.17 for directed graphs, in general. For instance, consider the graph in Figure 3 with sink s and edge (u, v) of multiplicity 2. The configuration $c = 1 \cdot u + 0 \cdot v$ is stable, is not recurrent, and has no FSC.




Figure 3. Path graph with an extra edge.

Theorem 7.20. *Let $G = (V, E, s)$ be a sandpile graph with no selfish vertices. Then a stable sandpile on G is recurrent if and only if it has no forbidden subconfigurations.*

Proof. Let c be a stable sandpile on G . We have just seen that if c is recurrent, then it has no FSC. For the converse, suppose that c has no FSC, and let \mathbf{b} be the burning sandpile for G . By Theorem 7.6 (2), the burning script for \mathbf{b} is $\sigma_{\mathbf{b}} = \chi_{\tilde{V}}$, and thus for each non-sink vertex v ,

$$(7.1) \quad \mathbf{b}(v) = \text{outdeg}(v) - \text{indeg}_{\tilde{V}}(v).$$

To show that c is recurrent, we will show that every non-sink vertex fires as $\mathbf{b} + c$ stabilizes; we already know that each vertex fires at most once (Theorem 7.5).

Set $V_1 := \tilde{V}$. Since V_1 is not an FSC, the following set is nonempty:

$$B_1 = \{v \in V_1 : c(v) \geq \text{indeg}_{V_1}(v)\}.$$

By (7.1), B_1 contains exactly the unstable vertices of $\mathbf{b} + c$. Continue, recursively, as follows: as long as $V_i \neq B_i$, define

$$\begin{aligned} V_{i+1} &= V_i \setminus B_i = \tilde{V} \setminus (B_1 \cup \cdots \cup B_i), \\ B_{i+1} &= \{v \in V_{i+1} : c(v) \geq \text{indeg}_{V_{i+1}}(v)\}. \end{aligned}$$

Now, B_i is empty if and only if V_i is an FSC. So given our assumptions, no B_i is empty. The recursion is finite, producing a partition B_1, \dots, B_k of \tilde{V} . For $i \geq 2$, each B_i is the set of non-sink vertices that become unstable after firing $B_1 \cup \cdots \cup B_{i-1}$ from $\mathbf{b} + c$. In detail, since $\mathbf{b} = \tilde{L}\chi_{\tilde{V}}$, the configuration obtained from $\mathbf{b} + c$ after firing the these vertices is

$$\mathbf{b} + c - \tilde{L}\chi_{B_1 \cup \cdots \cup B_{i-1}} = \tilde{L}\chi_{\tilde{V}} + c - \tilde{L}\chi_{B_1 \cup \cdots \cup B_{i-1}} = \tilde{L}\chi_{V_i} + c.$$

The amount of sand on $v \in V_i$ in this configuration is $\text{outdeg}(v) - \text{indeg}_{V_i}(v) + c(v)$, which is unstable exactly when $c(v) \geq \text{indeg}_{V_i}(c)$, i.e., when $v \in B_i$. Proceeding, we see that ordering the elements of each B_i arbitrarily, the concatenation $B_1 \cdots B_k$ is a legal firing sequence for $\mathbf{b} + c$ consisting of all the non-sink vertices. \square

7.5. Dhar's burning algorithm for recurrent.

The word ‘‘burning’’ used in this chapter comes from the original version of a test for recurrence due to Dhar. In his words ([36]):

Given a configuration [on an undirected graph], at first all the sites are considered unburnt. Then, burn each site whose height is larger than the number of its unburnt neighbors. This process is repeated recursively, until no further sites can be burnt. Then, if all the sites have been burnt, the original configuration was recurrent, whereas if some unburnt sites

were left, then the original configuration was transient, and the remaining sites form an FSC.

... We also see that if there are no greedy sites, the burning test is a necessary and sufficient test for recurrence even for unsymmetrical [i.e., directed] graphs.

A minor technical point concerning the above quote: Dhar is assuming a firing rule by which a vertex becomes unstable when the amount of sand on it is strictly greater than its outdegree. So we would replace his “larger than the number of its unburnt neighbors” by “at least as large as the number of its unburnt neighbors”.

The reader should consider how Dhar’s algorithm is expressed in the proof of Theorem 7.20. Burning sites based on a comparison of heights and unburnt neighbors is another way of thinking about vertices firing in the stabilization of a configuration after adding the burning sandpile. Here is an implementation of Dhar’s algorithm—generalized to work for all sandpile graphs—in light of Theorem 7.5 and the proof of Theorem 7.20:

Algorithm 6 Dhar’s burning algorithm for recurrences.

- 1: INPUT: A sandpile c on a sandpile graph G .
 - 2: – Compute the burning sandpile, \mathbf{b} (Remark 7.7).
 - 3: – Stabilize $b + c$, and set W equal to the set of unfired non-sink vertices. (The firing script for the stabilization is bounded above by the burning script, $\sigma_{\mathbf{b}}$.)
 - 4: OUTPUT: Return W . If $W = \emptyset$, then c is recurrent. Otherwise, c is not recurrent, and further, if G has no selfish vertices, then $c|_W$ is an FSC.
-

Remark 7.21.

- We sometimes take *Dhar’s burning algorithm* to mean Theorem 7.5 (2) or (3), which are also tests for recurrence but with no reference to FSCs.
- Problem 7.6 asks for an implementation of the burning algorithm that more closely follows the description in the quotation (cf. Algorithm 7 in Chapter 9).

Exercise 7.22. For an undirected graph, explain how Dhar’s algorithm for recurrences is the “dual” of the version given in Chapter 3 for superstables (which is then applied to computing q -reduced divisors).

Example 7.23. There are sandpile graphs with selfish vertices yet, nonetheless, a stable sandpile on the graph is recurrent if and only if it has no FSC. For instance, let G be any directed acyclic graph with a selfish vertex. Each stable configuration on G is recurrent and hence has no FSC.

Exercise 7.24. Show that for the graph in Figure 3, a stable sandpile is recurrent if and only if it has no FSC, even though v is selfish.


Notes

The burning script algorithm (Remark 7.7) for directed graphs comes from the work of Speer ([84]). For the original version of Dhar’s algorithm, see [34], [36], and [37].

Problems for Chapter 7

7.1. Let $G = (V, E, s)$ be a sandpile graph with no selfish vertices, so $\text{indeg}_{\tilde{V}}(v) \leq \text{outdeg}(v)$ for all $v \in \tilde{V}$. Show that it is possible to form a new graph G' by starting with G and adding or removing edges emanating from the sink to achieve $\text{indeg}_V(v) = \text{outdeg}(v)$ for all $v \in \tilde{V}$ where the indegree is with respect to all of the vertices now. Hence, G' has the same vertex set as G and the same edges except possibly for edges of the form (s, v) . Explain why G' is actually Eulerian and the mapping $c \mapsto c$ gives an isomorphism of sandpile groups, $\mathcal{S}(G) \approx \mathcal{S}(G')$.

7.2. Consider the following sandpile graph with sink s :



- Find the (minimal) burning sandpile and its script.
- Apply Dhar's algorithm to the sandpile $c = (1, 1, 1, 1, 1, 1, 1)$ to find a forbidden subconfiguration (and thus demonstrate that c is not recurrent).

7.3. Let $G = (V, E, s)$ be an undirected, simple sandpile graph. (Simple means each edge has weight 1.) Show that $c_{\max} - 1_{\tilde{V}}$ is not recurrent.

7.4. Let e be an edge of an undirected sandpile graph G . For each of the following graphs G' (i) describe (with proof) the relationship between superstable on G and on G' , and (ii) describe (with proof) the corresponding isomorphism of sandpile groups, $\mathcal{S}(G) \approx \mathcal{S}(G')$.

- Suppose e is a *bridge*, i.e., deleting e disconnects G . Let $G' := G/e$, the graph obtained by contracting e , i.e., by deleting e and identifying its vertices.
- Suppose e is a loop, and let $G' := G \setminus e$, the graph obtained from G by deleting e .

7.5. Let G be the $m \times n$ sandpile grid graph, described in Section 6.6, and let \vec{k} denote the sandpile on G having k grains of sand on each vertex. Use the theory of forbidden subconfigurations to show the following:

- $\vec{1}$ is not recurrent.
- $\vec{2}$ is recurrent.

7.6. Give pseudocode for an implementation of Dhar's burning algorithm that closely follows the statement of the algorithm in the quotation at the beginning of Section 7.5. (Algorithm 7 in Chapter 9 is a variant that produces a bijection between recurrents and spanning trees.)

Threshold density

In Section 6.1, we discussed an experiment where grains of sand are dropped randomly onto the vertices of a sandpile graph, with time allowed for stabilization between each additional grain. We saw that only certain stable sandpiles appear more than once, which led us to the definition of recurrent sandpiles and the subsequent theory of the sandpile group.

One could also consider the same experiment conducted on a graph with no sink vertex. Drop a grain of sand on a randomly chosen vertex, stabilize, and repeat. This time, since there is no sink, enough sand will eventually accumulate so that the system cannot be stabilized: no matter how many times we fire, there will always be an unstable vertex. We will call the first unstabilizable state reached in this manner the *threshold state*. Since this is a probabilistic system, the threshold state may change on each run of the experiment.

Example 8.1. Let $G = C_3$ be a triangle, and start at the state $(0, 0, 0)$. Suppose we happen to randomly pick the first vertex three times in a row. The system then evolves as follows:

$$(0, 0, 0) \xrightarrow{(v_1)} (1, 0, 0) \xrightarrow{(v_1)} (2, 0, 0)^\circ = (0, 1, 1) \xrightarrow{(v_1)} (1, 1, 1).$$

(We use the notation (v_1) instead of v_1 to distinguish adding a grain at v_1 and attempting to stabilize from simply firing v_1 .) From $(1, 1, 1)$, no matter where a grain of sand is dropped, the system will be unstabilizable. For instance, if the second vertex is chosen, we arrive at the threshold state, $(1, 2, 1)$.

Exercise 8.2. Let $G = C_3$, as above, and start at the state $(0, 0, 0)$. If vertices are chosen uniformly at random, how many grains of sand do you expect to drop before the threshold state is reached? (Draw out all possible routes to a threshold state. What is the average number of steps in these routes?)

One can further generalize this sinkless version of the experiment by trying to take into account the starting state of the system. For instance, how many steps will it take to reach threshold in the previous example if we start at $(1, 0, 0)$ or if

we start at $(-10, 0, 0)$? For the latter, we need to imagine “negative” amounts of sand, or perhaps switch our metaphor back to dollars, as in the first part of this book.

In order to study the dynamics behind these types of experiments, we formalize them as examples of Markov chains. Our ultimate goal is Theorem 8.48, the threshold density theorem, due to Levine ([66]), characterizing the expected number of grains of sand per vertex at the threshold state in the limit as the amount of sand in the starting state goes to $-\infty$. Along the way, we will also see in Theorem 8.29 that the inverse of the reduced Laplacian of the graph encodes the expected number of firings in the stabilization of a randomly chosen recurrent configuration caused by dropping a grain of sand at a specific vertex.

8.1. Markov Chains

Definition 8.3. A *finite Markov chain* consists of the following data:

- (1) A finite set of *states*, Ω .
- (2) A function $P: \Omega \times \Omega \rightarrow [0, 1]$ with the property that for all $x \in \Omega$, the function $P(x, \cdot)$ is a probability distribution on Ω :

$$\sum_{y \in \Omega} P(x, y) = 1.$$

The function P is called the *transition matrix* of the Markov chain.

- (3) A sequence of random variables (X_0, X_1, X_2, \dots) satisfying the *law of the chain*: for all $t \geq 0$,

$$\mathbb{P}(X_{t+1} = y \mid X_t = x) = P(x, y).$$

We interpret a Markov chain as an evolving system, with the random variable X_t giving the state of the system at time t . The system evolves in discrete time steps via an update rule defined by the transition matrix: if the system is in state x at time t , then it will move to state y at the next time $t + 1$ with probability $P(x, y)$. Note that the transition probabilities only depend on the current state, not the previous history of the system or the time—this is the *Markov property*.

It is often helpful to visualize a finite Markov chain as a directed graph with the states as vertices and a directed edge from x to y if and only if $P(x, y) > 0$. Then the evolution of the Markov chain may be thought of as a random walk on the associated directed graph, where at each time step the system moves from its current state to a neighboring state by choosing an outgoing edge according to the probabilities given by P .

Example 8.4. Suppose that $\Omega = \{A, B, C, D, E, F\}$ and P is given by the following matrix, with rows and columns listed in alphabetical order:

$$P = \begin{matrix} & A & B & C & D & E & F \\ \begin{matrix} A \\ B \\ C \\ D \\ E \\ F \end{matrix} & \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0.7 & 0.3 & 0 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix} \end{matrix}.$$

The associated directed graph is pictured in Figure 1. When the chain is at B , it will move to C with probability 0.7 and to D with probability 0.3. When the chain is at C , it will move to either D or E with equal probability.




Figure 1. Directed graph associated to a Markov chain on the 6 states $\{A, B, C, D, E, F\}$.

The next example shows that the experiment from Section 6.1 is a finite Markov chain.

Definition 8.5 (Abelian sandpile model). Let $G = (V, E, s)$ be a sandpile graph, and choose any probability distribution $\alpha: V \rightarrow [0, 1]$ on the vertices of G . Assume that $\alpha(v) > 0$ for all $v \in V$, so there is a positive probability of dropping a grain of sand on every vertex (including the sink). The *abelian sandpile* is the Markov chain with state space $\Omega = \text{Stab}(G)$, the finite set of stable sandpiles, and transition matrix

$$P(c, c') = \sum_{v \in V : (c+v)^\circ = c'} \alpha(v).$$

If there is no vertex v such that $(c + v)^\circ = c'$, then $P(c, c') = 0$. Moreover, we define $(c + s)^\circ = c$ for all stable c . Starting with a stable sandpile c_0 , define a sequence of random variables by $c_{t+1} = (c_t + v_t)^\circ$ where the vertices v_0, v_1, v_2, \dots are independent random draws from α . The sequence (c_0, c_1, c_2, \dots) clearly satisfies the law of the chain. **Notation:** We write $c_{t-1} \xrightarrow{(v_t)} c_t$, using parentheses, to distinguish the evolution of the chain from our earlier notation for firing a vertex, $c \xrightarrow{v} c - \tilde{L}v$.

Example 8.6. For example, consider the abelian sandpile model on the triangle C_3 for which the probabilities of picking the first, the second, and the sink vertices are α_1, α_2 , and α_3 , respectively. The states are $\text{Stab}(C_3) = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$, only 3 of which are recurrent. The associated directed graph for the Markov chain appears in Figure 2.




Figure 2. The abelian sandpile model for the triangle graph.

Suppose that $(\Omega, P, (X_t))$ is a finite Markov chain. We are interested in the long term behavior, and a basic question is the following: suppose that $\pi_0: \Omega \rightarrow [0, 1]$ is a probability distribution on the set of states, describing probabilities for various initial conditions of the system. What is the probability distribution π_t for the state of the system at some later time t ? Note that for each $y \in \Omega$, we have

$$\begin{aligned} \pi_1(y) &= \mathbb{P}(X_1 = y) = \sum_{x \in \Omega} \mathbb{P}(X_1 = y \mid X_0 = x) \mathbb{P}(X_0 = x) \\ &= \sum_{x \in \Omega} P(x, y) \pi_0(x) \\ &= (\pi_0 P)(y), \end{aligned}$$

where we view π_1 and π_0 as row vectors and P as a matrix. By induction, it follows that for all $t \geq 1$:

$$\begin{aligned} \pi_t(y) &= \mathbb{P}(X_t = y) = \sum_{x \in \Omega} \mathbb{P}(X_t = y \mid X_{t-1} = x) \mathbb{P}(X_{t-1} = x) \\ &= \sum_{x \in \Omega} P(x, y) \pi_{t-1}(x) \\ &= (\pi_{t-1} P)(y) \\ &= (\pi_0 P^{t-1} P)(y) \\ &= (\pi_0 P^t)(y). \end{aligned}$$

Thus, the probability distribution π_t is obtained from the initial probability distribution π_0 via right-multiplication by P^t , the t -th power of the transition matrix. Note that the probability of the chain moving from state x to state y in exactly n steps is given by the matrix element $P^n(x, y)$.

Definition 8.7. Suppose that x and y are states of a Markov chain $(\Omega, P, (X_t))$. We say that y is *accessible* from x if there is a positive probability of moving from x to y in a finite number of steps: there exists $n \geq 0$ such that $P^n(x, y) > 0$. If in addition, x is accessible from y , then we say that x and y *communicate*. The Markov chain is *irreducible* if every state communicates with every other state.

Exercise 8.8. Show that communication is an equivalence relation on the set of states in a Markov chain.

In the Markov chain of Figure 1, the communicating classes are $\{A\}$, $\{B\}$, $\{C\}$, $\{D, E, F\}$. Note that D, E, F are accessible from all states, C is accessible from A and B , while A and B are each accessible only from themselves. Only states D, E , and F are essential in the sense of the next definition.

Definition 8.9. A state x of a Markov chain is *essential* if the communicating class of x consists of all states y that are accessible from x . States that are not essential are called *inessential*.

Exercise 8.10. Show that being essential is a property of communicating classes in a Markov chain: if x and y communicate, then x is essential if and only if y is essential.

Exercise 8.11. In the context of the abelian sandpile model, show that there is a unique essential communicating class, represented by c_{\max} .

In the experiment of Section 6.1, we saw that some sandpiles—the recurrent—appeared many times, while others appeared at most once. The next definition formalizes this distinction for a general Markov chain.

Definition 8.12. Suppose that $(\Omega, P, (X_t))$ is a Markov chain. A state $x \in \Omega$ is *recurrent* if, starting from x , the chain will return to x with probability 1. A non-recurrent state is called *transient*.

Exercise 8.13. Show that recurrence is a property of communicating classes in a Markov chain: if x and y communicate, then x is recurrent if and only if y is recurrent.

Note that if x is a recurrent state of a Markov chain, then starting from x , the chain will return to x infinitely many times: each time the chain is in state x , it will return to x with probability 1. On the other hand, if x is transient, then starting from x , the chain will return to x only finitely many times: each time the system is in state x , there is a positive probability that it will never return to x . In particular, every finite Markov chain has at least one recurrent element—otherwise each of the states would be visited only finitely many times, leaving nowhere for the chain to go! Moreover, this implies that no matter where a finite Markov chain starts, it will reach a recurrent state after finitely many steps.

Proposition 8.14. *Suppose that $(\Omega, P, (X_t))$ is a finite Markov chain. Then a state $x \in \Omega$ is essential if and only if it is recurrent. In particular, every finite Markov chain has at least one essential state. Moreover, if z is an inessential state, then there exists an essential state x that is accessible from z .*

Proof. First suppose that x is recurrent. Then if y is accessible from x , it must be that x is accessible from y , else the chain could never return to x once it reached y . Thus, x is essential. Note that this direction does not require that the set of states be finite.

Now suppose that x is essential. Then the restriction of P to the communicating class $[x]$ defines an irreducible finite Markov chain on $[x]$, which must have at

least one recurrent element. By irreducibility, every element of $[x]$ is recurrent, including x .

Finally, suppose that z is inessential. Then it is also transient, so that if the chain starts at $X_0 = z$, it will reach a recurrent/essential state x after finitely many steps. Thus, there exists an $n > 0$ such that $P^n(z, x) > 0$, so that x is accessible from z . \square

Example 8.15. Suppose that $G = (V, E, s)$ is a sandpile graph. We wish to show that the recurrent sandpiles (in the sense of Definition 6.20) are exactly the recurrent states of the abelian sandpile model Markov chain. First note that the recurrent sandpiles (in the sense of Chapter 6) form the communicating class $[c_{\max}]$, which is the unique essential class in the chain. By the previous proposition, $[c_{\max}]$ is also the unique recurrent class.

Return now to a general finite Markov chain $(\Omega, P, (X_t))$ and the probability distributions $\pi_t = \pi_0 P$. Do these probability distributions converge to a limiting distribution, π ? We will later (in Theorem 8.27) state conditions that guarantee the existence of a limiting distribution, and we will prove its uniqueness (i.e., independence of the initial distribution π_0) in Proposition 8.22. For now, we simply assume the existence of the limit $\pi = \lim_{t \rightarrow \infty} \pi_t$, and investigate its properties. Taking the limit of the identity $\pi_t = \pi_{t-1} P$ immediately yields the equation $\pi = \pi P$, showing that if the system begins in the probability distribution π , it will stay there forever. For this reason, a distribution π satisfying $\pi = \pi P$ is called a *stationary distribution* for the Markov chain.

Proposition 8.16. *If π is a stationary distribution of the finite Markov chain $(\Omega, P, (X_t))$, then $\pi(x) = 0$ for all inessential states x .*

Proof. Since the chain is finite, it must have at least one essential communicating class \mathcal{C} by Proposition 8.14. The restriction of P to any such \mathcal{C} defines an irreducible Markov chain $(\mathcal{C}, P, (X_t))$. We compute

$$\begin{aligned} \sum_{y \in \mathcal{C}} (\pi P)(y) &= \sum_{y \in \mathcal{C}} \sum_{x \in \Omega} \pi(x) P(x, y) \\ &= \sum_{y \in \mathcal{C}} \left(\sum_{x \in \mathcal{C}} \pi(x) P(x, y) + \sum_{x \notin \mathcal{C}} \pi(x) P(x, y) \right) \\ &= \sum_{x \in \mathcal{C}} \pi(x) \sum_{y \in \mathcal{C}} P(x, y) + \sum_{y \in \mathcal{C}} \sum_{x \notin \mathcal{C}} \pi(x) P(x, y) \\ &= \sum_{x \in \mathcal{C}} \pi(x) + \sum_{y \in \mathcal{C}} \sum_{x \notin \mathcal{C}} \pi(x) P(x, y), \end{aligned}$$

where we use the fact that $\sum_{y \in \mathcal{C}} P(x, y) = 1$ for any $x \in \mathcal{C}$. Since π is stationary, it follows that $\sum_{y \in \mathcal{C}} (\pi P)(y) = \sum_{x \in \mathcal{C}} \pi(x)$, so that $\pi(x) P(x, y) = 0$ for all $y \in \mathcal{C}$ and $x \notin \mathcal{C}$.

Now suppose that x_0 is an inessential state. Then by Proposition 8.14, there exists an essential state y that is accessible from x_0 . Therefore, there is a finite sequence $x_0, x_1, \dots, x_{n-1}, x_n = y$, with $P(x_{k-1}, x_k) > 0$ for all k and x_k inessential for $0 \leq k < n$. Setting $\mathcal{C} = [y]$ in the previous paragraph, we see that $\pi(x_{n-1})P(x_{n-1}, y) = 0$, which implies that $\pi(x_{n-1}) = 0$. Moreover, note that if $\pi(x_k) = 0$ for some k , then by stationarity we have

$$0 = \pi(x_k) = (\pi P)(x_k) = \sum_{z \in \Omega} \pi(z)P(z, x_k),$$

so that $\pi(z)P(z, x_k) = 0$ for all states z . Since $P(x_{k-1}, x_k) > 0$ by assumption, this shows that $\pi(x_{k-1}) = 0$. By backwards induction along the sequence starting with x_{n-1} , we see that $\pi(x_0) = 0$ as required. \square

In light of this proposition, we may restrict attention to the essential states of a finite Markov chain for the purposes of studying stationary distributions.

For the abelian sandpile model, we have seen that there is a unique essential communicating class $[c_{\max}]$ consisting of the recurrent sandpiles. Moreover, from Chapter 6, we know that $[c_{\max}] = \mathcal{S}(G)$ is an abelian group under the operation of stable addition. Hence, the restriction of the abelian sandpile model to the essential states yields an irreducible Markov chain on the sandpile group. More precisely, this is an example of a *random walk on a group*.

Definition 8.17. Let H be a finite group and $\gamma: H \rightarrow [0, 1]$ a probability distribution on H . Then the *random walk on H with respect to γ* is the Markov chain on $\Omega = H$ defined by the transition matrix

$$P(h, h') := \gamma(h'h^{-1}).$$

Note that this does define a transition matrix, since for all $h \in H$ we have

$$\sum_{h' \in H} P(h, h') = \sum_{h' \in H} \gamma(h'h^{-1}) = \sum_{h' \in H} \gamma(h') = 1.$$

Starting at a group element $X_0 = h_0 \in H$, define a sequence of random variables by $X_{t+1} = h_t X_t$, where the group elements h_0, h_1, h_2, \dots are independent random draws from γ . In words: at each time t pick a random element $h_t \in H$ according to the distribution γ , and then update the system by multiplying the current state on the left by h_t .

Exercise 8.18. When is a random walk on a group irreducible?

Example 8.19. We now show that the abelian sandpile model (restricted to the essential states) is a random walk on the sandpile group. For every non-sink vertex $v \in \tilde{V}$, denote by $r_v \in \mathcal{S}(G)$ the unique recurrent sandpile that is equivalent to v modulo the reduced Laplacian lattice $\tilde{\mathcal{L}}$. For the sink vertex s , define $r_s := e$, the identity of $\mathcal{S}(G)$. Then transfer the probability distribution α on V to a probability distribution on $\mathcal{S}(G)$ by defining

$$\gamma(c) := \sum_{v \in V : r_v = c} \alpha(v).$$


The transition matrix for the random walk on $\mathcal{S}(G)$ with respect to γ is given by $P(c, c') = \gamma(c' \otimes c^{-1})$ for all $c, c' \in \mathcal{S}(G)$.

Now recall the update rule for the abelian sandpile model: at each time t , choose a random vertex $v_t \in V$ according to α and update the system by dropping a grain of sand on v_t and stabilizing: $c_{t+1} = (c_t + v_t)^\circ$. But in terms of the sandpile group we have

$$(c_t + v_t)^\circ = c_t \otimes r_{v_t} = r_{v_t} \otimes c_t,$$

which is the update rule for the random walk on $\mathcal{S}(G)$ with respect to the probability distribution γ .

Exercise 8.20. Consider the graph G with three undirected edges and two directed edges pictured below:



Take s as the sink vertex, and order the remaining vertices v_1, v_2, v_3 .

(1) Show that

$$\mathcal{S}(G) = \{(1, 0, 1), (1, 0, 0), (1, 1, 1), (0, 0, 1), (1, 1, 0), (0, 1, 1)\}$$

with $(1, 0, 1)$ the identity element.

(2) Show that $r_{v_2} = (1, 1, 1)$, and hence, using the notation from Example 8.19, the probability distribution γ is given by the following table:

recurrent	$(1, 0, 1)$	$(1, 0, 0)$	$(1, 1, 1)$	$(0, 0, 1)$	$(1, 1, 0)$	$(0, 1, 1)$
γ	$\alpha(s)$	$\alpha(v_1)$	$\alpha(v_2)$	$\alpha(v_3)$	0	0

(3) Compute the transition matrix for the abelian sandpile model on G restricted to the recurrents, and draw the corresponding directed graph.

The next proposition shows that the uniform distribution is a stationary distribution for any random walk on a group.

Proposition 8.21. *Let H be a finite group, and choose any probability distribution $\gamma: H \rightarrow [0, 1]$. Denote by $u: H \rightarrow [0, 1]$ the uniform distribution on H , defined as $u(h) = \frac{1}{|H|}$ for all $h \in H$. Then u is a stationary distribution for the random walk on H with respect to γ .*

Proof. We just need to check that $uP = u$:

$$(uP)(h) = \sum_{k \in H} u(k^{-1}h)P(k^{-1}h, h) = \frac{1}{|H|} \sum_{k \in H} \gamma(k) = \frac{1}{|H|} = u(h).$$

□

Returning to the abelian sandpile model, we have established the existence of a stationary distribution: the uniform distribution on recurrent sandpiles. But the question remains: is this the unique stationary distribution, and if so, will the chain always converge to the uniform distribution, independent of the initial

distribution π_0 ? As the next proposition shows, the irreducibility of a Markov chain guarantees the uniqueness of its stationary state.

Proposition 8.22. *Suppose that $(\Omega, P, (X_t))$ is an irreducible finite Markov chain and that π and μ are stationary distributions. Then $\pi = \mu$.*

Proof. The stationary condition $\pi = \pi P$ shows that the row rank of the matrix $P - I$ is at most $|\Omega| - 1$. Since π and μ are both probability distributions, in order to show that $\pi = \mu$, we just need to show that the row rank is exactly $|\Omega| - 1$. Since row rank equals column rank, it will suffice (by rank-nullity) to show that the kernel of $P - I$ is one-dimensional. In fact, we will show that the kernel of $P - I$ consists of the constant vectors. First, note that the all one's column vector is contained in the kernel: for all $x \in \Omega$,

$$(P\vec{1})(x) = \sum_{y \in \Omega} P(x, y) = 1 \implies (P - I)\vec{1} = 0.$$

Now suppose that a function $f: \Omega \rightarrow \mathbb{R}$ is in the kernel of $P - I$, where we view f as a column vector. Choose a state x where f attains a maximum, and let $y \in \Omega$ be an arbitrary state. We wish to show that $f(y) = f(x)$. By irreducibility, there exists a sequence $x = x_0, x_1, \dots, x_{n-1}, x_n = y$ with $P(x_{k-1}, x_k) > 0$ for all k . Note that

$$(Pf)(x) = \sum_{z \in \Omega} P(x, z)f(z) \leq f(x) \sum_{z \in \Omega} P(x, z) = f(x),$$

with equality if and only if $f(z) = f(x)$ for all $z \in \Omega$ such that $P(x, z) > 0$. Since $(Pf)(x) = f(x)$ by assumption, the equality obtains, so that we have $f(x_1) = f(x)$. But then the same argument applied to the computation of $(Pf)(x_1)$ shows that $f(x_2) = f(x_1) = f(x)$. By induction, we see that $f(y) = f(x)$ as claimed. \square

Exercise 8.23. Show that if π and μ are stationary distributions for a finite Markov chain with a unique essential communicating class, then $\pi = \mu$.

Definition 8.24. Let $(\Omega, P, (X_t))$ be a finite Markov chain. For any state x , define $\mathcal{T}(x) := \{n \geq 1 : P^n(x, x) > 0\}$. This is the set of positive times when it is possible for the chain to return to x , starting at $X_0 = x$. The *period* of the state x is defined to be the greatest common divisor of the set $\mathcal{T}(x)$, provided $\mathcal{T}(x)$ is nonempty, otherwise, the period is not defined.

In the Markov chain of Figure 1, we have $\mathcal{T}(A) = \mathcal{T}(B) = \mathcal{T}(C) = \emptyset$, while $\mathcal{T}(D) = \mathcal{T}(E) = \mathcal{T}(F) = \{3k : k \geq 1\}$. Hence, states D, E , and F each have period 3, while A, B, C have no period.

Exercise 8.25. Show that if x and y communicate, then they have the same period. Hence, we may speak of the period of a communicating class.

Definition 8.26. Suppose that $(\Omega, P, (X_t))$ is an irreducible Markov chain. We say that it is *aperiodic* if every state $x \in \Omega$ has period 1. By the previous exercise, this is equivalent to the existence of a state with period 1.

Theorem 8.27. *Suppose that $(\Omega, P, (X_t))$ is an irreducible finite Markov chain. Then there exists a unique stationary distribution π satisfying $\pi = \pi P$. If the chain is also aperiodic, then given any initial distribution π_0 on Ω , the chain converges to the stationary distribution: $\pi = \lim_{t \rightarrow \infty} \pi_0 P^t$.*

Proof. See Proposition 1.14 and Theorem 4.9 in [65]. □

Recall that for the abelian sandpile model, we assumed that the probability distribution α on V satisfied $\alpha(v) > 0$ for all $v \in V$. In particular, $\alpha(s) > 0$, so there is a positive probability of dropping a grain of sand on the sink at each time step. But $(c+s)^\circ = c$ for all stable c , since we ignore sand on the sink. It follows that each stable sandpile has period 1, so that the abelian sandpile model is aperiodic. Hence, if π_0 is an arbitrary initial distribution on $\mathcal{S}(G)$, we see that $\lim_{t \rightarrow \infty} \pi_0 P^t = u$, the uniform distribution on the sandpile group. Moreover, since c_{\max} is accessible from every stable sandpile, if we start the chain at any $c_0 \in \text{Stab}(G)$, it will reach the set of recurrents in finite time. From that point on, the chain will never leave the set of recurrents, and it will converge to the uniform distribution on $\mathcal{S}(G)$. Thus, in the long-run, the abelian sandpile model spends equal amounts of time in each recurrent sandpile, independent of the initial sandpile or the probability distribution α on V . To formally summarize this discussion:

Corollary 8.28. *The abelian sandpile model is a finite aperiodic Markov chain. If restricted to the sandpile group, it is also irreducible. In either case, the stationary distribution is the uniform distribution on the sandpile group.*

Knowing the long-run, stationary distribution of the abelian sandpile model allows us to answer questions such as the following. For any recurrent sandpile c and any pair of vertices $v, w \in V$, let $n(v, w; c)$ denote the number of times that vertex w topples in the stabilization of $c+v$. For fixed v and w , what is the expected value of $n(v, w; \cdot)$ when the system is in the stationary (uniform) distribution?

Theorem 8.29. *Let $G = (V, E, s)$ be a sandpile graph, and let u denote the uniform distribution on the sandpile group, $\mathcal{S}(G)$, which is the stationary distribution for the abelian sandpile model on G . Then in the stationary distribution, the expected number of topplings at $w \in \tilde{V}$ induced by the addition of a single grain of sand at $v \in \tilde{V}$ is given by the wv matrix element of the inverse of the reduced Laplacian:*

$$\frac{1}{|\mathcal{S}(G)|} \sum_{c \in \mathcal{S}(G)} n(v, w; c) = (\tilde{L}^{-1})_{wv}.$$

Proof. Define a $|\tilde{V}| \times |\tilde{V}|$ matrix N by setting the wv entry equal to the average toppling number:

$$N_{vw} := \frac{1}{|\mathcal{S}(G)|} \sum_{c \in \mathcal{S}(G)} n(v, w; c).$$

We wish to show that $N\tilde{L}^t = I$, the identity matrix. For any recurrent sandpile c , we have

$$(c+v)^\circ(w) = c(w) + v(w) - \sum_{z \in \tilde{V}} n(v, z; c) \tilde{L}_{wz},$$

where v is thought of as a configuration, (hence, $v(w)$ is 1 if $w = v$ and 0, otherwise). Now average over the recurrent sandpiles:

$$\begin{aligned} \frac{1}{|\mathcal{S}(G)|} \sum_{c \in \mathcal{S}(G)} (c + v)^\circ(w) &= \frac{1}{|\mathcal{S}(G)|} \left(\sum_{c \in \mathcal{S}(G)} c(w) \right) + v(w) \\ &\quad - \frac{1}{|\mathcal{S}(G)|} \sum_{z \in \tilde{V}} \sum_{c \in \mathcal{S}(G)} n(v, z; c) \tilde{L}_{wz}. \end{aligned}$$

Note that (again denoting the recurrent equivalent to v by $r_v \in \mathcal{S}(G)$),

$$\sum_{c \in \mathcal{S}(G)} (c + v)^\circ(w) = \sum_{c \in \mathcal{S}(G)} (c \circledast r_v)(w) = \sum_{c \in \mathcal{S}(G)} c(w).$$

Hence, the first sums on each side of the average cancel, and we are left with

$$\begin{aligned} v(w) &= \sum_{z \in \tilde{V}} \frac{1}{|\mathcal{S}(G)|} \sum_{c \in \mathcal{S}(G)} n(v, z; c) \tilde{L}_{wz} \\ &= (N\tilde{L}^t)_{vw} \end{aligned}$$

This shows that the matrix elements of the inverse transpose of the reduced Laplacian are given by the average toppling numbers, as claimed. \square

8.2. The fixed-energy sandpile

We now consider the abelian sandpile but without a sink vertex. At each step, a grain of sand is dropped on a random vertex. If the vertex is still stable, the chain continues as usual; otherwise, we attempt to stabilize by repeatedly firing unstable vertices. We can imagine that at first, when there is just a small amount of sand on the graph, short sequences of topplings suffice to stabilize the system. However, since there is no sink, as time goes on, sand accumulates and the system tends to take longer to settle down. Eventually it reaches a threshold at which stabilization is no longer possible—the system is “alive”.

To formalize these ideas, recall from Part 1 that a *divisor* on the graph $G = (V, E)$ is just an element of the free abelian group $\text{Div}(G) := \mathbb{Z}V$. For consistency, in this chapter we will interpret a divisor as assigning an amount of sand, rather than dollars, to each vertex. A negative amount of sand may be interpreted as a ‘hole’ that must be filled before a positive amount of sand can accumulate.

If $D \in \text{Div}(G)$, then we have the usual definitions of stable and unstable vertices for D and the toppling or firing of vertices (Definition 5.21). The toppling of an unstable vertex is called a *legal* toppling. Although we allow negative sand in D , a legal toppling does not create or deepen any holes. Recall that D is *linearly equivalent* to D' , written $D \sim D'$, if D' is obtainable from D through a sequence of (not necessarily legal) topplings or reverse topplings. Equivalently, there exists a firing script $\sigma: V \rightarrow \mathbb{Z}$ such that $D' = D - L\sigma$, where L is the Laplacian of G , in which case we write $D \xrightarrow{\sigma} D'$.

To attempt to stabilize D , we sequentially topple unstable vertices:

$$D \xrightarrow{v_1} D' \xrightarrow{v_2} D'' \xrightarrow{v_3} \dots$$

The resulting chain of topplings is called an *avalanche*. An avalanche either ends in a stable divisor, in which case we say D is *stabilizable*, or it goes on forever, in which case, we say D is *unstabilizable* or *alive*. A finite avalanche has a corresponding firing script $\sigma: V \rightarrow \mathbb{Z}$ recording the number of times each vertex has fired.

To address the question of existence and uniqueness of stabilizations, we generalize Proposition 5.27 from Part 1 to the case of directed graphs. The key idea is the least action principle for divisors, whose proof is just as in the earlier least action principle for configurations, Theorem 6.7.

Theorem 8.30 (Least action principle). *Let G be a directed graph, and let $D \in \text{Div}(G)$. Suppose $\sigma, \tau \geq 0$ are firing scripts such that σ arises from a sequence of legal topplings and that $D \xrightarrow{\tau} D'$ with D' stable. Then $\sigma \leq \tau$.*

As an immediate corollary we get the following:

Corollary 8.31. *Let $D \in \text{Div}(G)$. If D is stabilizable, then every avalanche for D is finite and has the same firing script.*

In this chapter, the words “alive” and “unstabilizable” mean there is no sequence of *legal* vertex-topplings that leads to a stable divisor. “Stabilizable” means some—and hence, by the preceding corollary, every—sequence of legal topplings leads to a (linearly equivalent) stable divisor. This usage differs from that in Part 1, in which all graphs are undirected and where “alive” means that every linearly equivalent divisor is unstable. For directed graphs, there is a difference between these two definitions (cf. Problem 8.2). However, for a certain class of graphs, which includes undirected graphs and, more generally, the Eulerian graphs with which we are primarily concerned in this chapter, the following proposition shows there is no difference. We say a directed graph is *strongly connected* if for all vertices u, v , there is a directed path from u to v . Equivalently, each vertex is globally accessible.

Proposition 8.32. *If G is strongly connected, then $D \in \text{Div}(G)$ is stabilizable if and only if there is a stable divisor linearly equivalent to D . Equivalently, D is alive if and only if every linearly equivalent divisor is unstable.*

Proof. If D is stabilizable, then a sequence of (legal) topplings leads to a linearly equivalent stable divisor.

For the converse, suppose D' is a stable divisor linearly equivalent to D . Since G is strongly connected, for each vertex v , there is at least one spanning tree rooted at v . So by Theorem 9.14, the kernel of the Laplacian of G is spanned by an integer vector ρ whose coordinates are all positive. Given any firing script τ such that $D \xrightarrow{\tau} D'$, we may assume $\tau \geq 0$ by adding multiples of ρ , if necessary. The least action principle then implies D is stabilizable. \square

We are now ready to define the *fixed-energy sandpile*, a Markov chain on G . The name comes from the physics literature in which the total number of grains of sand is a measure of the energy of the system, and topplings model energy transfer.

Thus, a fixed-energy sandpile is a Markov chain in which each state has a fixed amount of “energy” which cannot dissipate through topplings.

Choose a probability distribution $\alpha: V \rightarrow [0, 1]$ (as before, we assume that $\alpha(v) > 0$ for all $v \in V$). For the set of states, we take $\Omega := \text{Div}(G)$. The transition function $P: \Omega \times \Omega \rightarrow [0, 1]$ is defined by

$$P(D, D') := \begin{cases} \alpha(v) & \text{if } D' = D + v \text{ and } D' \text{ is alive} \\ \sum_{v \in V : (D+v)^\circ = D'} \alpha(v) & \text{otherwise.} \end{cases}$$

For each $v \in V$, define an operator a_v on $\text{Div}(G)$ by

$$a_v D := \begin{cases} D + v & \text{if } D + v \text{ is alive} \\ (D + v)^\circ & \text{otherwise.} \end{cases}$$

The sequence of random variables (D_t) is defined as follows: starting with an initial divisor D_0 , define

$$D_{t+1} := a_{v_{t+1}} D_t$$

where the vertices v_1, v_2, v_3, \dots are independent random draws from α . In words: the chain evolves from an initial divisor by randomly dropping single grains of sand on V and stabilizing if possible. **Notation:** as with the abelian sandpile, we write

$$D_{t-1} \xrightarrow{(v_t)} D_t,$$

with parentheses around v_t , to denote the evolution of the chain in order to avoid confusion with our earlier notation for firing a vertex: $D \xrightarrow{v} D - Lv$.

Note that although the set of states $\Omega = \text{Div}(G)$ is countably infinite, the transition function P is locally finite in the sense that for every state D , there exist at most $|V|$ other states D' such that $P(D, D') > 0$. Hence, for each D , the sum $\sum_{D' \in \Omega} P(D, D') = 1$ has only finitely many nonzero terms.

8.2.1. Threshold density. Since the amount of sand in the fixed-energy sandpile continually increases, there are no essential or recurrent states, and hence there is no stationary distribution. Instead, we are interested in the *threshold* for the fixed-energy sandpile, defined as the random time

$$\tau := \tau(D_0) = \min\{t \geq 0 : D_t \text{ is alive}\}.$$

The probability that τ assumes a certain value t is

$$\mathbb{P}_{D_0}(\tau = t) = \sum \prod_{i=1}^t \alpha(v_i),$$

where the sum is over all strings $v_1 \cdots v_t$ of vertices such that D_t is the first alive divisor in the corresponding sequence of states,

$$\begin{array}{ccccccc} D_0 & \xrightarrow{(v_1)} & D_1 & \xrightarrow{(v_2)} & D_2 & \xrightarrow{(v_3)} & \cdots & \xrightarrow{(v_t)} & D_t. \\ & & \parallel & & \parallel & & & & \\ & & a_{v_1} D_0 & & a_{v_2} D_1 & & & & \end{array}$$

Starting from an initial state D_0 , which may be highly stable (e.g., $D_0 \ll 0$), our system evolves toward an unstabilizable *threshold state* D_τ , obtained from a last stable state $D_{\tau-1}$ by randomly adding a grain of sand. The probability D_τ is a given divisor D is $\mathbb{P}_{D_0}(D_\tau = D) = \sum \prod_{i=1}^t \alpha(v_i)$, where the sum is over all strings $v_1 \dots v_t$ of all lengths $t \geq 1$ such that $D_0 \xrightarrow{(v_1)} D_1 \xrightarrow{(v_2)} \dots \xrightarrow{(v_t)} D_t$ with $D_t = D$ and such that D is the first alive divisor in the sequence.

The random vertex v_τ that causes the system to first reach a threshold state is called the *epicenter*. In detail, the probability a particular vertex v is the epicenter is $\mathbb{P}_{D_0}(v_\tau = v) = \sum \prod_{i=1}^t \alpha(v_i)$, this time the sum being taken over all strings $v_1 \dots v_t$ for all $t \geq 1$ ending with $v_t = v$ and such that the corresponding sequence of states first reaches threshold at D_t .

We can now define the main statistic of interest and precisely state the questions motivating this chapter:

Definition 8.33. The *threshold density* of the fixed-energy sandpile with initial state D_0 is the expected number of grains of sand per vertex in the threshold state:

$$\zeta_\tau(D_0) := \mathbb{E}_{D_0} \frac{\deg(D_\tau)}{|V|}.$$

Questions:

- What is the threshold density for a given system? Or, what amounts to the same thing: how much sand is expected on the graph at threshold?
- What is $\mathbb{P}_{D_0}(\deg D_\tau = k)$ for each k ?
- What is $\mathbb{P}_{D_0}(v_\tau = v)$ for each $v \in V$?

We are able to answer these questions for the class of Eulerian graphs, which includes undirected graphs (see Appendix A, Definition A.28 and Proposition A.29). Note that, for a vertex v in an Eulerian graph G , we can define $\deg_G(v) := \text{indeg}_G(v) = \text{outdeg}_G(v)$.

8.2.2. Comparison with the abelian sandpile. To answer our questions, we first need to understand the relationship between the abelian sandpile and the fixed-energy sandpile.

If G is Eulerian, then by Theorem 7.6, its burning script is $\chi_{\bar{v}}$, the characteristic function of the non-sink vertices, and hence its burning sandpile is obtained by firing the sink vertex. So one way to create an alive divisor is to start with a recurrent sandpile c and add just enough sand to the sink so that the sink is unstable. By the burning algorithm (Theorem 7.5), toppling the sink then induces an avalanche in which every non-sink vertex fires exactly once, yielding c again, then repeating.

Definition 8.34. Let G be an Eulerian sandpile graph, and let s be a vertex of G . A *basic alive divisor* (with respect to s) or simply a *basic divisor* is a divisor of the form $c + \deg_G(s)s$ for some recurrent $c \in \mathcal{S}(G, s)$. The collection of basic alive divisors is denoted $\mathcal{B}(G, s)$ or just $\mathcal{B}(G)$ if s is clear from context.

In what follows, by relabeling recurrents with their corresponding basic alive divisors, we often identify the abelian sandpile with a Markov chain on $\mathcal{B}(G)$. In

detail, if $B = c + \deg_G(s)s \in \mathcal{B}(G)$ and $v \in V$ is drawn at random, then the transition from B is

$$c + \deg_G(s)s \xrightarrow{(v)} (c + v)^\circ + \deg_G(s)s \in \mathcal{B}(G).$$

Lemma 8.35 (Basic decomposition). *Let $G = (V, E)$ be an Eulerian graph and $D \in \text{Div}(G)$. For each $s \in V$, there is a unique triple (B, e, σ) , where $B \in \mathcal{B}(G, s)$ is a basic alive divisor, $e \in \mathbb{Z}$, and $\sigma: V \rightarrow \mathbb{Z}$ is a firing script with $\sigma(s) = 0$, such that*

$$D = B + es - L\sigma,$$

where L is the Laplacian of G . Moreover, D is alive if and only if $e \geq 0$.

Proof. Let $c \in \mathcal{S}(G, s)$ be the unique recurrent sandpile equivalent to the configuration $D|_{\tilde{V}}$ modulo the reduced Laplacian lattice $\tilde{\mathcal{L}}$ with respect to the chosen sink, s . This means that (as configurations on $\tilde{V} := V \setminus \{s\}$)

$$D|_{\tilde{V}} = c - \tilde{L}\sigma$$

for some firing script $\sigma: \tilde{V} \rightarrow \mathbb{Z}$. This firing script is unique since \tilde{L} is invertible. Extend σ to a script on all of V by letting $\sigma(s) = 0$, and define the basic alive divisor $B := c + \deg_G(s)s$. Since D and $D - L\sigma$ are equal except possibly at the sink, we have

$$D = B + es - L\sigma,$$

where $e = \deg(D) - \deg(B)$ since $\deg(L\sigma) = 0$. The uniqueness of the decomposition follows immediately from the uniqueness of the recurrent sandpile c and the toppling script σ .

Next, note that $B + es$ is alive if and only if $e \geq 0$ by Theorem 7.5. The same holds for D by Proposition 8.32 since $D \sim B + es$ and G is Eulerian, hence strongly connected. So e is the *excess* sand, over that required for D to be alive. \square

Definition 8.36. Using the notation of Lemma 8.35, let $B_s(D) := B$ and let $R_s(D) := c$ to define mappings

$$B_s: \text{Div}(G) \rightarrow \mathcal{B}(G, s) \quad \text{and} \quad R_s: \text{Div}(G) \rightarrow \mathcal{S}(G, s)$$

for each $s \in V$.

Thus, the basic decomposition of $D \in \text{Div}(G)$ with respect to s has the form

$$D = B_s(D) + es - L\sigma = (R_s(D) + \deg_G(s)s) + es - L\sigma,$$

Note that $R_s(D) = R_s(B)$.

Now suppose we drop a grain of sand on the vertex v and then stabilize $D + v$ (if possible, otherwise do nothing) as well as stabilize $R_s(D) + v$ with respect to a chosen sink s (which is always possible, yielding another recurrent). Lemma 8.40 below implies that this process commutes with the mapping R_s . That is, the update rules for the fixed-energy sandpile and the abelian sandpile model are intertwined:

$$R_s(a_v D) = (R_s(D) + v)^\circ.$$

For the purposes of the upcoming lemma and for later results, we will need to keep track of the amount of sand that goes into the sink s in each step of the

abelian sandpile Markov chain. Recall that for each non-sink vertex v , we defined r_v to be the unique recurrent sandpile equivalent to v modulo the reduced Laplacian lattice \tilde{L} and defined r_s to be the identity of the sandpile group.

Definition 8.37. Let G be a sandpile graph with sink s , and let $c \in \mathcal{S}(G)$ with corresponding basic alive divisor $B = c + \deg_G(s)s$. Fixing a vertex v , let $\tilde{c} := r_v^{-1} \otimes c$ so that \tilde{c} is the unique recurrent sandpile such that $(\tilde{c} + v)^\circ = c$. The *burst size* is the number of grains of sand that fall into the sink as the abelian sandpile moves from state \tilde{c} to state c upon the addition of a grain of sand at v :

$$\beta_v(B) := \beta_v(c) := \deg(\tilde{c}) - \deg(c) + 1.$$

In particular, $\beta_s(B) = \beta_s(c) = 1$.

Thus, $\beta_v((c + v)^\circ)$ is the amount of sand that falls into the sink as the chain moves from state c to $(c + v)^\circ$.

Exercise 8.38. Consider the complete graph K_4 with ordered list of vertices s, v_1, v_2, v_3 and sink s .

- (1) Show that $r_{v_1}^{-1} = (1, 2, 2) = v_1 + 2v_2 + 2v_3$.
- (2) Verify the burst sizes in Table 1.

recurrents	β_s	β_{v_1}	β_{v_2}	β_{v_3}
(0,1,2)	1	1	2	3
(0,2,2)	1	1	0	0
(1,1,2)	1	0	0	3
(1,2,2)	1	0	0	0
(2,2,2)	1	0	0	0

Table 1. Up to symmetry, all recurrents on K_4 and their burst sizes.

It will be useful to have formal notation for the firing script of a stabilization:

Definition 8.39. The *odometer* function,

$$\begin{aligned} \text{odo}: \text{Config}(G) &\rightarrow \tilde{\mathcal{M}}(G) \\ c &\mapsto \tilde{L}^{-1}(c - c^\circ), \end{aligned}$$

assigns to each configuration c the firing script for its stabilization. Thus, $\text{odo}(c)$ records how many times each vertex fires as $c \rightsquigarrow c^\circ$ through legal vertex firings.

As usual, we consider each configuration as a divisor with no sand on the sink, in other words, through the natural inclusion $\mathbb{Z}\tilde{V} \subset \mathbb{Z}V$. Every script $\sigma : \tilde{V} \rightarrow \mathbb{Z}$, such as that given by the odometer function, is naturally considered as a script $\sigma : V \rightarrow \mathbb{Z}$ by letting $\sigma(s) = 0$.

Lemma 8.40. Fix $v \in V$. Suppose that $D = B + es - L(\sigma)$ is the basic decomposition of a divisor D where $R_s(D) =: c \in \mathcal{S}(G)$ so that $B_s(D) = B = c + \deg_G(s)s$ is the corresponding basic divisor. Then the basic decomposition of $a_v D$ is

$$a_v D = (c + v)^\circ + (e + \deg_G(s) + \beta)s$$

$$-L(\sigma - \text{odo}(c+v) + \mu - \mu(s)),$$

where $\beta := \beta_v((c+v)^\circ)$ is the burst size, μ is the toppling script for the stabilization of $D+v$ as a divisor ($\mu = 0$ if $D+v$ is alive), and $\mu(s)$ is the constant script $\mu(s) \cdot \vec{1}$.

Proof. We have

$$\begin{aligned} a_v D &= D + v - L\mu \\ &= c + v + (e + \deg_G(s))s - L(\sigma + \mu) \\ &= c + v - L \text{odo}(c+v) + (e + \deg_G(s))s \\ &\quad - L(\sigma + \mu - \text{odo}(c+v)). \end{aligned}$$

In the stabilization $c+v \rightarrow (c+v)^\circ$, the amount of sand going into the sink is $\beta = \beta_v((c+v)^\circ)$ (with $\beta = 1$ if $v = s$). Therefore,

$$\begin{aligned} a_v D &= (c+v)^\circ + (e + \deg_G(s) + \beta)s - L(\sigma + \mu - \text{odo}(c+v)) \\ &= (c+v)^\circ + (e + \deg_G(s) + \beta)s \\ &\quad - L(\sigma - \text{odo}(c+v) + \mu - \mu(s)) \end{aligned}$$

where in the last step we have subtracted the constant function $\mu(s)$, which is in the kernel of L , in order to ensure that the firing script vanishes at the sink. \square

Now suppose that we run the fixed-energy sandpile starting at the initial state D_0 and that we simultaneously run the abelian sandpile starting at state $c_0 = R_s(D_0)$. At each time step, we make a random draw from the probability distribution α on V and use the result to update both Markov chains. The previous lemma implies that for every time t ,

$$R_s(D_{t+1}) = R_s(a_{v_t} D_t) = (R_s(D_t) + v_t)^\circ = (c_t + v_t)^\circ = c_{t+1}.$$

Hence, the evolution of the abelian sandpile Markov chain is completely determined by the evolution of the fixed-energy sandpile Markov chain.

8.2.3. Stationary density. Fixing any sink vertex s , the threshold state of the fixed-energy sandpile will be linearly equivalent to a basic alive divisor $B \in \mathcal{B}(G)$ with excess sand $e \geq 0$ at the sink: $B + es = c + (\deg_G(s) + e)s$ where $c \in \mathcal{S}(G)$ is some recurrent for the abelian sandpile. Since the abelian sandpile Markov chain evolves towards the uniform distribution among the recurrents, we might guess that the average density of the basic divisors is relevant to understanding the threshold density of the fixed-energy sandpile.

Definition 8.41. The *stationary density* of the abelian sandpile model is

$$\begin{aligned} \zeta_{\text{st}} &:= \frac{1}{|\mathcal{S}(G)|} \sum_{B \in \mathcal{B}(G)} \frac{\deg B}{|V|} \\ &= \frac{\deg_G(s)}{|V|} + \frac{1}{|\mathcal{S}(G)|} \sum_{c \in \mathcal{S}(G)} \frac{\deg(c)}{|V|}. \end{aligned}$$

Now think of the abelian sandpile model (restricted to the sandpile group) as a Markov chain on $\mathcal{B}(G)$ by relabeling each recurrent sandpile by its corresponding basic divisor. The stationary density is then the average amount of sand per vertex in the stationary distribution of the chain.

When G is Eulerian, the following theorem shows that the basic stationary density is independent of the choice of sink vertex. So in that case, the basic stationary density is really a property of the graph; we may refer to it simply as the *stationary density* and write $\zeta_{\text{st}} = \zeta_{\text{st}}(G)$.

Theorem 8.42. *If G is Eulerian, then for all $n \geq 0$ the following quantities are independent of the choice of sink vertex s :*

- (1) $|\{c \in \mathbb{Z}\tilde{V} : c \text{ is superstable, } \deg(c) = n\}|$,
- (2) $|\mathcal{S}(G, s)|$,
- (3) $|\{B \in \mathcal{B}(G, s) : \deg(B) = n\}|$.

Proof. Part (1) is an immediate consequence of Merino's Theorem and its generalization to Eulerian graphs. For a statement and proof of Merino's Theorem, see Theorem 14.18. The theorem was generalized to Eulerian graphs in [79] (and then extended to strongly connected graphs in [25]).

By Theorem 7.12, c is recurrent if and only if $c_{\text{max}} - c$ is superstable, so for each sink s we have a bijection between superstables and recurrences. By (1), the total number of superstables is sink-independent, hence so is the total number of recurrences, which is (2).

Finally, for a fixed sink s , the number of recurrences of degree k is equal to the number of superstables of degree $m_s - k$, where

$$m_s := \deg(c_{\text{max}}) = \sum_{v \in \tilde{V}} (\deg_G(v) - 1) = \sum_{v \in V} (\deg_G(v) - 1) - \deg_G(s) + 1.$$

It follows that $m_s + \deg_G(s)$ is independent of the sink s , so for all $n \geq 0$, the following quantity is also sink-independent:

$$|\{c \in \mathcal{S}(G, s) : \deg(c) + \deg_G(s) = n\}|.$$

Since basic divisors have the form $c + \deg_G(s)s$, part 3 follows. □

Remark 8.43. If G is an undirected multigraph, then Merino's Theorem (Theorem 14.18) leads to a formula for the stationary density in terms of a specialization of the Tutte polynomial for G . A combinatorial interpretation of the stationary density in terms of spanning trees and spanning "unicycles" of the graph then follows. These ideas then lead to a closed form and asymptotics for the stationary density of the complete graph K_n in terms of Ramanujan's Q -function. For details, see Problems 14.10 and 14.11.

The threshold density $\zeta_\tau(D_0)$ depends on the initial state, D_0 . For instance, if D_0 is already alive, then $\zeta_\tau(D_0) = \deg(D_0)/|V|$. The case where D_0 is stable, e.g., $D_0 = 0$, is more interesting. Problem 8.5 shows that if G is a tree, then the threshold density starting from any stable divisor actually coincides with the

stationary density. Can the same be said for an arbitrary graph? We consider some examples.

Example 8.44. The recurrences on the cycle graph C_n consist of the maximal stable sandpile, c_{\max} , having one grain of sand on each non-sink vertex, and the $n - 1$ sandpiles of the form $c_v := c_{\max} - v$ where v is a non-sink vertex. Adding 2 grains of sand at the sink to each of these, we see that there is one basic divisor of degree $\deg(c_{\max}) + 2 = n + 1$, and $n - 1$ with degree $\deg(c_v) + 2 = n$. Therefore, the stationary density is

$$\zeta_{\text{st}} = \frac{1}{n} \cdot \frac{(n+1) + (n-1)n}{n} = \frac{n^2 + 1}{n^2} = 1 + \frac{1}{n^2}.$$

To compute the threshold density of the fixed-energy sandpile, note that since a divisor on C_n is stable if and only if the number of grains of sand on each vertex is at most one, the only stable divisor of degree n is the all ones divisor, $1_V := \sum_{v \in V} v$. Similarly, a stable divisor of degree $n - 1$ must have the form $1_V - v$ for some vertex v , i.e., the all ones divisor but with no sand on v . Start the sandpile at an initial stable state D_0 . If $\deg(D_0) = n$, then $D_0 = 1_V$, and the chain reaches threshold at the next step at a divisor of degree $n + 1$. Otherwise, the system eventually evolves to a divisor $D = 1_V - v$ of degree $n - 1$. From that point, if vertices are chosen uniformly at random, with probability $(n - 1)/n$ the next step results in an alive divisor, and with probability $1/n$, the next step is the stable divisor 1_V , followed by a threshold state. So starting at any stable state with degree at most $n - 1$, the expected amount of sand at threshold—the degree of the threshold divisor—is

$$\frac{n-1}{n} \cdot n + \frac{1}{n} \cdot (n+1) = \frac{n^2 + 1}{n}.$$

In sum, starting at a stable state D_0 ,

$$\zeta_{\tau}(D_0) = \begin{cases} 1 + \frac{1}{n} & \text{if } D_0 = 1_V, \\ 1 + \frac{1}{n^2} = \zeta_{\text{st}} & \text{if } D_0 \preceq 1_V. \end{cases}$$

As a special case, note that $\lim_{n \rightarrow \infty} \zeta_{\tau}(0) = \lim_{n \rightarrow \infty} \zeta_{\text{st}}(C_n) = 1$.

Example 8.45. Let B_3 be the *banana graph* consisting of three edges joining two vertices u and v . Figure 3 shows that fixed-energy sandpile on B_3 starting at $(0, 0)$ and evolving up to threshold. The divisor $au + bv$ is denoted (a, b) . Choosing the second vertex, v , as the sink, the basic divisors for B_3 are $(0, 3)$, $(1, 3)$, and $(2, 3)$. So the stationary density is

$$\zeta_{\text{st}} = \frac{1}{3} \left(\frac{3 + 4 + 5}{2} \right) = 2.$$

Underneath each state in Figure 3 is listed the number of paths to that state from the starting state, $(0, 0)$. These numbers are easily calculated by recursion: to find the number of paths to a state, add the corresponding numbers for its immediately preceding states. Assuming that each vertex is equally likely to be chosen at each step, we can then calculate the probability of reaching each threshold


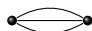


Figure 3. The fixed-energy sandpile Markov chain for the graph  with initial state (0,0). The small number appearing just below each state is the number of paths to that state from (0,0). Threshold states are in bold. See Example 8.45.

state. For instance, the probability the threshold state is (3,1) is $3/2^4$. These probabilities then determine the threshold density:

$$\zeta_\tau((0,0)) = \frac{1}{2} \left(2 \cdot \frac{1}{2^3} \cdot 3 + 2 \cdot \frac{3}{2^4} \cdot 4 + 2 \cdot \frac{6}{2^5} \cdot 5 \right) = \frac{33}{16}.$$

So in this case, the stationary density is a bit smaller than the threshold density.

The key to understanding the discrepancy is to consider the dependence of the threshold density on the initial state. Using the stationary density as an estimate of threshold density assumes that the probability the threshold state has degree d equals the proportion of basic divisors of degree d . For B_3 , this means that the threshold degree would be 3, 4, or 5, each with probability $1/3$. From the path counts in Figure 3, the actual probabilities of these degrees for the chain starting at (0,0) are $1/4$, $3/8$, and $3/8$, respectively—skewed towards higher threshold degrees compared with the stationary density estimate. Starting at (2,2), the effect is even more pronounced: the corresponding threshold probabilities are then 0, 0, and 1.

Exercise 8.46. Use Figure 3 to verify the entries in Table 2.

What would happen if the chain started at $(-100, -100)$? Eventually, the system evolves to a state that is nonnegative, and from that point, the behavior is accounted for by Figure 3. The table below shows the result of 1000 trials of a computer simulation in which the Markov chain is started at $(-100, -100)$ and allowed to evolve to a nonnegative divisor; the number of times each nonnegative stable divisor is reached is recorded:

divisor	(0,0)	(1,0)	(0,1)	(2,0)	(1,1)	(0,2)	(2,1)	(1,2)	(2,2)
frequency	310	174	186	158	0	172	0	0	0

	(0, 0)	(1, 0), (0, 1)	(2, 0), (0, 2)	(1, 1)	(2, 1), (1, 2)	(2, 2)
3	$\frac{2}{8}$	$\frac{2}{8}$	$\frac{4}{8}$	0	0	0
4	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{2}{8}$	$\frac{4}{8}$	$\frac{4}{8}$	0
5	$\frac{3}{8}$	$\frac{3}{8}$	$\frac{2}{8}$	$\frac{4}{8}$	$\frac{4}{8}$	1
ζ_τ	$\frac{33}{16}$	$\frac{33}{16}$	$\frac{30}{16}$	$\frac{36}{16}$	$\frac{36}{16}$	$\frac{40}{16}$

Table 2. The first row lists the possible nonnegative stable initial states for the fixed-energy sandpile on B_3 . The columns list the corresponding threshold densities and probabilities for the various threshold degrees.

Exercise 8.47. Convince yourself that the divisors $(1, 1)$, $(2, 1)$, $(1, 2)$, and $(2, 2)$ will never occur as the first nonnegative divisor in the chain started at $(-100, -100)$.

Suppose, contrary to the table, that the Markov chain always first reached a nonnegative state at $(0, 0)$, $(1, 0)$, or $(0, 1)$. Then according to Table 2, we would have equality between $\zeta_\tau((0, 0))$ and $\zeta_\tau((-100, -100))$. Instead, the data show there is also the possibility of reaching either $(2, 0)$ or $(0, 2)$. Having reached one of these states, the threshold degree is biased more towards 3, and the expected density, $30/16$, is below the stationary density.

Using Table 2 and the results from the simulation, we get the following estimate for $\zeta_\tau(-100, -100)$:

$$\left((310) \frac{33}{16} + (174 + 186) \frac{33}{16} + (158 + 172) \frac{30}{16} \right) \frac{1}{1000} = 2.000625,$$

which, magically, is close to the stationary density of 2.

Problem 8.10 considers the n -banana graph, B_n , defined by connecting two vertices by n edges, and requires the reader to show that $\lim_{n \rightarrow \infty} \zeta_\tau((0, 0))/\zeta_{\text{st}} = 4/3$.

8.3. The threshold density theorem

The paper on which this chapter is based, [66], starts with the question: “How much memory does a critical system retain of its pre-critical past?” The above two examples show that when the fixed-energy sandpile Markov chain reaches threshold—its critical state—the density of sand may depend on the initial state: the chain has some memory of its starting point. Once one realizes this dependence, it becomes natural to ask what would happen if the chain were allowed to mix for a long time before reaching threshold, say by starting at a state D_0 far from threshold. To help answer that question, we will compare the fixed-energy sandpile of divisors (having no sink) with the abelian sandpile of recurrents (having chosen a sink vertex).

In Figure 4, the left-hand side depicts divisors D_t in the fixed energy sandpile evolving to threshold. As we have seen, these divisors determine a corresponding evolution of recurrents c_t in the abelian sandpile, depicted on the right. Over time, the c_t approach the uniform distribution on the set of recurrents (Corollary 8.28),

and thus, the basic divisors B_t appearing on the left also approach the uniform distribution. This might lead one to believe that each basic divisor is equally likely to appear as B_τ in the limit as $\deg(D_0) \rightarrow -\infty$. It would then follow that in the limit the threshold density is the average density of the basic alive divisors, i.e., the stationary density.




Figure 4. Evolution to threshold. Divisors and their linearly equivalent basic decompositions evolving to threshold are on the left, and their corresponding recurrences are on the right: writing $B_t = c_t + \deg_G(s)s$ with $c_t \in \mathcal{S}(G)$, we have $R_s(D_t) = R_s(B_t) = c_t$.

Theorem 8.48 (Threshold density theorem). *Let $G = (V, E)$ be an Eulerian digraph. Then in the limit as $\deg(D_0) \rightarrow -\infty$, the threshold density of the fixed-energy sandpile on G converges to the stationary density of G :*

$$\zeta_\tau(D_0) \rightarrow \zeta_{\text{st}} \quad \text{as } \deg(D_0) \rightarrow -\infty,$$

i.e., for all $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that $|\zeta_\tau(D_0) - \zeta_{\text{st}}| < \epsilon$ whenever $\deg(D_0) < -N$.

We are not quite ready to prove the threshold density theorem, for the reasoning just employed to motivate it is faulty! The basic divisors appearing at threshold are not uniformly distributed. To understand this, consider the basic decomposition of a divisor in the chain before threshold,

$$D_{t-1} \sim B_{t-1} + e_{t-1}s = (c_{t-1} + \deg_G(s)s) + e_{t-1}s$$

with $e_{t-1} < 0$. In the next step in the abelian sandpile, $c_{t-1} \xrightarrow{(v_t)} c_t$, the amount of sand going into the sink is the burst size, $\beta_{v_t}(B_t) = \beta_{v_t}(c_t)$. In the fixed-energy sandpile, we have then

$$D_t \sim B_t + e_t s = B_t + (e_{t-1} + \beta_{v_t}(B_t))s,$$

and threshold is reached if and only if

$$e_{t-1} + \beta_{v_t}(B_t) \geq 0.$$

Thus, a basic divisor with large burst sizes is more likely to appear at threshold.

Example 8.49. Consider the 3-banana graph of Example 8.45 consisting of two vertices, u and v , connected by three edges. Take $s := v$ as the sink. The basic divisors and their burst sizes with respect to the two vertices are given below:

basic divisor	burst sizes for u and s , respectively
$B := (0, 3)$	3, 1
$B' := (1, 3)$	0, 1
$B'' := (2, 3)$	0, 1.

If the first vertex, u , is the epicenter, then (evolving from a stable initial state) there are three possibilities for the basic decomposition of the threshold divisor: $B + 0 \cdot s$, $B + 1 \cdot s$, and $B + 2 \cdot s$, whereas if s is the epicenter, the possibilities are $B + 0 \cdot s$, $B' + 0 \cdot s$, and $B'' + 0 \cdot s$. As we will see later, it turns out that if vertices are chosen with equal probabilities, these six possibilities are equally likely, making B four times more likely to occur at threshold than either B' or B'' . Nevertheless, in accordance with the threshold density theorem, averaging these six possibilities gives a threshold density of 2, which happens to equal the stationary density!

A curious marathon. We've just seen that as the fixed-energy sandpile runs in tandem with the abelian sandpile, the basic divisors and their corresponding recurrences approach the uniform distribution. However, the basic divisor at the time of threshold depends on burst sizes, which are functions of both the basic divisors and the vertices. We construct a new Markov chain to keep track of these data.

Definition 8.50. Fix a sink vertex s , and consider the directed graph $M(G, s)$ with vertex set $\mathcal{V} := V \times \mathcal{B}(G)$ and directed edges

$$\mathcal{E} := \{((v', B'), (v, B)) : (v', B'), (v, B) \in \mathcal{V} \text{ and } B' \xrightarrow{(v)} B\}.$$

Define the *length* of each edge by the burst size

$$\ell((v', B'), (v, B)) = \beta_v(B) = \deg(B') - \deg(B) + 1.$$

Next, define the corresponding *vertex-divisor Markov chain* corresponding to $M(G, s)$ to have state space \mathcal{V} and transition probabilities

$$P((v', B'), (v, B)) = \begin{cases} \alpha(v) & \text{if } B' \xrightarrow{(v)} B, \\ 0 & \text{otherwise.} \end{cases}$$

Since we are assuming that there is a positive probability of choosing each vertex in the fixed-energy sandpile ($\alpha(v) > 0$ for all $v \in V$), it is straightforward to check that the vertex-divisor chain is irreducible, with stationary distribution π given by $\pi((v, B)) = \frac{\alpha(v)}{|\mathcal{S}(G)|}$ for all $(v, B) \in M(G, s)$ (cf. Problem 8.12).

Note that $M(G, s)$ is Eulerian with each vertex having degree $|V|$. In particular, for each vertex-divisor pair (v, B) , there is a unique basic divisor B' such that $B' \xrightarrow{(v)} B$, and thus an edge $((v', B'), (v, B))$ for each $v' \in V$. Hence, the length of each incoming edge is $\beta_v(B)$.

Example 8.51. Returning to the banana graph B_3 in Example 8.49, Figure 5 displays the graph $M(B_3, s)$ with edges labeled by burst size.




Figure 5. The vertex-divisor digraph $M(B_3, s)$ with edges labeled by burst size.

As suggested in [66], we now imagine a “Markov” marathon in a town whose street map is $M(G, s)$. The intersections in the town are the vertex-divisor pairs forming the vertices of $M(G, s)$. From each intersection (v', B') there is one outgoing (one-way) street for each vertex v of G . The length of the road from (v', B') to (v, B) is the burst size $\beta_v(B)$. The route for the marathon is determined randomly for each runner in accordance with the Markov chain corresponding to $M(G, s)$.

To start the marathon, an integer total distance d is decided upon, and runners are assigned to starting intersections in whatever manner. During the race, when a runner arrives at an intersection, the out-going road to their next intersection is chosen at random according to the distribution α for the vertices of the abelian sandpile on G . A runner finishes at the first intersection at which their cumulative distance C_d is d or greater. At that point, the runner records the pair $(r, C_d - d)$ consisting of the finishing road r (an edge of $M(G, s)$) and the excess distance.

What can we expect for each runner’s record $(r, C_d - d)$ in the limit as $d \rightarrow \infty$? Start with the excess distance, $e_d = C_d - d$. If $r = ((v', B'), (v, B))$, we must have $0 \leq C_d - d < \beta_v(B)$. As we will soon see, it turns out that in the limit, given that the runner has stopped at (v, B) , each of these $\beta_v(B)$ possibilities for the excess distance is equally likely. (The loops of length 1 in Figure 5 may suggest why this is true.) And next for the insight which unlocks the threshold density theorem: in the limit, the triple (v, B, e) where e is any of the possible excess distances at (v, B) occurs with probability $\alpha(v)/|\mathcal{S}(G)|$. For instance, if α is the uniform distribution, then all of these triples are equally likely.

Markov renewal. To formalize the marathon just described, let $(\Omega, P, (X_t))$ be an irreducible finite Markov chain with stationary distribution π , and let $\mathcal{E} := \{(x, y) : P(x, y) > 0\}$ be the edges in the chain’s corresponding digraph. Let $\ell: \mathcal{E} \rightarrow \mathbb{N}$ be a length function on the edges. The length of a path p is $\lambda(p) = \sum_{\epsilon \in p} \ell(\epsilon)$ where the sum is taken over the edges $\epsilon \in \mathcal{E}$ of p . We assume that ℓ is aperiodic in the sense that the gcd of the lengths of the closed paths is one. This

means that closed paths of every sufficiently large length may occur. For each $d \in \mathbb{N}$ consider the random time $\tau_d := \min\{t : \lambda_t \geq d\}$ where $\lambda_t := \sum_{i=1}^t \ell((X_{i-1}, X_i))$.

Theorem 8.52. (Markov renewal theorem) *Let $x_0, x, y \in \Omega$ be any states, and let $e \in \mathbb{N}$. Then, as $d \rightarrow \infty$,*

$$\mathbb{P}_{x_0}\{(X_{\tau_d-1}, X_{\tau_d}, \lambda_{\tau_d} - d) = (x, y, e)\} \rightarrow \begin{cases} \frac{1}{Z} \pi(x)P(x, y) & \text{if } 0 \leq e < \ell(x, y), \\ 0 & \text{otherwise.} \end{cases}$$

with normalization constant $Z = \sum_{(x,y) \in \mathcal{E}} \pi(x)P(x, y)\ell(x, y)$.

A proof of the above theorem is beyond the scope of this text. The result appears as Proposition 11 in [66] where it is noted to be a special case of a more general Markov renewal theorem due to Kestent ([54]).

We now apply our Markov renewal theorem to the vertex-divisor Markov chain of Definition 8.50.

Theorem 8.53. *Let G be an Eulerian graph, and choose a probability distribution $\alpha: V \rightarrow [0, 1]$ such that $\alpha(v) > 0$ for all $v \in V$. Let (D_t) be the fixed-energy sandpile on G with threshold time $\tau = \tau(D_0)$ and epicenter v_τ . Suppose the basic decomposition of the threshold divisor with respect to $s \in V$ is*

$$D_\tau = B_\tau + e_\tau s - L\sigma_\tau.$$

As $\deg(D_0) \rightarrow -\infty$, the joint probability distribution of the triple (v_τ, B_τ, e_τ) converges:

$$\mathbb{P}_{D_0}\{(v_\tau, B_\tau, e_\tau) = (v, B, e)\} \rightarrow \begin{cases} \frac{\alpha(v)}{|\mathcal{S}(G)|} & \text{if } 0 \leq e < \beta_v(B), \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Consider the basic decomposition $D_t = B_t + e_t s - L\sigma_t$. Since D_t is alive if and only if $e_t \geq 0$ (Lemma 8.35), it follows that the threshold time may be written as

$$\tau = \min\{t \geq 0 : e_t \geq 0\}.$$

But by Lemma 8.40,

$$e_t = e_{t-1} + \beta_{v_t}(B_t) = e_{t-1} + \ell((v_{t-1}, B_{t-1}), (v_t, B_t)).$$

Therefore, $e_t = e_0 + \lambda_t$ for all $t > 0$, so that

$$\tau = \min\{t \geq 0 : \lambda_t \geq -e_0\}.$$

The length function ℓ for the vertex-divisor chain is aperiodic since for each basic divisor B we have $\ell((s, B), (s, B)) = 1$. So the Markov renewal theorem applies and, letting $d := -e_0$, it says that in the limit as $e_0 \rightarrow -\infty$,

$$\begin{aligned} & \mathbb{P}_{D_0}\{(v_{\tau-1}, B_{\tau-1}), (v_\tau, B_\tau), e_\tau = ((v', B'), (v, B), e)\} \\ & \rightarrow \begin{cases} \frac{1}{Z} \frac{\alpha(v')}{|\mathcal{S}(G)|} \alpha(v) & \text{if } B' \xrightarrow{(v)} B \text{ and } 0 \leq e < \beta_v(B), \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Here, the normalization constant Z is given by

$$\begin{aligned} Z &= \sum_{((v', B'), (v, B)) \in \mathcal{E}} \frac{\alpha(v')}{|\mathcal{S}(G)|} \alpha(v) \beta_v(B) \\ &= \sum_{v, v' \in V} \alpha(v) \alpha(v') \frac{1}{|\mathcal{S}(G)|} \sum_{B \in \mathcal{B}(G)} \beta_v(B). \end{aligned}$$


But the inner average of burst sizes must be one by the conservation of sand. Indeed, for any vertex $v \in V$, the operator $c \mapsto c \otimes r_v := (c + v)^\circ$ permutes the elements of the sandpile group, $\mathcal{S}(G)$, and hence,

$$\begin{aligned} \sum_{c \in \mathcal{S}(G)} \beta_v((c + v)^\circ) &= \sum_{c \in \mathcal{S}(G)} \deg(c) - \sum_{c \in \mathcal{S}(G)} \deg(c \otimes r_v) + |\mathcal{S}(G)| \\ &= |\mathcal{S}(G)|. \end{aligned}$$

Substituting, we see that $Z = \sum_{v, v' \in V} \alpha(v) \alpha(v') = (\sum_{v \in V} \alpha(v))^2 = 1$. Our theorem now follows by summing the earlier convergence statement over the pairs (v', B') and noting that $e_0 \rightarrow -\infty$ if and only if $\deg D_0 \rightarrow -\infty$. \square

To appreciate the clarity the previous theorem brings to the nature of the threshold state, the reader is encouraged to try the following exercise.

Exercise 8.54. Let G be the path graph on three vertices but with one doubled edge:



Take α to be the uniform distribution on the vertices. There are seven possible threshold divisors (assuming a stable initial state), which we group below according to linear equivalence classes:

$$\begin{aligned} (0, 3, 0) &\sim (0, 2, 1) \sim (2, 1, 0) \sim (2, 2, -1) \\ (1, 3, 0) &\sim (1, 2, 1) \\ (2, 2, 0). \end{aligned}$$

For the following exercises, take “at threshold” to mean “at threshold in the limit as $D_0 \rightarrow -\infty$ ”. For each of the three choices of sink vertex:

- (1) Compute the two basic divisors and their burst sizes with respect to all three vertices.
- (2) For each of the seven possible threshold divisors D , write $D \sim B + ev$ where B and e come from the basic decomposition. (There are really just three cases—one for each divisor class of a threshold divisor.)
- (3) Find the six 3-tuples of the form (v, B, e) where v is a vertex, B is a basic divisor, and e is a possible excess value (i.e., a nonnegative integer less than $\beta_v(B)$). Since α is the uniform distribution, Theorem 8.53 says that each of these tuples is equally likely to occur at threshold.

- (4) Using the above 3-tuples, compute the probabilities of each of the two basic divisors occurring in the basic decomposition at threshold. Use these probabilities to show that the threshold density is the stationary density (in the limit).
- (5) Again using the 3-tuples, check that each vertex occurs at threshold with the same probability.

You should find that in the case of one of the vertices, one basic divisor is five times more likely to occur than the other at threshold, and in the case of the other two vertices, one basic divisor is twice as likely to occur.

Why can't the seven possible threshold divisors be equally likely to occur at threshold?

Corollary 8.55. *Let $s, v \in V$, let $B \in \mathcal{B}(G, s)$, and let $e, n \in \mathbb{N}$. We have the following in the limit as $\deg(D_0) \rightarrow -\infty$:*

(1)

$$\mathbb{P}_{D_0}\{B_\tau = B\} \rightarrow \frac{1}{|\mathcal{S}(G)|} \sum_{v \in V} \alpha(v) \beta_v(B),$$

(2)

$$\mathbb{P}_{D_0}\{v_\tau = v\} \rightarrow \alpha(v),$$

(3)

$$\mathbb{P}_{D_0}\{e_\tau = e \mid v_\tau = v, B_\tau = B\} \rightarrow \begin{cases} \frac{1}{|\beta_v(B)|} & \text{if } 0 \leq e < \beta_v(B), \\ 0 & \text{otherwise,} \end{cases}$$

(4)

$$\mathbb{P}_{D_0}\{\deg(D_\tau) = n\} \rightarrow \frac{|\{B \in \mathcal{B}(G, s) : \deg(B) = n\}|}{|\mathcal{S}(G)|}.$$

Remark 8.56. In words, the first three parts of this corollary say that in the long run:

- (1) The probability of a given basic divisor at threshold is proportional to the average burst size of that divisor, with proportionality constant $1/|\mathcal{S}(G)|$.
- (2) The probability of a given vertex being the epicenter is just the probability of drawing that vertex. The paper [66, p. 1007] suggests this result is an instance of a general principal:

... in a system driven slowly to criticality from a highly subcritical initial state, stress is distributed uniformly in the sense that the probability of triggering a system-spanning avalanche by applying additional stress does not depend on where the additional stress is applied.

- (3) Given the epicenter, threshold divisor, and choice of sink vertex, each possibly occurring amount of excess sand—above that needed to be alive—in the basic decomposition of the divisor is equally likely.

Proof of Corollary 8.55. Part (1) follows from summing over v and e in Theorem 8.53.

For Part (2), we first consider the case where the epicenter is the sink. A divisor with basic decomposition $B + es - L\sigma$ is stabilizable if and only if $e < 0$. So the only way for the epicenter to be the sink is if $e_\tau = 0$. (Recall that $\beta_s(B) = 1$ since adding a grain of sand to the sink causes no toppling in the abelian sandpile.) Therefore,

$$(8.1) \quad \mathbb{P}_{D_0}\{v_\tau = s, B_\tau = B\} = \mathbb{P}_{D_0}\{(v_\tau, B_\tau, m_\tau) = (s, B, 0)\} \rightarrow \frac{\alpha(s)}{|\mathcal{S}(G)|}$$

by Theorem 8.53. Part (2) follows by taking v to be the sink and summing over B .

Part (3) is an immediate consequence of Theorem 8.53.

For Part (4), we need to consider basic decompositions with respect to various sinks $s \in V$. So for a given choice of sink s , write the basic decomposition of a divisor as

$$D_t = B_t^s + e_t^s s - L(\sigma_t^s).$$

Then we have $\deg(D_t) = \deg(B_t^s) + e_t^s$ for all t . As mentioned in the proof of Part (2), the only way for the epicenter to be the sink, $v_\tau = s$, is for $e_\tau^s = 0$, so that $\deg(D_\tau) = \deg(B_\tau^s)$. But then conditioning on $v_\tau = s$ yields

$$\begin{aligned} \mathbb{P}_{D_0}\{\deg(D_\tau) = n\} &= \sum_{s \in V} \mathbb{P}_{D_0}\{v_\tau = s, \deg(D_\tau) = n\} \\ &= \sum_{s \in V} \mathbb{P}_{D_0}\{v_\tau = s, \deg(B_\tau^s) = n\}. \end{aligned}$$

By (8.1), above, in the proof Part (2), this sum converges as $\deg(D_0) \rightarrow -\infty$ to

$$\sum_{s \in V} \frac{\alpha(s)}{|\mathcal{S}(G)|} |\{B \in \mathcal{B}(G, s) : \deg(B) = n\}|,$$

Our result follows since $|\{B \in \mathcal{B}(G, s) : \deg(B) = n\}|$ is independent of s by Theorem 8.42. \square

8.3.1. Proof of the threshold density theorem. Corollary 8.55 (4) immediately implies Theorem 8.48, which states that the limit of the threshold density as $\deg(D_0) \rightarrow -\infty$ is the stationary density:

$$\zeta_\tau(D_0) \rightarrow \zeta_{\text{st}}.$$

Indeed, we have

$$\begin{aligned} \zeta_\tau(D_0) &= \mathbb{E}_{D_0} \frac{\deg(D_\tau)}{|V|} \\ &= \sum_{n \geq 0} \mathbb{P}_{D_0}\{\deg(D_\tau) = n\} \frac{n}{|V|} \\ &\rightarrow \sum_{n \geq 0} \frac{|\{B \in \mathcal{B}(G, s) : \deg(B) = n\}|}{|\mathcal{S}(G)|} \frac{n}{|V|} \\ &= \frac{1}{|\mathcal{S}(G)|} \sum_{B \in \mathcal{B}(G, s)} \frac{\deg(B)}{|V|} \\ &= \zeta_{\text{st}}. \end{aligned}$$

Note that the sum over n is finite since there is an upper bound for the degree of recurrent sandpiles, hence an upper bound for the degree of basic alive divisors.

Notes

This chapter is an exposition of the main result in the paper *Threshold State and a Conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle* ([66]), by Lionel Levine. The reader is encouraged to consult that work for further context and history. We recommend [65] as a compatible reference for the underlying theory of Markov chains.

Problems for Chapter 8

8.1. Suppose that $(\Omega, P, (X_t))$ is a finite Markov chain and $f: \Omega \rightarrow \mathbb{R}$ is a function on the set of states. Thinking of P as a matrix and f as a column vector, show that the entries of Pf are the expectation values for f at the next time step. That is, for all states $x \in \Omega$,

$$Pf(x) = \mathbb{E}_{X_0=x}\{f(X_1)\}.$$

8.2. Show that Proposition 8.32 is false without the assumption of strong connectivity by giving an example of a sandpile graph G with linearly equivalent divisors $D \sim D'$ for which D is alive (unstabilizable through a sequence of legal topplings) and D' is stable.

8.3. By Corollary 5.24, a divisor on an undirected graph is minimally alive if and only if its degree equals the number of edges of the graph. Give an example of an Eulerian graph for which this result does not hold.

8.4. Let G be a sandpile graph with sink s . Let $c \in \mathcal{S}(G)$ and $v \in V$. Show the following inequalities for burst sizes:

- (a) $\beta_v(c_{\max}) \leq \beta_v(c)$
- (b) $\beta_v((c_{\max} + v)^\circ) \geq \beta_v(c)$.

8.5. Let T be a tree with n vertices, and consider the fixed-energy sandpile on T .

- (a) Show that the stationary density of T is $\zeta = 1 - 1/n$.
- (b) Show that for every stable divisor D , the threshold density of the fixed-energy sandpile with initial state D is equal to the stationary density: $\zeta_r(D) = \zeta$.

8.6. Let G be the diamond graph (formed by removing an edge from the complete graph on 4 vertices). Consider the fixed-energy sandpile on G with the uniform distribution on draws from the vertices.

- (a) Compute the stationary density of G directly from the definition.
- (b) Compute the threshold density of G starting with $D_0 = \vec{0}$ using the method of Example 8.45.

8.7. Let G be the house graph (Figure 7 in the Problems section for Chapter 2). Find a stable divisor D and distinct vertices u and v such that $D(u) < 0$ but $D + v$ is alive.

8.8. Let G be a loopless undirected graph of genus g .

- (a) Let c be a recurrent sandpile on G with respect to the sink vertex s . Prove the following upper bound for burst sizes:

$$\beta_v(c) \leq g + 1,$$

for all vertices v .

- (b) Let D be a stable divisor on G , and suppose there exists some $v \in V$ such that $F = D + v$ is alive. Thus, F is a possible threshold divisor for the fixed-energy sandpile on G . Prove that $D(u) \geq \deg_G(u) - g - 1$ for all $u \in V$. (See Problem 8.7 for an example where the inequality is sharp.) Thus, every divisor

that is the threshold divisor for the fixed-energy sandpile on G can arise from the chain with initial condition

$$D_0 = \sum_{v \in V} (\deg_G(v) - g - 1)v = K + (1 - g)1_V$$

where K is the canonical divisor and $1_V = \sum_{v \in V} v$.

8.9. Consider the fixed-energy sandpile on K_4 in which draws from the vertices are made uniformly at random. According to Chapter 11, Section 11.1, the recurrent sandpiles on K_4 are the 16 coordinate-wise permutations of the following:

$$(0, 1, 2), (0, 2, 2), (1, 1, 2), (1, 2, 2), (2, 2, 2).$$

Exploiting symmetry will help with the following calculations.

- Compute the stationary density of K_4 directly from the definition.
- Compute the threshold density of K_4 starting with $D_0 = \vec{0}$ using the method of Example 8.45. (Warning: this is a fairly long calculation—it might be best done by dividing the work among a group of people.)

8.10. Let B_n be the n -banana graph consisting of two vertices u, v joined by n edges. Example 8.45 considers the case where $n = 3$. Figure 3 should be helpful for this problem.

- Show that the stationary density of B_n is

$$\zeta_{\text{st}}(B_n) = \frac{3}{4}n - \frac{1}{4}.$$

- Show that the threshold density for initial state $(0, 0)$ is

$$\zeta_{\tau}((0, 0)) = \frac{n}{2^n} \sum_{k=0}^{n-1} \frac{1}{2^k} \binom{n+k}{k}.$$

- From the preceding, show

$$\zeta_{\tau}((0, 0)) = n \left(1 - \frac{\binom{2n}{n}}{4^n} \right).$$

- Using Stirling's approximation, show

$$\zeta_{\tau}((0, 0)) \sim n \left(1 - \frac{1}{\sqrt{\pi n}} \right),$$

where $f(n) \sim g(n)$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$.

- Show

$$\lim_{n \rightarrow \infty} \frac{\zeta_{\tau}((0, 0))}{\zeta_{\text{st}}(B_n)} = \frac{4}{3}.$$

8.11. We say a divisor D is a *potential threshold divisor* if it may arise at the threshold divisor for the fixed-energy sandpile with respect to some stable initial state. This problem shows that not every potential threshold divisor is a basic alive divisor with respect to some vertex. Let G be the graph pictured in Figure 6. The sandpile group for G has order 3, and thus there are 3 basic alive divisors for each of the 5 vertices of G , accounting for 15 threshold divisors. Find the four remaining threshold divisors (two of which are nonnegative).




Figure 6. Graph for Problem 8.11.

8.12. Show that the vertex-divisor Markov chain of Definition 8.50 is irreducible and has stationary density π given by $\pi((v, B)) = \frac{\alpha(v)}{|\mathcal{S}(G)|}$ for all pairs (v, B) .

8.13. Let B_n be the n -banana graph consisting of two vertices v, s joined by n edges. Consider the fixed-energy sandpile on B_n with the uniform distribution $\alpha(v) = \alpha(s) = 1/2$ on its vertices.

- (a) Describe all potential threshold divisors (using the terminology of Problem 8.11) and their basic decompositions.
- (b) For each basic divisor B , compute $\mathbb{P}(B_\tau = B)$ in the limit as $\deg(D_0) \rightarrow -\infty$.
- (c) Prove that, unlike in the case of Exercise 8.54, each potential threshold divisor is equally likely in the limit.

Part 3

Topics

Trees

This chapter provides two proofs of the matrix-tree theorem, which counts the number of spanning trees of a graph using the Laplacian. The version of the theorem presented here applies to directed graphs, for which the appropriate notion of spanning tree is a *rooted spanning tree*—a subgraph for which each vertex has a unique directed path to a selected *root* vertex. For an undirected graph, we may instead think of a spanning tree as a connected acyclic subgraph containing all of the vertices.

For us, an important consequence of the matrix-tree theorem is that the number of elements in the sandpile group of a graph is the number of spanning trees of the graph (rooted at the sink). For example, Figure 1 displays the eight spanning trees of the diamond graph, a graph whose sandpile group is isomorphic to \mathbb{Z}_8 . Section 9.2 presents several corollaries, including the use of rooted spanning trees to compute the kernel of the Laplacian of a directed sandpile graph.

The remaining sections of the chapter discuss tree bijections and describe the remarkable rotor-routing process which provides an action of the sandpile group on the set of spanning trees.




Figure 1. Spanning trees of the diamond graph.


9.1. The matrix-tree theorem

Let $G = (V, E)$ be a directed graph. Recall that for us this has been taken to mean that V is a finite set of vertices, E is a finite multiset of directed edges, and G is (weakly) connected. For convenience, however, in this chapter we drop the last condition and make no connectedness assumption. A *subgraph* of G is a directed graph whose vertices and directed edges form sub(multi)sets of V and E , respectively.

Definition 9.1. A (*directed*) *spanning tree* of G rooted at $s \in V$ is a subgraph T such that for all $v \in V$, there exists a unique directed path in T from v to s . The vertex s is the *root* or *sink* of the tree.

If T is a directed spanning tree of G rooted at s , then (cf. Proposition A.26): (i) T contains *all* of the vertices of G (hence, the word “spanning”); (ii) T contains no directed cycles; and (iii) for all vertices v of G , the outdegree of v in T is 0 if $v = s$, and is 1, otherwise. In particular, T contains no multiple edges.


Example 9.2. The graph pictured below has three directed edges and one undirected edge:



The determinant of its reduced Laplacian with respect to s is

$$\det \begin{pmatrix} 3 & 0 \\ -2 & 1 \end{pmatrix} = 3,$$

which is the number of spanning trees rooted at s , as shown below:



Note that second two trees are different since the multiple edges of the form (v_1, v_2) are counted as distinct.

Let L be the Laplacian matrix of G relative to an ordering of the vertices, v_1, \dots, v_n . For each $k \in \{1, \dots, n\}$, let $L^{(k)}$ denote the $(n-1) \times (n-1)$ matrix formed by removing the k -th row and column of L . This is the *reduced Laplacian* for G with respect to v_k .

Theorem 9.3 (Matrix-tree). *The determinant of $L^{(k)}$ is the number of spanning trees of G rooted at v_k .*

Proof. Since a permutation of the vertices induces a corresponding permutation of the rows and columns of L , it suffices to consider the case $k = n$. For ease

of notation, we write $\tilde{L} := L^{(n)}$. Letting a_{ij} denote the number of times (v_i, v_j) appears as an edge of G , we have

$$\tilde{L} = \begin{pmatrix} \sum_{i \neq 1} a_{1i} & -a_{21} & -a_{31} & \cdots & -a_{n-1,1} \\ -a_{12} & \sum_{i \neq 2} a_{2i} & -a_{32} & \cdots & -a_{n-1,2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -a_{1,n-1} & -a_{2,n-1} & -a_{3,n-1} & \cdots & \sum_{i \neq n-1} a_{n-1,i} \end{pmatrix}$$

where $\sum_{i \neq k} a_{ki}$ denotes the sum over $i \in \{1, \dots, n\} \setminus \{k\}$. Each column encodes the rule for reverse-firing the corresponding vertex.

Let \mathfrak{S}_{n-1} be the permutation group on $\{1, \dots, n-1\}$. Recall that the *sign* of $\sigma \in \mathfrak{S}_{n-1}$ is $\text{sgn}(\sigma) := (-1)^t$ where t is the number of factors in any expression for σ as a product of transpositions—it records whether an even or odd number of swaps is required to create the permutation σ . Let $\text{Fix}(\sigma)$ be the set of *fixed points* of σ :

$$\text{Fix}(\sigma) := \{i \in \{1, \dots, n-1\} : \sigma(i) = i\}.$$


Then

$$(9.1) \quad \det \tilde{L} = \sum_{\sigma \in \mathfrak{S}_{n-1}} \text{sgn}(\sigma) \tilde{L}_{\sigma(1),1} \cdots \tilde{L}_{\sigma(n-1),n-1},$$

where

$$\tilde{L}_{\sigma(k),k} = \begin{cases} \sum_{i \neq k} a_{k,i} & \text{if } k \in \text{Fix}(\sigma) \\ -a_{k,\sigma(k)} & \text{otherwise.} \end{cases}$$

The idea now is to expand (9.1) into signed monomials in the a_{ij} and to think of each monomial as a directed graph by identifying a_{ij} with the directed edge (v_i, v_j) labeled with the number of times this edge appears in G , i.e., with a_{ij} , itself:



We then show that after cancellation due to the signs of the permutations, the remaining monomials correspond exactly to the distinct spanning trees rooted at v_n . Each monomial itself—a product of various a_{ij} —is an integer which counts the number of times its corresponding spanning tree occurs as a spanning tree of G . (Recall that since G may have repeated edges, a spanning tree may occur more than once.)

We pause now for an extended example. For those readers wishing to skip ahead, the end of the example is marked with a line.

Example 9.4. Take $n = 10$ and $\sigma = (2, 7)(3, 5, 9) \in \mathfrak{S}_9$. The set of fixed points is $\text{Fix}(\sigma) = \{1, 4, 6, 8\}$ and $\text{sgn}(\sigma) = \text{sgn}((2, 7)) \text{sgn}((3, 5, 9)) = (-1) \cdot 1 = -1$. The term in the expansion of $\det \tilde{L}$ corresponding to σ is

$$\begin{aligned} & \text{sgn}(\sigma) \tilde{L}_{\sigma(1),1} \tilde{L}_{\sigma(2),2} \cdots \tilde{L}_{\sigma(9),9} \\ &= (-1)(a_{1,2} + \cdots + a_{1,10})(-a_{2,7})(-a_{3,5})(a_{4,1} + \cdots + a_{4,10})' \\ & \quad \cdot (-a_{5,10})(a_{6,1} + \cdots + a_{6,10})'(-a_{7,2})(a_{8,1} + \cdots + a_{8,10})'(-a_{9,3}), \end{aligned}$$

where the prime symbol on a factor indicates the term of the form $a_{i,i}$ should be omitted from the enclosed sum. Continuing,

$$\begin{aligned}
 &= (-1) \left[\overbrace{(a_{1,2} + \dots)}^{\sigma(1)=1} \overbrace{(a_{4,1} + \dots)'}^{\sigma(4)=4} \overbrace{(a_{6,1} + \dots)'}^{\sigma(6)=6} \overbrace{(a_{8,1} + \dots)'}^{\sigma(8)=8} \right] \\
 &\quad \cdot \left[\underbrace{(-a_{2,7})(-a_{7,2})}_{(2,7)} \underbrace{(-a_{3,5})(-a_{5,9})(-a_{9,3})}_{(3,5,9)} \right].
 \end{aligned}$$

Question: Which monomials, identified with directed graphs, appear in the expansion of the above?

Answer: For each fixed point i of σ , we get a choice of any edge of the form

$\begin{matrix} & a_{ij} & \\ & \bullet & \\ v_i & \longrightarrow & v_j \end{matrix}$ where $j \in \{1, \dots, 10\}$ and $j \neq i$. For each non-trivial cycle of σ , there is only one choice:




Figure 2 considers three monomials coming from the expansion of the term in $\det \tilde{L}$ corresponding to σ . Each monomial m corresponds to a directed graph G_m . Column F shows the part of G_m that comes from choices for the fixed points of σ , and column C shows the part that comes from the nontrivial cycles. Note that, as in example (c), these two parts may share vertices. There may be an edge connecting a fixed point vertex to a cycle vertex in G_m . Example (b) shows that it is not necessary for v_{10} to occur in G_m . In general, v_{10} does not appear if and only if each vertex in G_m has a path to a directed cycle (since the outdegree for each non-root vertex is 1).

Finally, it is important to determine the sign of each monomial corresponding to σ in the expansion of $\det \tilde{L}$. The sign is determined by $\text{sgn } \sigma$ and by the number of factors of the form $-a_{ij}$ that go into the calculation of the monomial. With these two considerations in mind, it is straightforward to see that the resulting sign is $(-1)^{\# \text{ non-trivial cycles of } \sigma}$. For instance, consider the monomial in example (a) in Figure 2. It appears in the expansion of $\det \tilde{L}$ in the term

$$\text{sgn}((2,7)(3,5,9)) \underbrace{a_{1,10} a_{4,8} a_{6,4} a_{8,6}}_{\text{Fix}(\sigma) = \{1, 4, 6, 8\}} \underbrace{(-a_{2,7})(-a_{7,2})}_{(2,7)} \underbrace{(-a_{3,5})(-a_{5,9})(-a_{9,3})}_{(3,5,9)}.$$

Each cycle of σ ultimately contributes a factor of -1 :

$$\begin{aligned}
 \text{sgn}((2,7))(-a_{2,7})(-a_{7,2}) &= -1 \cdot a_{2,7} a_{7,2} \\
 \text{sgn}((3,5,9))(-a_{3,5})(-a_{5,9})(-a_{9,3}) &= -1 \cdot a_{3,5} a_{5,9} a_{9,3}.
 \end{aligned}$$

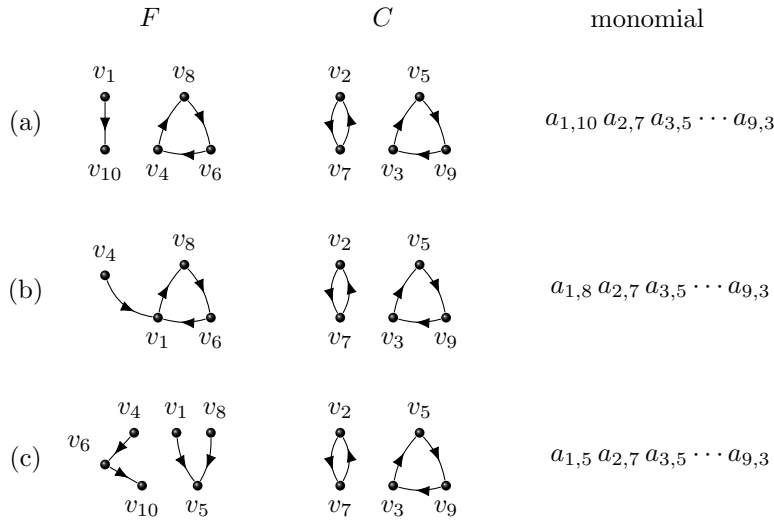


Figure 2. Monomials and corresponding graphs for Example 9.4.

We now return to the proof. The monomials in the expansion of (9.1) correspond exactly with signed, weighted, ordered pairs (F, C) of graphs F and C formed as follows:

- (1) Choose a subset $X \subseteq \{1, \dots, n - 1\}$ (representing the fixed points of some σ).
- (2) Make any loopless, directed (not necessarily connected) graph F with vertices $\{1, \dots, n\}$ such that

$$\text{outdeg}_F(i) = \begin{cases} 1 & \text{if } i \in X \\ 0 & \text{if } i \notin X. \end{cases}$$

- (3) Let C be any vertex-disjoint union of directed cycles of length at least 2 (i.e., no loops) such that C contains all of the vertices $\{1, \dots, n - 1\} \setminus X$.

Each of these ordered pairs of graphs (F, C) is associated with an element of \mathfrak{S}_{n-1} , with the vertices of outdegree one in F determining the fixed points and with C determining the cycles. In general, this is a many-to-one relationship, given the choices in step (2). The *weight* of (F, C) , denoted $\text{wt}(F, C)$ is the product of its labels—those a_{ij} such that (v_i, v_j) occurs in either F or C —multiplied by $(-1)^\gamma$ where γ is the number of cycles in C . For instance, for each of the three examples in Figure 2, the number of cycles in C is 2, so the weight is just the listed monomial, without a sign change. With this notion of weight, it then follows in general that

$$\det \tilde{L} = \sum_{(F,C)} \text{wt}(F, C).$$

Let Ω denote the set of ordered pairs (F, C) , constructed as above, but such that either F or C contains a directed cycle. We show that the monomials corresponding to elements of Ω cancel in pairs in the expansion of (9.1) by constructing a *sign reversing transposition* on Ω . Given $(F, C) \in \Omega$, pick the cycle γ of the disjoint

union $F \sqcup C$ with the vertex of smallest index. Then if γ is in F , move it to C , and vice versa. Formally, if the cycle is in F , define $F' = F \setminus \{\gamma\}$ and $C' = C \cup \{\gamma\}$, and if it is in C , define $F' = F \cup \{\gamma\}$ and $C' = C \setminus \{\gamma\}$. This defines a transposition $(F, C) \mapsto (F', C')$ on Ω such that $\text{wt}(F, C) = -\text{wt}(F', C')$ since the number of cycles in C differs from the number of cycles in C' by one. See Figure 3 for an example. It follows that in the sum $\sum \text{wt}(F, C)$, terms paired by the transposition




Figure 3. Sign reversing transposition.

cancel, leaving only those terms $\text{wt}(F, C)$ for which the transposition is undefined, i.e., those (F, C) such that the graph $F \sqcup C$ contains no cycles. In this case, the corresponding permutation is the identity permutation, $C = \emptyset$, and F is a spanning tree rooted at v_n . The weight, $\text{wt}(F, C)$, counts the number of times this spanning tree occurs as a spanning tree of G due to G having multiple edges. \square

9.1.1. Matrix-tree through deletion and contraction. This section presents a second proof of the matrix-tree theorem. If $e = (u, v)$ is a non-loop edge of $G = (V, E)$, define $G - e$ to be the graph formed from G by removing e (i.e., decrease its multiplicity by 1). It has the same vertex set as G , and like G , it is not necessarily connected. We say $G - e$ is formed from G by *deletion* of e . Next, define $G \nearrow e$ to be the graph formed from G by removing all edges with tail u , i.e., all edges of the form (u, w) with $w \in V$, and then identifying the vertices u and v , naming the resulting vertex v . We say $G \nearrow e$ is formed from G by *contraction* of e .¹ See Figure 4 for an example.

The following counts are evident for spanning trees rooted at $v = e^+$:

$$\# \left(\begin{array}{c} \text{spanning trees} \\ \text{of } G - e \end{array} \right) = \# \left(\begin{array}{c} \text{spanning trees of } G \text{ not} \\ \text{containing } e \end{array} \right),$$

and

$$\# \left(\begin{array}{c} \text{spanning trees} \\ \text{of } G \nearrow e \end{array} \right) = \# \left(\begin{array}{c} \text{spanning trees of } G \\ \text{containing } e \end{array} \right).$$

Letting κ denote the number of spanning trees rooted at v , we have

$$(9.2) \quad \kappa(G) = \kappa(G - e) + \kappa(G \nearrow e).$$

¹We use the notation $G \nearrow e$ to distinguish it from the usual contraction, G/e , for undirected graphs, formed by simply contracting the edge e and not removing additional edges (cf. Chapter 14).




Figure 4. Deletion and contraction of an edge. Undirected edges represent a pair of oppositely oriented directed edges.

We now present a proof of the matrix-tree theorem based on deletion and contraction.

Proof. The proof goes by induction on the number of non-loop edges. To avoid trivialities, assume that G has at least two vertices. Fix a root vertex v , and consider all rooted trees and reduced Laplacians with respect to v .

If G has no non-loop edges, then both $\det \tilde{L}_G$ and the number of spanning trees are 0. More generally, if there is no non-loop edge into the root, v , and hence no rooted spanning trees, then the sum of the rows of \tilde{L}_G is zero, hence $\det \tilde{L}_G = 0$. So the theorem holds in that case. Otherwise, we may assume there is an edge $e := (u, v)$ with $u \neq v$. For the sake of writing down matrices, order the vertices so that u and v appear first and second, respectively. Comparing Laplacians (in block form), we have

$$L_G = L_{G-e} + \left(\begin{array}{cc|c} 1 & 0 & 0 \\ -1 & 0 & 0 \\ \hline 0 & 0 & 0 \end{array} \right).$$

Hence, for reduced Laplacians:

$$\tilde{L}_G = \tilde{L}_{G-e} + \left(\begin{array}{c|c} 1 & 0 \\ \hline 0 & 0 \end{array} \right).$$

The matrix $\tilde{L}_G \gamma_e$ is formed by removing the rows and columns corresponding to u and to v from L_G , which is the same as removing those same rows and columns from L_{G-e} . Take determinants in the equation displayed above and use multilinearity to expand the right-hand side to see that

$$\det \tilde{L}_G = \det \tilde{L}_{G-e} + \det \tilde{L}_G \gamma_e.$$

The result follows from (9.2) and induction. \square

Remark 9.5. Another proof of the matrix-tree theorem in the undirected case is based on writing the reduced Laplacian as the product of the reduced vertex-edge incidence matrix with its dual, then expanding the determinant of the product using the Binet-Cauchy theorem. For further details, see Chapter 9 of [87], for instance.

Remark 9.6 (Generalized weighted version of matrix-tree). In our first proof of the matrix-tree theorem, we found it useful to think of an edge (v_i, v_j) which appears with multiplicity a_{ij} as a single edge whose *weight* is a_{ij} . We can then define the *weight*, $\text{wt}(T)$, of a rooted spanning tree T to be the product of the weights of its edges. These weights appeared as monomials in the expansion of $\det \tilde{L}$ in the proof. The weight of T then counts the number of times T appears as a rooted spanning tree in G . We can generalize the matrix-tree theorem by recognizing that, in fact, in either of our proofs, the weights a_{ij} could be replaced by arbitrary elements in a commutative ring. For instance, one may assume that the a_{ij} are indeterminates in a polynomial ring. The matrix-tree theorem that says that

$$\det \tilde{L} = \sum_T \text{wt}(T)$$

where the sum is over the rooted spanning trees of a directed edge-weighted graph G .

9.2. Consequences of the matrix-tree theorem

We now obtain several corollaries of the matrix-tree theorem. We first list a few immediate corollaries, then present another determinantal formula for counting trees, and finally consider the kernel of the Laplacian of a (directed) sandpile graph.

Corollary 9.7. *Let G be a sandpile graph. Then the order of the sandpile group of G is the number of directed spanning trees of G rooted at the sink.*

Proof. The result follows immediately from Proposition 6.33. \square

A *tree on n labeled vertices* is a connected undirected graph with n labeled vertices and no cycles.

Corollary 9.8 (Cayley's formula). *The number of trees on n labeled vertices is n^{n-2} .*

Proof. Problem 9.5. \square

Let $L^{(ij)}$ denote the matrix obtained from the Laplacian L of G by removing the i -th row and j -th column.

Corollary 9.9. *The ij -th cofactor, $(-1)^{i+j} \det L^{(ij)}$, of L is the number of directed spanning trees rooted at the j -th vertex.*

Proof. We have $(-1)^{i+j} \det L^{(ij)} = \det L^{(jj)}$ since the sum of the rows of L is the zero vector (Problem 9.3). The result then follows from the matrix-tree theorem. \square

Corollary 9.10. *Let G be a directed graph with n vertices. Suppose the Laplacian matrix of G has eigenvalues μ_1, \dots, μ_n with $\mu_n = 0$. For each vertex v , let κ_v be the number of directed spanning trees of G rooted at v . Then,*

$$\mu_1 \cdots \mu_{n-1} = \sum_v \kappa_v.$$

In other words, the product of these eigenvalues is the total number of rooted trees.

Proof. We may assume the vertex set is $1, \dots, n$, with the i -th column of L corresponding to vertex i . First note that since L is singular, a zero eigenvalue μ_n always exists. Factoring the characteristic polynomial of L , we have

$$\det(L - I_n x) = (\mu_1 - x) \cdots (\mu_n - x).$$

We calculate the coefficient of x in this expression in two ways. Since $\mu_n = 0$, by expanding the right-hand side, we see the coefficient is $-\mu_1 \cdots \mu_{n-1}$. Now consider the left-hand side. For each i , let r_i denote the i -th row of L , and let e_i denote the i -th standard basis vector. Then

$$\det(L - I_n x) = \det(r_1 - e_1 x, \dots, r_n - e_n x).$$

By multilinearity of the determinant, letting \tilde{L}_i denote the reduced Laplacian with respect to vertex i , the coefficient of x is

$$\sum_{i=1}^n \det(r_1, \dots, r_{i-1}, -e_i, r_{i+1}, \dots, r_n) = - \sum_{i=1}^n \det(\tilde{L}_i) = - \sum_{i=1}^n \kappa_i. \quad \square$$

Remark 9.11. If G is undirected, or more generally, if G is Eulerian (cf. Definition A.28 and Problem 9.4), then the number of spanning trees rooted at a vertex is independent of the particular vertex. Calling this number κ , Corollary 9.10 says in this case that

$$\kappa = \frac{\mu_1 \cdots \mu_{n-1}}{n}.$$

9.2.1. Another tree-counting formula. For an Eulerian multigraph (e.g., an undirected graph), the number of directed trees rooted at any vertex is independent of the vertex. In this case, we have a matrix-tree-like formula that does not depend on a choice of vertex.

Proposition 9.12. *Let G be an Eulerian multigraph on n vertices. Let L be the Laplacian matrix of G , and let J be the $n \times n$ matrix with all entries equal to 1. Then the number of spanning trees of G rooted at any vertex is*

$$\det(L + J)/n^2.$$

Proof. Problem 9.7. \square

Exercise 9.13. Use Proposition 9.12 to quickly compute the number of spanning trees (and hence the size of the sandpile group) of the complete graph, K_n .

9.2.2. Kernel of the Laplacian. Let $G = (V, E, s)$ be a sandpile graph with sink s , thus ensuring that G has at least one rooted spanning tree. For each $v \in V$, let $\theta(v)$ be the number of spanning trees of G rooted at v , thus defining a vector $\theta \in \mathbb{Z}^{|V|}$. Then let $\gamma := \gcd\{\theta(v)\}_{v \in V}$, and define $\tau := \theta/\gamma$.

Theorem 9.14. *The kernel of the Laplacian of G is generated by τ .*

Proof. Order the vertices of G as v_1, \dots, v_n with $v_n = s$. Define $L^{(ij)}$ and the cofactors $C_{ij} := (-1)^{i+j} \det L^{(ij)}$, as above. Since the sum of the rows of L is zero, Problem 9.3 says that if we fix j , then C_{ij} is independent of i . Now, fix i and calculate the determinant of L by expanding along its i -th row:


$$\begin{aligned} 0 = \det L &= \sum_{j=1}^n L_{ij} C_{ij} \\ &= \sum_{j=1}^n L_{ij} C_{jj} && \text{(independence of } C_{ij} \text{ on } i) \\ &= \sum_{j=1}^n L_{ij} \theta(v_j) && \text{(matrix-tree theorem).} \end{aligned}$$

Thus, $\tau \in \ker L$.

Since G has at least one directed spanning tree into s , the matrix tree theorem says that L has at least one nonzero $(n-1) \times (n-1)$ minor; hence, the rank of L is at least $n-1$. Since the rows of L add up to the zero vector, its rank is at most $n-1$. Therefore, $\ker L$ consists of all integer multiples of a single nonzero integer vector. Since $\theta \in \ker L$, it follows that τ generates the kernel. \square

Example 9.15.

- (1) If G is undirected, or more generally, Eulerian, the number of trees rooted at each vertex is independent of the vertex (cf. Corollary 12.3). Hence, $\gamma = 1$, and $\ker L$ is spanned by $(1, \dots, 1)$, as we saw earlier by other means.
- (2) The graph



has q trees into u and p trees into v . Hence, the kernel of its Laplacian is

$$\ker L = \ker \begin{pmatrix} p & -q \\ -p & q \end{pmatrix} = \left\{ k \cdot \frac{1}{\gcd(p,q)} (q, p) \in \mathbb{Z}^2 : k \in \mathbb{Z} \right\}.$$

Using the notation of Proposition 9.12, we have

$$\det(L + J)/2^2 = \frac{p+q}{2}.$$

Our graph is Eulerian if and only if $p = q$, in which case, $\det(L + J)/2^2 = p$ gives the number of spanning trees rooted into either vertex.

9.3. Tree bijections

As a consequence of the matrix-tree theorem, Corollary 9.7 says that the set of recurrent sandpiles and the set of spanning trees of a sandpile graph are equinumerous. In this section, we refine Dhar’s burning algorithm (Algorithm 6, Section 7.5) to provide an explicit bijection between these sets. In Part 1, we saw a “dual” version of Dhar’s algorithm which tested for superstability. Correspondingly, there are various refinements of Dhar’s algorithm providing bijections between superstable and spanning trees. We will discuss two of these with the property that under the bijections, the degree of a sandpile corresponds to an interesting property of its associated tree.

Since the recurrents form a group, as do the superstable, each of these bijections induces a group structure on the set of spanning trees. One could hope that there is a *natural* group structure on the spanning trees, independent of these bijections. Figure 1 from the beginning of the chapter shows this hope would be unfounded: each tree in that figure is paired with a distinct symmetric tree; so there can be no natural choice for the identity element. We could instead ask if there is a natural action of the sandpile group on the set of spanning trees. Later in this section, we shall see the elegant and astonishing rotor-router model provides a free and transitive group action, although it requires adding a bit more structure to the graph. The rotor-router model has the advantage of applying to arbitrary directed sandpile graphs whereas the bijections we present based on Dhar’s algorithm apply only to undirected multigraphs. Finally, we use what we have learned to address the problem of choosing a random spanning tree.

9.3.1. Dhar’s tree bijection algorithm. Let G be an undirected sandpile graph with sink s . Multiple edges are allowed, but to avoid trivialities, we assume G has no loops. Let $c \in \mathcal{S}(G)$ be a recurrent sandpile. Firing the sink vertex adds sand to its neighboring vertices and produces the sandpile $b + c$ where b is the burning sandpile of Chapter 7. By Dhar’s burning algorithm, $b + c$ stabilizes back to c with each vertex firing exactly once. Dhar and Majumdar ([38]) imagine this process as a fire spreading along edges, burning vertices in its path. At time 0, we burn (fire) the sink, causing a subset B_1 of its neighboring vertices to become unstable. At time 1, we burn (fire) these vertices, causing a new set of vertices B_2 to become unstable. We burn these at time 2, and so on, until every vertex is burnt. Each vertex $u \in B_i$ becomes unstable after firing its neighbors in B_{i-1} . The trick to forming a spanning tree is to choose one such neighbor for each u . But how does one make a choice? The answer in [38] is to initially make an arbitrary choice, once and for all, for an ordering of the edges of G . Having made that choice, we create a spanning tree corresponding to each c as follows: For each $u \in B_i$, let N_u be its set of neighbors in B_{i-1} . Just after firing B_{i-1} , say there are $k \geq \deg_G(u)$ grains of sand on u . If u is barely unstable, i.e., if $k = \deg_G(u)$, choose the first edge in N_u to be part of the tree. If $k = \deg_G(u) + 1$, choose the second edge, instead. In general, choose the $(k - \deg_G(u) + 1)$ -th edge to be part of the tree.

Example 9.16. The top of Figure 5 depicts a graph G with edge ordering e_1, \dots, e_9 . Along the bottom, we illustrate the process just described for associating a spanning tree to the recurrent $c = (1, 3, 3, 2)$.




Figure 5. Dhar's tree bijection via vertex firings (cf. Example 9.16).

After firing the sink, s , the vertices v_1 and v_2 are unstable. There is only one edge connecting s to v_1 ; so it becomes part of the spanning tree we are constructing. On the other hand, both e_2 and e_3 connect s to v_2 , which means we need to make a choice. Since v_2 has 5 grains of sand, which is one more than needed to make it unstable, we choose the second of the two, e_3 . After firing v_1 and v_2 , the vertex v_3 becomes unstable with 5 grains of sand, which is just enough to make it unstable. Therefore, of the two edges, e_4 and e_5 , connecting the vertices v_1 and v_2 to v_3 , we choose the first, e_4 . Then, firing v_3 results in 6 grains of sand on v_4 . That's two more than needed to make v_4 unstable, so e_9 is selected for the spanning tree. Finally, firing v_4 returns us to the original configuration.

Algorithm 7 is an implementation of Dhar's tree bijection that closely follows the original description in [38]. At each stage of the algorithm, the sum of the burnt and unburnt neighbors of a vertex equals its outdegree. This allows the algorithm to go forward by counting unburnt neighbors rather than firing vertices, modifying the starting configuration, and testing for instability, as in Example 9.16.

By Proposition 9.19, below, Algorithms 7 and 8 provide a bijection between recurrents and spanning trees.

Example 9.17. Let G be the graph in Figure 5 with the displayed edge-ordering. Figure 6 illustrates Algorithm 7 applied to the recurrent $c = (0, 3, 4, 1)$ on G .

Exercise 9.18. Again consider the graph G from Figure 5 with the given ordering of edges.

- (1) Apply Dhar's tree bijection algorithm to associate spanning trees with each of the following recurrents: $(1, 1, 4, 2)$, $(1, 0, 4, 2)$, $(1, 1, 4, 0)$, and $(0, 3, 4, 0)$.

Algorithm 7 Dhar's tree bijection algorithm (recurrent \rightarrow spanning tree).

```

1: INPUT:
     $G = (V, E)$  – undirected loopless multigraph with fixed ordering for  $E$ 
     $s \in V$  – sink vertex
     $c \in \mathcal{S}(G)$ 
2: OUTPUT: tree – spanning tree of  $G$ 
3: initialization:
    tree =  $\emptyset$ , unburnt =  $\tilde{V} := V \setminus \{s\}$ , newly_burnt =  $\{s\}$ 
4: while unburnt  $\neq \emptyset$  do
5:   burnable =  $\emptyset$ 
6:   unburnt_nbrs =  $\{u \in \text{unburnt} : u \text{ a neighbor of a vertex in newly\_burnt}\}$ 
7:   for  $u \in \text{unburnt\_nbrs}$  do
8:      $\xi = \#$  edges connecting  $u$  to a vertex in unburnt
9:     if  $c(u) \geq \xi$  then
10:      add  $u$  to burnable
11:       $e = [(c(u) - \xi) + 1]$ -th edge connecting  $u$  to a vertex in newly_burnt
12:      add  $e$  to tree
13:   remove burnable from unburnt
14:   newly_burnt = burnable
15: return tree

```

Algorithm 8 Dhar's tree bijection algorithm (spanning tree \rightarrow recurrent).

```

1: INPUT:
     $G, s$  – as in Algorithm 7
    tree – spanning tree of  $G$ 
2: OUTPUT:  $c$  – recurrent sandpile on  $G$ 
3: initialization:
     $c = 0$ -sandpile, unburnt =  $\tilde{V}$ , newly_burnt =  $\{s\}$ 
4: while unburnt  $\neq \emptyset$  do
5:   burnable =  $\{u \in \text{unburnt} : \exists uv \in \text{tree with } v \in \text{newly\_burnt}\}$ 
6:   for  $u \in \text{burnable}$  do
7:      $\xi = \#$  edges in  $G$  connecting  $u$  to a vertex in unburnt
8:      $F = \{e_1, e_2, \dots\}$  = edges connecting  $u$  to a vertex in newly_burnt, in order
9:      $c(u) = \xi + (\ell - 1)$  for the unique index  $\ell$  such that  $e_\ell \in F \cap \text{tree}$ .
10:  remove burnable from unburnt
11:  newly_burnt = burnable
12: return  $c \in \mathcal{S}(G)$ 

```

(2) Apply Dhar's tree bijection algorithm to associate recurrents with each of the following spanning trees: $\{e_1, e_2, e_5, e_9\}$, $\{e_1, e_2, e_4, e_7\}$, $\{e_3, e_4, e_6, e_8\}$, and $\{e_2, e_4, e_5, e_6\}$.

Proposition 9.19. *Algorithm 7 defines a bijection from recurrent sandpiles to spanning trees with inverse defined by Algorithm 8.*

Proof. For each stable sandpile c on G , let $B_0 = B_0(c) := \{s\}$ and, for $i \geq 1$, let $B_i = B_i(c)$ be the vertices of G that are unstable after firing $B_0 \cup \dots \cup B_{i-1}$ from c . Firing B_0 produces the sandpile $b + c$, where b is the burning sandpile. By




Figure 6. Another example of Dhar’s tree bijection algorithm (cf. Example 9.17).

Theorems 7.5 and 7.6, the B_i are disjoint, and the sandpile c is recurrent if and only if $\cup_{i \geq 0} B_i = V$.

For each tree T of G containing s , let $C_0 = C_0(T) = \{s\}$, and, for $i \geq 1$, let $C_i = C_i(T)$ be vertices of G that are connected by an edge of T to $C_0 \cup \dots \cup C_{i-1}$ but are not contained in the set $C_0 \cup \dots \cup C_{i-1}$. In other words, $C_i(T)$ is the set of vertices at distance i from s in T . We have that T is a spanning tree if and only if $\cup_{i \geq 0} C_i = V$.

Consider Algorithm 7 with input $c \in \mathcal{S}(G)$. The effect of the i -th iteration of the while-loop is to add a unique edge to **tree** for each $u \in B_i$ which connects u to a vertex in B_{i-1} . Thus, each element of B_i has a path in **tree** back to s , and cycles are never formed. Since c is recurrent, $\cup_{i \geq 0} B_i = V$, and hence the algorithm terminates with **tree** a spanning tree of G . Thus, the algorithm determines a mapping $\tau: \mathcal{S}(G) \rightarrow \mathcal{T}(G)$ where $\mathcal{T}(G)$ denotes the set of spanning trees of G .

Next, consider Algorithm 8 with input $T \in \mathcal{T}(G)$. Since T spans G , the algorithm terminates, and its output is a sandpile which we will denote by $\rho(T)$. For each $u \in V$ and subset $W \subseteq V$, let $n(u, W)$ be the number of edges of G connecting u to a point in W . For each $i \geq 1$, let $W_i = \cup_{j=0}^{i-1} C_j(T)$. For $u \in \tilde{V}$, there is a unique $i \geq 1$ such that $u \in C_i(T)$, and

$$\rho(T)(u) = n(u, W_i^c) + (\ell - 1).$$

where $1 \leq \ell \leq n(u, C_{i-1}(T))$. Since $\deg_G(u) = n(u, W_i) + n(u, W_i^c)$, we have

$$\rho(T)(u) = \deg_G(u) + (\ell - n(u, W_i)) - 1,$$

and hence, $\rho(c)$ is a stable sandpile.

We now claim $B_i(\rho(T)) = C_i(T)$ for all i . We have $B_0(\rho(T)) = C_0(T) = \{s\}$. Let $i \geq 1$, and by induction suppose that $B_j(\rho(T)) = C_j(T)$ for $0 \leq j \leq i - 1$. Take $u \in C_i(T)$. Firing $W_i = \cup_{j=0}^{i-1} C_j(T) = \cup_{j=0}^{i-1} B_j(\rho(T))$ adds $n(u, W_i)$ grains of sand to u , and hence, having started at $\rho(T)$, the net amount of sand on u is

$$\begin{aligned} \rho(T)(u) + n(u, W_i) &= n(u, W_i^c) + (\ell - 1) + n(u, W_i) \\ &= \deg_G(u) + (\ell - 1) \\ &\geq \deg_G(u). \end{aligned}$$

Thus, u is unstable. This shows that $B_i(\rho(T)) \supseteq C_i(T)$.

For the opposite inclusion, take $v \in B_i(\rho(T))$. There is a unique k such that $v \in C_k(T)$. If $k \neq i$, then since the $B_j(\rho(T))$ are disjoint and $B_j(\rho(T)) = C_j(T)$

for $0 \leq j \leq i - 1$, it follows that $k > i$. By definition,

$$\rho(T)(v) = n(v, W_k^c) + (m - 1)$$

where $1 \leq m \leq n(v, C_{k-1}(T))$. But then, after firing W_i from $\rho(T)$, the amount of sand on v is

$$\begin{aligned} \rho(T)(v) + n(v, W_i) &= \rho(T)(v) + n(v, W_k) - n(v, W_k \setminus W_i) \\ &= \deg_G(v) + m - 1 - n(v, W_k \setminus W_i). \end{aligned}$$

Since $C_{k-1}(T) \subseteq W_k \setminus W_i$, we have $m \leq n(v, C_{k-1}(T)) \leq n(v, W_k \setminus W_i)$. So it follows that $\rho(T)(v) + n(v, W_i)$ is stable. That's a contradiction since $v \in B_i(\rho(T))$ means that after firing $W_i = \cup_{j=0}^{i-1} B_j(\rho(T))$, the resulting sandpile is unstable at v . Therefore, it must be that $k = i$.

We have shown by induction that $B_i(\rho(T)) = C_i(T)$ for all i . One consequence, since $\cup_{i \geq 0} C_i(T) = V$, is that $\rho(T)$ is recurrent. Therefore, Algorithm 8 defines a mapping $\rho: \mathcal{T}(G) \rightarrow \mathcal{S}(G)$.

Our next goal is to show that τ and ρ are inverses of each other. First, take $c \in \mathcal{S}(G)$, and consider $\rho(\tau(c))$. Directly from the construction of $\tau(c)$, it follows that $B_i(c) = C_i(\tau(c))$ for all $i \geq 0$, i.e., letting $W_i = \cup_{j=0}^{i-1} B_j(c)$, the $C_i(\tau(c))$ are exactly those vertices that become unstable after firing W_i . Let $u \in \hat{V}$ and select the unique $i \geq 1$ for which $u \in B_i(c)$. Let $F = \{e_1, e_2, \dots\}$ be the ordered list of edges connecting u to a vertex in $B_{i-1}(c)$. Then by line 11 of Algorithm 7, we have $e_\ell \in \tau(c)$ where $\ell = c(u) - n(u, W_i^c) + 1$. So by line 9 of Algorithm 8,

$$\rho(\tau(c))(u) = n(u, W_i^c) + (\ell - 1) = c(u).$$

Hence, $\rho(\tau(c)) = c$ for all $c \in \mathcal{S}(G)$.

Now let $T \in \mathcal{T}(G)$, and consider $\tau(\rho(T))$. We saw above that $B_i(\rho(T)) = C_i(T)$ for all $i \geq 0$. Let $e \in T$. Then there exists a unique $i \geq 1$ such that $e = \{u, v\}$ with $u \in C_i(T)$ and $v \in C_{i-1}(T) \subseteq W_i = \cup_{j=0}^{i-1} C_j(T)$. Say e is the ℓ -th edge connecting u to a vertex in $C_{i-1}(T)$. Then

$$\rho(T)(u) = n(u, W_i^c) + (\ell - 1).$$

Now apply Algorithm 7 to $\rho(T)$. We have $u \in B_i(\rho(T))$ and $v \in B_{i-1}(\rho(T))$. So during the i -th iteration of the while-loop, the vertex u is considered as an element of `unburnt_nbrs` with $\xi = n(u, W_i^c)$. Edge $e = e_\ell$ is added to `tree` since at line 11,

$$(\rho(T)(u) - \xi) + 1 = \ell.$$

Hence, $e \in \tau(\rho(T))$ and it follows that $\tau(\rho(T)) \supseteq T$. Since every spanning tree of G has the same cardinality, $\tau(\rho(T)) = T$. □

9.3.2. Tree bijections and external activity. In Part 1, Section 3.4, we presented Dhar's algorithm for testing whether a configuration is superstable. There, a configuration c is imagined as an assignment of $c(v)$ firefighters to each vertex v . The sink vertex is lit, and fire spreads along incident edges. If at any point in the burning process, the number of burning edges incident on a vertex is strictly greater than the number of firefighters there, the firefighters abandon the vertex and the vertex is burnt. Fire can then spread along edges incident to that vertex.

The configuration c is superstable if and only if in the end, each vertex has been burnt.

Just as we saw for recurrent, a more controlled burning process yields a tree bijection algorithm matching superstables with spanning trees. Here, we describe a slight variation on a superstable-tree algorithm from [8]. Again fix an ordering of the edges of G once and for all. Assign firefighters to vertices according to a superstable c , and light the sink. At each step in the algorithm, list the unburnt edges incident on burnt vertices in order: e_1, \dots, e_k . Set each of the edges on fire, one at a time, in *reverse order* (from largest to smallest), until a first edge $e_i = \{u, v\}$ is reached which joins a burnt vertex u to an unburnt vertex v such that the number of already-burnt edges incident on v is equal to $c(v)$. Light e_i (which then overwhelms the firefighters at v) and mark e_i . Now repeat. The algorithm halts when all of the vertices are burnt. At that point the set of marked edges forms a spanning tree. The formal description of the algorithm in terms of pseudocode is left to the reader.




Figure 7. Superstable-tree bijection algorithm (cf. Example 9.20). All red edges are burnt. The dotted red edges are burnt but not part of the spanning tree.

Example 9.20. Figure 7 illustrates the superstable-tree bijection algorithm. We use the notation and edge-ordering displayed in Figure 5. The sink is lit, and fire spreads along the largest incident edge, e_3 , to vertex v_2 . There are no firefighters at v_2 , so v_2 is burnt, and e_3 becomes part of the spanning tree we are constructing. In the next step, we examine the edges incident to one of two burnt vertices and to an unburnt vertex: e_1, e_5 , and e_6 . Of these, $e_6 = \{v_2, v_4\}$ is the largest, so it is burnt. There are two firefighters at v_4 , and thus, v_4 is not burnt. We burn the next smallest edge, $e_5 = \{v_2, v_3\}$, and v_3 is burnt since there are no firefighters protecting it. The edge e_5 is added to the tree. Continue: in the next step, we consider edges e_1, e_4, e_7, e_8 , and e_9 . Burning edges e_9 then e_8 overwhelms the firefighters as v_4 and e_8 is added to the tree. In the last step, we burn edges e_4 and e_1 which burns v_1 , and e_1 is added to the tree. At this point, all vertices are burnt so the algorithm halts. The tree spans G .

External activity. In addition to being a bijection, the algorithm just presented has an extra, intriguing, property. Having fixed an ordering of the edges of G , there is a statistic $e(T) \in \mathbb{N}$ called the *external activity* associated to each spanning tree T of G . An edge e is *externally active* in T if it is not in T and if it is the smallest edge in the unique cycle formed by adding e to T . Then $e(T)$ is the number edges of that are externally active in T . External activity will be discussed at length later, in Section 14.5. We can now state the main theorem:

Theorem 9.21 ([8, 26]). *The superstable-tree algorithm gives a bijection between superstables and spanning trees of G . After running the algorithm on a superstable c , the number of unburnt edges is the external activity of the resulting spanning tree, T , and we have*

$$(9.3) \quad e(T) = g - \deg(c),$$


where $g = |E| - |V| + 1$ is the genus of G .

After running the algorithm on a superstable c , note that each non-sink vertex v accounts for $c(v) + 1$ burnt edges—just the number needed to overwhelm the firefighters at v . This means that the number of burnt edges is $\deg(c) + |\tilde{V}|$. So if we know that the number of unburnt edges is equal to the external activity of the resulting spanning tree, equation (9.3) follows.

Example 9.22. In Example 9.20, the two unburnt edges are e_2 and e_7 . Adding e_2 to the spanning tree, T , gives a cycle with two edges, e_2 and e_3 . Since e_2 is smaller than e_3 in our ordering, it is externally active in T . The story is similar with e_7 , but with no other edges. For instance, adding e_6 to T produces a cycle with tree edges e_5 and e_8 , but e_5 is smaller than e_6 . So e_6 is not externally active.

Exercise 9.23.

- (1) Consider the graph G of Figure 5 with its edge ordering. Let T be the spanning tree of G pictured in red below:



- (a) Which edges are externally active?
 (b) Find the superstable c corresponding to T under our superstable-tree algorithm and verify that equation (9.3) holds in this instance.
- (2) Describe a general algorithm that takes spanning trees to superstables and is inverse to the superstable-tree algorithm presented in this section.

Exercise 9.24. (Recurrent-superstable duality.)

- (1) Describe Algorithm 7 in terms of a controlled burning of edges, analogous to our superstable-tree bijection algorithm.
- (2) A controlled burning of edges can be thought of as a controlled firing of vertices: when an edge burns, a grain of sand is deposited on its newly burnt vertex. Consider Algorithm 7 when it reaches line 9. The number of unburnt vertices adjacent to u is ξ . Let β be the number of burnt adjacent vertices. Then $c(u) \geq \xi$ if and only if

$$c(u) + \beta \geq \beta + \xi = \deg_G(u).$$

Thus, u is an unstable vertex in c after it receives β additional grains of sand, i.e., after β adjacent edges burn. Recall the duality between recurrents and

superstables given by $c \mapsto c_{\max} - c$ (Theorem 7.12). Show that $c(u) \geq \xi$ if and only if $(c_{\max} - c)(u) > \beta$ (i.e., exactly when the firefighters at u are overwhelmed).

- (3) Modify Algorithm 7 so that the tree T it matches to a recurrent c is the same as that matched to the superstable $c_{\max} - c$ returned by the superstable-tree algorithm, and then state a version of Theorem 9.21 (including equation (9.3)) in terms of recurrences.

9.3.3. Tree bijections and tree inversions. The basis for our tree bijection algorithms has been Theorem 7.5: starting with a recurrent sandpile c and firing the sink, the resulting configuration stabilizes back to c , and in the process, each vertex fires exactly once. Our controlled burning of vertices and edges is an expression of these vertex firings. In the previous section, the rule for burning edges produced a bijection relating two statistics: the degree of a superstable and the external activity of its corresponding tree. In this section, we modify the burning rule so that the resulting bijection involves a different tree statistic called the κ -inversion number (Definition 9.29). The results described in this section appear in [78], to which the reader is referred for details and proofs.

In this section G will be a simple, connected, undirected graph with vertex set $V = \{0, \dots, n\}$. Choose any $s \in V$ as the sink vertex, and let $\tilde{V} := V \setminus \{s\}$. (In all of our examples, we will take $s = 0$.) To describe our new tree-bijection, take a superstable configuration c and again think of it as an assignment of firefighters to vertices. We again light the sink vertex and proceed by burning vertices and edges in a controlled manner. This time, though, the spreading is determined by a *depth-first search* of the vertices of G . In detail: at the beginning of each step of the algorithm, there will be a currently *active* burnt vertex i (initially $i = s$). To find the next edge to burn, find the maximal (in numerical order) vertex j such that: (i) j is unburnt, and (ii) $\{i, j\}$ is an unburnt edge. (We consider the case there is no such j below.) Burn the edge $e := \{i, j\}$. If the number of burnt edges incident on j is now greater than $c(j)$, then the firefighters at j are overwhelmed and abandon the vertex. In that case, e is added to the tree we are constructing, and j is burnt, becoming the active vertex for the next step in the algorithm. If not, the next step of the algorithm proceeds with i again the active vertex.

If there are no vertices j adjacent to i meeting the two criteria specified above, the algorithm proceeds by recursively backtracking: the vertex $i' \neq i$ that was active just before i became active is set as the active vertex, and the algorithm proceeds as before. (The vertex i' will be the unique vertex adjacent to i in the tree built so far.)

In any event, the algorithm halts as soon as all vertices are burnt, returning a spanning tree of G . A precise description is provided in Algorithm 9, displayed below.

Example 9.25. Figure 8 considers the depth-first search algorithm for the superstable $c = (1, 0, 2, 1)$ on the complete graph K_5 . The graph and its vertex labels appear in the bottom right of the figure. The sink vertex, 0, is lit and its neighbors are probed in reverse numerical order. Firefighters protect vertices 4 and 3, but there are none at vertex 2. So the edges connecting 0 to these vertices are

Algorithm 9 Depth-first search burning algorithm.

```

1: INPUT:
     $G = (V, E)$  – simple undirected graph with  $V = \{0, \dots, n\}$ 
     $s \in V$  – sink vertex
     $c \in \mathbb{N}^{\tilde{V}}$  – sandpile on  $G$  with respect to  $s$ 
2: OUTPUT: tree – tree of  $G$ , a spanning tree iff  $c$  is superstable
3: initialization:
    burnt_vertices =  $\{s\}$ ,  burnt_edges =  $\emptyset$ ,  tree =  $\emptyset$ 
4: execute DFS_FROM( $s$ )
5: return tree

DFS_FROM
6: function DFS_FROM( $i$ )
7:   for all  $j$  adjacent to  $i$  in  $G$ , from largest numerical value to smallest do
8:     if  $j \notin$  burnt_vertices then
9:       if  $c(j) = 0$  then
10:        append  $j$  to burnt_vertices
11:        append  $(i, j)$  to tree
12:        DFS_FROM( $j$ )
13:       else
14:         $c(j) = c(j) - 1$ 
15:        append  $(i, j)$  to burnt_edges

```

burnt. The vertex 2 is burnt and the edge $\{0, 2\}$ becomes part of the tree. Vertex 2 becomes the active vertex, and the algorithm continues.

The largest unburnt neighbor of 2 is vertex 4, and the single firefighter at vertex 4 is already occupied with the burnt edge $\{0, 4\}$. So that firefighter is overwhelmed. The edge $\{2, 4\}$ is burnt and added to the tree. The vertex 4 is newly burnt and becomes the active vertex. The unburnt edges joining 4 to unburnt





Figure 8. The depth-first search algorithm for the superstable $c = (1, 0, 2, 1)$. See Example 9.25 for details. The red rectangle marks the active vertex, all red edges are burnt, and solid red edges are tree-edges.

neighbors are $\{3, 4\}$ and $\{1, 4\}$. There are sufficient firefighters to protect vertices 1 and 3. So these edges are burnt but not vertices 1 and 3. We then backtrack: vertex 2 again becomes the active vertex, and the algorithm continues from there to completion. In the figure, the edges of the resulting spanning tree are shown in solid red. Dotted red edges are edges that were burnt but did not become part of the tree.

Tree inversions. Let T be a tree with vertices $V = \{0, \dots, n\}$, and pick a root/sink vertex $s \in V$. If $i, j \in V$ and i is on the unique path from j to s , then i is an *ancestor* of j , and j is a *descendant* of i . If i is an ancestor of j and $\{i, j\}$ is an edge of T , we say i is the *parent* of j and j is a child of i . Each non-root vertex of T has a unique parent, but vertices may have many children.

Definition 9.26. An *inversion* of the rooted tree T is an ordered pair (i, j) of vertices such that: (i) $i > j$, (ii) i is not the root vertex², and (iii) i is an ancestor of j . The number of inversions of T is the *inversion number* for T .

Example 9.27. In Figure 8, the depth-first burning algorithm produces the following tree rooted at vertex 0:



This tree has two inversions, $(2, 1)$ and $(3, 1)$, so its inversion number is 2.

Exercise 9.28.

- (1) Find all sixteen trees with vertices $\{0, 1, 2, 3\}$, or equivalently, the sixteen spanning trees of the complete graph K_4 on this vertex set.
- (2) Describe the sixteen superstable sandpiles on K_4 .
- (3) Let τ_k denote the number of these trees with inversion number k , and let h_k denote the number of superstables on K_4 with degree k . Verify the following table:

k	0	1	2	3	
τ_k	6	6	3	1	.
h_k	1	3	6	6	

Note that the inversion and degree counts in the table in Exercise 9.28 are the same but in reverse order. This is indicative of a general phenomenon first proved by Kreweras ([64]): Let $g = n(n-1)/2$, the genus of the complete graph K_{n+1} . Let τ_k denote the number of spanning trees on $V = \{0, \dots, n\}$ with inversion number k , and let h_k be the number of superstables on K_{n+1} of degree k . Then,

$$(9.4) \quad \tau_k = h_{g-k}$$

for $0 \leq k \leq g$.

²We will usually choose 0 as the root vertex, in which case condition (i) implies condition (ii).

To try to generalize Kreweras' formula, fix an arbitrary graph G of the type considered in this section. Let τ_k be its number of spanning trees with inversion number k , let h_k be the number of superstable of G with degree k , and let $g = |E| - |V| + 1$ be its genus. It turns out that it is *sometimes* the case that equation (9.4) continues to hold (cf. Problem 9.9), but not always. For a case where it does not hold consider the house graph as pictured in Figure 9 with its 11 spanning trees. The genus of the house graph is 2. The inversion numbers and degree counts for this graph are:

k	0	1	2	3	
τ_k	4	3	3	1	.
h_k	1	4	6	0	

The problem is that our notion of an inversion of a spanning tree does not take into account the structure of G . The appropriate generalization is:

Definition 9.29. Let G be a graph (simple, connected, undirected, with the vertex set $\{0, \dots, n\}$ and fixed root/sink vertex s), and let T be a spanning tree rooted at s . An inversion (i, j) of T is a κ -inversion if the parent of i is adjacent to j in G . The κ -inversion number, $\kappa(T) = \kappa(G, T)$, is the number of κ -inversions of T .




Figure 9. The 11 superstable sandpiles on the house graph, G , with their κ -inversion numbers and corresponding spanning trees (via the depth-first search burning algorithm).

Exercise 9.30. Verify the κ inversion numbers in Figure 9. For instance, the inversion $(3, 2)$ of the spanning tree $\overset{0}{\cdot} \overset{1}{\cdot} \overset{3}{\cdot} \overset{2}{\cdot} \overset{4}{\cdot}$ is not a κ -inversion since the parent of 3 in the tree is 1, and $\{1, 2\}$ is not an edge of G . Hence, the inversion number of T is 1, but $\kappa(T) = 0$.


We can now state the main theorem of this section relating the depth-first search bijection to κ -inversions in the same way that Theorem 9.21 relates the superstable-tree bijection of the previous section to external activity:

Theorem 9.31 ([78]). *The depth-first search algorithm, Algorithm 9, is a bijection between superstables and spanning trees of G . If T is the spanning tree corresponding to a superstable c , then*

$$\kappa(T) = g - \deg(c).$$

Exercise 9.32.

- (1) Pick any superstable on the house graph G of Figure 9. Use the depth-first search algorithm to find its corresponding tree, and verify that the formula in Theorem 9.31 holds.
- (2) Let W be the following graph with spanning tree T in blue:



Let vertex 0 be the root/sink. Find the superstable c that is in bijection with T via the depth-first search algorithm, and verify the formula in Theorem 9.31 holds in this case.

9.3.4. Rotor-routers. Let $G = (V, E)$ be a sandpile graph, possibly directed, with sink s . As usual, we may consider an undirected edge as a pair of oppositely oriented edges. We endow G with some extra structure: at each non-sink vertex v , fix a cyclic ordering³ of the edges e with tail $e^- = v$. Now imagine a rotor mechanism attached to v pointing along one of these edges (see the blue arrow in Figure 10). The rotor can be activated if there is sand on v , in which case it first spins to point along the next edge in the cyclic ordering, then sends a grain of sand along that edge. If the edge points to the sink, the grain of sand is lost. The collection of rotors (one for each non-sink vertex v) forms a “machine” on G called a *rotor-router*. It is a tool for routing sand along the edges of the graph. A couple of natural questions: Is it always possible to route all of the sand into the sink? If the answer is ‘yes’, then will every activation-order of the rotors eventually direct all of the sand into the sink?

³In the case of a non-simple graph, if there are d distinct edges directed from v to w , then each must occur separately in the cyclic ordering, but they need not be consecutive.

If the rotor at a particular vertex rotates through a complete cycle, back to its original position, the corresponding redistribution of sand exactly simulates a vertex firing in the sandpile model. The results stated below will expand upon this connection between the rotor-router model and the sandpile model.

To discuss the rotor-router model more precisely, define a *rotor configuration* to be a function $\rho: \tilde{V} \rightarrow E$ such that $\rho(v)$ is an edge emanating from v . Denote the edge following $\rho(v)$ in the cyclic ordering at v by $\rho(v)_{\text{next}}$. A *state* of the rotor-router is a pair (c, ρ) consisting of a sandpile and a rotor configuration. A *rotor firing* at $v \in \tilde{V}$ on a state (c, ρ) such that $c(v) > 0$ produces a new state as follows:

- (1) Rotate the rotor, replacing $\rho(v)$ by $e := \rho(v)_{\text{next}}$.
- (2) Send a grain of sand along e , replacing c by $c - v + e^+$.

So $(c, \rho) \xrightarrow{v} (c - v + \rho(v)_{\text{next}}^+, \rho(v)_{\text{next}})$. A rotor firing is depicted in Figure 10.




Figure 10. A rotor firing at the vertex v .

A rotor-router state is *stabilized* by performing operations until all of the sand has been routed to the sink. We now state several results for the rotor-router model without proof; for details, see [56].

Theorem 9.33. *Let (c, ρ) be a rotor-router state.*

- (1) *The state (c, ρ) has a stabilization: by repeatedly applying rotor-router operations, all of the sand is eventually routed to the sink.*
- (2) *The stabilization of (c, ρ) is independent of the order of rotor firings, as is the number of times each rotor is fired at each vertex in the stabilization process.*

Let $c(\rho)$ denote the rotor configuration for the stabilization of the rotor-router state (c, ρ) . The next proposition shows that the sandpile group $\mathcal{S}(G)$ acts freely on the set of rotor-router configurations.

Proposition 9.34. *Let a and b be sandpiles, and let ρ be a rotor configuration. Then,*

- (1) $(a + b)(\rho) = a(b(\rho))$,
- (2) $a(\rho) = b(\rho)$ if and only if $a = b \bmod \tilde{\mathcal{L}}$, where $\tilde{\mathcal{L}}$ is the reduced Laplacian lattice for G .

A rotor configuration ρ is *acyclic* if the set of edges $\{\rho(v) : v \in \tilde{V}\}$ has no directed cycles, i.e., if the rotors form a directed spanning tree into the sink.

Proposition 9.35. *A rotor configuration ρ is acyclic if and only if it is reachable from every rotor configuration: for any rotor configuration ρ' there exists a sandpile c such that $c(\rho') = \rho$. We say that the acyclic rotors are the recurrent rotors under the action of the sandpile group.*

These two propositions immediately imply:

Theorem 9.36. *Let \mathcal{A} denote the set of acyclic rotor configurations (i.e., the set of directed spanning trees into s on G). Then the sandpile group for G acts freely and transitively on \mathcal{A} by the action*

$$\begin{aligned} \mathcal{S}(G) \times \mathcal{A} &\rightarrow \mathcal{A} \\ (c, \rho) &\mapsto c(\rho). \end{aligned}$$

Figure 11 illustrates what happens when a recurrent sandpile c is placed on a directed spanning tree and then all of the sand is rotor-routed to the sink. In light of Proposition 9.34, the resulting rotor configuration would be the same for any sandpile equal to c modulo $\tilde{\mathcal{L}}$, i.e., equal up to vertex firings or reverse firings.

Theorem 9.36 implies that after having made an arbitrary choice of tree to serve as the identity element, the set of directed spanning trees of G naturally forms a group isomorphic to the sandpile group. Indeed, if we fix a directed spanning tree for G and think of it as an acyclic rotor configuration ρ_0 , then the mapping

$$\begin{aligned} \mathcal{S}(G) &\rightarrow \mathcal{A} \\ c &\mapsto c(\rho_0) \end{aligned}$$

is a bijection endowing \mathcal{A} with the group structure of the sandpile group. In this way, the set of spanning trees is a group that has forgotten its identity element. Using language introduced in Section 3.5, the set of spanning trees is a *torsor* for the sandpile group, just as the sets $\text{Pic}^d(G)$ are torsors for the Jacobian group as described in Part 1.

9.3.5. Random trees. A bijection between $\mathcal{S}(G)$ and the set of spanning trees of G provides a method for choosing a random spanning tree ([8]). Letting s_1, \dots, s_k be the invariant factors of $\mathcal{S}(G)$ and $\tilde{\mathcal{L}}$ the image of the reduced Laplacian, \tilde{L} , the idea is to use the string of isomorphisms:

$$\prod_{i=1}^k \mathbb{Z}_{s_i} \xrightarrow{\sim} \mathbb{Z}^{n-1} / \tilde{\mathcal{L}} \xrightarrow{\sim} \mathcal{S}(G).$$

The first isomorphism comes from the computation of the Smith normal form of \tilde{L} , and the second maps an element to the unique equivalent recurrent modulo $\tilde{\mathcal{L}}$. Choosing a random element of $\prod_{i=1}^k \mathbb{Z}_{s_i}$ (easy) and using these isomorphisms produces an element of $\mathcal{S}(G)$, and then the tree bijection determines a tree.

In detail, use the methods of Section 2.4.2 to find matrices P and Q , invertible over the integers, such that $P\tilde{L}Q = D$, where D is a diagonal matrix with the invariant factors of \tilde{L} along its diagonal. Let j_1, \dots, j_k be the indices for the columns of D containing the invariant factors for $\mathcal{S}(G)$ (the diagonal entries of \tilde{L} that are




Figure 11. Rotor-routing. Edges at each vertex are cyclically ordered counterclockwise.

greater than 1). Let R be the $(n - 1) \times k$ matrix consisting of columns j_1, \dots, j_k of P^{-1} . Then

$$\begin{aligned} \phi: \prod_{i=1}^k \mathbb{Z}_{s_i} &\rightarrow \mathbb{Z}^{n-1} / \tilde{\mathcal{L}} \\ (a_1, \dots, a_k) &\mapsto R(a_1, \dots, a_k)^t. \end{aligned}$$

is a well-defined isomorphism (cf. Problem 9.12). In particular, the columns of R define generators for $\mathbb{Z}^{n-1} / \tilde{\mathcal{L}}$. Let $c_1, \dots, c_k \in \mathcal{S}(G)$ be the recurrents equivalent to these generators (cf. Theorem 6.28).

To generate a random tree: (i) pick random integers $a_j \in \{0, \dots, s_j - 1\}$ for $1 \leq j \leq k$, and (ii) define the sandpile $c := \sum_{j=1}^k a_j c_j$, (iii) stabilize c to get $c^\circ \in \mathcal{S}(G)$, and finally, (iv) use a fixed tree bijection to find the tree corresponding to c° .

In practice, it is probably easier to work with superstable sandpiles rather than recurrent. In that case, we take c_1, \dots, c_k to be superstable representatives for the elements in $\mathbb{Z}^{n-1}/\tilde{\mathcal{L}}$ corresponding to the columns of R . If G is undirected, one might use Algorithm 3 for a fast implementation of this step. Then, again in the undirected case, use any form of Dhar's bijection between superstables and spanning trees. For particulars, including a runtime analysis, see [8].

Dhar ([38]) suggested another way in which the sandpile model may be used to generate random spanning trees: In Chapter 8 we considered the abelian sandpile model as a Markov chain. The chain evolves from the current state—a recurrent sandpile—by dropping a grain of sand on a random vertex and stabilizing. We saw that the stationary distribution for the chain is the uniform distribution on the set of recurrents (Corollary 8.28). Combining this method of generating a chain of recurrents with a tree bijection yields a method for sampling from the set of spanning trees.

Notes

The first proof of the directed version of the matrix-tree theorem is due to Tutte in [89] using a deletion-contraction argument. The proof presented here using a sign reversing transposition comes from the beautiful paper by Zeilberger, *A combinatorial approach to matrix algebra* [95]. Proposition 9.12 is due to Temperley in [88]. Each of the tree-bijections in Section 9.3 also constitute proofs of a version of the matrix-tree theorem. For more on the history of this theorem, see [86], p. 68, and for a generalization to higher dimensions, see Theorem 15.24.

The free transitive action of the sandpile group on the set of spanning trees rooted at the sink via rotor-routing appears in the paper ([56], 2008) by Holroyd et al. Ellenberg ([43], 2011) asked whether this action is independent of the chosen sink, and if not, whether there a nice class of graphs for which it is. That question was answered, for the case of undirected graphs, by Chan et al. ([23], 2015): the action is independent of the sink precisely in the case of planar graphs. By Corollary 13.23, the sandpile group of a planar graph embedded in the plane is naturally isomorphic to its dual graph. And it turns out that the rotor-router action on spanning trees is compatible with this duality ([24, 9]).

Another free transitive action of $\mathcal{S}(G)$ on the set of spanning trees comes from the theory of “break divisors”. Using the language of Part 1 of this book, there is a natural free transitive action of $\text{Pic}^0(G)$ on $\text{Pic}^g(G)$ via $(D, D') \mapsto D + D'$. Baker et al. ([9, 94]) have described canonical bijections between $\text{Pic}^g(G)$ and the set of spanning trees of G where g is the genus of G . The action of $\mathcal{S}(G)$ then comes from the isomorphism $\mathcal{S}(G) \approx \text{Jac}(G) = \text{Pic}^0(G)$.

Problems for Chapter 9

9.1. Prove that an undirected sandpile graph with sandpile group isomorphic to \mathbb{Z}_2 must have a repeated edge (i.e., an edge of multiplicity greater than one.) Describe all such graphs.

9.2. Find all undirected, loopless graphs with Jacobian group isomorphic to \mathbb{Z}_4 . (Do not count graphs obtained from others by attaching a tree at a vertex. In other words, only consider graphs with no vertices of degree 1.) Why are you sure that you haven't missed any such graphs?

9.3. Let M be an $n \times n$ matrix, and suppose the row vectors of M sum to the zero vector. Let $M^{(ij)}$ be the matrix obtained by removing the i -th row and j -th column from M . Then

$$(-1)^{i+j} \det M^{(ij)} = \det M^{(jj)}.$$

9.4. Let $G = (V, E)$ be an Eulerian graph (Definition A.28). Use the matrix-tree theorem, Problem 9.3, and Theorem 9.14 to show that the number of spanning trees rooted at $v \in V$ is independent of v and thus that the kernel of the Laplacian of G is $\mathbb{Z} \cdot \vec{1}$. (Hint: Consider the row and column sums of the Laplacian of G . In fact, by Corollary 12.3, the sandpile group of an Eulerian graph is, up to isomorphism of groups, independent of the choice of sink, but do not use that result here.)


9.5.

(a) Show by induction the following formula for an $m \times m$ determinant:

$$\det \begin{pmatrix} x & y & y & \cdots & y \\ y & x & y & \cdots & y \\ \vdots & & \ddots & & \vdots \\ y & y & y & \cdots & x \end{pmatrix} = (x - y)^{m-1} (x + (m - 1)y).$$

(b) Use the matrix-tree theorem to prove Cayley's formula: the number of trees on n labeled vertices is n^{n-2} . (Note that "tree" in this case means a spanning tree of K_n , the complete (undirected) graph on n vertices. The labels are mentioned to distinguish between isomorphic trees, i.e., trees isomorphic as graphs.)

9.6. Find all directed spanning trees into s in the following graph, checking for agreement with the matrix-tree theorem. Note that two of the edges have weight 2.



9.7. Let G be an Eulerian multigraph (Definition A.28) with n vertices, and let L be its Laplacian matrix. Let J be the $n \times n$ matrix whose entries are all 1s. Prove Proposition 9.12 using the hints below.

Let M be any $n \times n$ matrix, and let $M^{(ij)}$ be the matrix obtained from M by removing its i -th row and j -th column. The *adjugate* of M in the $n \times n$ matrix, $\text{adj}(M)$, defined by $\text{adj}(M)_{ij} := (-1)^{i+j} \det M^{(ji)}$; so the adjugate is the transpose of the cofactor matrix. Some well-known properties of the adjugate are: (i) $M \text{adj}(M) = (\det M)I_n$, and (ii) if N is another $n \times n$ matrix, $\text{adj}(MN) = \text{adj}(N) \text{adj}(M)$.

(a) Prove that

$$(nI_n - J)(L + J) = nL.$$


(b) Prove that $\text{adj}(nI - J) = n^{n-2}J$.

(c) Prove that $n^{-2} \det(L + J)$ is the number of spanning trees of G .

(d) Show that, in general, $n^{-2} \det(L + J)$ does not count directed spanning trees in the case of a directed multigraph.

9.8. Let $K_{m,n}$ be the *complete bipartite graph*. Its vertex set is the disjoint union of two sets, U and V , of sizes m and n , respectively. The edge set consists of all pairs $\{u, v\}$ such that $u \in U$ and $v \in V$. Show that the number of spanning trees of $K_{m,n}$ is $n^{m-1}m^{n-1}$.

9.9. Consider the diamond graph G with sink vertex 0:




(a) Find the spanning trees corresponding to each of the 8 superstable of G under the superstable-tree bijection presented in Section 9.3.2. Verify that Theorem 9.21 holds in each case.

(b) Find the spanning trees corresponding to each of the 8 superstable of G under the depth-first search algorithm (Algorithm 9). Verify that Theorem 9.31 holds in each case. In addition, show that the inversion number for each tree is equal to its κ -inversion number. (Thus, there is no difference between inversions and κ -inversions for G . This turns out to be the case for any *threshold graph*, of which G is an example (cf. [78])).

9.10. Figure 11 starts with the configuration $(2, 1, 0)$ and an initial tree. By redrawing Figure 11 starting, instead, with the identity configuration, show that the final tree is the same as the initial tree.


9.11. Consider the following graph G with sink s :



- Find all recurrent configurations on G , indicating the identity configuration.
- Use the matrix-tree theorem to determine the number of directed spanning trees directed into the sink.
- Choose a generator c for the sandpile group of G . The rotor-router action, $T \mapsto c(T)$, permutes the spanning trees of G . Describe this permutation by drawing a directed graph with the spanning trees of G as vertices and edges $(T, c(T))$.

9.12.

- Prove that the mapping ϕ described in Section 9.3.5 is a well-defined isomorphism.
- Let G be the graph pictured below:



- Fixing v_0 as the sink, let \tilde{L} be the reduced Laplacian of G . Compute matrices P and Q , invertible over the integers, such that $P\tilde{L}Q = D$ where $D = \text{diag}(1, 1, 2, 10)$ is the Smith normal form of \tilde{L} .
- Describe the mapping $\phi: \mathbb{Z}_2 \times \mathbb{Z}_{10} \rightarrow \mathbb{Z}^4 / \tilde{\mathcal{L}}$ defined in Section 9.3.5. (Your answer will depend on P , which is not uniquely defined.)
- Find generators for $\mathcal{S}(G)$ by identifying recurrents equivalent to $\phi(1, 0)$ and $\phi(0, 1)$.
- Using the version of Dhar's tree bijection implemented by Algorithm 7, find the spanning tree associated with $(1, 2) \in \mathbb{Z}_2 \times \mathbb{Z}_{10}$ using the method described in Section 9.3.5. For the sake of the algorithm, use lexicographic edge ordering: $v_0v_1, v_0v_2, v_0v_3, v_0v_4, v_1v_2, v_1v_3, v_1v_4$.

Harmonic morphisms

In Part 1 we considered the dollar game and related algebraic structures on a fixed graph $G = (V, E)$. In this chapter we study mappings between graphs, and we identify a special class of mappings (called *harmonic*) that behave well with respect to the formation of the Picard and Jacobian groups. In Section 10.2 we continue the analogy with Riemann surfaces described in Section 5.3 by viewing harmonic mappings between graphs as discrete versions of holomorphic mappings between surfaces. Finally, in Section 10.3 we provide an interpretation of harmonic mappings in terms of *household-solutions* to the dollar game. As in Part 1, by a graph we mean a finite, connected, undirected multigraph without loop edges.

10.1. Morphisms between graphs

We begin by defining the general notion of a *morphism* between graphs (see Figure 1).

Definition 10.1. Suppose that $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ are graphs. A *graph morphism*, $\phi: G_1 \rightarrow G_2$, is a set-function $\phi: V_1 \cup E_1 \rightarrow V_2 \cup E_2$ taking vertices to vertices and preserving incidence (although perhaps contracting edges to vertices):

$$\begin{aligned} \phi(v) &\in V_2 \quad \text{for all vertices } v \in V_1, \\ \phi(e) &= \begin{cases} \phi(v)\phi(w) \in E_2 & \text{if } e = vw \in E_1 \text{ and } \phi(v) \neq \phi(w) \\ \phi(v) \in V_2 & \text{if } e = vw \in E_1 \text{ and } \phi(v) = \phi(w). \end{cases} \end{aligned}$$

Any graph morphism $\phi: G_1 \rightarrow G_2$ induces a homomorphism $\phi_*: \text{Div}(G_1) \rightarrow \text{Div}(G_2)$ on the corresponding groups of divisors in the natural way:

$$\phi_* \left(\sum a_v v \right) = \sum a_v \phi(v).$$

Note that ϕ_* preserves the degree of divisors: $\deg(\phi_*(D)) = \deg(D)$. We would like this homomorphism to descend to a homomorphism between Picard and Jacobian groups. Since $\text{Pic}(G) = \text{Div}(G)/\text{Prin}(G)$, this will happen exactly when ϕ_* sends




Figure 1. A graph morphism given by vertical projection. The edges incident to v_1 are sent to the edge h_1h_2 , the edges incident to v_3 are sent to h_2h_3 , and the edge v_2v_4 is contracted to the vertex h_2 .

principal divisors on G_1 to principal divisors on G_2 . As we will show in Proposition 10.4, the *harmonic* condition defined below guarantees that principal divisors are preserved in this fashion.

Definition 10.2. A graph morphism $\phi: G_1 \rightarrow G_2$ is *harmonic* at a vertex $v \in V_1$ if the following quantity is independent of the choice of edge $e' \in E_2$ incident to $\phi(v) \in V_2$:

$$\begin{aligned} m_v &:= |\{e = vw \in E_1 : \phi(e) = e'\}| \\ &= \# \text{ of pre-images of } e' \text{ incident to } v. \end{aligned}$$

The quantity $m_v \geq 0$ is called the *horizontal multiplicity* of ϕ at v . We say that ϕ is *harmonic* if it is harmonic at all vertices $v \in V_1$.

Figure 2 shows examples of harmonic and non-harmonic morphisms.

Remark 10.3. Consider the harmonic morphism $\phi: G_1 \rightarrow G_2$ pictured on the right in Figure 2. To create the mapping in steps, first imagine stacking four horizontal, disjoint copies of the path graph G_2 parallel above G_2 . Pinch together the left-hand endpoints of each of the copies to form the vertex v_1 . Do the same to the right-hand endpoints to form v_3 . Next, take the bottom three copies of G_1 , and pinch together their middle vertices to form v_3 . Finally, add the vertical edge joining v_2 and v_4 . This description generalizes. Nonconstant harmonic mappings may be formed by (1) creating a graph G_1 from several disjoint copies of G_2 by pinching together subsets of corresponding vertices and adding vertical edges, (2) projecting G_1 down to G_2 by sending vertices of G_1 down to their counterparts in G_2 . Figure 4 indicates a further possibility in which an added vertical edge is subdivided. In fact, one could replace a single vertical edge joining vertices u and v with any graph containing vertices u and v and declaring all of the edges of that graph to be vertical for the sake of the mapping.

Proposition 10.4. Suppose that $\phi: G_1 \rightarrow G_2$ is a harmonic morphism, and that $\text{div}(f)$ is a principal divisor on G_1 . Then $\phi_*(\text{div}(f))$ is a principal divisor on G_2 . It follows that if $D \sim D'$ are linearly equivalent divisors on G_1 , then their images




Figure 2. The morphism on the left is not harmonic at v_4 since the edge h_1h_2 has three preimages incident to v_4 while h_2h_3 has only one. The morphism on the right is harmonic, with horizontal multiplicities $m_{v_1} = m_{v_3} = 4, m_{v_2} = 1,$ and $m_{v_4} = 3.$

$\phi_*(D) \sim \phi_*(D')$ are linearly equivalent on G_2 . Hence, ϕ induces homomorphisms of the Picard and Jacobian groups:

$$\phi_* : \text{Pic}(G_1) \rightarrow \text{Pic}(G_2) \quad \text{and} \quad \phi_* : \text{Jac}(G_1) \rightarrow \text{Jac}(G_2).$$

Proof. For any vertex $v \in V_1$, set $h = \phi(v)$ and apply the homomorphism ϕ_* to the principal divisor obtained by firing v :

$$\begin{aligned} \phi_*(\text{div}(\chi_v)) &= \phi_* \left(\sum_{vw \in E_1} (v - w) \right) \\ &= \sum_{vw \in E_1} (\phi(v) - \phi(w)) \\ &= \sum_{k \in V_2} \left(\sum_{vw \in E_1: \phi(w)=k} (h - k) \right) \\ &= \sum_{hk \in E_2} m_v (h - k) \\ &= \text{div}(m_v \chi_h). \end{aligned}$$

Since the principal divisors $\text{div}(\chi_v)$ generate the subgroup $\text{Prin}(G_1) \subset \text{Div}(G_1)$, it follows that $\phi_*(\text{Prin}(G_1)) \subset \text{Prin}(G_2)$. Two divisors are linearly equivalent if and only if their difference is principal, so ϕ_* preserves linear equivalence. \square

To explain the name *harmonic* for this special class of morphisms, recall the following notion from Definition 3.24: a function $f: V \rightarrow A$ from the vertices of a graph G to an abelian group A is called *harmonic* if its value at each vertex v is equal to the average of its values at the neighbors of v :


$$\text{deg}_G(v)f(v) = \sum_{vw \in E} f(w).$$

As shown in Proposition 3.27, the Abel-Jacobi mapping $S_g: V \rightarrow \text{Jac}(G)$ is universal for harmonic functions in the sense that the set of harmonic functions $f: V \rightarrow A$

such that $f(q) = 0$ is in bijection with the set of homomorphisms $\rho: \text{Jac}(G) \rightarrow A$ via the mapping $\rho \mapsto \rho \circ S_q$.

Proposition 10.5. *Let $\phi: G_1 \rightarrow G_2$ be a morphism of graphs, and consider the induced homomorphism $\phi_*: \text{Div}(G_1) \rightarrow \text{Div}(G_2)$ on divisor groups. Then ϕ_* induces a homomorphism on Jacobian groups if and only if the morphism ϕ pulls back harmonic functions on G_2 to harmonic functions on G_1 : for all abelian groups A , if $f: V_2 \rightarrow A$ is harmonic, then $f \circ \phi: V_1 \rightarrow A$ is also harmonic.*


Proof. First suppose that ϕ_* induces a homomorphism on Jacobian groups. The diagram below illustrates the network of mappings involved in the ensuing argument.



Pick a vertex $q_1 \in V_1$ and set $q_2 = \phi(q_1) \in V_2$. Suppose that $f: V_2 \rightarrow A$ is harmonic on G_2 ; we wish to show that $\tilde{f} := f \circ \phi: V_1 \rightarrow A$ is harmonic on G_1 . Since constant functions are harmonic, we may replace f by $f - f(q_2)$ and \tilde{f} by $\tilde{f} - \tilde{f}(q_1)$ and thus assume that $f(q_2) = \tilde{f}(q_1) = 0$. By Proposition 3.27, there exists a homomorphism $\rho: \text{Jac}(G_2) \rightarrow A$ such that $f = \rho \circ S_{q_2}$. Consider the homomorphism $\tilde{\rho} := \rho \circ \phi_*: \text{Jac}(G_1) \rightarrow A$. By Exercise 3.26, the function $\tilde{\rho} \circ S_{q_1}$ is harmonic on G_1 . But this function is just \tilde{f} :

$$\begin{aligned}
 (\tilde{\rho} \circ S_{q_1})(v) &= \rho(\phi_*([v - q_1])) \\
 &= \rho([\phi(v) - q_2]) \\
 &= (\rho \circ S_{q_2})(\phi(v)) \\
 &= (f \circ \phi)(v) \\
 &= \tilde{f}(v).
 \end{aligned}$$

Conversely, suppose that ϕ pulls back harmonic functions on G_2 to harmonic functions on G_1 , and consider the following commutative diagram:



The function $S_{q_2} \circ \phi: V_1 \rightarrow \text{Jac}(G_2)$ is harmonic on G_1 , being the pullback of the Abel-Jacobi map S_{q_2} on G_2 . By Proposition 3.27, there exists a homomorphism $\tau: \text{Jac}(G_1) \rightarrow \text{Jac}(G_2)$ such that $\tau \circ S_{q_1} = S_{q_2} \circ \phi$. Hence, for any vertex $v \in V_1$, we have

$$[\phi(v) - \phi(q_1)] = [\phi(v) - q_2] = S_{q_2}(\phi(v)) = \tau(S_{q_1}(v)) = \tau([v - q_1]).$$

It follows that the equivalence class of $\phi_*(v - q_1) = \phi(v) - \phi(q_1)$ in $\text{Jac}(G_2)$ depends only on the equivalence class of $v - q_1$ in $\text{Jac}(G_1)$, which implies that $\phi_*: \text{Jac}(G_1) \rightarrow \text{Jac}(G_2)$ is well-defined. \square

Corollary 10.6. *If $\phi: G_1 \rightarrow G_2$ is a harmonic morphism, then ϕ pulls back harmonic functions on G_2 to harmonic functions on G_1*

Proof. Since ϕ is harmonic, it induces a homomorphism ϕ_* on Jacobian groups by Proposition 10.4. \square

As the next exercise shows, the harmonic property is a stronger condition on graph morphisms than the property of pulling back harmonic functions to harmonic functions (equivalently, inducing a homomorphism on Jacobian groups). As we will see, harmonic morphisms have many good properties in addition to the preservation of harmonic functions under pullback, and they have an especially nice interpretation in terms of the dollar game which we describe in Section 10.3.

Exercise 10.7. Show that the graph inclusion morphism displayed in Figure 3 is not harmonic but induces an isomorphism on Jacobian groups.




Figure 3. A non-harmonic graph morphism that induces an isomorphism on Jacobian groups.

As a trivial example of a harmonic morphism, suppose that G_1 and G_2 are arbitrary graphs, and pick a vertex $h \in G_2$. Then define $\phi: G_1 \rightarrow G_2$ by $\phi(v) = \phi(e) = h$ for all $v \in V_1$ and $e \in E_1$. The morphism ϕ is clearly harmonic, since $m_v = 0$ for all $v \in V_1$. Such morphisms are called *constant*, since they collapse the entire graph G_1 onto the single vertex h . A less drastic phenomenon occurs when $m_v = 0$ for some, but not all vertices of G_1 . In that case, ϕ sends each vertex v with $m_v = 0$ to the same place as all of its neighbors, contracting the connecting edges (see Figure 4). Thus, ϕ collapses entire neighborhoods in G_1 to vertices of G_2 ,

and we call such morphisms *degenerate*. Since we will often want to exclude such degeneracies, we make the following definition.

Definition 10.8. Let $\phi: G_1 \rightarrow G_2$ be a harmonic morphism. Then ϕ is *non-degenerate* at a vertex $v \in V_1$ if the horizontal multiplicity $m_v > 0$. We say that ϕ is *non-degenerate* if it is non-degenerate at all vertices of G_1 .




Figure 4. This harmonic morphism is degenerate at the vertex v_5 .

Exercise 10.9. Prove that if $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism, then ϕ is surjective on both vertices and edges.

Exercise 10.10. Show that if $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism, then $\phi_*: \text{Div}(G_1) \rightarrow \text{Div}(G_2)$ is surjective, so that the induced maps on the Picard and Jacobian groups are also surjective.

Definition 10.11. Suppose that $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism. The *degree* of ϕ is defined as the number of pre-images of any edge $e' \in E_2$:

$$\text{deg}(\phi) := |\phi^{-1}(e')|.$$

The degree is well-defined by the next proposition.

Proposition 10.12. *The degree of a non-constant harmonic morphism is well-defined, i.e., the previous definition does not depend on the choice of edge $e' \in E_2$. Moreover, for any vertex $h \in V_2$, we have*

$$\text{deg}(\phi) = \sum_{v \in \phi^{-1}(h)} m_v.$$

Proof. Suppose that $e' = hk \in E_2$. Then

$$\begin{aligned} |\phi^{-1}(e')| &= \sum_{e \in E_1: \phi(e)=e'} 1 \\ &= \sum_{v \in \phi^{-1}(h)} \left(\sum_{vw \in E_1: \phi(vw)=e'} 1 \right) \\ &= \sum_{v \in \phi^{-1}(h)} m_v. \end{aligned}$$

This shows that the definition yields the same number for all choices of e' incident to the vertex $h \in V_2$. Since the graph G_2 is connected, it follows that the degree is well-defined. \square

Definition 10.13. Suppose that $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism. For each vertex $v \in V_1$, define the *vertical multiplicity* of ϕ at v to be

$$\text{vert}(v) := \# \text{ of edges incident to } v \text{ that are contracted by } \phi.$$

Example 10.14. The harmonic morphism shown in Figure 4 has degree 4, since every edge of the target graph has exactly 4 preimages. The horizontal and vertical multiplicities are as follows:

$$\begin{aligned} m_{v_1} &= 4 & , & & \text{vert}(v_1) &= 0 \\ m_{v_2} &= 1 & , & & \text{vert}(v_2) &= 1 \\ m_{v_3} &= 4 & , & & \text{vert}(v_3) &= 0 \\ m_{v_4} &= 3 & , & & \text{vert}(v_4) &= 1 \\ m_{v_5} &= 0 & , & & \text{vert}(v_5) &= 2. \end{aligned}$$

Exercise 10.15. Suppose that $\phi: G_1 \rightarrow G_2$ is a harmonic morphism and $\phi(v) = h$ for $v \in V_1$ and $h \in V_2$. Show that

$$\deg_{G_1}(v) = \deg_{G_2}(h)m_v + \text{vert}(v).$$

So far we have been studying the *push-forward* homomorphisms ϕ_* (on the divisor, Picard, and Jacobian groups) induced by a harmonic morphism ϕ . Note that the push-forward goes in the same direction as the original morphism ϕ . That is,

$$\phi: G_1 \rightarrow G_2 \quad \Longrightarrow \quad \phi_*: \text{Pic}(G_1) \rightarrow \text{Pic}(G_2).$$

There is also a natural *pull-back* homomorphism ϕ^* going in the opposite direction.

Definition 10.16. Let $\phi: G_1 \rightarrow G_2$ be a harmonic morphism. Then define $\phi^*: \text{Div}(G_2) \rightarrow \text{Div}(G_1)$ by

$$\phi^*(D') = \sum_{v \in V_1} m_v D'(\phi(v)) v \quad \text{for all } D' \in \text{Div}(G_2).$$

Exercise 10.17. Show that if $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism, then $\deg(\phi^*(D')) = \deg(\phi) \deg(D')$ for all $D' \in \text{Div}(G_2)$. Moreover, show that the composition $\phi_* \circ \phi^*$ is simply multiplication by $\deg(\phi)$ on $\text{Div}(G_2)$.

Proposition 10.18. *Suppose that $\phi: G_1 \rightarrow G_2$ is harmonic. Then ϕ^* sends principal divisors on G_2 to principal divisors on G_1 , thus preserving linear equivalence. Hence, ϕ induces homomorphisms of the Picard and Jacobian groups:*

$$\phi^*: \text{Pic}(G_2) \rightarrow \text{Pic}(G_1) \quad \text{and} \quad \phi^*: \text{Jac}(G_2) \rightarrow \text{Jac}(G_1).$$

Proof. Let $h \in V_2$ be any vertex. We begin by computing the pull-back of the degree-one divisor h :

$$\phi^*(h) = \sum_{v \in \phi^{-1}(h)} m_v v.$$

Now suppose that k is a neighbor of h in G_2 , so that we may fix a particular edge $e' = hk \in E_2$. Pulling back the degree-zero divisor $h - k$ yields

$$\begin{aligned} \phi^*(h - k) &= \phi^*(h) - \phi^*(k) \\ &= \sum_{v \in \phi^{-1}(h)} m_v v - \sum_{w \in \phi^{-1}(k)} m_w w \\ &= \sum_{v \in \phi^{-1}(h)} \sum_{vw \in E_1: \phi(vw) = e'} v - \sum_{w \in \phi^{-1}(k)} \sum_{vw \in E_1: \phi(vw) = e'} w \\ &= \sum_{vw \in E_1: \phi(vw) = e'} (v - w). \end{aligned}$$

Finally, consider $\text{div}(\chi_h) = \sum_{hk \in E_2} (h - k)$, the principal divisor obtained by firing h once. Pulling back yields

$$\begin{aligned} \phi^*(\text{div}(\chi_h)) &= \sum_{e' = hk \in E_2} \phi^*(h - k) \\ &= \sum_{e' = hk \in E_2} \sum_{vw \in E_1: \phi(vw) = e'} (v - w) \\ &= \sum_{v \in \phi^{-1}(h)} \sum_{vw \in E_1: \phi(w) \neq h} (v - w) \\ &= \sum_{v \in \phi^{-1}(h)} \sum_{vw \in E_1} (v - w) \\ &= \sum_{v \in \phi^{-1}(h)} \text{div}(\chi_v) \in \text{Prin}(G_1). \end{aligned}$$

Note that the fourth line above is obtained from the third by the inclusion of contracted edges of the form vv' for $v, v' \in \phi^{-1}(h)$. This is permissible because each such edge appears twice in the double sum, contributing the two terms $v - v'$ and $v' - v$ of opposite signs. Since the principal divisors $\text{div}(\chi_h)$ for $h \in V_2$ generate the subgroup $\text{Prin}(G_2)$, it follows that ϕ^* preserves linear equivalence as claimed. \square

As a counterpart to Exercise 10.10, we have the following injectivity result for the pullback homomorphism.

Proposition 10.19. *If $\phi: G_1 \rightarrow G_2$ is harmonic and non-constant, then the pullback $\phi^*: \text{Jac}(G_2) \rightarrow \text{Jac}(G_1)$ is injective.*

Proof. Suppose D is a degree-zero divisor on G_2 with $F := \phi^*(D) \sim 0$, i.e., the dollar game F on G_1 is winnable. We need to show D is winnable on G_2 . We accomplish this through a modified version of the greedy algorithm from Section 3.1 to win both F and D in tandem, as follows. For as long as there exists a vertex u of G_1 with $F(u) < 0$, let $v := \phi(u)$, and perform the borrowing operations

$$F \xrightarrow{-\phi^{-1}(v)} F', \quad D \xrightarrow{-v} D'.$$

As detailed at the end of the proof of Proposition 10.18, $F' = \phi^*(D')$. Replace F and D by the linearly equivalent F' and D' , respectively, and repeat.

Restricting our attention to F , if ϕ is non-degenerate, then our algorithm is an ordinary implementation of the greedy algorithm. Hence, eventually F is replaced by $0 \in \text{Div}(G_1)$, at which point it follows that D is replaced by $0 \in \text{Div}(G_2)$, and we are done.

So suppose that ϕ is degenerate and that at some point in our algorithm we need to borrow at a fiber $S := \phi^{-1}(v)$ containing a point of multiplicity 0. There will be $u \in S$ with $F(u) < 0$, and hence, $D(v) < 0$. At each $w \in S$, we have $F(w) \leq 0$ with equality exactly when $m_w = 0$. The ordinary greedy algorithm does not allow borrowing at S , since some of its members are not in debt. However, there is a way to perform a sequence of borrowings, allowed by the ordinary greedy algorithm, whose net effect is borrowing at S : repeatedly, until $S = \emptyset$, (i) let T be the in-debt vertices of S , (ii) borrow at T , and (iii) replace S by $S \setminus T$. After thus borrowing from all vertices in $S = \phi^{-1}(v)$, we still have $F(w) = 0$ for all degenerate vertices $w \in S$ since $\text{outdeg}_S(w) = 0$. So for the purpose of our algorithm, we may safely ignore vertices at which ϕ is degenerate. Our algorithm eventually halts, winning both dollar games simultaneously. \square

We can now state the *Riemann-Hurwitz formula for graphs*, which provides the relationship between the genera of graphs connected by a harmonic morphism.

Theorem 10.20. *Suppose that $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism between graphs of genera g_1 and g_2 , respectively. Then*

$$2g_1 - 2 = \text{deg}(\phi)(2g_2 - 2) + \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v)).$$

Proof. This result will follow from an investigation of the relation between the canonical divisor on G_1 and the pullback of the canonical divisor on G_2 (see Definition 5.7). Denote the canonical divisors by K_1 and K_2 , respectively. We compute

$$\begin{aligned} \phi^*(K_2) &= \sum_{h \in V_2} (\text{deg}_{G_2}(h) - 2)\phi^*(h) \\ &= \sum_{h \in V_2} (\text{deg}_{G_2}(h) - 2) \sum_{v \in \phi^{-1}(h)} m_v v \\ &= \sum_{v \in V_1} \text{deg}_{G_2}(\phi(v))m_v v - \sum_{v \in V_1} 2m_v v \\ &= \sum_{v \in V_1} (\text{deg}_{G_1}(v) - \text{vert}(v))v - \sum_{v \in V_1} 2m_v v \\ &= \sum_{v \in V_1} (\text{deg}_{G_1}(v) - 2)v - \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v))v \\ &= K_1 - \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v))v. \end{aligned}$$

Rearranging and taking degrees now yields the result:

$$\begin{aligned} 2g_1 - 2 &= \text{deg}(K_1) \\ &= \text{deg}(\phi^*(K_2)) + \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v)) \end{aligned}$$

$$\begin{aligned}
&= \deg(\phi) \deg(K_2) + \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v)) \\
&= \deg(\phi)(2g_2 - 2) + \sum_{v \in V_1} (2(m_v - 1) + \text{vert}(v)).
\end{aligned}$$

□

Exercise 10.21. Show that if $\phi: G_1 \rightarrow G_2$ is a non-constant harmonic morphism, then $g_1 \geq g_2$, with equality if and only if ϕ has degree 1. If ϕ is non-degenerate, show that $g_1 = g_2$ if and only if ϕ is an isomorphism of graphs.

10.2. Branched coverings of Riemann surfaces

As you might guess from the name, the Riemann-Hurwitz formula for graphs (Theorem 10.20) is a discrete version of a result about Riemann surfaces, initially used by Bernhard Riemann in 1857 and proved by Adolf Hurwitz in 1891. In order to state this classical result, we need to extend the discussion of Section 5.3 to include holomorphic mappings between Riemann surfaces and the attendant concepts of branching and ramification.

In Example 5.20, we viewed a meromorphic function f on a Riemann surface S as a mapping $f: S \rightarrow \mathbb{P}^1 = \mathbb{C} \cup \{\infty\}$ to the Riemann sphere. This is a special case of a *holomorphic mapping* $\phi: S_1 \rightarrow S_2$ between Riemann surfaces. The general idea is the following: near each point $p \in S_1$, an arbitrary mapping $\phi: S_1 \rightarrow S_2$ may be expressed as $z = \phi(w)$, where w is a local coordinate near p and z is a local coordinate near $\phi(p)$. We say that the mapping ϕ is *holomorphic* if (for all points $p \in S_1$) the resulting function $z = \phi(w)$ is complex-differentiable.

Using complex analysis, one may establish the following basic facts about any non-constant holomorphic mapping $\phi: S_1 \rightarrow S_2$:

- ϕ is surjective.
- There exists a positive integer $d \geq 1$ such that all but finitely many points $q \in S_2$ have exactly d pre-images under ϕ ; the integer d is called the *degree* of the mapping ϕ .
- The finitely many points q_1, q_2, \dots, q_b for which $\phi^{-1}(q_i)$ has cardinality less than d are called the *branch points* of ϕ .
- Consider a point $p \in S_1$ and its image $q = \phi(p) \in S_2$. Then there exist local coordinates w near p and z near q such that ϕ has the form $z = w^m$ for some integer $m \geq 1$; the integer $m = m_p$ is called the *ramification index* of p , and p is a *ramification point* if $m_p > 1$.
- For any point $q \in S_2$, we have $d = \sum_{p \in \phi^{-1}(q)} m_p$. In particular, there are only finitely many ramification points for ϕ , and $q \in S_2$ is a branch point if and only if $\phi^{-1}(q)$ contains a ramification point.

To visualize the phenomenon of ramification, first consider a point $q \in S_2$ that is *not* a branch point. Then q has a small neighborhood V such that $\phi^{-1}(V) = U_1 \sqcup U_2 \sqcup \dots \sqcup U_d$, a disjoint union of d neighborhoods U_j —picture them as a “stack of pancakes” mapping down to V as in Figure 5. On the other hand, if q is a

branch point, then some of these pancakes “come together” in groups, one for each preimage of q , so that $\phi^{-1}(V) = U_1 \sqcup \cdots \sqcup U_n$, where $n < d$ is the number of points in $\phi^{-1}(q)$. If U_j corresponds to the point p_j with ramification index m_j , then U_j maps to V under the m_j -to-1 power mapping $w \mapsto w^{m_j}$.




Figure 5. The picture on the left shows the local behavior of a degree-6 holomorphic mapping near a non-branch point q . The picture on the right shows a possibility for the local picture when q is a branch point: p_1 has ramification index 3, p_2 has ramification index 2, and p_3 is unramified. In each picture, the 6 red points upstairs together comprise the preimage of the red point downstairs. The dotted lines indicate schematically how the 6 pancakes on the left come together to form only 3 on the right.

In this way, the horizontal ramification indices m_v for harmonic graph morphisms are analogues of the ramification indices m_p for mappings of Riemann surfaces (cf. Remark 10.3). For instance, in Figure 4, the vertex v_4 has horizontal ramification index 3, indicating that near v_4 the mapping is 3-to-1 (ignoring the vertical edge). By contrast, the vertex v_2 is horizontally unramified and nearby the mapping is 1-to-1. Note that there is no analogue of vertical ramification in the setting of Riemann surfaces.

We are now able to state the classical Riemann-Hurwitz formula. As in the graph-theoretic case, one may prove this result by comparing the canonical divisor of S_1 with the pullback of the canonical divisor of S_2 .

Theorem 10.22 (Riemann-Hurwitz). *Suppose that $\phi: S_1 \rightarrow S_2$ is a non-constant holomorphic mapping between Riemann surfaces of genera g_1 and g_2 , respectively. Then*

$$2g_1 - 2 = \deg(\phi)(2g_2 - 2) + \sum_{p \in S_1} (m_p - 1).$$

The Riemann-Hurwitz formula has many important consequences for the theory of Riemann surfaces. For instance, it may be used to establish the famous Hurwitz genus bound for the automorphism group of a Riemann surface S : if $g = g(S) \geq 2$, then S has at most $84(g - 1)$ automorphisms. In fact, there is a graph-analogue of this result, proved via the the graph-theoretic Riemann-Hurwitz formula (see Problem 10.5). It concerns so-called *harmonic group actions* on finite graphs, and may be stated as follows: if G is a finite graph of genus $g \geq 2$ and Γ is a finite group acting harmonically on G , then $|\Gamma| \leq 6(g - 1)$. For further details, see [30, 31, 32].

10.3. Household-solutions to the dollar game

Recall the initial discussion of the dollar game in Chapter 1, where we thought of the multigraph $G_1 = (V_1, E_1)$ as a community, and a divisor $D \in \text{Div}(G_1)$ as a distribution of wealth. In this section we interpret a non-constant harmonic morphism $\phi: G_1 \rightarrow G_2$ as the creation of *households* whose members pool their money in an attempt to solve their debt problems. The harmonic condition for ϕ plays the role of a *fairness property* which ensures that the lending moves available to the households reflect the structure of the original community represented by G_1 . The details of this interpretation are presented in the following paragraphs.

First of all, note that by Exercise 10.9, the non-constant harmonic morphism ϕ is surjective. Hence, the collection of sets

$$H := \{\phi^{-1}(h) \cap V_1 : h \in V_2\}$$

forms a partition of the vertex set V_1 by non-empty subsets. Identifying each vertex $h \in V_2$ with its set of pre-images, we may thus think of the vertices of G_2 as households in the original community G_1 . In these terms, we may think of the morphism ϕ as given on vertices by

$$\phi(v) = \text{unique household in } H \text{ containing } v.$$

Moreover, the definition of a graph morphism guarantees that every edge between individuals in G_1 is sent to an edge between households in G_2 except that edges between members of the same household are forgotten in the passage to G_2 .

Roughly speaking, the harmonic condition ensures that each individual v , when joining her household, reduces the total strength of her bond with each other household in a fair manner. In detail, for any vertex $v \in V_1$, consider the corresponding household $h = \phi(v)$. Then select any edge $e' \in E_2$ incident to $\phi(v)$ in G_2 , say $e' = hk$. By the definition of harmonic morphism, there are exactly m_v pre-images of the edge e' incident to v in G_1 , and each of these has the form vw for some $w \in \phi^{-1}(k)$. Thus, edges incident to v are identified in groups of size m_v with edges incident to $\phi(v)$. This justifies the fairness interpretation: each individual v , when joining a household, reduces the strength of her relationship with all *other* households by the *same* factor m_v . Note in particular that if there are d edges between $\phi(v)$ and k in G_2 , then there must be exactly dm_v edges in G connecting v to members of the household k .

Now suppose that D is a divisor on G_1 . When the individuals in G_1 form households according to the harmonic morphism ϕ , they pool their money/debt.

This pooling is exactly what is described by the push-forward homomorphism on divisor groups: if $D = \sum a_v v \in \text{Div}(G_1)$, then

$$\phi_*(D) = \sum_{v \in V_1} a_v \phi(v) = \sum_{h \in V_2} \left(\sum_{v \in \phi^{-1}(h)} a_v \right) h.$$

In this way, the dollar game on G_1 starting with D naturally induces the dollar game on G_2 starting with $\phi_*(D)$. Note that the divisor $\phi_*(D) \in \text{Div}(G_2)$ has the same degree as D .

An effective divisor linearly equivalent to $\phi_*(D)$ on G_2 will be called a *household-solution* to the original dollar game starting with D on G_1 . The next proposition justifies the hope that forming households may help the community G_1 solve its debt problems:

Proposition 10.23. *Let $\phi: G_1 \rightarrow G_2$ be a non-constant harmonic morphism. Let r_1 and r_2 denote the rank functions for the graphs G_1 and G_2 , respectively. Then for all divisors $D \in \text{Div}(G_1)$, we have the inequality*

$$r_2(\phi_*(D)) \geq r_1(D).$$

Proof. Set $R = r_1(D)$. We need to show that $\phi_*(D) - E'$ is winnable for all effective divisors $E' \in \text{Div}(G_2)$ of degree R . So let E' be any such effective divisor of degree R on G_2 . By the surjectivity of the push-forward ϕ_* , there exists an effective divisor $E \in \text{Div}(G_1)$, also of degree R , such that $\phi_*(E) = E'$. Since the rank of D is R , there exists an effective divisor $F \in \text{Div}(G_1)$ such that $D - E \sim F$. But ϕ_* preserve linear equivalence, so $\phi_*(D) - E' = \phi_*(D - E) \sim \phi_*(F) \geq 0$. Hence, $\phi_*(D) - E'$ is winnable as claimed. \square

If we drop the requirement that ϕ be non-constant, then the entire community G_1 could form into a single household. The corresponding harmonic morphism $\phi: G \rightarrow \star$ is the constant map onto the single vertex \star , thought of as a graph with no edges. The induced map ϕ_* on divisors may be identified with the degree map $\text{deg}: \text{Div}(G) \rightarrow \text{Div}(\star) \simeq \mathbb{Z}$. Hence, by forming a giant single household, the community G automatically solves the induced dollar game on \star starting with $\phi_*(D)$ whenever $\text{deg}(D) \geq 0$. But this household-solution isn't very satisfying, which explains why we have restricted our attention to non-constant harmonic morphisms.

For a slightly less trivial household-solution, pick a vertex $q \in V_1$ and consider the harmonic morphism $\phi: G_1 \rightarrow P$ onto the path graph consisting of one edge e' connecting two vertices h and k :

$$\phi(v) = \begin{cases} h & \text{if } v = q, \\ k & \text{if } v \neq q \end{cases}, \quad \phi(e) = \begin{cases} e' & \text{if } e = qv \text{ for some } v \in V_1, \\ k & \text{otherwise.} \end{cases}$$

Now there are two households, one containing only q , and the other containing everyone else. The harmonic morphism ϕ is non-constant, sending all edges incident to q in G_1 to the unique edge e' in P and contracting all other edges to the vertex k . For any divisor $D \in \text{Div}(G_1)$ of nonnegative degree, we have $\text{deg}(\phi_*(D)) = \text{deg}(D) \geq 0$, which means that the total amount of money on the two vertices of P is nonnegative. If one of these vertices is in debt, then the other

can perform lending moves until both are debt-free. Thus, the harmonic morphism ϕ yields a household-solution to all nonnegative dollar games on G_1 . But this household-solution is quite drastic in most cases: all individuals in $V \setminus \{q\}$ have forgotten their bonds to each other, and only connections with q are remembered in P . Note that unless every vertex in G_1 is connected to q , the morphism ϕ will be degenerate. Hence, we may avoid these types of household-solutions by considering only non-degenerate harmonic morphisms. In terms of households, the requirement that ϕ be non-degenerate corresponds to the demand that no individual in G_1 joins the same household as all of her neighbors. We thus think of non-degenerate harmonic morphisms as remembering enough of the structure of the original community G_1 to provide interesting household-solutions to dollar games on G_1 .

Suppose that $\phi: G_1 \rightarrow G_2$ is a non-degenerate harmonic morphism between graphs of genus g_1 and g_2 , respectively. By the Riemann-Hurwitz Theorem 10.20, we know that $g_2 \leq g_1$. This should help in the search for household solutions to the dollar game since Corollary 4.9 guarantees winnability when the degree is larger than the genus: the dollar game on G_2 starting with $\phi_*(D)$ is winnable if $g_2 \leq \deg(\phi_*(D)) = \deg(D)$. So if we want to find household solutions to the dollar game, we should look for harmonic morphisms with g_2 as small as possible.

In particular, for $g_2 = 0$ we obtain non-degenerate household-solutions to all dollar games on G_1 of nonnegative degree. Hence we would like to find non-degenerate harmonic morphisms $\phi: G_1 \rightarrow T$ where the target graph T is a tree. As mentioned at the end of Section 5.3, these morphisms are the graph-analogues of branched coverings of the projective line \mathbb{P}^1 for Riemann surfaces. The next theorem extends this analogy by demonstrating that harmonic morphisms to trees are essentially the same as complete linear systems of rank 1.

Theorem 10.24. *Suppose that G is d -edge connected with $|V(G)| > d \geq 1$. Then the set of degree- d non-degenerate harmonic morphisms $\phi: G \rightarrow T$ such that T is a tree is in natural bijection with the collection of complete linear systems of degree d and rank 1 on G .*

The proof of this theorem requires a few definitions and lemmas.

Definition 10.25. Let $D \in \text{Div}(G)$. A complete linear system $|D|$ partitions the graph G , if $V(G)$ is the disjoint union of the supports of the divisors in $|D|$.

Definition 10.26. A complete linear system $|D|$ is called a g_d^r if $\deg(D) = d$ and $r(D) = r$.

Lemma 10.27. *If G is d -edge connected, then any g_d^r on G with $r \geq 1$ partitions G .*

Proof. Suppose that $|D|$ is a g_d^r on G , with $r \geq 1$. Then for any $v \in V(G)$, the linear system $|D-v|$ is nonempty, so there exists an effective divisor E with $D \sim v+E$. This shows that the supports of the divisors in $|D|$ cover $V(G)$. For the disjointness of the supports, we need to show that E is the unique effective divisor linearly equivalent to $D-v$. So suppose that $E' \sim E$ with E' effective. Applying the Abel-Jacobi map $S_q^{(d-1)}: \text{Div}_+^{(d-1)}(G) \rightarrow \text{Jac}(G)$, we see that $S_q^{(d-1)}(E) = S_q^{(d-1)}(E')$. But by

Proposition 3.28, the d -edge connectivity of G implies that $S_q^{(d-1)}$ is injective, so that $E = E'$ as required. \square

Lemma 10.28. *Suppose that G is d -edge connected, with $|V(G)| > d$. Then every divisor of degree d on G has rank at most 1.*

Proof. Suppose that D has degree d and $r(D) = r \geq 1$. Fix a vertex $v \in V(G)$. Then by the previous lemma, there is a unique effective divisor $E \in |D|$ with v in its support. Since $|V(G)| > d$ and $\deg(E) = d$, there exists a vertex $w \in V(G)$ outside the support of E . We claim that $|D - v - w| = \emptyset$, which shows that $r(D) = 1$ as claimed. For this, suppose on the contrary that F is effective and $F \sim D - v - w$. Then $F + v + w$ contains v in its support and is linearly equivalent to D . Hence $F + v + w = E$ by the uniqueness of E . But this contradicts the choice of vertex w . \square

Proof of Theorem 10.24 Suppose that $|D|$ is a g_d^1 on G . We wish to construct a tree T together with a harmonic morphism $\phi: G \rightarrow T$. For the vertex set, we define $V(T) = |D|$. By Lemma 10.27, $|D|$ partitions G , so for each vertex $v \in V(G)$, there is a unique effective divisor $E_v \in |D|$ with v in its support. Hence, sending v to E_v defines a surjective map from $V(G)$ to $V(T)$. We wish to choose the edge-set $E(T)$ in such a way that T is a tree and that the map on vertices described above extends to a harmonic morphism of graphs.

Suppose that $E \neq F$ are distinct effective divisors in $|D|$. Then there exists a non-constant function $f \in \mathcal{M}(G)$ such that $E - F = \text{div}(f)$. Set $M(f) = \{x \in V(G) : f(x) = \max(f)\}$. Then for any $x \in M(f)$ we have

$$\begin{aligned} E(x) &= \text{div}(f)(x) + F(x) \\ &\geq \text{div}(f)(x) \\ &= \sum_{xy \in E(G)} (f(x) - f(y)) \\ &\geq \text{outdeg}_{M(f)}(x). \end{aligned}$$

Let $\delta(M(f))$ denote the set of edges in G that connect a vertex in $M(f)$ to a vertex outside of $M(f)$. Then

$$|\delta(M(f))| = \sum_{x \in M(f)} \text{outdeg}_{M(f)}(x) \leq \deg(E) = d.$$

But removing the edge-set $\delta(M(f))$ disconnects the graph G , so by d -edge connectivity, we must have $|\delta(M(f))| = d$. This implies that all of the preceding inequalities are actually equalities. In particular, we must have $E(x) = \text{outdeg}_{M(f)}(x)$ for all $x \in M(f)$, and $\text{supp}(E) \subset M(f)$. Moreover, if $xy \in \delta(M(f))$ with $x \in M(f)$, then $f(x) - f(y) = 1$. A similar argument shows that the support of F is contained in $m(f) := \{y \in V(G) : f(y) = \min(f)\}$, and if $yz \in \delta(m(f))$ with $y \in m(f)$, then $f(z) - f(y) = 1$.

Now suppose that there exists an edge $e = vw$ connecting the support of E to the support of F . It follows that $e \in \delta(M(f)) \cap \delta(m(f))$, and that $f(v) - f(w) = 1$. But $v \in M(f)$ and $w \in m(f)$, so we see that f takes only two values, and $V(G) = M(f) \cup m(f)$. This implies that $\delta(M(f)) = \delta(m(f))$. Further, if $e' = xy \in \delta(M(f))$

is an edge with $x \in M(f)$, then $E(x) = \text{outdeg}_{M(f)}(x) \geq 1$, so x is in the support of E . But $y \in m(f)$, so

$$F(y) = E(y) - \text{div}(f)(y) = \sum_{yz \in E(G)} (f(z) - f(y)) = \text{outdeg}_{m(f)}(y) \geq 1,$$

so y is in the support of F . Thus, the d edges in $\delta(M(f)) = \delta(m(f))$ are exactly the edges in $E(G)$ connecting the support of E to the support of F .

We are now ready to define the edge-set of the graph T : if $E, F \in |D| = V(T)$, then $EF \in E(T)$ if and only if there is an edge of G connecting the support of E to the support of F . Our work above has shown that if one such edge exists in $E(G)$, then there are exactly d such edges.

Define a graph morphism $\phi: G \rightarrow T$ as follows:

$$\begin{aligned} \phi(v) &= E_v \quad \text{for all } v \in V(G) \\ \phi(vw) &= \begin{cases} E_v E_w & \text{if } E_v \neq E_w \\ E_v & \text{if } E_v = E_w \end{cases} \quad \text{for all } e = vw \in E(G). \end{aligned}$$

At this point, we know that ϕ is surjective, and that every edge of T has exactly d pre-images. Hence, if ϕ is harmonic, then it has degree d as required.

To show that ϕ is harmonic, let $v \in V(G)$ be arbitrary, and set $E = \phi(v)$. Choose an edge $EF \in E(T)$ incident to the vertex E . The number of edges $e \in E(G)$ incident to v such that $\phi(e) = EF$ is the number of edges from v to the support of F . Call this number $m(v, F)$. Using the notation from above, if $E - F = \text{div}(f)$, then

$$E(v) = \text{outdeg}_{M(f)}(v) \geq m(v, F).$$

Summing over the support of E , we find that

$$d = \deg(E) \geq \sum_{v \in \text{supp}(E)} m(v, F) = d,$$

since there are exactly d edges connecting the support of E to the support of F . It follows that $E(v) = m(v, F)$ for all $v \in \text{supp}(E)$. In particular, for fixed v , the number $m(v, F) = E(v)$ is independent of the choice of edge $EF \in E(T)$ incident to $\phi(v) = E$. Hence, $\phi: G \rightarrow T$ is harmonic, with horizontal multiplicities given by $m_v = E_v(v)$ for all $v \in V(G)$. This implies that ϕ is non-degenerate, since for all $v \in V(G)$, we have $m_v = E_v(v) \geq 1$.

It remains to show that T is a tree. For this, it suffices to show that any two vertices of T are linearly equivalent as degree-one divisors on T (cf. Problem 4.3). So suppose $F \in |D|$ is a vertex of T , and note that

$$\phi^*(F) = \sum_{v \in \text{supp}(F)} F(v)v = F \sim D.$$

Hence, we see that for any two vertices $E, F \in |D|$, we have $\phi^*(E) \sim \phi^*(F)$. By the injectivity of ϕ^* as a map of Picard groups, this implies that $E \sim F$ as degree-one divisors on T .

Thus, to any g_d^1 on G , we have associated a non-degenerate, degree- d harmonic morphism $\phi: G \rightarrow T$, where T is a tree.

For the other direction, suppose that $\phi: G \rightarrow T$ is any non-degenerate, degree- d harmonic morphism to a tree. We wish to find a g_d^1 on G that will yield ϕ via the construction described above. For this, choose any vertex $t \in V(T)$, and set $D := \phi^*(t) \in \text{Div}(G)$, an effective divisor of degree d . To show that D has rank 1, first note that $r(D) \leq 1$ by Lemma 10.28. To show that $r(D) = 1$, let $v \in V(G)$ be arbitrary. Then since T is a tree, $\phi(v) \in V(T)$ is linearly equivalent to t . But then $\phi^*(\phi(v))$ is linearly equivalent to $\phi^*(t) = D$ on G . But $\phi^*(\phi(v))(v) = m_v v \geq v$ since ϕ is non-degenerate. It follows that $\phi^*(\phi(v)) - v$ is effective, and

$$D - v \sim \phi^*(\phi(v)) - v.$$

Thus, $|D - v|$ is nonempty for all $v \in V(G)$, so $r(D) = 1$ as claimed. Hence, $|D|$ is a g_d^1 on G . Moreover, we have shown that $\phi^*(s) \in |D|$ for every $s \in V(T)$. But these must account for the entire complete linear system since $|D|$ partitions the graph G by Lemma 10.27. Hence, applying the construction described in the first part of this proof to $|D|$ will produce the given tree T together with the harmonic morphism $\phi: G \rightarrow T$.

□

Notes

- (1) This chapter is based on the paper [7] by Baker and Norine. In particular, the Riemann-Hurwitz formula appears as Theorem 2.14 of that paper.
- (2) Our discussion of the relation between harmonic morphisms and harmonic functions differs in its focus from that found in Section 2.2 of [7], where the authors characterize harmonic morphisms to simple graphs (i.e., those without multiple edges) in terms of a *local* preservation of harmonic functions under pullback.
- (3) Proposition 10.19 appears as Theorem 4.13 of [7], with a different (and substantially more difficult) proof. As observed by Baker and Norine, the corresponding injectivity result for Riemann surfaces is false in general.
- (4) Theorem 10.24 is due to the first author and Avi Steiner. It generalizes Theorem 5.12 of [7], which is essentially the case $d = 2$.

Problems for Chapter 10

10.1. Using Remark 10.3, construct all non-constant harmonic morphisms of degree 2 to the triangle graph C_3 with at most 2 vertical edges. Are any of your morphisms degenerate?

10.2. Consider the following composition of graph morphisms:

$$G_1 \xrightarrow{\phi_1} G_2 \xrightarrow{\phi_2} G_3.$$

- Show that if ϕ_1 and ϕ_2 are harmonic, then so is their composition $\phi_2 \circ \phi_1$.
- Show that if $\phi_2 \circ \phi_1$ is harmonic and ϕ_1 is harmonic and nonconstant, then ϕ_2 is harmonic.
- Prove or provide a counterexample: if $\phi_2 \circ \phi_1$ is harmonic and ϕ_2 is harmonic and nonconstant, then ϕ_1 is harmonic.

10.3. Find all g_3^1 's on the complete graph K_4 and explicitly describe the corresponding harmonic morphisms to trees.

10.4. A graph G of genus $g \geq 2$ is called *hyperelliptic* if it has a g_2^1 .

- Use Riemann-Roch to show that every graph of genus $g = 2$ is hyperelliptic, with g_2^1 given by the canonical divisor.
- Use Clifford's Theorem to show that if $g(G) \geq 2$ and D is a divisor of degree 2 on G , then $r(D) \leq 1$.
- Suppose that $\phi: G \rightarrow G'$ is a non-constant harmonic morphism and that $g(G') \geq 2$. Show that if G is hyperelliptic, then G' is hyperelliptic.
- Show that the banana graph B_n is hyperelliptic for $n \geq 3$.
- Suppose that $G \neq B_n$ is a hyperelliptic graph. Use Proposition 3.28 to show that the edge-connectivity of G is at most 2.
- Show that the g_2^1 on a hyperelliptic graph is unique. (Hint: if D and D' are two degree-2 divisors of rank 1, apply Riemann-Roch to the divisors $E = D + (g-2)D'$ and $E' = (g-1)D'$ and conclude that both are linearly equivalent to the canonical divisor. You will find Exercise 5.4 helpful.)

For more about hyperelliptic graphs, including the study of their Weierstrass points (cf. Problem 5.9), see [7].

10.5. Let $\Gamma \subset \text{Aut}(G)$ be a group of automorphisms of the connected graph G . The group Γ thus acts on the vertices and on the edges of G . Define the *quotient graph* G/Γ as follows:

$$V(G/\Gamma) = \{\Gamma v : v \in V(G)\} = \Gamma\text{-orbits of vertices of } G$$

$$E(G/\Gamma) = \{\Gamma e : e = vw \in E(G) \text{ and } \Gamma v \neq \Gamma w\}.$$

$$= \Gamma\text{-orbits of edges of } G, \text{ with orbits yielding loop-edges removed.}$$

Define a graph morphism $\phi_\Gamma: G \rightarrow G/\Gamma$ by

$$\begin{aligned} \phi_\Gamma(v) &= \Gamma v && \text{for all } v \in V(G); \\ \phi_\Gamma(e) &= \begin{cases} \Gamma e & \text{if } e = vw \text{ and } \Gamma v \neq \Gamma w, \\ \Gamma v & \text{if } e = vw \text{ and } \Gamma v = \Gamma w. \end{cases} \end{aligned}$$

- (a) Show by explicit example that the morphism ϕ_Γ need not be harmonic. (Suggestion: consider a \mathbb{Z}_2 -action on the banana graph B_3 .)
- (b) We say that Γ *acts harmonically* on G if for all subgroups $\Delta < \Gamma$, the quotient morphism $\phi_\Delta: G \rightarrow G/\Delta$ is harmonic. Show that Γ acts harmonically on G if and only if for every vertex v , the stabilizer subgroup $\Gamma_v := \{\gamma \in \Gamma : \gamma v = v\}$ acts freely on the set of edges incident to v :

if $\gamma \neq \text{id}_\Gamma$ and $\gamma v = v$, then $\gamma e \neq e$ for all edges $e = vw$.

For the remainder of this problem, suppose that Γ acts harmonically on G .

- (c) Show that ϕ_Γ has degree $|\Gamma|$, and for each vertex $v \in V(G)$, the horizontal ramification index m_v is equal to $|\Gamma_v|$, the size of the stabilizer subgroup. Also, show that m_v divides the vertical multiplicity $\text{vert}(v)$ for each vertex v , and define $n_v := \text{vert}(v)/m_v$.
- (d) Apply the Riemann-Hurwitz formula to ϕ_Γ to obtain:

$$2g(G) - 2 = |\Gamma|(2g(G/\Gamma) - 2 + R).$$

Here, $R = \sum_{x \in V(G/\Gamma)} \left(2\left(1 - \frac{1}{m_x}\right) + n_x\right)$, where $m_x := m_v$ and $n_x := n_v$ for any choice of vertex v such that $\phi_\Gamma(v) = x$.

- (e) Show that if $R > 2$, then in fact $R \geq 7/3$.
- (f) Now suppose that $g(G) \geq 2$. First show that if $g(G/\Gamma) \geq 1$, then $|\Gamma| \leq 2g(G) - 2$. Finally, show that if $g(G/\Gamma) = 0$, then $R > 2$ and $|\Gamma| \leq 6(g - 1)$. Thus, the maximal size of a harmonic group action on a graph of genus $g \geq 2$ is $6(g - 1)$.
- (g) Find a group of order 6 acting harmonically on the genus-2 banana graph B_3 .

For more about harmonic group actions and genus bounds, see [30, 31, 32].

Divisors on complete graphs

11.1. Parking functions

There is a nice description of the superstable configurations on K_{n+1} , the complete graph on $n + 1$ vertices, in terms of a protocol for parking cars. Suppose there is a line of n cars, C_1, \dots, C_n traveling down a street with C_1 in the lead. Further along that street, there is a line of n parking spaces labeled, in order, $1, \dots, n$. The driver of each car has a preferred parking space. We list these preferences as a vector $p = (p_1, \dots, p_n)$ where p_i is the preference for C_i . The protocol is that the driver of C_i will drive to parking space p_i , ignoring the state of any previous parking spaces. If space p_i is empty, car C_i parks there. If it is full, then C_i parks in the next available space. Figure 1 gives three examples.

If p is a permutation of the vector $(1, \dots, n)$, then there is a unique parking space for each car, and each car C_i will end up in its preferred space. On the other hand, suppose p is the constant vector $(1, 1, \dots, 1)$. Then car C_1 will drive to space 1 and park; car C_2 will find space 1 filled and drive on to 2, the next available space. In the end, each C_i parks in space i . Only C_1 gets its preferred spot.

Not every list of parking preferences p allows every car to park. For instance, consider the constant vector $p = (n, n, \dots, n)$. Car C_1 parks in space n . Next, C_2 drives past the empty parking spaces $1, \dots, n - 1$ to its preferred space n but finds it filled. The protocol says C_2 should drive on and take the next available space. However, there are no more available spaces. In fact, only C_1 can park with this p . Those parking preferences p that allow every car to park are called *parking functions* of length n .

Exercise 11.1.

(1) Which of the following lists of parking preferences are parking functions? For each that is, find the resulting assignment of cars to parking spaces.

- | | |
|-----------------------|-----------------------|
| (a) $(3, 1, 3, 1, 4)$ | (b) $(2, 3, 2, 4)$ |
| (c) $(2, 1, 3, 2)$ | (d) $(4, 3, 1, 3, 4)$ |




Figure 1. Three examples of parking functions. In each case, the cars C_1, C_2, C_3 drive across the page from right-to-left to parking spots labeled 1, 2, 3. The parking preferences for each car are listed in order above the arrows.

(2) Let $X := \{(p_1, p_2, p_3) \in \mathbb{Z}^3 : 1 \leq p_i \leq 3\}$. What is the probability that an element of X chosen uniformly at random is a parking function?

The list of parking preferences (2, 3, 2, 4) in Exercise 11.1, 1 (b), has no driver preferring parking space 1. That means all four cars need to park in the three remaining spaces, 2, 3, 4, which is impossible. A similar problem arises in 1 (d): if the preferences are (4, 3, 1, 3, 4), then there are four cars competing for the three parking spaces 3, 4, 5.

Let p be a list of preferences, and let the cars park according to p . If p is not a parking function, then some of the cars are not able to park in spaces $1, \dots, n$. Suppose we send these cars to a special overflow parking lot. So now everyone has a space to park, and p is a parking function exactly when no car ends up parked in the overflow lot, i.e., exactly when all the spaces $1, \dots, n$ are filled. Note that space 1 is filled exactly when at least one car prefers space 1. Next note that spaces 1 and 2 are both filled exactly when space 1 is filled *and* at least two cars prefer spaces numbered at most 2, taking into account the possibility that a car preferring space 1 is forced to park in space 2, instead. Continuing this line of thought proves the following result.

Proposition 11.2. *Let $p = (p_1, \dots, p_n) \in \mathbb{Z}^n$ with $1 \leq p_i \leq n$ for all i . Then p is a parking function if and only if for each $j = 1, \dots, n$,*

$$|\{i : p_i \leq j\}| \geq j.$$

For $p, q \in \mathbb{Z}^n$ write $q \leq p$ if $q_i \leq p_i$ for all i . A *maximal parking function* is a parking function p maximal with respect to \leq , i.e., with the property that if $p \leq q$ for some parking function q , then $p = q$. Let $\bar{1} = (1, \dots, 1)$. We have the following immediate corollary to Proposition 11.2:

Corollary 11.3. *Suppose that $p = (p_1, \dots, p_n)$ is a parking function.*

- (1) *Then so is $(p_{\pi(1)}, \dots, p_{\pi(n)})$ for any permutation π of the indices.*
- (2) *If $\vec{1} \leq q \leq p$, then q is a parking function.*
- (3) *The maximal parking functions are exactly the $n!$ vectors obtained by permuting the components of $(1, \dots, n)$.*

Corollary 11.3 provides an easy way to determine whether a given integer vector q is a parking function. First, sort the components of q to obtain the vector \tilde{q} with $\tilde{q}_i \leq \tilde{q}_{i+1}$ for all i . Then q is a parking function if and only if $\vec{1} \leq \tilde{q} \leq (1, \dots, n)$. To find all parking functions, start with the maximal parking function $p = (1, \dots, n)$; next list all vectors q such that $\vec{1} \leq q \leq p$ and q is increasing, i.e., $q_1 \leq \dots \leq q_n$; finally, take all vectors obtained by permuting the components of these increasing parking functions.

Proposition 11.4. *There are $(n + 1)^{n-1}$ parking functions of length n .*

Proof. Problem 11.1. □

Theorem 11.5. *Identifying configurations on K_{n+1} with elements of \mathbb{N}^n , as usual, the superstable are exactly*

$$p - \vec{1} = (p_1 - 1, \dots, p_n - 1)$$

as p ranges over all parking functions of length n .

Proof. Fix $q \in K_{n+1}$, and consider configurations on the remaining vertices $\tilde{V} = \{v_1, \dots, v_n\}$. Let c be a nonnegative configuration on K_{n+1} , and define $p = c + \vec{1}$. If c is not superstable, there exists a nonempty set $S \subseteq \tilde{V}$ that can be legally fired. This means that for each $v \in S$,

$$c(v) \geq \text{outdeg}_S(v) = n + 1 - |S|.$$

Letting $j := n + 1 - |S|$,

$$\{v_i : p_i \leq j\} = \{v_i : c(v_i) < j\} \subseteq \tilde{V} \setminus S.$$

It follows that

$$|\{i : p_i \leq j\}| \leq |\{\tilde{V} \setminus S\}| = j - 1 < j.$$

So by Proposition 11.2, p is not a parking function.

Conversely, if p is not a parking function, by Proposition 11.2, there exists $j \in \{1, \dots, n\}$ such that $|\{i : p_i \leq j\}| < j$. Let

$$T := \{v_i : p_i \leq j\} = \{v \in \tilde{V} : c(v) < j\},$$

and define $S := \tilde{V} \setminus T = \{v_i : p_i > j\} = \{v \in \tilde{V} : c(v) \geq j\}$. Then for each $v \in S$, we have

$$\text{outdeg}_S(v) = n + 1 - |S| = |T| + 1 \leq j \leq c(v).$$

Thus, c is not superstable since S can be legally fired. □

Corollary 11.6. *A configuration $c = (c_0, \dots, c_{n-1})$ on K^{n+1} is superstable if and only if after sorting the components so that $c_0 \leq \dots \leq c_{n-1}$, we have $c_i \leq i$ for all i . That is,*

$$c \leq (0, 1, \dots, n-1).$$

The maximal superstables (with respect to \leq) on K_{n+1} are the $n!$ configurations obtained from $(0, 1, \dots, n-1)$ by permuting components.

Corollary 11.7. $|\text{Jac}(K_{n+1})| = (n+1)^{n-1}$.

Proof. Immediate from Proposition 11.4. □

Remark 11.8. By the matrix-tree theorem, Corollary 11.7 provides a proof of Cayley's formula which states that the number of trees on $n+1$ labeled vertices is $(n+1)^{n-1}$.

11.2. Computing ranks on complete graphs

Although the general problem of computing the rank of a divisor on a graph is difficult, for certain classes of graphs there are efficient algorithms. In this section, we describe an algorithm due to Robert Cori and Yvan Le Borgne ([27]) for computing the rank of divisors on the complete graph K_{n+1} with vertices v_0, v_1, \dots, v_n . Fix $q = v_0$. Note that the symmetric group S_{n+1} acts on $\text{Div}(K_{n+1})$ by permuting the coefficients, and two divisors differing by a permutation have complete linear systems of the same size. It follows that the rank function is constant on the S_{n+1} -orbits in $\text{Div}(K_{n+1})$.

By Theorem 11.5, the superstable configurations on K_{n+1} are in bijection with the parking functions of length n . In particular, if c is a superstable on K_{n+1} , then there exists an index $1 \leq i \leq n$ such that $c(v_i) = 0$. Every divisor D is linearly equivalent to a unique q -reduced divisor of the form $c + kq$ with c superstable. Moreover, $c + kq$ is winnable if and only if $k \geq 0$. The idea of the Cori-Le Borgne algorithm is to iterate the following process until the divisor becomes unwinnable: replace $c + kq$ by $c + kq - v_i$ for $c(v_i) = 0$, and then compute the corresponding q -reduced divisor. The number of iterations is then one more than the rank of D .

Algorithm 10 Cori-Le Borgne algorithm.

- 1: INPUT: a divisor $D \in \text{Div}(K_{n+1})$.
 - 2: OUTPUT: the rank of D .
 - 3: **initialization:** $R = -1$
 - 4: compute the q -reduced divisor $c + kq \sim D$
 - 5: **while** $k \geq 0$ **do**
 - 6: choose $v_i \in \tilde{V}$ such that $c(v_i) = 0$
 - 7: $c + kq = q$ -reduced divisor $\sim c - v_i + kq$
 - 8: $R = R + 1$
 - 9: return R
-

The proof of validity for this algorithm depends on the following lemma.

Lemma 11.9. *Suppose $D \in \text{Div}(K_{n+1})$ is an effective divisor, and that $D(v_i) = 0$ for some $1 \leq i \leq n$. Then there exists an effective divisor $E \in \text{Div}(K_{n+1})$ of degree $r(D) + 1$ such that $E(v_i) > 0$ and $|D - E| = \emptyset$.*

Proof. Let $F \in \text{Div}(K_{n+1})$ be an effective divisor of degree $r(D) + 1$ such that $|D - F| = \emptyset$. If $F(v_i) > 0$ then we are done, so assume that $F(v_i) = 0$. By assumption, $D - F$ is not effective, so we may choose a vertex $v_j \neq v_i$ such that $D(v_j) - F(v_j) =: -a < 0$. Note that $v_j = v_0 = q$ is allowed. Setting $A = D - (F - av_j)$, we have $A(v_i) = D(v_i) - F(v_i) = 0$ and $A(v_j) = D(v_j) - F(v_j) + a = 0$. Thus, if $\tau \in S_{n+1}$ is the transposition that switches i and j , then $\tau A = A$. Define $E = F - av_j + av_i$, which is effective since $E(v_j) = F(v_j) - a = D(v_j) \geq 0$. Moreover, $E(v_i) = F(v_i) + a = a > 0$, and $\deg(E) = \deg(F) = r(D) + 1$. But we also have

$$\begin{aligned} D - E &= D - F + av_j - av_i \\ &= A - av_i \\ &= \tau(A - av_j) \\ &= \tau(D - F). \end{aligned}$$

It follows that $|D - E| = |\tau(D - F)| = \tau|D - F| = \emptyset$, as required. □

Proof of validity for algorithm 10. The algorithm must terminate in at most $\deg(D) + 1$ iterations since the degree of the divisor decreases by 1 in each iteration of the while-loop. To prove that the returned value R is equal to the rank, we proceed by induction on $r(D)$. For a base case, if $r(D) = -1$, then D is unwinnable, the while-loop never executes, and the algorithm returns $R = -1 = r(D)$. Now suppose that the algorithm is valid for all divisors of rank at most $m \geq -1$, and suppose that $r(D) = m + 1$. If $c + kq \sim D$ is the q -reduced divisor linearly equivalent to D , then c is superstable, so there exists $v \in \tilde{V}$ such that $c(v) = 0$. Some such v will be subtracted from $c + kq$ during the first run of the while-loop, and R will be incremented to $R = 0$. By the lemma, there exists an effective divisor E of degree $m + 2$ such that $E(v) > 0$ and $|D - E| = \emptyset$. Hence, the divisor $D - v$ has rank at most m , since subtracting the effective divisor $E - v$ of degree $m + 1$ yields a divisor with an empty linear system. It follows that $r(D - v) = m$, since subtracting v from D can decrease the rank by at most 1 (Exercise 5.3). By the induction hypothesis, the algorithm will run for $m + 1$ additional steps before terminating, when it will return $R = m + 1 = r(D)$. □

The potentially time-consuming part of Algorithm 10 occurs in the computation of the q -reduced divisor linearly equivalent to $c - v_i + kq$ in step 7. However, it turns out that the reduction may be accomplished through a single lending move by q followed by a single set-firing, which we now describe.

For notational purposes, let $c_i := c(v_{i+1})$ for all i so that $c = (c_0, \dots, c_{n-1})$. Next, assume that c is *increasing*, i.e., $c_0 \leq \dots \leq c_{n-1}$. This entails no loss of generality: at the initialization stage of the algorithm, we can add a step that sorts the components of c so that c is increasing. By symmetry of K_{n+1} , the rank does not change. Similarly, each time through the while-loop, after the q -reduction in

step 7, we can sort again (and, luckily, it turns out that the sorting required is especially easy).

Let's follow c through one iteration of the while-loop. For the following, recall that by Corollary 11.6, an increasing configuration c' on K_{n+1} is superstable if and only if $c' \leq (0, 1, \dots, n - 1)$. Since c is sorted and superstable at the beginning of the loop, we have $c_0 = 0$, so we may choose v_1 at step 6. Subtracting 1 from $c(v_1) = c_0$, then firing q transforms c as follows:

$$\begin{aligned} c = (0, c_1, \dots, c_{n-1}) &\rightarrow (-1, c_1, \dots, c_{n-1}) \\ &\rightarrow (0, c_1 + 1, \dots, c_{n-1} + 1) =: c'. \end{aligned}$$

Next, we want to superstabilize c' through a single set-firing.

If $c_i = i$ for some $i = 1, \dots, n - 1$, then let ℓ be the smallest such i , and fire the set $\{v_{\ell+1}, v_{\ell+2}, \dots, v_n\}$:

$$\begin{aligned} c' &= (0, c_1 + 1, \dots, c_{\ell-1} + 1, \overbrace{c_{\ell} + 1, \dots, c_{n-1} + 1}^{\text{fire}}) \\ &\quad \downarrow \\ &= (n - \ell, c_1 + n - \ell + 1, \dots, c_{\ell-1} + n - \ell + 1, c_{\ell} - \ell, \dots, c_{n-1} - \ell). \end{aligned}$$

Each of the vertices $v_{\ell+1}, \dots, v_{n-1}$ loses a dollar to each of v_1, \dots, v_{ℓ} and to q .

Cyclically permute components to get the increasing configuration

$$(11.1) \quad \tilde{c} := (c_{\ell} - \ell, \dots, c_{n-1} - \ell, n - \ell, c_1 + n - \ell + 1, \dots, c_{\ell-1} + n - \ell + 1).$$

Then $\tilde{c} \leq (0, 1, \dots, n - 1)$, hence superstable.

Otherwise, if $c_i < i$ for $i = 1, \dots, n - 1$, then $c' \leq (0, 1, \dots, n - 1)$. So it is already superstable. We set

$$\tilde{c} := c' = (0, c_1 + 1, \dots, c_{n-1} + 1),$$

which may be interpreted as the special case of equation (11.1) in which $\ell = n$ and the empty set of vertices is fired. In either case, vertex q loses n dollars when it fires, but then gains $n - \ell$ during superstabilization.

Thus, in one trip through the while loop, we might as well replace $c + kq$ with $\tilde{c} + (k - \ell)q$. Algorithm 11 incorporates this idea.

11.2.1. Dyck paths. A *balanced string of parentheses of order n* is a character string of length $2n$ consisting of n open parentheses, (, and n close parentheses,), such that as the string is read from left-to-right, at no time do more close parentheses appear than open parentheses. For example, the string $((()())()$ is balanced of order $n = 4$, and $(())()$ is not balanced.

Exercise 11.10. Find the five balanced strings of parentheses of order 3.

A *Dyck path of order n* is a walk in \mathbb{Z}^2 starting at $(0, 0)$ and ending at (n, n) where (i) each step in the walk is either north, adding $(0, 1)$, or east, adding $(1, 0)$, and (ii) at each lattice point (i, j) reached in the walk, $i \leq j$. Interpreting each step north as an open parenthesis and each step east as a closed parenthesis gives a bijection between the sets of balanced strings of parentheses and Dyck paths of order n . There is also a bijection between these sets and the set of increasing

Algorithm 11 Cori-Le Borgne algorithm, version 2.

```

1: INPUT: a divisor  $D \in \text{Div}(K_{n+1})$ .
2: OUTPUT: the rank of  $D$ .
3: initialization:  $R = -1$ 
4: compute the  $q$ -reduced divisor  $c + kq \sim D$ 
5: sort  $c$  so that  $c = (c_0, \dots, c_{n-1})$  with  $0 = c_0 \leq \dots \leq c_{n-1} < n$ 
6: while  $k \geq 0$  do
7:    $R = R + 1$ 
8:   find smallest  $\ell > 0$  such that  $c_\ell = \ell$ ,
9:   or let  $\ell = n$  if no such index exists
10:   $k = k - \ell$ 
11:  if  $k \geq 0$  then
12:    # superstabilize  $c$  after firing  $q$ 
13:     $c_0 = n - \ell$ 
14:     $(c_1, \dots, c_{\ell-1}) = (c_1 + n - \ell + 1, \dots, c_{\ell-1} + n - \ell + 1)$ 
15:     $(c_\ell, \dots, c_{n-1}) = (c_\ell - \ell, \dots, c_{n-1} - \ell)$ 
16:    # rotate  $c$ 
17:     $c = (c_\ell, \dots, c_{n-1}, c_0, \dots, c_{\ell-1})$ 
18:  return  $R$ 

```

superstables on K_{n+1} . Given a superstable $c = (c_0, \dots, c_{n-1})$ with $c_i \leq c_{i+1}$ for all i , create a Dyck path as follows: take c_0 steps east and then one step north. Then for $i = 0, \dots, n-2$, take $c_{i+1} - c_i$ steps east followed by one step north. Finish the path by taking as many steps east as necessary to reach (n, n) . See Figure 2 for an example.

Exercise 11.11. For each of the five balanced strings of parentheses of order 3, make a drawing as in Figure 2, showing the corresponding Dyck path and shading boxes to indicate the corresponding increasing superstable on K_4 .




Figure 2. The Dyck path of order 4 corresponding to the balanced string of parentheses $((())())$. The increasing superstable $(0, 0, 1, 3)$ on K_5 is encoded as the number of shaded squares in each row.

Our goal now is to visualize the Cori-Le Borgne algorithm in terms of Dyck paths. For concreteness, we take $n = 8$ and consider a divisor on K_9 . Recall that we have fixed an ordering of the vertices, v_0, \dots, v_8 with $q = v_0$. Take $D = c + kq$ where

$$c = 3v_4 + 3v_5 + 3v_6 + 5v_7 + 7v_8 = (0, 0, 0, 3, 3, 3, 5, 7)$$

and k is not yet specified.

Figure 3 depicts one iteration of the while-loop in Algorithm 11, transforming c into the configuration \tilde{c} as described starting on page 220. The number of shaded boxes in each row encodes the current state of c . From the initial state of c shown

on the left in (i), one can see that once the corresponding Dyck path leaves the lattice point $(0, 0)$, it next hits the diagonal at $(3, 3)$. Thus, $\ell = 3$ for this time through the while-loop.

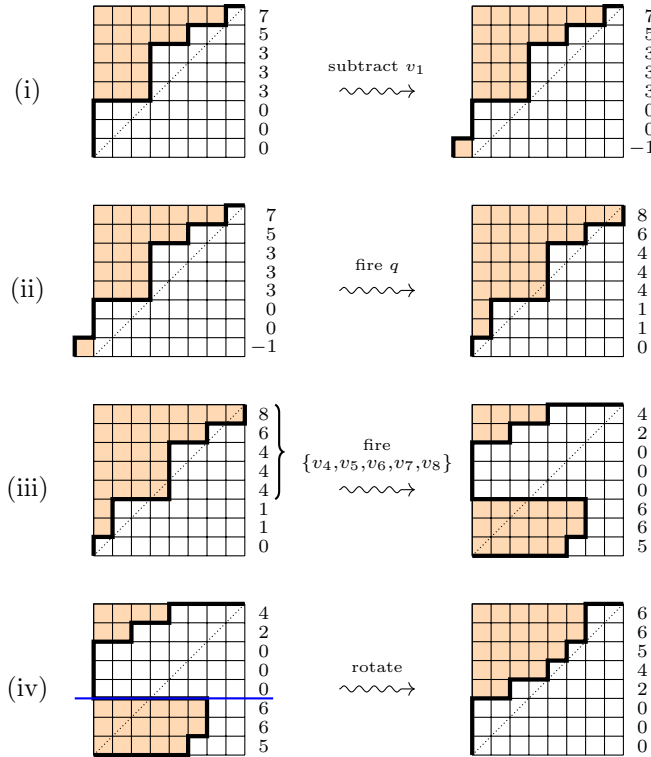


Figure 3. One iteration of the while-loop in Algorithm 11.

In (i), c is replaced by $c - v_1$. Firing q in step (ii) adds 1 to each component of c ; the fact that c is then not superstable can be seen from the two shaded boxes that jut over the diagonal. To superstabilize, since $\ell = 3$, we fire the set of vertices $S := \{v_4, v_5, v_6, v_7, v_8\}$. There are edges connecting these vertices to the four remaining vertices, including q , so they each lose \$4. Meanwhile, the remaining vertices each gain \$5, one from each vertex in S . Finally, in (iv), the configuration is rotated to obtain an increasing superstable configuration denoted earlier as \tilde{c} .

Figure 4 summarizes the transformation of c into \tilde{c} . It also suggests another way of thinking about the transformation. Consider the area between each Dyck path and the diagonal. We have divided it into two parts, blue and yellow, determined by where the Dyck path for c first reconnects with the diagonal, at (ℓ, ℓ) where $\ell = 3$. What is the effect of the while-loop on these regions? Their positions along the diagonal are swapped, but in the process, the blue region changes. If you tilt your head counter-clockwise, you might think of the two regions as mountains. The blue mountain has split into two: a tiny mountain consisting of half a square and a “sunken” version of the original. (Un-tilting your head, the sunken version of the blue mountain is obtained by scooting the blue region one unit to the right, then




Figure 4. Visualizing one iteration of the while-loop of Algorithm 11 through the action on the regions between the Dyck path and the diagonal.

excising everything under the diagonal.) The shape of the yellow mountain does not change. Figure 5 shows the next iteration of the while-loop. This time, $\ell = 4$. Again the blue mountain splits off a tiny blue mountain and sinks. The yellow mountain is unchanged.




Figure 5. Next iteration (see Figure 4).

Exercise 11.12. For each of the following, repeatedly iterate the while-loop of Algorithm 11, drawing pictures as in Figure 4. After enough iterations, a stable state it reached. Describe it.

- (1) Continue with $(0, 0, 0, 2, 4, 5, 6, 6)$, from Figure 4.
- (2) $c = (0, 0, 0, 0)$ on K_5 .
- (3) Try a couple of your own increasing superstables on complete graphs.

Now for the important question: Is there a way to read off the rank of $D = c + kq$ from these diagrams? We will assume that D has been processed so that c is an increasing superstable. Since c is superstable, $r(D) \geq 0$ if and only if $k \geq 0$. Each iteration of the while-loop essentially replaces D with the divisor $\tilde{c} + (k - \ell)q$, decreasing the rank by 1. Thus, the rank of D is one less than the number of times the while-loop is executed, and this number is completely determined by the succession of values attained by ℓ . Recall the geometric meaning of ℓ in our diagrams: (ℓ, ℓ) is where the Dyck path first reconnects to the diagonal after leaving $(0, 0)$. So we can find the rank of D by paying attention to these special points!

Consider our previous example, in which $c = (0, 0, 0, 3, 3, 3, 5, 7)$. Assuming $k \geq 0$, the while-loop of Algorithm 11 is executed. The resulting transformation, $c \rightsquigarrow \tilde{c}$ is again portrayed in Figure 6 but this time with numbers along the diagonals. In this iteration, we have $\ell = 3$, and hence q loses \$3. Thus, if $k \in \{0, 1, 2\}$, the




Figure 6. First iteration of the while-loop. Numbered diagonals.

while-loop will only be performed once; so $r(D) = 0$. If $k \geq 3$, the loop is executed again. Repeating the same argument, this time using the diagram on the right in Figure 6, this next iteration results in a loss to q of an additional \$4, hence, a net loss of \$7. So if $k \in \{3, 4, 5, 6\}$, no more iterations will occur, and since there are two iterations, $r(D) = 1$. Using this diagram to look ahead, one may see that q loses \$1 in both the third and fourth iterations, then loses \$2 in the fifth. So if $k = 7$, then $r(D) = 2$; if $k = 8$, then $r(D) = 3$; and if $k = 9$ or 10 , then $r(D) = 4$.

Notice if k is large enough, each corner in the original Dyck path eventually corresponds with an ℓ -value. Recalling the “sinking mountain” analogy, with repeated iterations each corner eventually sinks to the point where it touches the diagonal. However, also recall that with each iteration, the “blue mountain” spawns an additional tiny mountain, and hence an additional corner is formed. Figure 7 shows the




Figure 7. Fifth iteration.

fifth iteration of the while-loop. The numbers in bold signify the sum of ℓ -values encountered so far. It would take a long time to determine how $r(D)$ changes as k increases if it were necessary to repeatedly draw these diagrams representing each iteration. Happily, there is an alternative, which we describe next.

11.2.2. Cori-Le Borgne diagrams. Consider Figure 6 illustrating the first iteration of the while-loop in Algorithm 11. The diagonal boxes are numbered sequentially with those appearing in “corners” of the Dyck path typeset in boldface. After the first iteration, 3 dollars are lost from q , as one can see from where the Dyck path on the left first hits the diagonal (not counting the origin). To perform the next iteration, one would start with the Dyck path pictured on the right in the figure. It meets the diagonal at $(4, 4)$. So in the second iteration, 4 dollars are lost from q . The net loss is then \$7: the next number after 3 that appears in bold. In




Figure 8. Cori-Le Borgne diagram for the divisor $D = (0, 0, 0, 3, 3, 3, 5, 7) + kq$.

Figure 7, one sees from the boldface numbers that after the fifth iteration of the algorithm, \$11 dollars have been lost from q . Our algorithm continues until q goes into debt, and the number of iterations then determines the rank of the divisor. For example, if q started with less than \$11, then the rank of our divisor would be less than 5.

The *Cori-Le Borgne diagram* of a divisor keeps track of the sequence of net losses—the boldface numbers—produced by the iterations of the while-loop in Algorithm 11. This diagram for our running example appears in Figure 8. To describe it in general, let D be any divisor on K_{n+1} . We may assume that $D = c + kq$ where q is any chosen vertex and c is a superstable with respect to q . List the components of c in nondecreasing order, then draw the Dyck path for c . The Dyck path sits in a square grid with side length equal to n , the number of components of c . Extend this grid infinitely to the left, keeping the same height (cf. Figure 8). Starting at the lower-left endpoint of the Dyck path, just left of the Dyck path, itself, list the natural numbers in the grid squares along diagonals parallel to the diagonal for the Dyck path. If a number lies to the left of the Dyck path, write it in bold face. This includes those numbers lying to the left of the original square containing the Dyck path.

To determine the rank of $D = c + kq$ for each value of k from the Cori-LeBorgne diagram for D , let b be the number of bold values less than or equal to k . Then $r(D) = b - 1$. For example, from Figure 8 we can fill in the table in Figure 9 giving the complete list of ranks for $D = (0, 0, 0, 3, 3, 3, 5, 7) + kq$ as a function of k . The Cori-Leborgne diagram tidily encodes the rank of $D = c + kq$ by turning the algorithm’s “rotation” into the wrap-around numbering and the “sinking” into the left-shift of the numbering along successive diagonals.

k	< 0	$0, 1, 2$	$3, 4, 5, 6$	7	8	$9, 10$	11	$12, 13$	> 13
$r(c + kq)$	-1	0	1	2	3	4	5	6	$k - 7$

Figure 9. The rank of $D = (0, 0, 0, 3, 3, 3, 5, 7) + kq$ as a function of k .

Exercise 11.13. Convince yourself that the bold numbers in a Cori-Leborgne diagram record the net loss from q at each iteration of the algorithm's while-loop, as claimed.

Problems for Chapter 11

11.1. Consider a variation of the protocol for parking cars described in §11.1. There are still n cars, C_1, \dots, C_n , but this time there is one extra parking space, numbered $n + 1$, and the spaces are arranged in a circle. Car C_i prefers to park in space $p_i \in \{1, \dots, n + 1\}$. Other than that, the rules are essentially the same: each car in turn drives to its preferred spot and parks there if possible. Otherwise, it drives on to the next available spot. Since the spaces are arranged in a circle, each car will eventually park. Call these preference lists *circular parking functions* (terminology we learned from Matthias Beck).

- (a) After the cars park according to a given circular parking function, there is one empty parking space. Show that the number of circular parking functions that leave space i empty is the same as the number that leave space 1 empty, for each i .
- (b) Show that a circular parking function is an actual parking function if and only if it leaves space $n + 1$ empty.
- (c) Now conclude that the number of ordinary parking functions of length n is $(n + 1)^{n-1}$.

11.2. Draw a diagram as in Figure 8 in order to determine the rank of the divisor $D = (0, 0, 0, 1, 1, 4, 6, 6, 8, 22)$ on K_{10} .

11.3. Find the ranks of all divisors on K_4 . Let p_1, \dots, p_5 be the increasing parking functions of length 3, and define the superstable $c_i := p_i - \bar{1}$ for each i . Define $D_i := c_i - \deg(c_i)q$ for $i = 1, \dots, 5$. For each $d \in \mathbb{Z}$, the divisor classes of the $D_i + dq$ and of all the divisors obtained from them by permuting the components of the c_i are exactly the elements of $\text{Pic}^d(D)$. In other words, up to symmetry and linear equivalence, the divisors of degree d are exactly $D_i + dq$ for $i = 1, \dots, 5$. Make a table showing the rank of $D_i + dq$ as i and d vary. Use the Cori-Le Borgne diagrams of Section 11.2.2 to compute the ranks.

More about sandpiles

In this chapter we collect some additional topics about sandpiles. The setting is the same as Part 2: by a *graph* we mean a finite, connected, directed multigraph, possibly with loop edges.

12.1. Changing the sink

Consider the graph



with two edges from u to v and one from v to u . The sandpile group with respect to the sink vertex u is trivial while the sandpile group with respect to v has two elements. For a general graph, how does the sandpile group depend on the choice of sink?

Extending the definitions from Part 1 to directed graphs, we define *divisors* on a directed graph G to be elements of $\text{Div}(G) := \mathbb{Z}V$, the free abelian group on the vertices. Two divisors, D, D' are *linearly equivalent* if their difference lies in the image of the (full) Laplacian $L: \mathcal{M}(G) \rightarrow \text{Div}(G)$, in which case we write $D \sim D'$. The *degree* of a divisor D , denoted $\deg(D)$, is the sum of the coefficients of D . Just as before, we define the *Picard group*

$$\text{Pic}(G) := \text{Div}(G) / \text{im}(L) = \text{Div}(G) / \sim .$$

Even in the case of a directed graph, the entries in each column of the Laplacian sum to zero, so the image contains only degree-0 divisors. Therefore, it makes sense to define the Jacobian group,

$$\text{Jac}(G) := \text{Pic}^0(G) := \text{Div}^0(G) / \sim := \{[D] \in \text{Pic}(G) : \deg(D) = 0\}.$$

Note that the Jacobian group does not depend on the choice of a sink vertex.

Now suppose that G has a least one globally accessible vertex, and let $\det(\tilde{L})$ denote the reduced Laplacian with respect to one such vertex. Since the sum of the rows of L is $\vec{0}$ and $\det(\tilde{L}) \neq 0$ by the matrix-tree theorem 9.3, the rank of L is $|V| - 1$, and we may write

$$\ker(L) = \mathbb{Z}\tau$$

for some $\tau \in \mathcal{M}(G) = \mathbb{Z}^V$. In fact, we specified a generator τ explicitly in Section 9.2.2: for each vertex v , let $\theta(v)$ be the number of spanning trees rooted into v , and set $\gamma := \gcd\{\theta(v)\}$. Then $\tau = \theta/\gamma$.

Theorem 12.1. *Let $G = (V, E, s)$ be a sandpile graph with chosen sink s .*

(1) *There is a commutative diagram with exact rows*

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathbb{Z}^{\tilde{V}} & \xrightarrow{\tilde{L}} & \mathbb{Z}^{\tilde{V}} & \longrightarrow & \mathcal{S}(G, s) \longrightarrow 0 \\ & & \downarrow \iota & & \downarrow \varepsilon & & \downarrow \bar{\varepsilon} \\ 0 & \longrightarrow & \mathbb{Z}^V / \ker(L) & \xrightarrow{L} & \text{Div}^0(G) & \longrightarrow & \text{Jac}(G) \longrightarrow 0, \end{array}$$

where $\varepsilon(v) = v - s$ and $\iota(v) = [v] \in \mathbb{Z}^V / \ker(L)$ for all $v \in \tilde{V}$. (Here, as usual, we identify v with its characteristic function χ_v .)

(2) *For each $v \in \tilde{V}$, let n_v denote the number of edges in G directed from s to v . Let $\tilde{c} \in \mathcal{S}(G, s) \simeq \mathbb{Z}^{\tilde{V}} / \tilde{\mathcal{L}}$ be the unique recurrent such that*

$$\tilde{c} = - \sum_{v \in \tilde{V}} n_v v \pmod{\tilde{\mathcal{L}}}.$$

There is short exact sequence

$$0 \longrightarrow \mathbb{Z}_{\tau(s)} \xrightarrow{\alpha} \mathcal{S}(G, s) \xrightarrow{\bar{\varepsilon}} \text{Jac}(G) \longrightarrow 0,$$

where $\alpha(k) := k \cdot \tilde{c}$, and $\bar{\varepsilon}(c) := [c - \deg(c) s]$.

Proof. In part 1, the exactness of the rows is immediate. For commutativity,

$$\begin{aligned} \varepsilon \tilde{L}v &= \varepsilon \left(\text{outdeg}(v)v - \sum_{vw \in E: w \neq s} w \right) \\ &= \text{outdeg}(v)v - \sum_{vw \in E} w \\ &= Lv = L\iota v. \end{aligned}$$

The mapping $\bar{\varepsilon}: \mathcal{S}(G, s) \rightarrow \text{Jac}(G)$ is then induced by ε .

For part 2, apply the snake lemma (see Appendix B.2.5), using the fact that ε is invertible ($\varepsilon^{-1}(\sum_{v \in V} a_v v) = \sum_{v \in \tilde{V}} a_v v$):

$$\begin{array}{ccccccc}
 & & \ker \iota & & 0 & & \ker \bar{\varepsilon} \\
 & & \downarrow & & \downarrow & & \downarrow \\
 0 & \longrightarrow & \mathbb{Z}^{\tilde{V}} & \xrightarrow{\tilde{L}} & \mathbb{Z}^{\tilde{V}} & \longrightarrow & \mathcal{S}(G, s) \longrightarrow 0 \\
 & & \downarrow \iota & & \downarrow \varepsilon & & \downarrow \bar{\varepsilon} \\
 0 & \longrightarrow & \mathbb{Z}^V / \ker(L) & \xrightarrow{L} & \text{Div}^0(G) & \longrightarrow & \text{Jac}(G) \longrightarrow 0 \\
 & & \downarrow & & \downarrow & & \downarrow \\
 & & \text{cok } \iota & & 0 & & \text{cok } \bar{\varepsilon}
 \end{array}$$

By the snake lemma, there is an exact sequence


$$0 \longrightarrow \ker \iota \longrightarrow 0 \longrightarrow \ker \bar{\varepsilon} \longrightarrow \text{cok } \iota \longrightarrow 0 \longrightarrow \text{cok } \bar{\varepsilon} \longrightarrow 0.$$

Hence, $\ker \iota = 0$, $\text{cok } \bar{\varepsilon} = 0$, and $\ker \bar{\varepsilon} \simeq \text{cok } \iota$. Since every non-sink vertex is equivalent to 0 in $\text{cok } \iota$, we see that the cokernel is generated by $[s]$. But $[s]$ has order $\tau(s)$ in $\text{cok } \iota$, since $\ker L = \mathbb{Z}\tau$. It follows that $\ker \bar{\varepsilon} \simeq \text{cok } \iota \simeq \mathbb{Z}_{\tau(s)}$. All that remains is to describe the mapping $\alpha: \mathbb{Z}_{\tau(s)} \rightarrow \mathcal{S}(G, s)$ explicitly, and for this we just need to find a generator for the kernel of $\bar{\varepsilon}$. From the definition of the snake lemma mapping $\ker \bar{\varepsilon} \xrightarrow{\sim} \text{cok } \iota$ and the generator $[s]$ for $\text{cok } \iota$, we find the generator

$$c := \varepsilon^{-1}(L[s]) = \varepsilon^{-1}((\text{outdeg}(s) - (\#\text{loops at } s))s - \sum_{v \in \tilde{V}} n_v v) = - \sum_{v \in \tilde{V}} n_v v,$$

where n_v is the number of edges directed from s to v . Letting $\tilde{c} \in \mathcal{S}(G, s)$ denote the equivalent recurrent sandpile, we see that $\alpha: \mathbb{Z}_{\tau(s)} \rightarrow \mathcal{S}(G, s)$ is defined by $\alpha(k) = k \cdot \tilde{c}$ as claimed. \square

Exercise 12.2. Consider the following 2-vertex graph G :



with full Laplacian

$$L = \begin{pmatrix} 4 & -2 \\ -4 & 2 \end{pmatrix}.$$

There are 2 spanning trees rooted into u and 4 rooted into v , so $\tau = (1, 2)$ generates the kernel of L . The Jacobian of G is cyclic of order 2:

$$\text{Jac}(G) = \text{Div}^0(G) / \text{im}(L) = \mathbb{Z}(1, -1) / \mathbb{Z}(2, -2) = \mathbb{Z}_2.$$

First take $s = u$ for the sink, so that $\tilde{L} = (2)$ and $\mathcal{S}(G, u) = \{0, v\} \simeq \mathbb{Z}_2$. Using the notation from Theorem 12.1 (2), we have $\tilde{c} = 0$, the unique recurrent equivalent to $-4v$ modulo \tilde{L} . In this case, α is the zero map, and $\bar{\varepsilon}: \mathcal{S}(G, u) \rightarrow \text{Jac}(G)$ is an isomorphism, defined by $v \mapsto [v - u]$.

Now take $s = v$ for the sink, so $\tilde{L} = (4)$ and $\mathcal{S}(G, v) = \{0, u, 2u, 3u\} \simeq \mathbb{Z}_4$. Now we have $\tilde{c} = 2u$, and the map $\alpha: \mathbb{Z}_2 \rightarrow \mathcal{S}(G, v)$ is given by $1 \mapsto 2u$. The mapping $\bar{\varepsilon}: \mathcal{S}(G, v) \rightarrow \text{Jac}(G) = \mathbb{Z}_2$ defined by $u \mapsto [u - v]$ fits into the exact sequence

$$0 \longrightarrow \mathbb{Z}_2 \xrightarrow{\alpha} \mathcal{S}(G, v) \xrightarrow{\bar{\varepsilon}} \mathbb{Z}_2 \longrightarrow 0.$$

Note that this exact sequence is not split, since $\mathcal{S}(G, v) \simeq \mathbb{Z}_4 \not\cong \mathbb{Z}_2 \times \mathbb{Z}_2$.

In the case where G is undirected, we saw in Proposition 2.8 that the kernel of the Laplacian is generated over \mathbb{Z} by $\tau = (1, \dots, 1)$. So in that case, $\mathcal{S}(G) \simeq \text{Jac}(G)$ for all choices of sink vertex s . Since $\text{Jac}(G)$ does not depend on s , the sandpile group is determined up to isomorphism, independent of the choice of sink. Which directed graphs have this independence property?

The following corollary of Theorem 12.1 answers this question. (In its statement, we use the term *well-defined* to mean that the vertex chosen as the sink is globally accessible, i.e., from each vertex there is some directed path to the sink.)

Corollary 12.3. *The sandpile group $\mathcal{S}(G)$ is well-defined and independent of the choice of sink (up to isomorphism) if and only if G is Eulerian, in which case the mapping $\bar{\varepsilon}$ gives an isomorphism $\mathcal{S}(G) \approx \text{Jac}(G)$.*

Proof. If $\mathcal{S}(G)$ is well-defined and independent of choice of sink, then for all v , we have $\tau(v) > 0$, and by Theorem 12.1 (2),

$$|\mathcal{S}(G)|/\tau(v) = |\text{Jac}(G)|.$$

Since $\text{Jac}(G)$ is independent of v , it follows that τ is a constant vector, and hence, $\tau = (1, \dots, 1)$. But $\tau = (1, \dots, 1)$ is equivalent to $\text{indeg}(v) = \text{outdeg}(v)$ for all $v \in V$, which in turn is equivalent to G being Eulerian (Proposition A.29).

Conversely, if G is Eulerian, the existence of an Eulerian cycle implies every vertex is globally accessible. Further, $\text{indeg}(v) = \text{outdeg}(v)$ for all v ; so $\tau = (1, \dots, 1)$, and by Theorem 12.1 (2), $\mathcal{S}(G, v) \approx \text{Jac}(G)$ for all v . □

12.2. Minimal number of generators for $\mathcal{S}(G)$

Let $G = (V, E, s)$ be a sandpile graph, and let $\mu(G)$ denote the minimal number of generators for $\mathcal{S}(G)$. We first consider some bounds on $\mu(G)$, then consider what happens to $\mu(G)$ as edges are removed from G . The results presented here are rooted in work by Lorenzini ([70, 71, 72]) and collaboration with Hoppenfeld ([58]).

By the structure theorem for finitely generated abelian groups, discussed in Section 2.4 and Remark 2.34,

$$\mathcal{S}(G) \simeq \prod_{i=1}^{\mu(G)} \mathbb{Z}_{s_i}$$

for some integers $s_i > 1$ such that $s_i | s_{i+1}$ for all i . The s_i are the *invariant factors* of $\mathcal{S}(G)$ and are uniquely determined by $\mathcal{S}(G)$. Since $\mathcal{S}(G)$ is finite, its free part is trivial.

Since $\mathcal{S}(G) \simeq \mathbb{Z}\tilde{V}/\tilde{\mathcal{L}}$, the invariant factors of $\mathcal{S}(G)$ are exactly the invariant factors of \tilde{L} not equal to 1. Suppose G has n vertices. Then \tilde{L} is an $(n-1) \times (n-1)$ invertible matrix, and hence,

$$\mu(G) = n - 1 - t.$$

where t is the number of invariant factors of \tilde{L} equal to 1.

Recall from Section 2.4 that the i -th determinantal divisor, $d_i := d_i(\tilde{L})$, is the gcd of the $i \times i$ minors of \tilde{L} and turns out to be the product of the first i invariant factors of \tilde{L} . (Since \tilde{L} is invertible, none of its determinantal divisors are 0.) Therefore, the t in the above displayed equation is the maximal i such that $d_i = 1$.

Of course, $0 \leq \mu(G) \leq n-1$. By the matrix-tree theorem 9.3, we have $\mu(G) = 0$ if and only if G contains exactly one directed spanning tree rooted at the sink. Thus, if G is undirected, then $\mu(G) = 0$ if and only if G is a tree. In general, if G is an undirected simple graph, then \tilde{L} will contain an entry equal to ± 1 , and hence, $d_1 = 1$ and $\mu(G) \leq n - 2$. By Problem 2.6, we know that $\mathcal{S}(K_n) \simeq \mathbb{Z}_n^{n-2}$. In fact, if G is a subgraph of K_n , then $\mu(G) = n - 2$ if and only if $G = K_n$ (cf. Problem 12.2).

Example 12.4. Figure 1 displays all connected 5-vertex subgraphs of the complete graph K_5 and their invariant factors. Graphs on the same horizontal level have the same number of edges. A line connects a graph G on one level to a graph G' one level down if G' is obtained from G by removal of an edge. The line is black if the edge-removal does not change the number of invariant factors, blue if the number decreases, and red if it increases. Note that removal of an edge changes the number of invariant factors by at most one, a general fact explained by Theorem 12.6.

We now concentrate on how the minimal number of generators for the sandpile group changes when a graph is modified by adding or removing edges. The modification can be expressed in terms of a change in reduced Laplacians. In general, let M be any $m \times n$ integer matrix, and define $\mu(M)$ to be the minimal number of generators of $\text{cok}(M)$. Therefore, $\mu(M) = m - t$ where t is the number of invariant factors of M equal to 1. For a sandpile graph G with reduced Laplacian \tilde{L} , we have $\mu(G) = \mu(\tilde{L})$. We are interested in the following lemma in the case when M and M' are reduced Laplacians of graphs obtainable from each other by adding or removing edges.

Lemma 12.5. *Let M and M' be $m \times n$ integer matrices. Then*

$$|\mu(M) - \mu(M')| \leq \text{rk}(M - M').$$

Proof. For notation, let Λ_i be the collection of all \mathbb{Z} -linear combinations of $i \times i$ minors of M , let $d_i := d_i(M)$ be the i -th determinantal divisor of M , and let s_i be the i -th invariant factor of M . Therefore, $\Lambda_i = (d_i)$, the principal ideal consisting of all integer multiples of d_i , and by Theorem 2.33, $d_i = \prod_{j=1}^i s_j$. Let Λ'_i be the \mathbb{Z} -linear combinations of $i \times i$ minors of M' , and let d'_i and s'_i be the corresponding determinantal divisor and invariant factor for M' .

For each i , we will show that

$$(12.1) \quad \Lambda'_i \subseteq \Lambda_{i-r}$$

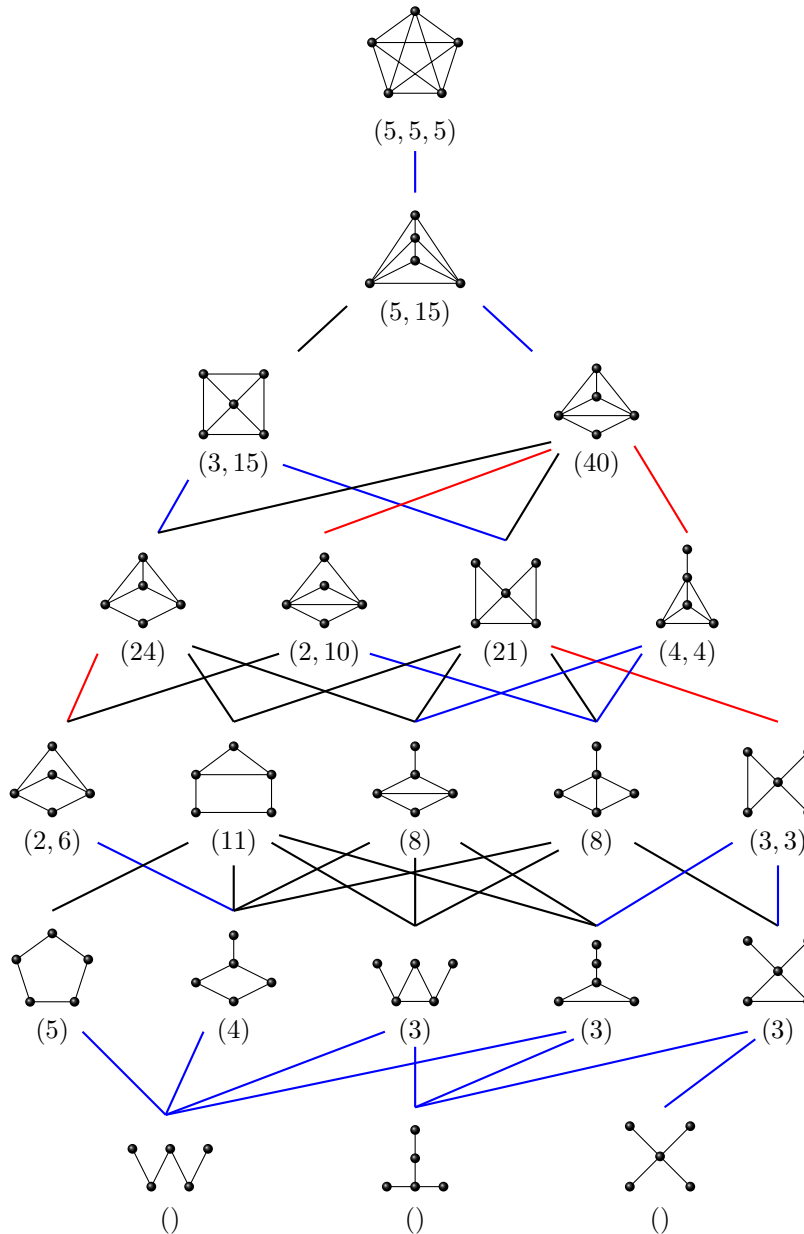


Figure 1. Invariant factors of the connected 5-vertex subgraphs of K_5 for Example 12.4. Lines between graphs indicate an edge removal, and colors indicate a change in the number of invariant factors: black = no change; blue = decrease; red = increase.

where $r := \text{rk}(M - M')$. (For this purpose, take $\Lambda_i = (1) = \mathbb{Z}$ and $d_i = 1$ if $i \leq 0$, and take $\Lambda_i = \{0\}$ if i is larger than the number of rows or columns of M , and similarly for M' .) To see that this implies the lemma, let t be the number of

invariant factors of M' equal to 1. Then $d'_t = 1$ and $\mu(M') = m - t$. Once (12.1) is established, it follows that $d_{t-r} = 1$, and thus, $\mu(M) \leq m - (t - r) = \mu(M)' + r$. By symmetry, we also have $\mu(M') \leq \mu(M) + r$, and the result follows.

To prove (12.1), take some i such that $1 \leq i \leq \min\{m, n\}$, the other cases being trivial. Replacing M' by an arbitrary $i \times i$ submatrix of M' and replacing M by its corresponding submatrix, we may assume M' and M are $i \times i$ matrices. Our task is to show $\det(M') \in \Lambda_{i-r}$.

Say the columns of M are u_1, \dots, u_i and the columns of $M' - M$ are v_1, \dots, v_i . Then

$$\det(M') = \det(u_1 + v_1, \dots, u_i + v_i).$$

Expanding the right-hand side of this equation using the fact that the determinant is multilinear and alternating, it follows that $\det(M')$ is a sum of terms of the form

$$(12.2) \quad \pm \det(u_{\ell_1}, \dots, u_{\ell_k}, v_{\mu_1}, \dots, v_{\mu_{i-k}}).$$

Writing each u_{ℓ_p} as a linear combination of the standard basis vectors, $\{e_j\}$, and expanding further, each of these terms is a \mathbb{Z} -linear combination of terms of the form

$$\det(e_{\nu_1}, \dots, e_{\nu_k}, v_{\mu_1}, \dots, v_{\mu_{i-k}}),$$

which is an $(i - k) \times (i - k)$ minor of the matrix with columns $v_{\mu_1}, \dots, v_{\mu_{i-k}}$, and thus, an $(i - k) \times (i - k)$ minor of $M' - M$. If $i - k > r = \text{rk}(M - M')$, then this minor is 0. Hence, we may assume $i - r \leq k$.

Next, go back to (12.2) and this time write the v_{μ_p} in terms of the standard basis vectors. Expand (12.2) to get a \mathbb{Z} -linear combination of $k \times k$ minors of M . In general, any $j \times j$ minor of a matrix is a linear combination of its $(j - 1) \times (j - 1)$ minors, as can be seen by computing a determinant by expanding along a row. In our case, since $i - r \leq k$, repeated application of this fact shows that $\det(M')$ is a linear combination of $(i - r) \times (i - r)$ minors of M , as required. \square

Theorem 12.6. *Let G and G' be sandpile graphs on the same vertex set with the same sink, and let U be a subset of the vertices of size $r + 1$. Suppose the edge sets of G and G' differ only by edges with both endpoints in U . Then*

$$|\mu(G) - \mu(G')| \leq r.$$

Proof. Let \tilde{L} and \tilde{L}' be the reduced Laplacians for G and G' , respectively. For each non-sink vertex v , let e_v be the v -th standard basis vector, and let $e_s := 0$ for the sink vertex, s . Then the column span of $\tilde{L} - \tilde{L}'$ is contained in $\text{Span}_{\mathbb{Z}}\{e_u - e_v : u, v \in U\}$, a free abelian group of rank r . Hence, $\text{rk}(\tilde{L} - \tilde{L}') \leq r$, so the result follows from Lemma 12.5. \square

Corollary 12.7. *Let G and G' be sandpile graphs on the same vertex set with the same sink.*

(1) *If G' is obtained from G by removal of a single edge or, more generally, by removal of any set of edges joining a fixed pair of vertices, then*

$$|\mu(G) - \mu(G')| \leq 1.$$

- (2) $\mu(G) \leq |\tilde{E}| - |V| + 1$ where \tilde{E} is the set of edges for the underlying simple undirected graph (e.g., $\tilde{E} = E$ if G is a simple undirected graph).
- (3) Let U be a subset of the vertices of size $r + 1$. Suppose that by removing only edges with both endpoints in U one may obtain a directed spanning tree of G rooted into the sink. Then, $\mu(G) \leq r$.

Proof. Part (1) is immediate from Theorem 12.6. By the matrix-tree theorem theorem 9.3, G has at least one spanning tree rooted into the sink. Choosing such a spanning tree T , note that it has $|V| - 1$ edges. Then part (2) follows from part (1) and the fact that the sandpile group of T is trivial. Similarly, part (3) follows from Theorem 12.6. □

As an application of the ideas presented above, we ask: How many edges must one remove from the complete graph K_n in order to obtain a sandpile graph with cyclic sandpile group? Since $\mathcal{S}(K_n) \simeq \mathbb{Z}_n^{n-2}$, we have $\mu(K_n) = n - 2$, so according to Corollary 12.7 (1), we must remove at least $n - 3$ edges. In fact, with proper choices, a cyclic sandpile group is always achievable by removing exactly $n - 3$ edges. One way of doing this is to remove a path graph from K_n . For an illustration of the following proposition, see Figure 2.

Proposition 12.8. *Let G be a sandpile graph obtained from the complete graph K_n by removing k (undirected) edges.*

- (1) *If $\mathcal{S}(G)$ is cyclic, then $k \geq n - 3$.*
- (2) *If G is obtained by removing a path graph of length $n - 3$, then $\mathcal{S}(G)$ is cyclic of order $nU_{n-3}(\frac{n-2}{2})$, where U_j is the j -th Chebyshev polynomial¹ of the second kind.*

Proof. We have already discussed part (1). For part (2), let $1, \dots, n$ be the vertices of K_n . By symmetry, we may assume G is obtained by removing the path graph on vertices $1, \dots, n - 2$, where vertices 1 and $n - 2$ have degree 1. The Laplacian for G is the $n \times n$ matrix

$$L = \begin{pmatrix} n-2 & 0 & -1 & -1 & \cdots & -1 & -1 & -1 & -1 \\ 0 & n-3 & 0 & -1 & \cdots & -1 & -1 & -1 & -1 \\ -1 & 0 & n-3 & 0 & \cdots & -1 & -1 & -1 & -1 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ -1 & -1 & -1 & -1 & \cdots & n-3 & 0 & -1 & -1 \\ -1 & -1 & -1 & -1 & \cdots & 0 & n-2 & -1 & -1 \\ -1 & -1 & -1 & -1 & \cdots & -1 & -1 & n-1 & -1 \\ -1 & -1 & -1 & -1 & \cdots & -1 & -1 & -1 & n-1 \end{pmatrix}$$

Choosing n as the sink vertex, drop the last row and column from L to obtain the reduced Laplacian, \tilde{L} , of G . To show $\mathcal{S}(G)$ is cyclic, it suffices to show the $(n - 2)$ -nd determinantal divisor for \tilde{L} is 1. To see this, start with the matrix \tilde{L} and perform the following operations, in order: (i) subtract the last row of \tilde{L} from all other rows, (ii) add all but the last column to the last column, and (iii) add the last row to the second-to-last row. Finally, dropping the first column and last row

¹The definition of Chebyshev polynomial follows the proof of this proposition.

of the resulting matrix leaves an $(n - 2) \times (n - 2)$ lower-triangular matrix with 1s on the diagonal, and hence determinant 1. Thus, G is cyclic.

To compute the size of $\mathcal{S}(G)$, let J be the $n \times n$ matrix with each entry equal to 1. Then by Proposition 9.12, the number of spanning trees, hence the size of $\mathcal{S}(G)$ is $\det(L + J)/n^2$. We have

$$L + J = \begin{pmatrix} n-1 & 1 & & & & & & & & & \\ & 1 & n-2 & 1 & & & & & & & \\ & & 1 & n-2 & 1 & & & & 0 & & \\ & & & \ddots & \ddots & & & & & & \\ & & & & 1 & n-2 & 1 & & & & \\ & 0 & & & & 1 & n-1 & 0 & & & \\ & & & & & & 0 & n & 0 & & \\ & & & & & & & 0 & n & & \end{pmatrix}.$$

Let A be the matrix obtained by dropping the last two rows and columns from $L + J$. Then

$$|\mathcal{S}(G)| = \det(L + J)/n^2 = \det(A).$$

The matrix A is closely related to a particular evaluation of a Chebyshev polynomial of the second kind:

$$U_{n-2}\left(\frac{n-2}{2}\right) = \det \begin{pmatrix} n-2 & 1 & & & & & & & & & \\ & 1 & n-2 & 1 & & & & & & & \\ & & 1 & n-2 & 1 & & & & & & \\ & & & \ddots & \ddots & & & & & & \\ & & & & 1 & n-2 & 1 & & & & \\ & 0 & & & & 1 & n-2 & 1 & & & \\ & & & & & & 1 & n-2 & & & \end{pmatrix}.$$

Letting r_i denote the i -th row of the above matrix, and letting e_i denote the i -th standard basis vector, we have

$$\begin{aligned} \det(A) &= \det(r_1 + e_1, r_2, \dots, r_{n-3}, r_{n-2} + e_{n-2}) \\ &= \det(r_1, r_2, \dots, r_{n-3}, r_{n-2}) + \det(e_1, r_2, \dots, r_{n-3}, r_{n-2}) \\ &\quad + \det(r_1, r_2, \dots, r_{n-3}, e_{n-2}) + \det(e_1, r_2, \dots, r_{n-3}, e_{n-2}) \\ &= U_{n-2}\left(\frac{n-2}{2}\right) + 2U_{n-3}\left(\frac{n-2}{2}\right) + U_{n-4}\left(\frac{n-2}{2}\right). \end{aligned}$$

Expanding the determinant defining Chebyshev polynomials of the second kind gives the identity $U_{n-2}(x) = 2xU_{n-3}(x) - U_{n-4}(x)$, from which it follows that $\det(A) = nU_{n-3}(\frac{n-2}{2})$, as desired. \square

Definition 12.9. A square matrix M is *tridiagonal* if $M_{ij} = 0$ whenever $|i - j| > 1$. The determinants of the following two $j \times j$ tridiagonal matrices are called the j -th




Figure 2. The complete K_9 after removing a path graph of length 6 (dashed). In accordance with Proposition 12.8, the sandpile group is cyclic of order $9 \cdot U_6(\frac{7}{2}) = 953433$, where $U_6(x) = 64x^6 - 80x^4 + 24x^2 - 1$ is a Chebyshev polynomial of the second kind.

Chebyshev polynomials of the first and second kind, respectively:

$$T_j(x) = \det \begin{pmatrix} x & 1 & & & & & & & \\ 1 & 2x & 1 & & & & & & \\ & & 1 & 2x & 1 & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2x & 1 & \\ & & & & & & & 1 & 2x \end{pmatrix}, \quad U_j(x) = \det \begin{pmatrix} 2x & 1 & & & & & & & \\ 1 & 2x & 1 & & & & & & \\ & & 1 & 2x & 1 & & & & \\ & & & & \ddots & & & & \\ & & & & & 1 & 2x & 1 & \\ & & & & & & & 1 & 2x \end{pmatrix}.$$

Equivalently, they are defined by the recurrences:

$$(12.3) \quad \begin{aligned} T_0(x) &= 1 \\ T_1(x) &= x \\ T_j(x) &= 2xT_{j-1}(x) - T_{j-2}(x) \quad \text{for } j \geq 2, \end{aligned}$$

and

$$(12.4) \quad \begin{aligned} U_0(x) &= 1 \\ U_1(x) &= 2x \\ U_j(x) &= 2xU_{j-1}(x) - U_{j-2}(x) \quad \text{for } j \geq 2. \end{aligned}$$

12.3. M-matrices

In this section we consider a generalization of the sandpile model. Define

$$Z_n := \{A \in M_{n \times n}(\mathbb{Z}) : A_{ij} \leq 0 \text{ for } i \neq j\},$$

the set of $n \times n$ integer matrices with nonpositive off-diagonal entries. Let $\tilde{S} := \{s_1, \dots, s_n\}$ be a collection of sites, and fix $A \in Z_n$. The idea now is to mimic the sandpile model, using A in place of the reduced Laplacian matrix.

Think of A as a mapping $\mathbb{Z}^{\tilde{S}} \rightarrow \mathbb{Z}^{\tilde{S}}$ where As_i is the i -th column of A . A configuration on \tilde{S} is an element $c \in \mathbb{Z}^{\tilde{S}} \simeq \mathbb{Z}^n$. Firing s_i from c results in the configuration $c - As_i$. A site $s_i \in \tilde{S}$ is unstable in c if $c(i) \geq A_{ii}$, and it is then legal to fire c . If $\sigma \in \mathbb{Z}^{\tilde{S}}$, the corresponding script-firing produces the configuration

$c - A\sigma$. Our previous definitions of *recurrent* and *superstable* configurations now make sense in this new setting.

Example 12.10. Let

$$A = \begin{pmatrix} 2 & -4 \\ -2 & 3 \end{pmatrix}.$$

The configuration $c = 2s_1 = (2, 0)$ is unstable since $c(s_1) = 2 \geq A_{11}$. Firing s_1 yields the stable configuration

$$c^\circ = \begin{pmatrix} 0 \\ 2 \end{pmatrix} = c - As_1 = \begin{pmatrix} 2 \\ 0 \end{pmatrix} - \begin{pmatrix} 2 & -4 \\ -2 & 3 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

Exercise 12.11. Is c° superstable?

Continuing the above example, consider the configuration $(0, 3)$. We have the following sequence of legal vertex firings:

$$(0, 3) \xrightarrow{s_2} (4, 0) \xrightarrow{s_1} (2, 2) \xrightarrow{s_1} (0, 4) \xrightarrow{s_2} (4, 1) \xrightarrow{s_1} \dots$$

This configuration will *never* stabilize.

Gabrielov [46], has shown that given any $A \in Z_n$, if a configuration has a stabilization (through legal vertex firings), then that stabilization and its firing script is unique, and the least action principle, Theorem 6.7, still holds. Further, whether a configuration will stabilize does not depend on the order of legal vertex firings.

Definition 12.12. A matrix $A \in Z_n$ is *avalanche finite* if every configuration $c \in \mathbb{Z}\tilde{S}$ has a stabilization.

For example, we have seen that the reduced Laplacian of a sandpile graph is always avalanche finite. It turns out avalanche finite matrices are well-known and useful not just in the context of sandpiles. They have applications in economics, game theory, probability and statistics, and finite-element analysis for partial differential equations. (See [46], [52], and [80] for references.) Avalanche finite matrices are exactly the nonsingular M -matrices which we now describe.

Definition 12.13. A matrix $A \in Z_n$ is an M -matrix if $A = rI - B$ for some matrix B with nonnegative entries and some real number $r \geq \rho(B)$, the maximum of the moduli of the eigenvalues of B .

Theorem 12.14 ([80]). *Let $A \in Z_n$. Then A is a nonsingular M -matrix if it satisfies any of the following equivalent conditions.*

- (1) *The matrix A is avalanche finite.*
- (2) *The real part of each eigenvalue of A is positive.*
- (3) *Every real eigenvalue of A is positive.*
- (4) *A^{-1} exists and its entries are nonnegative.*
- (5) *All of the principal minors of A are positive. (A principal minor is the determinant of a submatrix formed by removing a set, possibly empty, of rows and the corresponding set of columns, i.e., rows and columns with the same indices.)*

- (6) The matrix $A + rI$ is nonsingular for all $r \geq 0$.
- (7) There exists $\sigma \geq 0$ such that $A\sigma > 0$.
- (8) The matrix A does not reverse signs: $\sigma(s) \cdot ((A\sigma)(s)) > 0$ for all $s \in \tilde{S}$ such that $\sigma(s) > 0$.
- (9) $A\sigma \geq 0 \Rightarrow \sigma \geq 0$.
- (10) There exists $\sigma > 0$ with $A\sigma \geq 0$ such that if $(A\sigma)(s_{i_0}) = 0$, then there exist $1 \leq i_1 \leq \dots \leq i_r \leq n$ such that $A_{i_k i_{k+1}} \neq 0$ for $0 \leq k \leq r-1$ and $(A\sigma)(s_{i_r}) > 0$.

We have already seen some of these conditions for the case where A is the reduced Laplacian of a sandpile graph. Condition (1) is the existence of a stabilization for the sandpile model, Theorem 6.12. Theorem 8.29 implies the nonnegativity of the entries of \tilde{L}^{-1} , condition (4). Condition (9) is Lemma 7.14. The existence of a burning sandpile implies condition (10).

Given $A \in Z_n$, define the configuration c_{\max} by $c_{\max}(s) = A_{ss} - 1$ for all $s \in \tilde{S}$.

Theorem 12.15 ([52]). *Let A be a nonsingular M -matrix.*

- (1) (Existence and uniqueness.) *In each equivalence class of configurations modulo the image of A , there is a unique recurrent and a unique superstable configuration.*
- (2) (Duality.²) *A configuration c is recurrent if and only if $c_{\max} - c$ is superstable.*

Exercise 12.16. Let

$$A = \begin{pmatrix} 2 & -3 \\ -3 & 6 \end{pmatrix}.$$

- (1) Show that A is a nonsingular M -matrix.
- (2) Find all recurrents and superstables.
- (3) Was there a quick way of determining the number of recurrents before explicitly finding all of them?

Remark 12.17. (1) We have defined integer M -matrices. Instead, we could replace Z_n by the set of $n \times n$ real matrices with nonpositive off-diagonal entries, and then Definition 12.13 yields the collection of real M -matrices, for which Theorem 12.14 still holds. In this case, we consider real-valued configurations, but firing scripts are still integer-valued, so A is viewed as a mapping $\mathbb{Z}^{\tilde{S}} \rightarrow \mathbb{R}\tilde{S}$.

- (2) In fact, [80] provides 40 equivalent conditions for non-singular M -matrices, and Theorem 12.14 is only a sample.

12.4. Self-organized criticality

The *Gutenberg-Richter Law* ([51]) is an empirical relationship describing the frequency of earthquakes as a function of their *magnitude*, defined to be proportional to the logarithm of the maximal amplitude reading of a standard seismogram located at a standard distance from the quake epicenter. Fixing a suitable geographic region and sufficiently long time-scale, let $N(m)$ denote the number of earthquakes of

²In fact, our proof of (1 \Rightarrow 2) in Theorem 7.12 is lifted directly from [52].

magnitude m . Then the Gutenberg-Richter Law states that the logarithm of $N(m)$ decreases linearly with m :

$$\log_{10} N(m) \approx a - bm.$$

Matching this formula to data, the slope b is generally found to be approximately 1, whereas the intercept a depends on the overall level of seismic activity in the region.

Equivalently, setting $s := 10^m$ as a measure of the intensity of an earthquake, and writing $\tilde{N}(s) := N(\log_{10}(s))$ for the number of earthquakes of intensity s , we find the *power law*

$$\tilde{N}(s) \approx As^{-b},$$

where $A := 10^a$. This implies a lack of preferred scale for earthquakes in the following sense: $\tilde{N}(s)s^b = A$ does not depend on s . Recalling that $b \approx 1$, we may interpret this fact as follows: *over a sufficiently long time period, the total intensity produced by quakes of a particular intensity s is independent of s .* As stated in [91]: “Nature expends the same total intensity shaking the earth at one point on the Gutenberg-Richter scale as it does at any other point on that scale.” Power laws and the corresponding lack of preferred scale occur in many other contexts, including the frequency of words in natural languages (Zipf’s law), the population of cities, the distribution of incomes, and self-similar fractal patterns in natural formations such as coastlines.

In an influential 1987 paper [4], Bak, Tang, and Wiesenfeld introduced their notion of *self-organized criticality* (SOC), which they offered as a general mechanism to explain the ubiquity of power-laws and scale-free structures in natural and social phenomena. They studied toy models of slowly-driven dissipative physical systems such as the abelian sandpile in the expectation that these models would display SOC. Here is Dhar writing about sandpiles and self-organized criticality 10 years later ([35]):

The sandpile model was proposed as a paradigm of self-organized criticality (SOC). It is certainly the simplest, and best understood, *theorist’s* model of SOC: it is a non-equilibrium system, driven at a slow steady rate, with local threshold relaxation rules, which in the steady state shows relaxation events in bursts of a wide range of sizes, and long-range spatio-temporal correlations, obtained without fine-tuning of any control parameters.

Many of the features listed by Dhar should be familiar to the reader from Part 2:

- *non-equilibrium system*: sand is regularly being added to the sandpile, as well as lost to the sink;
- *driven at a slow steady rate*: the additional grains come one-by-one, with time allowed between grains for stabilization to occur through toppling (avalanches);
- *local threshold relaxation rules*: a vertex topples only when it becomes unstable, and the stability threshold is the degree of the vertex, determined by its immediate neighbors;
- *steady state*: a random walk on the recurrent sandpiles (Corollary 8.28);

- *relaxation events in bursts of a wide range of sizes*: adding a single grain of sand to a recurrent sandpile may immediately result in another recurrent (avalanche size zero) or it may result in an unstable sandpile that must topple many times before stabilizing (a large avalanche).

It remains to explain the final two features: *long-range spatio-temporal correlations* and *without fine-tuning of any control parameters*. For this, we must briefly and informally describe the phenomenon of continuous thermodynamic phase transitions. So consider a ferromagnetic substance (like iron), which you might imagine³ as a regular lattice of atoms, each of which has a physical quantity called *spin* that may be in one of two states: *up* or *down*. It is energetically favorable for neighboring atoms to have the same spin rather than opposite, so at low temperatures the spins will tend to align with each other—at equilibrium the system will fluctuate around either the all-up or the all-down state. At high temperatures, thermal fluctuations will disrupt the spin-alignment, so at equilibrium the system will fluctuate around a state of average spin zero, and the correlation between spins will be extremely short-range. More precisely, there will be a temperature dependent *correlation length* $\xi(T)$ that sets the scale for the exponential decay of correlation as a function of separation r via $\exp(-r/\xi(T))$: nearby atoms tend to fluctuate together, while atoms separated by more than the correlation distance fluctuate independently. Starting at high temperatures and slowly cooling, the system will pass through a *critical temperature* T_c which separates the two regimes. The correlation length $\xi(T)$ diverges to infinity as $T \rightarrow T_c$, while the correlation function changes from exponential decay to a power law $r^{-\tau}$ for some *critical exponent* τ .

The appearance of the power law correlation function $r^{-\tau}$ at the critical temperature T_c indicates that there is no characteristic length scale for the system, so that clusters of spin-aligned atoms of all sizes appear, leading to a fractal-like spatial structure. Similarly, there is no characteristic time scale for the formation and persistence of spin-aligned regions, so that cluster lifetimes of all durations will appear. These are the long-range spatio-temporal correlations mentioned by Dhar in the quote above. Note the essential fact that this critical behavior is dependent on fine-tuning the control parameter of temperature to the critical value T_c .

In contrast, the hallmark of SOC is that such long-range correlations occur *without* the fine-tuning of any parameter—the system organizes itself into a critical state. In the abelian sandpile model, the steady state is expected to be critical in the sense of displaying a power law distribution of avalanche sizes and durations. These expectations for avalanche statistics are born out by numerical simulations, and in some cases there are rigorous results indicating the presence of power laws—see, for example, [12]. Overall, the idea of self-organized criticality has been extremely influential in a variety of fields, and it has also generated much controversy—for a survey, see [92].

³We are roughly describing the famous *Ising model* of magnetism. For further details, see, e.g., [83, Section 8.1 and Chapter 12].

Problems for Chapter 12

12.1. Let G be Eulerian, and let $\mathcal{S}(G, s)$ and $\mathcal{S}(G, s')$ denote the sandpile group of G with respect to sink vertices, $s, s' \in V$, respectively. Describe an isomorphism $\mathcal{S}(G, s) \simeq \mathcal{S}(G, s')$.

12.2. Suppose that G is a subgraph of K_n but not equal to K_n . Show that $\mu(G)$, the minimal number of generators for G , is strictly less than $\mu(K_n) = n - 2$.

12.3. Let C_m be the cycle graph with vertices u_0, \dots, u_{m-1} and let C_n be the cycle graph with vertices v_0, \dots, v_{n-1} . Let $G_{m,n}$ be the graph obtained by identifying the vertices u_0 and v_0 and adding the edge $\{u_1, v_{n-1}\}$. See Figure 3 for $G_{4,6}$.

- (a) Show that $\mathcal{S}(G)$ is cyclic of order $3mn - m - n$.
- (b) Removing the edge $\{u_1, v_{n-1}\}$ results in a graph with sandpile group isomorphic to $\mathbb{Z}_m \times \mathbb{Z}_n$. Thus, removing this edge causes the minimal number of generators of the sandpile group to *increase* by 1 if and only if $\gcd(m, n) \neq 1$. Find a graph G with a sequence of k edges e_1, \dots, e_k such that (i) $G_i := G \setminus \{e_1, \dots, e_i\}$ is a sandpile graph for $i = 0, \dots, k$ (with $G_0 := G$) and (ii) $\mu(G_{i+1}) = \mu(G_i) + 1$ for all i , i.e., the minimal number of generators for the sandpile group increases with each edge removal.




Figure 3. $G_{4,6}$ for Problem 12.3

12.4. (See [46], Example 1.17).

- (a) Let $A \in Z_n$, i.e., A is an integer matrix with nonpositive entries off of the main diagonal. Furthermore, assume that the diagonal entries of A are nonnegative. Show that if $n \leq 3$, then A is avalanche finite if and only if $\det(A) > 0$.
- (b) Consider the matrix with determinant $16 > 0$,

$$A = \begin{pmatrix} 1 & -3 & -1 & 0 \\ -3 & 1 & -1 & 0 \\ -1 & -1 & 1 & -2 \\ 0 & 0 & -2 & 1 \end{pmatrix}.$$

Show that A is not avalanche finite by exhibiting a configuration that will not stabilize.

Cycles and cuts

In this chapter, we relate the sandpile group to the algebraic theory of cycles and cuts. Unless stated otherwise, $G = (V, E)$ will denote an undirected multigraph, not necessarily connected, and possibly containing loop edges. See Appendix A for background and terminology from graph theory.

13.1. Cycles, cuts, and the sandpile group

In order to develop the algebraic theory of cycles and cuts, we need to choose an orientation \mathcal{O} for the undirected multigraph G and consider the directed multigraph (G, \mathcal{O}) . However, we will see that the main results are independent of the choice of orientation. If the vertices of G are ordered as v_1, \dots, v_n , then we will generally choose the *standard orientation*, which assigns the directed edge (v_i, v_j) to the undirected edge $\{v_i, v_j\}$ whenever $i < j$.

Let $\mathbb{Z}E$ be the free abelian group on the undirected edges of G . In the case where E is a multiset, copies of edges are treated as distinct in $\mathbb{Z}E$. For example, if G is the banana graph B_2 consisting of two vertices connected by two edges, then $\mathbb{Z}E \simeq \mathbb{Z}^2$. As usual, if $g = \sum_{e \in E} a_e e$ is an element of $\mathbb{Z}E$, then the *support* of g , denoted $\text{supp}(g)$, is the set of edges e for which the coefficient a_e is nonzero.

The orientation \mathcal{O} allows us to define the *boundary* of an edge $e \in E$ as $\partial e := e^+ - e^- \in \mathbb{Z}V$. Extending linearly defines the *boundary map*,

$$\partial: \mathbb{Z}E \rightarrow \mathbb{Z}V.$$

Fixing an ordering of the vertices and of the edges realizes ∂ as a matrix called the *oriented incidence matrix* whose rows are indexed by the vertices and whose columns are indexed by the edges. See Figure 1 for an example. The next exercise reveals the connection between the oriented incidence matrix and the Laplacian of the multigraph G .

Exercise 13.1. Let ∂^t denote the transpose of the boundary mapping. Show that $\partial\partial^t = L$, the Laplacian matrix.




Figure 1. The oriented incidence matrix with respect to the standard orientation. Rows and columns are labeled with vertex indices.

The cycle space. Consider a cycle $C = u_0, e_1, u_1, e_2, \dots, e_k, u_k$ in the undirected graph G . The *sign* of an edge $e \in E$ with respect to C and the orientation \mathcal{O} is $\sigma(e, C) = 1$ if $C = v, e, v$ is a loop at a vertex v , and otherwise

$$\sigma(e, C) = \begin{cases} 1 & \text{if } e^- = u_i \text{ and } e^+ = u_{i+1} \text{ for some } i, \\ -1 & \text{if } e^+ = u_i \text{ and } e^- = u_{i+1} \text{ for some } i, \\ 0 & \text{otherwise (} e \text{ does not occur in } C\text{).} \end{cases}$$

We then identify C with the formal sum $\sum_{e \in E} \sigma(e, C)e \in \mathbb{Z}E$. For notational convenience, if $e = uv$, we denote $-e$ by vu .

Example 13.2. For the graph in Figure 1 (with the standard orientation), the cycle

$$C = v_1, \{v_1, v_2\}, v_2, \{v_2, v_3\}, v_3, \{v_1, v_3\}, v_1$$

is identified with

$$C = \overline{12} + \overline{23} - \overline{13} = \overline{12} + \overline{23} + \overline{31} \in \mathbb{Z}E$$

where $\overline{12} := v_1v_2$, etc. This cycle is shown in red in Figure 3.

Definition 13.3. The (*integral*) *cycle space*, $\mathcal{C} \subset \mathbb{Z}E$, is the \mathbb{Z} -span of all cycles.

Example 13.4. Let G be the oriented graph pictured in Figure 2. The cycle space is isomorphic to \mathbb{Z}^2 with basis $e_1 - e_2, e_3$.




Figure 2. $\mathcal{C} \simeq \mathbb{Z}^2$.

The cut space. A *directed cut* of G is an ordered partition of the vertices into two nonempty parts. For each nonempty $U \subsetneq V$, we get the directed cut, (U, U^c) . The *cut-set* corresponding to U , denoted c_U^* , is the collection of edges with one vertex in U and the other in the complement, U^c . For each $e \in E$, define the *sign* of e in c_U^* with respect to the orientation \mathcal{O} by

$$\sigma(e, c_U^*) = \begin{cases} 1 & \text{if } e^- \in U \text{ and } e^+ \in U^c, \\ -1 & \text{if } e^+ \in U \text{ and } e^- \in U^c, \\ 0 & \text{otherwise (} e \text{ does not occur in } c_U^*\text{).} \end{cases}$$

We identify the cut-set c_U^* with the formal sum $\sum_{e \in E} \sigma(e, c_U^*)e \in \mathbb{Z}E$. Thus, for instance, $c_{U^c}^* = -c_U^*$. If G is not connected, there will be empty cut-sets, identified with $0 \in \mathbb{Z}E$. A *vertex cut* is the cut-set corresponding to a single vertex, $U = \{v\}$, and we write c_v^* for c_U^* in that case. A minimal nonempty cut-set with respect to inclusion is called a *bond*. For example, the cut-set $c_{\{v_2, v_3\}}^*$ in Example 13.5 is not a bond.

Example 13.5. For the graph in Figure 1 (with the standard orientation), the cut-set corresponding to $\{v_2, v_3\}$ is

$$\begin{aligned} c_{\{v_2, v_3\}}^* &= -\overline{12} - \overline{13} + \overline{24} + \overline{34} \\ &= \overline{21} + \overline{31} + \overline{24} + \overline{34} \in \mathbb{Z}E. \end{aligned}$$

It is shown in blue in Figure 3.




Figure 3. The cycle C from Example 13.2 (left) and the cut-set $c_{\{v_2, v_3\}}^*$ from Example 13.5 (right).

Definition 13.6. The (*integral*) *cut space*, $\mathcal{C}^* \subset \mathbb{Z}E$, is the \mathbb{Z} -span of all cut-sets.

Exercise 13.7. If U is a nonempty subset of $V(G)$, the *subgraph of G induced by U* , denoted $G[U]$, is the graph with vertex set U and edge multiset consisting of those edges with both ends in U . If G is connected, show that the cut-set corresponding to a nonempty set $U \subsetneq V(G)$ is a bond if and only if $G[U]$ and $G[U^c]$ are connected. If G is not connected, show that its bonds are exactly the bonds of its connected components.

Bases for cycle and cut spaces. Fix a spanning forest F for G , and for notational purposes, identify F with its set of edges. Let $F^c := E \setminus F$.

Exercise 13.8. Show that for each $e \in F^c$, the graph with edges $F \cup \{e\}$ has a unique cycle, c_e , such that $\sigma(e, c_e) = 1$. (This holds even if e is a loop.)

Pick $e \in F$. The forest F is a disjoint union of spanning trees of the connected components of G , and one of these spanning trees, say T , contains e . Removing e disconnects T into two connected components T^- and T^+ where e^- is contained in T^- . Let U be the vertices of T^- . Define the cut-set $c_e^* := c_U^*$, and note that $\sigma(e, c_e^*) = 1$.

Exercise 13.9. Show that the cut-sets c_e^* are bonds.

Theorem 13.10.

- (1) The kernel of the boundary mapping is the cycle space: $\ker \partial = \mathcal{C}$.
- (2) Let F be a spanning forest of G . Then $\{c_e : e \in F^c\}$ is a \mathbb{Z} -basis for \mathcal{C} and $\{c_e^* : e \in F\}$ is a \mathbb{Z} -basis for \mathcal{C}^* .
- (3) $\text{rank}_{\mathbb{Z}} \mathcal{C} = |E| - |V| + \kappa$, and $\text{rank}_{\mathbb{Z}} \mathcal{C}^* = |V| - \kappa$ where κ is the number of connected components of G .
- (4) $\mathcal{C} = (\mathcal{C}^*)^\perp := \{f \in \mathbb{Z}E : \langle f, g \rangle = 0 \text{ for all } g \in \mathcal{C}^*\}$ where $\langle \cdot, \cdot \rangle$ is defined for $e, e' \in E$ by

$$\langle e, e' \rangle := \delta(e, e') = \begin{cases} 1 & \text{if } e = e', \\ 0 & \text{if } e \neq e' \end{cases}$$

and extended linearly for arbitrary pairs in $\mathbb{Z}E$.

- (5) If G is connected, then the following sequence is exact:

$$0 \longrightarrow \mathcal{C} \longrightarrow \mathbb{Z}E \xrightarrow{\partial} \mathbb{Z}V \xrightarrow{\text{deg}} \mathbb{Z} \longrightarrow 0.$$

Proof. It is clear that $\mathcal{C} \subseteq \ker \partial$. For the opposite inclusion, consider an arbitrary $f = \sum_{e \in E} a_e e \in \mathbb{Z}E$. Fix a spanning forest F , and define

$$g := f - \sum_{e \in F^c} a_e c_e.$$

Then $\partial f = \partial g$ and $\text{supp}(g) \subseteq F$. If $g \neq 0$, then the union of the edges in $\text{supp}(g)$ is a subforest of F with at least one edge; choose a leaf vertex v in this subforest. Then $v \in \text{supp}(\partial g)$, and hence, $\partial f = \partial g \neq 0$. So if $f \in \ker \partial$, then $g = 0$, i.e.,

$$f = \sum_{e \in F^c} a_e c_e \in \mathcal{C}.$$

This proves part 1.

For part 2, we have just seen that $\{c_e : e \in F^c\}$ spans $\ker \partial$. These elements are linearly independent since $c_e \cap F^c = \{e\}$. To see that $\{c_e^* : e \in F\}$ is a basis for \mathcal{C}^* , first note that each cut-set is a linear combination of vertex cuts:

Exercise 13.11. Show that for each nonempty $U \subsetneq V$,

$$c_U^* = \sum_{v \in U} c_v^*.$$

Hence, the vertex cuts span \mathcal{C}^* . However, each vertex cut is a linear combination of the c_e^* :

Exercise 13.12. Show that for each $v \in V$,

$$c_v^* = \sum_{e \in F: e^- = v} c_e^* - \sum_{e \in F: e^+ = v} c_e^*.$$

(One way to proceed: First argue that we may assume G is connected. Then analyze the above expression in terms of the components of F after removing all edges incident on v .)

Hence, $\{c_e^* : e \in F\}$ spans the cut space. Linear independence of the c_e^* follows from the fact that $c_e^* \cap F = \{e\}$.

For part 3, first note that we have just shown that both the cycle and cut space have \mathbb{Z} -bases, hence it makes sense to consider their ranks (see Theorem 2.23). Since the number of edges in a tree is one less than the number of its vertices, we have

$$\text{rank}_{\mathbb{Z}} \mathcal{C}^* = |F| = |V| - \kappa,$$

and

$$\text{rank}_{\mathbb{Z}} \mathcal{C} = |F^c| = |E| - |F| = |E| - |V| + \kappa.$$

For part 4, let $f = \sum_{e \in E} a_e e$. For each $v \in V$,

$$\langle f, c_v^* \rangle = \sum_{e: e^- = v} a_e - \sum_{e: e^+ = v} a_e,$$

which is the negative of the coefficient of v in ∂f . Thus, $\langle f, c_v^* \rangle = 0$ for all $v \in V$ if and only if $f \in \ker \partial = \mathcal{C}$. Since the vertex cuts span \mathcal{C}^* , the result follows.

Part 5 is Problem 13.1. □

Remark 13.13. The number $\text{rank}_{\mathbb{Z}} \mathcal{C}$ is known as the *cycle rank* or *cyclomatic number* of G . If G is connected, Theorem 13.10 (3) shows that this number is what we called the *genus* of G in Part 1. In graph theory, the name *genus* is typically reserved to mean something else: the minimal number g such that the graph can be drawn on a sphere with g handles (an oriented compact surface of genus g) without edge crossings.

Relation to the sandpile group. The connection of cycles and cuts to sandpile groups is the fact that taking the boundary of a vertex cut corresponds to firing the vertex! Details appear below.

Proposition 13.14.

$$\partial(\mathcal{C}^*) = \mathcal{L} := \text{im } L.$$

Proof. For $v \in V$,

$$\begin{aligned} \partial c_v^* &= \partial \left(\sum_{e: e^- = v} e - \sum_{e: e^+ = v} e \right) \\ &= \sum_{e: e^- = v} (e^+ - e^-) - \sum_{e: e^+ = v} (e^+ - e^-) \\ &= -Lv. \end{aligned}$$

Since the c_v^* generate \mathcal{C}^* , it follows that $\partial(\mathcal{C}^*) = \mathcal{L}$. □

Corollary 13.15. *The boundary map induces an isomorphism*

$$\mathcal{E}(G) := \mathbb{Z}E / (\mathcal{C} + \mathcal{C}^*) \xrightarrow{\partial} \text{Div}^0(G) / \mathcal{L} = \text{Jac}(G) \simeq \mathcal{S}(G).$$

Proof. This result follows immediately from the previous proposition and Theorem 13.10 (5). The isomorphism $\text{Jac}(G) \simeq \mathcal{S}(G)$ sends $\sum_{v \in V} c(v)v$ to the recurrent equivalent to $\sum_{v \in \tilde{V}} c(v)v$ modulo $\tilde{\mathcal{L}}$, the inverse of $\tilde{\varepsilon}$ in Theorem 12.1 (2). [The $\tau(s)$ appearing in Theorem 12.1 is 1 since here we are assuming G is undirected, hence, Eulerian (cf. Corollary 12.3).] \square

13.2. Planar duality

This section involves properties of graph-embeddings in the plane and hence, necessarily, a certain amount of topology. We will rely on intuition for these aspects and refer the reader to [39] for a rigorous treatment.

Let $G = (V, E)$ be an undirected plane multigraph, not necessarily connected. By *plane multigraph* we mean that G is drawn in the plane with non-crossing edges. Removing G from the plane divides the plane into several connected components, one of which is unbounded. These are the *faces* of G , which we denote by F .

Definition 13.16. The *dual* of G , denoted $G^* = (V^*, E^*)$, is the graph with vertices $V^* := F$, and one edge e^* for each edge e of G . For $e \in E$, the corresponding edge $e^* \in E^*$ joins the faces on either side of e . (These two faces may coincide if e contains a leaf vertex, in which case e^* forms a loop.)

We think of the dual of G as a plane graph by picking a point in each face of G , including the unbounded face, to serve as the vertices for G^* , then connecting a pair of these vertices with an edge if their corresponding faces share an edge in G . We require that each edge e^* of G^* is drawn so that it crosses its corresponding edge, $e \in E$, exactly once and crosses no other edges. See Figure 4, ignoring the arrows for now.

We have not required that G be connected. However, G^* is always connected. Roughly, the reason is as follows. Let f be a bounded face and pick a point in its interior. Next, consider the infinitely many lines through that point. Fix one of them that misses the finitely many vertices of G . That line represents a path of edges in G^* connecting f to the unbounded face. Thus, every vertex of G^* is connected to the vertex representing the unbounded face of G .

Exercise 13.17.

- (1) Find the duals of the complete graphs K_3 and K_4 . What is remarkable about the relationship between K_4 and its dual?
- (2) Check that $(K_3^*)^* \simeq K_3$ and $(K_4^*)^* \simeq K_4$.
- (3) Give an example of a plane graph G that is not connected, and show $(G^*)^* \not\simeq G$. Compare the number of vertices, edges, and faces of G and G^* .

It turns out that $(G^*)^* \simeq G$ if and only if G is connected (cf. [39, Section 4.6]). (So one might object to calling G^* the *dual* if G is not connected.) By definition, the faces of G are always in bijection with the vertices of G^* . If G is connected, then by duality, the vertices of G are in bijection with the faces of G^* .

Fix the usual orientation of the plane. An orientation \mathcal{O} on G induces an orientation \mathcal{O}^* on G^* as follows. If e is an edge of G , let $(e^*)^- := \text{right}(e)$, the face

to the right of e when traveling from e^- to e^+ , and let $(e^*)^+ := \text{left}(e)$, the face to the left. Figure 4 gives an example. In that example, there are two edges of G joining the same vertices (an edge of multiplicity two), and we have chosen to give them opposite orientations, to show that it is not prohibited.




Figure 4. An oriented graph G in black and its dual G^* in blue.

The next exercise shows that something funny happens with orientations when taking the double-dual.

Exercise 13.18. Starting from the oriented plane graph G^* in Figure 4, find the orientation on $G = (G^*)^*$ induced by G^* . It should be opposite the original orientation.

The previous exercise suggests that we should really be thinking of G and G^* as sitting in different copies of \mathbb{R}^2 —copies with opposite orientations. We used a right-left rule, above, to induce an orientation from G to G^* . To think of G^* as sitting in the *same* oriented plane as G , use a left-right rule to induce an orientation from G^* to its dual, $(G^*)^* = G$.

Theorem 13.19. *Let G be an oriented plane graph, and let G^* have the induced orientation. Then the mapping on edges $e \mapsto e^*$ determines isomorphisms*

$$\mathcal{C}_G \simeq \mathcal{C}_{G^*}^*, \quad \mathcal{C}_G^* \simeq \mathcal{C}_{G^*}$$

between the cycle and cut spaces of G and G^ . These isomorphisms restrict to bijections between cycles and bonds.*

Sketch of a proof. Let X be a collection of edges of G with corresponding dual edges X^* in G^* . Let f and g be faces of G with corresponding vertices f^* and g^* of G^* . Suppose there is a path $f^*, e_1^*, f_2^*, \dots, f_k^*, e_k^*, g^*$ in the undirected graph G^* , using none of the edges in X^* , going from f^* to g^* . This path corresponds to a walk in \mathbb{R}^2 from a point in f to a point in g moving from face to face by transversally crossing the edges e_1, \dots, e_k of G and not meeting any edge of X .

By the Jordan curve theorem, removing any cycle in G divides the plane into two components: the inside and the outside. Any two faces f and g in the same component are connected by a path in \mathbb{R}^2 not meeting the cycle. By the previous

paragraph, this means that the vertices f^* and g^* are connected by a path in G^* . Hence, if U is the set of faces of G inside the cycle, then the induced subgraphs $G^*[U^*]$ and $G^*[(U^*)^c]$ are both connected. By Exercise 13.7, the edges connecting U^* to its complement form a bond. These are the edges dual to those in the original cycle.

Now suppose that X^* is the set of edges in a bond of G^* . Removing X^* divides the vertices of G^* into two connected components, so removing X from \mathbb{R}^2 disconnects the faces of G into two connected components. If X does not contain a cycle, then it forms a forest. However, a forest has only one face, which would contradict the disconnection just mentioned. So X must contain a cycle C . Minimality of the bond implies that $X = C$. \square

Exercise 13.20. Let G be a plane graph consisting of the disjoint union of a square and a triangle. Illustrate the bijection between cycles of G and bonds of G^* and between bonds of G and cycles of G^* given by Theorem 13.19.

Exercise 13.21. Let $G = (V, E)$ be a connected plane graph, and let F denote the set of faces of G . Show that Euler's formula,

$$|V| - |E| + |F| = 2,$$

is an immediate consequence of Theorems 13.10 and 13.19.

Remark 13.22. It turns out that the boundaries of the faces of a plane graph are cycles generating the graph's cycle space. Omitting any one of them gives a basis for the cycle space. This basis has the property that each edge appears at most twice (since each edge lies on at most two faces). *MacLane's planarity condition* says the existence of such a basis characterizes graphs that may be embedded in the plane—see [39, Section 4.5] for more details.

In Corollary 13.15, we defined $\mathcal{E}(G) := \mathbb{Z}E/(C + C^*)$ and showed that the boundary mapping from edges to vertices induces an isomorphism $\mathcal{E}(G) \simeq \mathcal{S}(G)$.

Corollary 13.23. *Let G be a connected plane graph. There is a commutative diagram of isomorphisms*

$$\begin{array}{ccc} \mathcal{E}(G) & \longrightarrow & \mathcal{E}(G^*) \\ \partial \downarrow & & \downarrow \partial^* \\ \mathcal{S}(G) & \longrightarrow & \mathcal{S}(G^*). \end{array}$$

The vertical maps are induced by the edge-vertex boundary mappings. The top horizontal mapping sends the class of an edge e to the class of its dual, e^ .*

In particular, the sandpile groups of G and G^ are isomorphic.*

Proof. Immediate from Corollary 13.15. \square

Example 13.24. Consider Figure 5, which shows an oriented plane graph G in black and its dual in blue with the induced orientation. Sink vertices s and h are fixed for G and G^* , respectively.

To illustrate the isomorphism $\mathcal{S}(G) \simeq \mathcal{S}(G^*)$ of Corollary 13.23, start with the recurrent $c_{\max} = 2u + 2v + w$ on G (see Figure 6). We represent this recurrent by




Figure 5. An oriented plane graph G with sink s (in black) and its dual G^* with sink h (in blue)

its corresponding element in $\text{Jac}(G)$ by using the sink vertex to create a divisor of degree zero:

$$c_{\max} - \deg(c_{\max})s = 2u + 2v + w - 5s.$$

The next step is to lift $c_{\max} - 5s$ to $\mathcal{E}(G) := \mathbb{Z}E/(\mathcal{C} + \mathcal{C}^*)$, that is, to find an element in $\partial^{-1}(c_{\max})$. To do this, it helps to think of elements of $\mathbb{Z}E$ as *flows* on G . Think of the edges of G as pipes. Let e be an oriented edge. Then, for example, $3e$ represents flow of three units of liquid from e^- to e^+ , which can be read off from the boundary mapping: $\partial(3e) = 3e^+ - 3e^-$. An element in $\ker \partial = \mathcal{C}$ then represents a flow in which there is no build-up of liquid at any vertex—flow is *conserved* at each vertex. These flows are sometimes called *circulations*. They are ignored by $\mathcal{E}(G)$.

In these terms, our problem is to find a flow on G such that the flow into each vertex is given by the coefficients of $c_{\max} - 5s$. One such flow is exhibited on the top right in Figure 6. To find that configuration, start at vertex s . To get -5 units of flow there, we arbitrarily start with 3 units of flow along su and 2 along sv . At that point, there is one surplus unit of flow into u . To correct for this, send one unit of flow from u to w , at which point, we are done: $\partial(3su + 2sv + uw) = c_{\max} - 5s$.

Exercise 13.25. Find a lifting of $c_{\max} - 5s$ that has 5 units of flow from s to u .

Next, dualize the flow, i.e., think of the flow along each edge $e \in E$ as a flow along the dual edge $e^* \in E^*$. Then apply the boundary mapping for $\mathcal{E}(G^*)$. This gives an element of $\text{Jac}(G^*)$ —in our case, $-f - g + 2h$. One way to find the corresponding recurrent in $\mathcal{S}(G^*)$ is to fire the sink of G^* enough times to enable the creation of an equivalent divisor whose restriction to the non-sink vertices is




Figure 6. The isomorphism of Corollary 13.23 sends the recurrent $2u + 2v + w$ of G to the recurrent of G^* equivalent to $-f - g$ modulo the reduced Laplacian of G^* , i.e., to $f + g$.

greater than c_{\max}^* , the maximal stable configuration for G^* . Then stabilize. In our case, firing the sink twice is enough and gives

$$-f - g \xrightarrow{2s} 3f + 3g \rightsquigarrow (3f + 3g)^\circ = f + g.$$

Exercise 13.26. Dualize the flow you found in Exercise 13.25, take the boundary, and show that the equivalent recurrent is still $f + g$.

Problems for Chapter 13

13.1. Prove Theorem 13.10 (5).

13.2. A graph is *planar* if it can be drawn in the plane without crossing edges. Show there exists a planar graph G with two embeddings in the plane, say as the graphs G_1 and G_2 , such that the duals of G_1 and G_2 are not isomorphic.

13.3. A planar graph G is *self-dual* if it is isomorphic to its dual, G^* .

(a) For each $n \geq 4$ find a self-dual planar graph with n vertices.

(b) Show that if G is self-dual, then $|E| = 2|V| - 2$.

13.4. In this problem, we provide an alternate computation of the sandpile group in terms of lattices and orthogonal projections. For more about this topic, including the connection with the Abel-Jacobi map from Section 3.5, see [2]. Let G be a connected undirected multigraph with a fixed orientation \mathcal{O} for its edges, and $\partial_{\mathbb{Q}}: \mathbb{Q}E \rightarrow \mathbb{Q}V$ denote the boundary mapping, now viewed as a linear transformation between vector spaces over \mathbb{Q} . We endow $\mathbb{Q}E$ with the inner product $\langle \cdot, \cdot \rangle$ as described in Theorem 13.10, for which the edges form an orthonormal basis. Set $K := \ker \partial_{\mathbb{Q}}$, and consider the orthogonal direct sum decomposition $\mathbb{Q}E = K \oplus K^{\perp}$, with orthogonal projection $\pi: \mathbb{Q}E \rightarrow K$ onto the kernel. The edge-group $\mathbb{Z}E \subset \mathbb{Q}E$ is a *lattice* in $\mathbb{Q}E$ in the sense of being a finitely generated subgroup of rank equal to the dimension of $\mathbb{Q}E$.

(a) Check that the integral cycle space \mathcal{C} is a lattice in K and the integral cut space \mathcal{C}^* is a lattice in K^{\perp} .

(b) Define the *dual lattice* $\mathcal{C}^{\#} := \{x \in K : \langle x, y \rangle \in \mathbb{Z} \text{ for all } y \in \mathcal{C}\}$. Show that $\pi(\mathbb{Z}E) = \mathcal{C}^{\#}$ as follows:

(i) Choose a spanning tree T for G , and let $\{c_e : e \in T^c\} = \{c_1, c_2, \dots, c_g\}$ denote the corresponding \mathbb{Z} -basis for \mathcal{C} (Theorem 13.10 (2)). Show that this may be extended to a \mathbb{Z} -basis $B = \{c_1, c_2, \dots, c_{|E|}\}$ for $\mathbb{Z}E$, hence a \mathbb{Q} -basis for $\mathbb{Q}E$.

(ii) Consider the dual basis $B' = \{c'_i\}$ for $\mathbb{Q}E$ defined by

$$\langle c'_i, c_j \rangle = \delta_{ij}.$$

Show that B' is actually a \mathbb{Z} -basis for $\mathbb{Z}E$, and that $\{\pi(c'_i) : i = 1, \dots, g\}$ is a \mathbb{Z} -basis for $\pi(\mathbb{Z}E)$.

(iii) Show that $\mathcal{C}^{\#} = \pi(\mathbb{Z}E)$ by writing an arbitrary element of the dual lattice as a \mathbb{Q} -linear combination of the basis elements $\pi(c'_1), \dots, \pi(c'_g)$ and then showing that the coefficients must be integers.

(c) Show that the map $\pi: \mathbb{Z}E \rightarrow \mathcal{C}^{\#}$ induces an isomorphism

$$\mathcal{E}(G) = \mathbb{Z}E / (\mathcal{C} + \mathcal{C}^*) \rightarrow \mathcal{C}^{\#} / \mathcal{C}.$$

By Corollary 13.15, it follows that $\mathcal{C}^{\#} / \mathcal{C} \simeq \mathcal{S}(G)$.

(d) Apply the preceding argument to the integral cut lattice to conclude that $(\mathcal{C}^*)^{\#} / \mathcal{C}^* \simeq \mathcal{S}(G)$.

Matroids and the Tutte polynomial

In the previous chapter we saw that the sandpile group is determined by a graph's cycles and cuts. In this chapter, we give a brief introduction to matroids—the abstract setting in which to analyze these structures. We present a theorem of Criel Merino showing that the coefficients of the Tutte polynomial of a graph encode the number of recurrents of each degree. As a further consequence, we provide a proof of Stanley's h -vector conjecture for the case of cographic matroids.

Unless otherwise specified, the graphs in this chapter are undirected multigraphs, possibly with loops.

14.1. Matroids

By way of motivation, let V be a finite-dimensional vector space over a finite field—so V is in fact a finite set. Recall the following properties of linear independence in V : (i) the empty set is linearly independent; (ii) if I is a linearly independent subset of V and $J \subseteq I$, then J is also linearly independent; (iii) if I, J are linearly independent subsets of V with $|I| > |J|$, then there exists $v \in I$ such that $J \cup \{v\}$ is still linearly independent in V . Matroids capture the essence of linear independence by abstracting these key properties.

Definition 14.1. A *matroid* is a pair $M = (E, \mathcal{I})$ consisting of a finite set E and a collection \mathcal{I} of subsets of E satisfying:

- (1) $\emptyset \in \mathcal{I}$;
- (2) \mathcal{I} is closed under taking subsets: if $I \in \mathcal{I}$ and $J \subseteq I$, then $J \in \mathcal{I}$;
- (3) the *exchange axiom*: if $I, J \in \mathcal{I}$ and $|I| > |J|$, then there exists $e \in I$ such that $J \cup \{e\} \in \mathcal{I}$.

The elements of \mathcal{I} are called the *independent sets* of the matroid, and E is called the *ground set* of M (or we say that M is a matroid on E). Two matroids are *isomorphic* if there is a bijection of their ground sets inducing a bijection of independent sets.

In what follows, the reader should check their intuition using the following three basic examples of matroids.

- *Uniform matroids.* Suppose $E = [n] := \{1, \dots, n\}$, and let $k \leq n$ be a natural number. Then $U_{k,n}$ is the matroid on E with every subset of size at most k taken as independent. Any matroid isomorphic to $U_{k,n}$ is called *k-uniform*.
- *Linear matroids.* Suppose E is a finite subset of a vector space, and let \mathcal{I} be the linearly independent subsets of E . Then $M = (E, \mathcal{I})$ is a matroid. A typical way for these to arise is for E to be the set of columns of a matrix.
- *Graphic matroids.* Let E be the edges of an undirected multigraph, and let \mathcal{I} be the edge-sets of forests (i.e., acyclic subgraphs) of the graph. Then $M = (E, \mathcal{I})$ is called the *cycle matroid* of G , denoted $M(G)$. Verifying the exchange property is part of Problem 14.1. Any matroid isomorphic to the cycle matroid of a graph is called a *graphic matroid*.

Vocabulary. Let $M = (E, \mathcal{I})$ be a matroid, and let $A \subseteq E$. The exchange property implies that every maximal independent subset of A under inclusion has the same cardinality. This cardinality is called the *rank* of A and is denoted $\text{rk}(A)$. An independent subset of size $\text{rk}(A)$ is a *basis* for A . The *rank of M* is $\text{rk}(E)$, and a *basis for M* is by definition a basis for E .

Exercise 14.2. Characterize the rank of a graphic matroid. What is a basis for a graphic matroid?

The *closure* of $A \subseteq E$ is

$$\text{cl}(A) := \{e \in E : \text{rk}(A \cup \{e\}) = \text{rk}(A)\}.$$

If $A = \text{cl}(A)$, then A is called *closed* or a *flat* or a *subspace*. If A is closed, then any subset of A whose closure is A is said to *span* A . A *hyperplane* of M is a maximal proper ($\neq E$) flat.

A subset of E that is not independent is *dependent*. A *circuit* of M is a minimal dependent set. A *cocircuit* is a minimal subset with the property that its intersection with every basis is nonempty.

Exercise 14.3. For a graphic matroid, show that circuits and cocircuits are exactly the cycles and bonds of the corresponding graph.

The *dual* of M , denoted M^* is the matroid on the same set E but whose bases are exactly the complements of the bases of M . The independent sets are subsets of the bases. Said another way: a subset of E is independent in M^* if and only if its complement spans M .

The *direct sum* of matroids $M = (E, \mathcal{I})$ and $N = (F, \mathcal{J})$, denoted $M \oplus N$, is the matroid whose ground set is the disjoint union of E and F and whose independent sets are disjoint unions of an independent set of M with an independent set of N or, equivalently, elements of $\mathcal{I} \times \mathcal{J}$. If $G \sqcup H$ is the disjoint union of graphs G and H , then $M(G \sqcup H) = M(G) \oplus M(H)$.

Deletion and contraction. A *loop* of M is an element $e \in E$ contained in no basis, i.e., $\text{rk}(\{e\}) = 0$. A *bridge*, also called a *coloop* or *isthmus*, is an element $e \in E$ contained in every basis, i.e., a loop of M^* .

If $e \in E$ is not a loop, we can form a new matroid by *contracting* e :

$$M/e := (E \setminus \{e\}, \{I \setminus \{e\} : e \in I \in \mathcal{I}\}),$$

the matroid formed from M by removing e from the ground set and from all its independent sets.

Dually, if $e \in E$ is not a bridge, we get a new matroid by *deleting* e :

$$M \setminus e := (E \setminus \{e\}, \{I \in \mathcal{I} : e \notin I\}),$$

the matroid formed from M by removing e from the ground set and removing all independent sets containing e .

Exercise 14.4. Show that if e is not a loop, then $(M/e)^* = M^* \setminus e$.

Let e be an edge in a graph G . If e is not a loop, define G/e to be the graph obtained from G by contracting e : remove e and identify its two vertices. Then $M(G/e) = M(G)/e$. If e is not a bridge, define $G \setminus e$ to be the graph obtained from G by deleting e . Then $M(G \setminus e) = M(G) \setminus e$.

It turns out that the dual of a graphic matroid is graphic if and only if the graph is *planar*, i.e., if and only if the graph has an embedding in the plane with non-crossing edges (cf. [74], Theorem 5.2.2). In that case, if the original graphic matroid is associated with the graph G , its dual is associated with G^* , the dual of G with respect to any planar embedding of G as described in Section 13.2.

There are many ways of characterizing matroids, including by properties of any of the following: bases, the rank function, dependent sets, circuits, cocircuits, flats, or hyperplanes. Problem 14.12 characterizes matroids as “subset systems” amenable to the greedy algorithm (so whenever you see the greedy algorithm, you might suspect there is a matroid lurking somewhere).

14.2. The Tutte polynomial

We now introduce a fundamental invariant of a matroid $M = (E, \mathcal{I})$.

Definition 14.5. The *Tutte polynomial* of a matroid M is

$$T(x, y) := T(M; x, y) := \sum_{A \subseteq E} (x - 1)^{\text{rk}(E) - \text{rk}(A)} (y - 1)^{|A| - \text{rk}(A)}.$$

The Tutte polynomial of an undirected multigraph is the Tutte polynomial of its cycle matroid: $T(G; x, y) := T(M(G); x, y)$.

Exercise 14.6. Compute the Tutte polynomial of $G = K_3$, a triangle, by computing the ranks of all subsets of the edge set.

Theorem 14.7 (Deletion-contraction). *The Tutte polynomial may be computed recursively as*

- (1) $T(\emptyset; x, y) = 1$ where \emptyset is the matroid with empty ground set.
- (2) If e is a loop, then $T(M; x, y) = yT(M \setminus e; x, y)$.


(3) If e is a bridge, then $T(M; x, y) = xT(M/e; x, y)$.

(4) If e is neither a loop nor a bridge, then

$$T(M; x, y) = T(M \setminus e; x, y) + T(M/e; x, y).$$

Proof. Problem 14.3. □

Example 14.8. Figure 1 illustrates the construction of the Tutte polynomial of the diamond graph G pictured at the top.



$$T(G; x, y) = x + 2x^2 + x^3 + (1 + 2x)y + y^2$$

Figure 1. The Tutte polynomial of G .

The idea is to delete and contract edges that are not bridges or loops (dashed in the figure) until graphs with only bridges and loops remain (at the bottom of the figure). A graph with a bridges and b loops is recorded as the monomial $x^a y^b$.

Exercise 14.9.

- (1) To appreciate a surprising aspect of Theorem 14.7, construct the Tutte polynomial of the graph in Figure 1 using a different sequence of edges to delete and contract.
- (2) Use deletion-contraction to verify the calculation made in Exercise 14.6.

Proposition 14.10. *Let $M = (E, \mathcal{I})$ be a matroid.*

- (1) $T(M; 1, 1)$ is the number of bases of M .
- (2) $T(M; 2, 1) = |\mathcal{I}|$, the number of independent sets of M .
- (3) $T(M; 1, 2)$ is the number of spanning sets of M (those sets containing a basis).
- (4) $T(M; 2, 2) = 2^{|E|}$.
- (5) $T(M^*; x, y) = T(M; y, x)$.

Proof. Problem 14.4. □

For the following proposition, we recall that a *coloring* of an undirected multigraph $G = (V, E)$ using λ colors is the assignment of one of λ distinct colors to each of the vertices of G . The coloring is *proper* if no two vertices sharing an edge have the same color. Let $\chi_G(\lambda)$ be the number of proper colorings of G using λ colors. The function $\chi_G(\lambda)$ is actually a polynomial in λ , called the *chromatic polynomial* of G . To see it's a polynomial, let c_i be the number of proper colorings of G using exactly i colors. Then

$$\chi_G(\lambda) = \sum_{i=0}^{|V|} c_i \binom{\lambda}{i}.$$

Proposition 14.11. *Let G be an undirected multigraph with κ connected components.*

- (1) $T(G; 1, 1)$ is the number of spanning forests of G .
- (2) $T(G; 2, 1)$ is the number of forests of G .
- (3) $T(G; 1, 2)$ is the number of spanning subgraphs of G , i.e., subgraphs whose edge sets contain all the vertices of G .
- (4) $T(G; 2, 0)$ is the number of acyclic orientations of G (cf. Chapter 4).
- (5) The chromatic polynomial of G is given by

$$\chi_G(\lambda) = (-1)^{|V|-\kappa} \lambda^\kappa T(G; 1 - \lambda, 0).$$

- (6) Fix $p \in (0, 1)$, and remove each edge of G independently with probability p . The probability the resulting graph has the same number of components as G (i.e., that no original connected component of G is disconnected) is

$$(1 - p)^{|V|-\kappa} p^{|E|-|V|+\kappa} T(G; 1, 1/p).$$

Proof. The proofs of parts (1)–(3) are immediate from Proposition 14.10. For parts (4)–(6), see [20, Section X.4]. □

The Tutte polynomial has the following universal property with respect to deletion and contraction.

Theorem 14.12. *Let f be any function from isomorphism classes of matroids to a commutative ring R satisfying the following:*

- (1) $f(M \oplus N) = f(M)f(N)$ for all matroids M and N ;

(2) There exist $a, b, c, d \in R$ such that for every matroid M ,

$$f(M) = \begin{cases} a f(M \setminus e) + b f(M/e) & \text{if } e \text{ is not a loop or bridge,} \\ c f(M/e) & \text{if } e \text{ is a bridge,} \\ d f(M \setminus e) & \text{if } e \text{ is a loop.} \end{cases}$$

Then, for every matroid M ,

$$f(M) = a^{\text{rk}(M^*)} b^{\text{rk}(M)} T(M; c/b, d/a).$$

In the case where a or b is not a unit in R , we interpret the expression on the right by treating them as indeterminates, expanding the Tutte polynomial according to its definition, and then canceling.

Proof. Use induction on $|E|$, the size of the ground set of the matroid M ; Problem 14.7. □

14.3. 2-isomorphisms

The following proposition says that the sandpile group is a matroid invariant.

Proposition 14.13. Let G and H be graphs (undirected with possible loops). Then

$$M(G) \simeq M(H) \implies \mathcal{S}(G) \simeq \mathcal{S}(H).$$

Proof. This is an immediate consequence of Corollary 13.15 and Exercise 14.3. □

So the question arises: when are the cycle matroids of two graphs isomorphic? To approach this question, we consider two operations on graphs. First, suppose $G_1 \sqcup G_2$ is the disjoint union of two graphs G_1 and G_2 . Identify a vertex in G_1 with a vertex in G_2 (i.e., glue the two graphs together at a vertex) to form a new graph $G_1 \vee G_2$ called a *one-point join* of G_1 and G_2 . The reader should spend a few moments verifying

$$M(G_1 \sqcup G_2) = M(G_1 \vee G_2).$$

Next consider *two-point joins* of the disjoint graphs G_1 and G_2 by choosing vertices u_1, v_1 in G_1 and u_2, v_2 in G_2 . We consider two gluings. The first identifies u_1 with u_2 and v_1 with v_2 . Call the resulting graph G . The second gluing identifies u_1 with v_2 and v_1 with u_2 . Call the result G' . We say that G' is formed from G by performing a *Whitney twist*. See Figure 2.




Figure 2. A Whitney twist.

Exercise 14.14. Show that if G' is a Whitney twist of G , then $M(G') = M(G)$.

Definition 14.15. Two undirected multigraphs are *2-isomorphic* if one can be obtained from the other by sequences of one-point joins and Whitney twists.

Theorem 14.16 (H. Whitney, 1933). *Let G and G' be undirected multigraphs. Then $M(G) \simeq M(G')$ if and only if G and G' are 2-isomorphic.*

In order to perform a one-point join on a graph, the graph must be disconnected. In order to perform a Whitney twist, there must be a pair of vertices whose removal disconnects the graph. In general, a graph is k -connected if there is no set of $k - 1$ vertices whose removal disconnects the graph.

Corollary 14.17. *If G is 3-connected and $M(G) \simeq M(G')$ for some graph G' , then $G \simeq G'$.*

14.4. Merino’s Theorem

In this section, we present Merino’s Theorem—a fundamental connection between sandpile theory and the Tutte polynomial. We start with an example. The superstacks of the diamond graph G are displayed in Figure 3, with ijk denoting the sandpile $iv_1 + jv_2 + kv_3$.



Figure 3. The superstacks of the diamond graph G .

How many are there of each degree? In Figure 1 we computed the Tutte polynomial of G :

$$T(G; x, y) = x + 2x^2 + x^3 + (1 + 2x)y + y^2.$$

Evaluating at $x = 1$ gives

$$T(G; 1, y) = 4 + 3y + y^2.$$

Merino’s Theorem says that the coefficients of $T(G; 1, y)$ record the number of superstacks of each degree (in reverse order): there is 1 superstack of degree 0, there are 3 of degree 1, and there are 4 of degree 2.

Let $G = (V, E)$ be a connected undirected multigraph, possibly with loops. Choose a sink vertex s . Earlier (see Remark 13.13), we defined the cycle rank of G to be $g = |E| - |V| + 1$. In Part 1 we called this number the *genus*, and we showed that when G has no loops, the maximal superstacks have degree g (Corollary 4.9). Note that each loop increases the cycle rank by 1, but does not affect the superstacks (Problem 7.4). It follows that if G has exactly l loops, then the maximal superstacks have degree $g - l \leq g$.

For $i = 0, \dots, g$, let h_i denote the number of superstacks of G of degree i . We call

$$h = h(G) = (h_0, \dots, h_g)$$

the h -vector of G . As a consequence of Theorem 14.18, below, we will see that the h -vector is independent of the choice of sink s .

Theorem 14.18 (Merino [73], 1997). *Let $T(x, y) = T(G; x, y)$ be the Tutte polynomial of the undirected multigraph G . Then*

$$T(1, y) = \sum_{i=0}^g h_{g-i} y^i.$$

Proof. Let g be the cycle rank of G , and fix s as the sink vertex for all graphs appearing below. The proof goes by induction on the number of edges. For these purposes we define $h_0 = 1$ when G has only the single vertex s (possibly with some loop edges).¹ Thus, in the case where G has no edges, we have $T(1, y) = 1 = h_0$, so the result holds.

Suppose that e is an edge of G incident to s . There are three cases to consider.

Case 1. Suppose e is a loop, and let $G' := G \setminus e$. It follows that $T(G; x, y) = yT(G \setminus e; x, y)$, the cycle rank of G' is $g' = g - 1$, and the superstable on G' are the same as on G (cf. Problem 7.4). Hence, $h' := h(G') = h$. Therefore, by induction and using the fact that $h_g = 0$:

$$T(G; 1, y) = yT(G'; 1, y) = y \sum_{i=0}^{g'} h_{g'-i} y^i = \sum_{i=0}^g h_{g-i} y^i.$$

Case 2. Suppose e is a bridge, and let $G' := G/e$. Then the circuit rank of G' is $g' = g$, and $h' := h(G') = h$ (cf. Problem 7.4). Since $T(G; x, y) = xT(G'; x, y)$, we have $T(G; 1, y) = T(G'; 1, y)$, and the result follows by induction.

Case 3. Suppose $e = \{s, v\}$ is neither a loop nor a bridge. Divide the superstable of G into two sets:

$$\begin{aligned} \mathcal{A} &:= \{c : c(v) = 0\} \\ \mathcal{B} &:= \{c : c(v) > 0\}. \end{aligned}$$

We will show that elements of \mathcal{A} are in bijection with superstable on G/e and elements of \mathcal{B} are in bijection with superstable on $G \setminus e$.

First consider \mathcal{A} . To form G/e , the vertices s and v are identified. Label the resulting vertex s , and think of the vertices of G/e as $V' := V \setminus \{v\}$. Let $\tilde{V} := V \setminus \{s\}$ denote the non-sink vertices of G , and let $\tilde{V}' := \tilde{V} \setminus \{v\} \subset \tilde{V}$ denote the non-sink vertices of G/e . Each configuration c on G with $c(v) = 0$ is naturally a configuration on G/e .

Let c be any sandpile on G with $c(v) = 0$. If it is legal to fire a nonempty set $W \subseteq \tilde{V}$ from c , then since $c(v) = 0$ and v is connected to s , it must be that $v \notin W$. Therefore, c is superstable if and only if there does not exist any nonempty legal firing-set $W \subseteq \tilde{V}'$. Further, if $w \in W \subseteq \tilde{V}'$, then $\text{outdeg}_W(w)$ is the same whether we are considering W as a set of vertices of G or of G/e . Therefore, c is superstable on G if and only if it is superstable on G/e . This shows that the elements of \mathcal{A}

¹This amounts to saying that the sandpile group of a 1-vertex sandpile graph is the trivial group, generated by the empty sandpile.

are in bijection with the superstable of G/e . Note that under this bijection, the degree of a superstable is preserved.

Now consider \mathcal{B} . If c is a sandpile on G with $c(v) > 0$, define

$$c^-(u) = \begin{cases} c(u) & \text{if } u \neq v \\ c(v) - 1 & \text{if } u = v. \end{cases}$$

Then c is superstable on G if and only if c^- is superstable on $G \setminus e$. Note that in this bijection, $\deg(c^-) = \deg(c) - 1$.

It follows from the above that if h' and h'' are the h -vectors for G/e and $G \setminus e$, respectively, then

$$h_i = h'_i + h''_{i-1}$$

for all i (setting $h''_{-1} = 0$ when $i = 0$). Recalling that the cycle rank of G/e is $g = g(G)$ and that the cycle rank of $G \setminus e$ is $g - 1$, it follows by induction that

$$T(G; 1, y) = T(G/e; 1, y) + T(G \setminus e; 1, y) = \sum_{i=0}^g h_{g-i} y^i.$$

□

Corollary 14.19.

- (1) The number of superstable of each degree is independent of the choice of sink vertex s .
- (2) Let $r_{i,s}$ be the number of recurrents of degree i for each i , and let $m_s := \deg(c_{\max})$ —these quantities generally depend on the sink s . Then the generating function for the recurrents by degree is

$$\sum_{i=0}^{m_s} r_{i,s} y^i = y^{m_s - g} \cdot T(G; 1, y).$$

- (3) $T(G; 1, 1) = |\mathcal{S}(G)|$.
- (4) If G has exactly ℓ loops, then

$$\frac{1}{y^\ell} \cdot T(G; 1, y) \Big|_{y=0}$$

is the number of maximal superstable (or minimal recurrents).

Proof. These all follow directly from Theorem 14.18. For part 2, use exercise 14.20 below and the fact that c is recurrent if and only if $c_{\max} - c$ is superstable (cf. Theorem 7.12). For part 4, note that y^ℓ is the maximal power of y dividing $T(G; x, y)$. □

Exercise 14.20. Show that $m_s := \deg(c_{\max}) = g + |E| - \deg(s)$ where s is the sink vertex. (Note that a loop contributes 2 to the degree of its vertex.)

The original version of Theorem 14.18 was defined in terms of “levels” of recurrents, where the *level* of a sandpile c is defined to be

$$\text{level}(c) := \deg(c) - |E| + \deg_G(s).$$

The generating function for recurrents by level is

$$P(y) := \sum_{i=0}^g \ell_i y^i,$$

where ℓ_i is the number of recurrents of level i .

Exercise 14.21. Show that $h_{g-i} = \ell_i$ and hence

$$P(y) = T(G; 1, y).$$

14.5. The Tutte polynomials of complete graphs

In Definition 14.5, we introduced the Tutte polynomial of a matroid via an explicit formula involving ranks of subsets of the ground set. We then observed (Theorem 14.7; Problem 14.3) that this definition is equivalent to one given by deletion and contraction. In the case of graphs, there is yet a third equivalent definition as a sum over spanning forests—in fact, this is Tutte’s original definition. Before presenting the spanning forest expansion, we need to briefly introduce two notions of *activity*.

Definition 14.22. Let G be an undirected multigraph, and fix a total ordering \prec on the edges of G . Suppose that F is a spanning forest of G .

- (1) An edge $e \in E(F)$ is *internally active* in F if $e \prec e'$ for all edges $e' \notin E(F)$ such that $(F \setminus e) \cup e'$ is a spanning forest of G . The *internal activity* $i(F)$ is the number of internally active edges in F .
- (2) An edge $e \notin E(F)$ is *externally active* in F if $e \prec e'$ for all edges $e' \in E(F)$ such that $(F \setminus e') \cup e$ is a spanning forest of G . The *external activity*² $e(F)$ is the number of externally active edges in F .

Figure 4 shows the internally and externally active edges of a spanning tree in a graph on 6 vertices.




Figure 4. The **thick** edges form a spanning tree with internally active edges $(0, 1)$ and $(1, 2)$ in **blue** and externally active edge $(1, 4)$ in **red**. We use the lexicographic edge ordering, where $(i, j) \prec (i', j')$ if $i < i'$ or $i = i'$ and $j < j'$.

²See Section 9.3.2 for a bijection between superstable and spanning trees relating the degree of a superstable to the external activity of its corresponding tree.

Theorem 14.23. *Let G be an undirected multigraph with an edge-ordering as in Definition 14.22. Then the Tutte polynomial of G may be computed as the following sum over spanning forests³ $F \subseteq G$:*

$$T(G; x, y) = \sum_{F \subseteq G} x^{i(F)} y^{e(F)}.$$

Proof. See [20], Section X.5, Theorem 10. □

In the remainder of this section, we use the spanning forest expansion to derive a recursion due to I. Pak for the Tutte polynomials of complete graphs—our exposition follows the short note [75].

Consider the complete graph K_{n+1} with vertex set $V = \{0, 1, 2, \dots, n\}$. The edges of K_{n+1} are given by ordered pairs (i, j) with $i < j$, and we use the lexicographic ordering, so that $(i, j) \prec (i', j')$ iff $i < i'$ or $i = i'$ and $j < j'$. For each spanning tree T , there is a unique path in T from vertex 0 to vertex 1; let $(0, j)$ be the first edge of this path. Removing the edge $(0, j)$ yields a disjoint union of two trees

$$T \setminus (0, j) = T' \sqcup T'',$$

where $0 \in V(T')$ and $j \in V(T'')$. Define the quantities $k(T) := |V(T'')|$ and $a(T) := |\{i \in V(T'') : i < j\}|$. Note that $1 \leq k(T) \leq n$ and $0 \leq a(T) \leq k(T) - 1$. Finally, set $A := V(T'') \setminus \{1\}$, which is a set of size $k(T) - 1$.

Now consider the collection $\mathcal{T}(k, a)$ of all spanning trees T of K_{n+1} such that $k(T) = k$ and $a(T) = a$. We claim that $\mathcal{T}(k, a)$ is in bijection with the set of triples (A, T', T'') with

- $A \subset \{2, 3, \dots, n\}$ a subset of size $k - 1$;
- T' a tree on $V(T') = \{0\} \cup (\{2, \dots, n\} \setminus A)$, a set of size $n - k + 1$;
- T'' a tree on $V(T'') = \{1\} \cup A$, a set of size k .

Indeed, the previous paragraph describes a recipe that produces such a triple from any spanning tree $T \in \mathcal{T}(k, a)$. But the inverse is easy to describe: given such a triple (A, T', T'') , list the elements of $A \cup \{1\} = V(T'')$ in increasing order, and set j to be the $(a + 1)$ st element of this list. Then adjoin the edge $(0, j)$ to obtain the tree $T = T' \cup (0, j) \cup T''$.

We now present two lemmas describing the activities of a spanning tree $T \in \mathcal{T}(k, a)$ with respect to the decomposition $T = T' \cup (0, j) \cup T''$.

Lemma 14.24. *The internal activity of a spanning tree $T \in \mathcal{T}(k, a)$ satisfies*

$$i(T) = i(T') + \delta_{a,0},$$

where $\delta_{a,0}$ is the Kronecker delta.

Proof. First note that every internally active edge of T has the form $(0, i)$ for some vertex i . This is because removing any other type of edge e from T yields two connected components T_0 and T_1 , with $0 \in V(T_0)$. Then adjoining any edge of the form $(0, l)$ with $l \in V(T_1)$ yields a spanning tree, and since $(0, l) \prec e$, this shows that e is not internally active. This observation immediately implies that no edge

³Recall that spanning forests are exactly spanning trees in the case where G is connected.

of T'' is internally active for T . Also, the special edge $(0, j)$ is internally active for T if and only if $a = 0$, since if $a \geq 1$, there exists $i < j$ such that $(T \setminus (0, j)) \cup (0, i)$ is a spanning tree.

Finally, consider an edge of T' of the form $(0, i)$. Then $T \setminus (0, i) = T_0 \sqcup T_1$, where T_0 is a tree containing 0 and T'' , and $T_1 \subset T'$ contains i . Now the only edges smaller than $(0, i)$ are of the form $(0, l)$ for $l < i$, and in order for $(T \setminus (0, i)) \cup (0, l)$ to be a spanning tree, we must have $l \in V(T_1) \subset V(T')$. It follows that $(0, i)$ is internally active for T if and only if it is internally active for T' . The claim follows. \square

Lemma 14.25. *The external activity of a spanning tree $T \in \mathcal{T}(k, a)$ satisfies*

$$e(T) = e(T') + e(T'') + a.$$

Proof. First note that every edge of the form $(0, i)$ with $i < j$ and $i \in V(T'')$ is externally active in T ; there are a such edges. But if (l, i) is any other type of edge with $0 < l \in V(T')$ and $i \in V(T'')$, then (l, i) is *not* externally active in T because of the existence of $(0, j)$. Now consider $i_1, i_2 \in V(T')$ but $(i_1, i_2) \notin E(T')$. Then $T \cup (i_1, i_2)$ contains a unique cycle in the subgraph $T' \cup (i_1, i_2)$, and in order to produce a spanning tree distinct from T , we must remove an edge of T' . This shows that (i_1, i_2) is externally active in T if and only if it is externally active in T' . The same argument holds for the subtree T'' , and the claim follows. \square

Theorem 14.26. *The Tutte polynomials of complete graphs satisfy the following recurrence:*

$$T(K_{n+1}; x, y) = \sum_{k=1}^n \binom{n-1}{k-1} (x + y + y^2 + \dots + y^{k-1}) T(K_k; 1, y) T(K_{n-k+1}; x, y).$$

Proof. We begin with the spanning tree expansion of the Tutte polynomial, where the sum is over spanning trees T of the complete graph K_{n+1} , and then we use the decomposition described above together with the preceding lemmas:

$$\begin{aligned} T(K_{n+1}; x, y) &= \sum_T x^{i(T)} y^{e(T)} \\ &= \sum_{k=1}^n \sum_{a=0}^{k-1} \sum_{T \in \mathcal{T}(k, a)} x^{i(T)} y^{e(T)} \\ &= \sum_{k=1}^n \sum_{a=0}^{k-1} \sum_{(A, T', T'')} x^{i(T') + \delta_{a,0}} y^{e(T') + e(T'') + a} \\ &= \sum_{k=1}^n \sum_{a=0}^{k-1} x^{\delta_{a,0}} y^a \sum_{(A, T', T'')} x^{i(T')} y^{e(T')} y^{e(T'')} \\ &= \sum_{k=1}^n \sum_{a=0}^{k-1} x^{\delta_{a,0}} y^a \binom{n-1}{k-1} \sum_{T''} y^{e(T'')} \sum_{T'} x^{i(T')} y^{e(T')} \\ &= \sum_{k=1}^n \binom{n-1}{k-1} \left(\sum_{a=0}^{k-1} x^{\delta_{a,0}} y^a \right) T(K_k; 1, y) T(K_{n-k+1}; x, y) \end{aligned}$$

$$= \sum_{k=1}^n \binom{n-1}{k-1} (x + y + y^2 + \cdots + y^{k-1}) T(K_k; 1, y) T(K_{n-k+1}; x, y).$$

Note that, starting in the 3rd line, we have used the bijection between the set of spanning trees $\mathcal{T}(k, a)$ and the set of triples (A, T', T'') discussed earlier. The binomial coefficient appears in the 5th line as the number of ways to choose the subset A . \square

Exercise 14.27. Note that $T(K_1; x, y) = 1$ and $T(K_2; x, y) = x$. Use the recursion of Theorem 14.26 to compute the Tutte polynomial of K_3 , and show that your answer agrees with that obtained by deletion and contraction. Then use the recursion to find the Tutte polynomial of K_4 .

Remark 14.28. In the proof of Theorem 14.26, we used the decomposition of a spanning tree T into two parts, achieved by removing the single edge $(0, j)$ beginning the path from 0 to 1 in T . Instead, we could remove *all* edges of T incident to 0, which would yield a decomposition of the form

$$T \setminus \bigcup_{i=1}^m (0, j_i) = T_1 \sqcup T_2 \sqcup \cdots \sqcup T_m,$$

where the T_i are disjoint trees. The collection of vertex sets $V(T_i)$ forms a partition of $\{1, 2, \dots, n\}$ into $m > 0$ parts. Setting $k_i := |V(T_i)|$ to be the size of the i th piece of the partition, a similar argument to the one provided above yields the following expression for the Tutte polynomial, where the sum is over all partitions of $\{1, 2, \dots, n\}$ into any number $m > 0$ parts.

$$T(K_{n+1}; x, y) = \sum_{V_1, V_2, \dots, V_m} \prod_{i=1}^m (x + y + y^2 + \cdots + y^{k_i-1}) T(K_{k_i}; 1, y).$$

This result appears as Theorem 14 of [47], which also contains the following more general result concerning the specialization $T(G; 1, y)$ for arbitrary connected multigraphs G .

Theorem 14.29 ([47], Theorem 10). *Let G be a connected multigraph, and fix a vertex v . For any subset $V_i \subseteq V(G) \setminus \{v\}$, let $G[V_i]$ denote the induced subgraph, consisting of the edges of G between vertices in V_i . Finally, let $\epsilon(V_i)$ denote the number of edges in G from v to V_i . Then*

$$T(G; 1, y) = \sum_{V_1, V_2, \dots, V_m} \prod_{i=1}^m (1 + y + y^2 + \cdots + y^{\epsilon(V_i)-1}) T(G[V_i]; 1, y),$$

where the sum is over all partitions of $V(G) \setminus \{v\}$ into $m > 0$ pieces V_i such that each $G[V_i]$ is connected.

For a deletion-contraction proof of Theorem 14.29, see [22].

14.6. The h -vector conjecture

The goal of this section is to state Stanley’s h -vector conjecture 14.36 and present Merino’s proof in the case of cographic matroids. In order to do so, we will need to briefly introduce three notions and the associated terminology: multicomplexes, simplicial complexes, and matroid complexes.

14.6.1. Multicomplexes. If c and c' are configurations on a sandpile graph, we have defined the relation $c' \leq c$ if $c'(v) \leq c(v)$ for all non-sink vertices v . This relation defines a partial ordering on every subset of configurations:

Definition 14.30. A *partially ordered set* (poset, for short) is a set P and a relation \leq such that for all $x, y, z \in P$,

- (1) $x \leq x$ (reflexivity);
- (2) if $x \leq y$ and $y \leq x$, then $x = y$ (antisymmetry);
- (3) if $x \leq y$ and $y \leq z$, then $x \leq z$ (transitivity).

For x, y in any poset P , we write $x < y$ if $x \leq y$ and $x \neq y$; we write $x \geq y$ if $y \leq x$; and so on. We say y *covers* x if $x < y$ and there is no $z \in P$ such that $x < z < y$.

It is convenient to think of a poset P in terms of its *Hasse diagram*: the graph with vertex set P in which two elements x, y are connected by an edge when y covers x . When drawn in the plane, we draw the edges so that the larger element is above the smaller.

Of special interest to us is the poset of superstables on a sandpile graph. Figures 5 and 6 display this poset for the cases of the diamond graph and K_5 , respectively.




Figure 5. The Hasse diagram for the superstables of the diamond graph (cf. Figure 3).




Figure 6. The Hasse diagram for the superstables of the complete graph K_5 .

The poset of superstables on a sandpile graph has several special properties. We recall the relevant terminology from the theory of posets. A subset I of a poset P is called an *order ideal* or *down-set* if whenever $a \in P$ and $b \in I$, then $a \leq b$ implies $a \in I$. An order ideal I is *generated by* $S \subseteq I$ if

$$I = \langle S \rangle := \{x \in P : \exists y \in S \text{ such that } x \leq y\}.$$

If S is finite, we say I is *finitely generated*. A *chain* in P of length k is a subset of elements $a_0 < a_1 < \cdots < a_k$. It is *saturated* if a_{i+1} covers a_i for all i . The poset P is *graded of rank m* if each maximal length chain has length m . In that case, the *rank function* is given by $\text{rk}(a) = 0$ if a is minimal, and $\text{rk}(b) = \text{rk}(a) + 1$ if b covers a .

Every subset of \mathbb{N}^n with our usual relation ($a \leq b$ if $a_i \leq b_i$ for all i) is a poset. A finitely generated order ideal in \mathbb{N}^n is called a *multicomplex*. A multicomplex is called *pure* if it is a graded poset. Thus, a pure multicomplex is one which has a generating set consisting of elements of the same degree (where $\text{deg}(a) := \sum_i a_i$). Figures 5 and 6 are Hasse diagrams of pure multicomplexes. Figure 7 is an example of a multicomplex that is not pure.




Figure 7. The (impure) multicomplex $\langle(2, 0), (1, 2)\rangle$.

Let $\mathcal{M} \subset \mathbb{N}^n$ be a multicomplex, and let the highest degree of an element of \mathcal{M} be d . The *degree sequence* for \mathcal{M} is the vector $h = (h_0, \dots, h_d)$ where h_i is the number of elements of \mathcal{M} of degree i . Any vector of the form (h_0, \dots, h_d) is called an *O -sequence* if it is the degree sequence of a multicomplex. An O -sequence is *pure* if it is the degree sequence of a pure multicomplex.

Now let G be a sandpile graph, possibly directed, with $n + 1$ vertices. Fixing an ordering of the vertices identifies configurations on G with elements of \mathbb{N}^n , and the set of superstable configurations forms a multicomplex. By the discussion just prior to Theorem 14.18, if G is undirected, the superstable configurations form a pure multicomplex of rank $g - l = |E| - |V| + 1 - \#(\text{loops})$. So, in that case the h -vector of G is a pure O -sequence.

Exercise 14.31. Give an example of a sandpile graph (necessarily directed) whose superstable configurations do not form a pure multicomplex.

14.6.2. Simplicial complexes. In this section we provide only the basic definition of a simplicial complex; see Chapter 15 for a fuller description including an introduction to simplicial homology.

An (abstract) *simplicial complex* Δ on a finite set S is a collection of subsets of S , closed under the operation of taking subsets. The elements of a simplicial complex Δ are called *faces*. An element $\sigma \in \Delta$ of cardinality $i + 1$ is called an *i -dimensional face* or an *i -face* of Δ . The empty set, \emptyset , is the unique face of dimension -1 . Faces of dimension 0, i.e., elements of S , are *vertices* and faces of dimension 1 are *edges*.

The maximal faces under inclusion are called *facets*. To describe a simplicial complex, it is often convenient to simply list its facets—the other faces are exactly determined as subsets. The *dimension* of Δ , denoted $\dim(\Delta)$, is defined to be the maximum of the dimensions of its faces. A simplicial complex is *pure* if each of its facets has dimension $\dim(\Delta)$.

Example 14.32. If $G = (V, E)$ is a simple connected graph (undirected with no multiple edges or loops), then G is the pure one-dimensional simplicial complex on V with E as its set of facets.

The *face-vector* or *f-vector* of Δ is $f(\Delta) = (f_{-1}, \dots, f_{d-1})$ where f_i is the number of faces of dimension i and $d := \dim(\Delta) + 1$ is the *rank* of Δ . Thus, $f_{-1} = 1$. The *face enumerator* of Δ is

$$f_{\Delta}(x) := x^d + f_0 x^{d-1} + \cdots + f_{d-1} := \sum_{i=0}^d f_{i-1} x^{d-i}.$$

The *h-vector* of Δ is the vector (h_0, \dots, h_d) defined by

$$h_{\Delta}(x) := h_0 x^d + h_1 x^{d-1} + \cdots + h_d = \sum_{i=0}^d h_{d-i} x^i := f_{\Delta}(x-1).$$

Example 14.33. Let Δ be the 2-dimensional simplicial complex on the set $[5] := \{1, 2, 3, 4, 5\}$ pictured in Figure 8. Its facets, $\{1, 2, 3\}$, $\{2, 3, 4\}$, $\{2, 4, 5\}$, $\{3, 4, 5\}$, determine the rest of the faces. Check that its *f-vector* is $(1, 5, 8, 4)$ and its *h-vector* is $(1, 2, 1, 0)$.




Figure 8. A pure simplicial complex of dimension 2.

Exercise 14.34. Compare coefficients in $h_{\Delta}(x) = f_{\Delta}(x-1)$ to show $h_0 = 1$ and $h_d = (-1)^{d-1} \tilde{\chi}(\Delta)$, where the alternating sum

$$\tilde{\chi}(\Delta) := -f_{-1} + f_0 - f_1 + \cdots + (-1)^{d-1} f_{d-1}$$

is the *reduced Euler characteristic* of Δ (see Problem 15.9). In general,

$$h_j = \sum_{i=0}^j (-1)^{j-i} \binom{d-i}{j-i} f_{i-1}.$$

14.6.3. Matroid complexes. If $M = (E, \mathcal{I})$ is a matroid, let $\Delta(M)$ denote the simplicial complex on E consisting of the independent sets. Simplicial complexes arising this way are called *matroid complexes*. The exchange axiom for matroids implies every matroid complex is pure.

Exercise 14.35. Is the simplicial complex in Figure 8 a matroid complex?

We now come to the h -vector conjecture:

Conjecture 14.36 (Stanley [85], 1977). *The h -vector of a matroid complex is a pure O -sequence.*

According to [19], published in 2012:

Most of the huge amount of work done on matroids over the last thirty-four years, involving ideas and techniques coming from several different disciplines, has in fact been motivated by that intriguing conjecture, which remains wide open today ...

The reader is encouraged to consult Stanley's influential paper, [85], as well as [19] for an appreciation of the context of this conjecture. Here, we present Merino's proof of the conjecture for the class of *cographic matroids*—those matroids that are matroid duals of graphic matroids.

Theorem 14.37 (Merino [73], 1997). *The h -vector of a cographic matroid is a pure O -sequence.*

Proof. We use the fact from [16] that for a matroid complex $\Delta = \Delta(M)$, the h -vector is encoded in the Tutte polynomial:

$$T(M; x, 1) = h_{\Delta}(x).$$

Suppose that M is cographic. Then there exists a graph G such that $M = M(G)^*$, the dual of the cycle matroid of G . By Proposition 14.10 (5) and the fact that $M^* = M(G)$

$$T(M; x, 1) = T(M^*; 1, x) = T(G; 1, x).$$

By Merino's Theorem (Theorem 14.18), we see that the h -vectors for M and G coincide. The h -sequence for G corresponds to the pure multicomplex formed by the superstable of G , and hence is a pure O -sequence. \square

Notes

Proposition 14.13 appears as Corollary 7.7 of D. Wagner's paper [90]. The formulation of Theorem 14.12 is adapted from V. Reiner's lecture notes [81], Proposition 28.

Problems for Chapter 14

14.1. Let $G = (V, E)$, and let \mathcal{I} be the collection of forests of G . Show that $M := (E, \mathcal{I})$ is a matroid.

14.2. Compute the Tutte polynomial for the k -uniform matroid, $U_{k,n}$.

14.3. Prove Theorem 14.7.

14.4.

(a) Prove Proposition 14.10, parts (1)–(4).

(b) Show that if $M = (E, \mathcal{I})$ is a matroid and $A \subseteq E$, then

$$\text{rk}_{M^*}(A) = |A| - \text{rk}_M(E) + \text{rk}_M(A^c).$$

(c) Prove Proposition 14.10, part 5.

14.5. Compute the Tutte polynomial for the diamond graph G in Figure 1 directly from Definition 14.5. Then verify that $T(G; 1, y) = 4 + 3y + y^2$. Identify the sets A in the definition that contribute to nonzero terms in $T(G; 1, y)$.

14.6. Use deletion and contraction to compute the Tutte polynomial for the dual of the graph G in Figure 1, and verify that it agrees with Proposition 14.10 (5).

14.7. Use induction to prove Theorem 14.12.

14.8. Use deletion and contraction to compute the Tutte polynomial for the cycle graph C_n , and then use it to verify Merino's theorem, Theorem 14.18, for C_n .

14.9. Verify Proposition 14.11 (6) in the following two cases by directly calculating the relevant probabilities and comparing with those given by the formula involving the Tutte polynomial:

(a) Let $G = B_k$ be the banana graph consisting of two vertices joined by k edges.

(b) Let G be the graph consisting of k loops at a single vertex.

14.10. (Unicycles and stationary density.) Let $G = (V, E)$ be a connected, undirected multigraph with Tutte polynomial $T(G; x, y)$. Define $t_G(y) := T(G; 1, y)$.

(a) Prove the following formula for the stationary density (Definition 8.41) of G :


$$\zeta_{\text{st}} = \frac{1}{|V|} \left(|E| + \frac{t'_G(1)}{t_G(1)} \right) = \frac{1}{|V|} (|E| + \ln(t_G(y))'|_{y=1}).$$

(b) Use the formula to compute the stationary density of the banana graph B_n .

(c) Show that $t'_G(1)$ is the number of *spanning unicycles* of G , where a *unicycle* is a subgraph having a single cycle. So a spanning unicycle is a subgraph obtained from a spanning tree by adding an edge, or equivalently, is a subgraph with $|V|$ vertices and $|V|$ edges. We therefore have the following formula for the stationary density:

$$\zeta_{\text{st}} = \frac{1}{|V|} \left(|E| + \frac{\# \text{ spanning unicycles}}{\# \text{ spanning trees}} \right).$$

(d) Let G be the graph pictured below:



Compute the threshold density of G in three ways:

- (i) directly from the definition by computing the recurrences of G ;
- (ii) using the formula given above involving $t_G(1)$ and $t'_G(1)$;
- (iii) and by describing the unicycles and spanning trees of G .

14.11. This problem assumes familiarity with basic generatingfunctionology (see [93], freely available online, for a nice introduction). The object is to establish a closed formula for the stationary density of the complete graph K_n .

For $n, k \in \mathbb{N}$ defining the *falling factorial*, $n^{\underline{k}} := n(n-1)\dots(n-k+1)$. Next, define the *Ramanujan Q -function* for $n \geq 1$ by

$$Q(n) := \sum_{k \geq 1} \frac{n^{\underline{k}}}{n^k}$$

$$= 1 + \frac{n-1}{n} + \frac{(n-1)(n-2)}{n^2} + \frac{(n-1)(n-2)(n-3)}{n^3} + \dots$$

Note that $Q(n)$ is a finite sum for each particular value of n . Our goal is to show that the stationary density of K_n is

$$\zeta_{\text{st}}(K_n) = \frac{1}{2} \left(Q(n) + n - 3 + \frac{1}{n} \right),$$

A *rooted spanning tree* of a graph is a spanning tree plus a choice of a “root” vertex. Let t_n and u_n be the number of rooted spanning trees and the number of spanning unicycles of K_n , respectively, and consider the corresponding exponential generating functions:

$$T(x) = \sum_{n \geq 1} t_n \frac{x^n}{n!} \quad \text{and} \quad U(x) = \sum_{n \geq 1} u_n \frac{x^n}{n!}.$$

By Cayley’s formula (Corollary 9.8), we have $t_n = n^{n-1}$, i.e., n times the number of spanning trees of K_n .

(a) Prove that

$$U(x) = \sum_{k \geq 3} \frac{T(x)^k}{2k} = \frac{1}{2} \left(\ln \left(\frac{1}{1-T(x)} \right) - T(x) - \frac{1}{2}T(x)^2 \right).$$

(Hint: one can form unicycles by grafting a rooted tree to each of the vertices of a cycle graph C_k with $k \geq 3$. Counting rotations and flips, there are $2k$ symmetries of C_k .)

(b) Show that

$$\frac{1}{1-T(x)} = \sum_{n \geq 0} n^n \frac{x^n}{n!}$$

and

$$\ln\left(\frac{1}{1-T(x)}\right) = \sum_{n \geq 0} n^{n-1} Q(n) \frac{x^n}{n!}.$$

- (c) Compute $n!$ times the n -th coefficient of $U(x)$ to show that the number of unicycles in K_n is $\frac{1}{2}n^{n-2}(nQ(n) - 2n + 1)$.
- (d) Use Problem 14.10 to compute the stated formula for $\zeta_{\text{st}}(K_n)$ in terms of the Q -function.
- (e) The function Q has the following asymptotic expression ([63]):

$$Q \sim \sqrt{\frac{\pi n}{2}} - \frac{1}{3} + \frac{1}{12} \sqrt{\frac{\pi}{2n}} - \frac{4}{135n} + O(n^{-3/2}).$$

Use this expression to show

$$\zeta_{\text{st}}(K_n) \sim \frac{n}{2} + \frac{\sqrt{2\pi n}}{4} - \frac{5}{3} + \frac{1}{48} \sqrt{\frac{2\pi}{n}} + \frac{131}{270n} + O(n^{-3/2}).$$

(It turns out that the approximation $\zeta_{\text{st}}(K_n) \approx \frac{n}{2} + \frac{\sqrt{2\pi n}}{4} - \frac{5}{3}$ is off by less than one percent for $n > 11$.)

Note: Our references for the connection between unicycles and Ramanujan's Q -function are [59] and [45], and we would like to acknowledge the assistance of Riley Thornton for the ideas presented here.

14.12. This problem shows that the exchange axiom for matroids is essentially related to the notion of a greedy algorithm.

Let $M = (E, \mathcal{I})$ where E is a finite set and \mathcal{I} is a nonempty collection of subsets of E , closed under inclusion. In particular, $\emptyset \in \mathcal{I}$. Any such M is called a *subset system*.

A *weight function* on M is any function of the form $\text{wt}: E \rightarrow \mathbb{R}$, and with respect to this function, define the *weight* of a subset A of E to be $\text{wt}(A) := \sum_{a \in A} \text{wt}(a)$. The *optimization problem* for M takes as input a weight function and asks for an element of \mathcal{I} with maximal weight. (The problem of *minimization* is equivalent by negating the weight function.)

One approach to the optimization problem is the *greedy algorithm*. Start with $A = \emptyset$. Then, as long as possible, choose an element of maximal weight among all $e \in E \setminus A$ such that $A \cup \{e\} \in \mathcal{I}$ and then add e to A . This certainly finds a maximal element of \mathcal{I} under inclusion, but does not guarantee that this maximal element has the maximal weight of all elements of \mathcal{I} .

- (a) Give an example of a subset system M that is not a matroid and exhibit a weight function for which the greedy algorithm does not produce a basis. (How about a minimal example?)
- (b) Show that a subset system M satisfies the matroid exchange axiom (and thus M is a matroid) if and only if for every weight function, the greedy algorithm produces an element of \mathcal{I} for M of maximal weight.
- (c) What changes if we only allow nonnegative weight functions?
- (d) The greedy algorithm in the special case of the cycle matroid of a graph is called *Kruskal's algorithm*. Use it to find minimal and maximal weight spanning trees

for the graph in Figure 9. For this purpose, the *weight* of a tree is defined as the sum of the weights of its edges.




Figure 9. Graph with weighted edges for Problem 14.12.

14.13. Find two connected simple (undirected and no multiple edges or loops) graphs that are not 2-isomorphic but have isomorphic sandpile groups.

14.14. Show that the exchange axiom for matroids is equivalent to the fact that matroid complexes are pure.

Higher dimensions

In this final chapter, we discuss higher-dimensional versions of many of the topics studied throughout the book. Graphs are replaced by simplicial complexes, and we begin by discussing their homology groups. We then generalize the Jacobian of a graph by defining the critical groups of a simplicial complex. We also introduce simplicial spanning trees and state the higher-dimensional matrix-tree theorem due to Duval, Klivans, and Martin. Finally, we conclude with a few remarks about higher-dimensional versions of the dollar game and the sandpile model.

15.1. Simplicial homology

We review the idea of a simplicial complex from Section 14.6.2 by way of an example.

Example 15.1. Figure 1 pictures a simplicial complex Δ on the set $[5] := \{1, 2, 3, 4, 5\}$:

$$\Delta := \{\emptyset, \bar{1}, \bar{2}, \bar{3}, \bar{4}, \bar{5}, \bar{12}, \bar{13}, \bar{23}, \bar{24}, \bar{34}, \bar{123}\},$$

writing, for instance, $\bar{23}$ to represent the set $\{2, 3\}$.




Figure 1. A 2-dimensional simplicial complex, Δ .

The sets of *faces* of each dimension are:

$$F_{-1} = \{\emptyset\} \qquad F_0 = \{\bar{1}, \bar{2}, \bar{3}, \bar{4}, \bar{5}\}$$

$$F_1 = \{\bar{12}, \bar{13}, \bar{23}, \bar{24}, \bar{34}\} \qquad F_2 = \{\bar{123}\}.$$

Its *facets* are $\overline{5}$, $\overline{24}$, $\overline{34}$, and $\overline{123}$. The dimension of Δ is 2, as determined by the facet $\overline{123}$. Since not all of the facets have the same dimension, Δ is not *pure*.

Let Δ be an arbitrary simplicial complex. By relabeling, if necessary, assume its vertices are $[n] := \{1, \dots, n\}$. For each i , let $F_i(\Delta)$ be the set of faces of dimension i , and define the *group of i -chains* to be the free abelian group with basis $F_i(\Delta)$:

$$C_i = C_i(\Delta) := \mathbb{Z}F_i(\Delta) := \left\{ \sum_{\sigma \in F_i(\Delta)} a_\sigma \sigma : a_\sigma \in \mathbb{Z} \right\}.$$

The *boundary* of $\sigma \in F_i(\Delta)$ is

$$\partial_i(\sigma) := \sum_{j \in \sigma} \text{sign}(j, \sigma) (\sigma \setminus j),$$

where $\text{sign}(j, \sigma) = (-1)^{k-1}$ if j is the k -th element of σ when the elements of σ are listed in order, and $\sigma \setminus j := \sigma \setminus \{j\}$. Extending linearly gives the *i -th boundary mapping*,

$$\partial_i: C_i(\Delta) \rightarrow C_{i-1}(\Delta).$$

If $i > n - 1$ or $i < -1$, then $C_i(\Delta) := 0$, and we define $\partial_i := 0$. We sometimes simply write ∂ for ∂_i if the dimension i is clear from context.

Example 15.2. Suppose $\sigma = \{1, 3, 4\} = \overline{134} \in \Delta$. Then $\sigma \in F_2(\Delta)$, and

$$\text{sign}(1, \sigma) = 1, \quad \text{sign}(3, \sigma) = -1, \quad \text{sign}(4, \sigma) = 1.$$

Therefore,

$$\partial(\sigma) = \partial_2(\overline{134}) = \overline{34} - \overline{14} + \overline{13}.$$

The (*augmented*) *chain complex* of Δ is the complex

$$0 \longrightarrow C_{n-1}(\Delta) \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_2} C_1(\Delta) \xrightarrow{\partial_1} C_0(\Delta) \xrightarrow{\partial_0} C_{-1}(\Delta) \longrightarrow 0.$$

The word *complex* here refers to the fact that $\partial^2 := \partial \circ \partial = 0$, i.e., for each i , we have $\partial_{i-1} \circ \partial_i = 0$ (cf. Problem 15.1).




Figure 2. Two boundary mapping examples. Notation: if $i < j$, then we write $i \longrightarrow j$ for \overline{ij} and $i \longleftarrow j$ for $-\overline{ij}$.

Figure 2 gives two examples of the application of a boundary mapping. Note that

$$\partial^2(\overline{12}) = \partial_0(\partial_1(\overline{12})) = \partial_0(\overline{2} - \overline{1}) = \emptyset - \emptyset = 0.$$

The reader is invited to verify $\partial^2(\overline{123}) = 0$.

Figure 3 shows the boundary of $\sigma = \overline{1234}$, the solid tetrahedron. Figure 4 helps to visualize the fact that $\partial^2(\sigma) = 0$. The orientations of the triangles may




Figure 3. ∂_3 for a solid tetrahedron. Notation: if $i < j < k$, then we write

$$\begin{array}{c} k \\ \triangle \\ i \quad j \end{array} \text{ for } \overline{ijk} \text{ and } \begin{array}{c} k \\ \triangle \\ i \quad j \end{array} \text{ for } -\overline{ijk}.$$

be thought of as inducing a “flow” along the edges of the triangles. These flows cancel to give a net flow of 0. This should remind you of Stokes’ theorem from multivariable calculus.




Figure 4. As seen in Figure 3, the boundary of a solid tetrahedron consists of oriented triangular facets.

Example 15.3. Let Δ be the simplicial complex on $[4]$ with facets $\overline{12}$, $\overline{3}$, and $\overline{4}$ pictured in Figure 5. The faces of each dimension are:

$$F_{-1}(\Delta) = \{\emptyset\}, \quad F_0(\Delta) = \{\overline{1}, \overline{2}, \overline{3}, \overline{4}\}, \quad F_1(\Delta) = \{\overline{12}\}.$$




Figure 5. Simplicial complex for Example 15.3.

Here is the chain complex for Δ :

$$\begin{array}{ccccccc}
 0 & \longrightarrow & C_1(\Delta) & \xrightarrow{\partial_1} & C_0(\Delta) & \xrightarrow{\partial_0} & C_{-1}(\Delta) \longrightarrow 0. \\
 & & & & \bar{1} & & \\
 & & & & \bar{2} & \swarrow & \\
 & & & & \bar{3} & \swarrow & \\
 & & & & \bar{4} & \swarrow & \\
 & & \bar{1}\bar{2} & \longmapsto & \bar{2} - \bar{1} & \longrightarrow & \emptyset
 \end{array}$$

In terms of matrices, the chain complex is given by

$$\begin{array}{ccccccc}
 & & & \bar{1}\bar{2} & & & \\
 & & & \bar{1} & \begin{pmatrix} -1 \\ 1 \\ 0 \\ 0 \end{pmatrix} & & \\
 & & & \bar{2} & & & \\
 & & & \bar{3} & & & \\
 & & & \bar{4} & & & \\
 0 & \longrightarrow & \mathbb{Z} & \xrightarrow{\partial_1} & \mathbb{Z}^4 & \xrightarrow{\partial_0} & \mathbb{Z} \longrightarrow 0. \\
 & & & & \emptyset & \begin{matrix} \bar{1} & \bar{2} & \bar{3} & \bar{4} \\ (1 & 1 & 1 & 1) \end{matrix} & &
 \end{array}$$

The sequence is not exact since $\text{rk}(\text{im } \partial_1) = \text{rk } \partial_1 = 1$, whereas by rank-nullity, $\text{rk}(\ker(\partial_0)) = 4 - \text{rk } \partial_0 = 3$.

Definition 15.4. For $i \in \mathbb{Z}$, the i -th (reduced) homology of Δ is the abelian group

$$\tilde{H}_i(\Delta) := \ker \partial_i / \text{im } \partial_{i+1}.$$

In particular, $\tilde{H}_{n-1}(\Delta) = \ker(\partial_{n-1})$, and $\tilde{H}_i(\Delta) = 0$ for $i > n-1$ or $i < 0$. Elements of $\ker \partial_i$ are called i -cycles and elements of $\text{im } \partial_{i+1}$ are called i -boundaries. The i -th (reduced) Betti number of Δ is the rank of the i -th homology group:

$$\tilde{\beta}_i(\Delta) := \text{rk } \tilde{H}_i(\Delta) = \text{rk}(\ker \partial_i) - \text{rk}(\partial_{i+1}).$$


Remark 15.5. To define ordinary (non-reduced) homology groups, $H_i(\Delta)$, and Betti numbers $\beta_i(\Delta)$, modify the chain complex by replacing $C_{-1}(\Delta)$ with 0 and ∂_0 with the zero mapping. The difference between homology and reduced homology is that $H_0(\Delta) \simeq \mathbb{Z} \oplus \tilde{H}_0(\Delta)$ and, thus, $\beta_0(\Delta) = \tilde{\beta}_0(\Delta) + 1$. All other homology groups and Betti numbers coincide. To see the motivation for using reduced homology in our context, see Example 15.12. From now on, we use “homology” to mean reduced homology.

In general, homology can be thought of as a measure of how close the chain complex is to being exact. In particular, $\tilde{H}_i(\Delta) = 0$ for all i if and only if the chain complex for Δ is exact. For the next several examples, we will explore how exactness relates to the topology of Δ .

The 0-th homology group measures “connectedness”. Write $i \sim j$ for vertices i and j in a simplicial complex Δ if $\bar{i}\bar{j} \in \Delta$. An equivalence class under the transitive closure of \sim is a *connected component* of Δ . By Problem 15.3, $\tilde{\beta}_0(\Delta)$ is one less than the number of connected components of Δ . For instance, for the simplicial complex Δ in Example 15.3,

$$\tilde{\beta}_0(\Delta) = \text{rk } \tilde{H}_0(\Delta) = \text{rk}(\ker \partial_0) - \text{rk}(\partial_1) = 3 - 1 = 2.$$

Example 15.6. The hollow triangle,




has chain complex

$$0 \longrightarrow \mathbb{Z}^3 \xrightarrow{\partial_1} \mathbb{Z}^3 \xrightarrow{\partial_0} \mathbb{Z} \longrightarrow 0.$$

$$\begin{matrix} & \bar{12} & \bar{13} & \bar{23} & & \\ \bar{1} & (-1 & -1 & 0 & & \\ \bar{2} & 1 & 0 & -1 & & \\ \bar{3} & 0 & 1 & 1 & & \end{matrix} \quad \begin{matrix} \bar{1} & \bar{2} & \bar{3} \\ \emptyset & (1 & 1 & 1) \end{matrix}$$

It is easy to see that $\text{rk}(\partial_1) = \text{rk}(\ker \partial_0) = 2$. It follows that $\tilde{\beta}_0(\Delta) = 0$, which could have been anticipated since Δ is connected. Since $\text{rk}(\partial_1) = 2$, rank-nullity says $\text{rk}(\ker \partial_1) = 1$, whereas $\partial_2 = 0$. Therefore, $\tilde{\beta}_1(\Delta) = \text{rk}(\ker \partial_1) - \text{rk}(\partial_2) = 1$. In fact, $\tilde{H}_1(\Delta)$ is generated by the 1-cycle



If we would add $\bar{123}$ to Δ to get a solid triangle, then the above cycle would be a boundary, and there would be no homology in any dimension. Similarly, a solid tetrahedron has no homology, and a hollow tetrahedron has homology only in dimension 2 (of rank 1). See Problem 15.4 for a generalization.

Exercise 15.7. Compute the Betti numbers for the simplicial complex formed by gluing two (hollow) triangles along an edge. Describe generators for the homology.

Example 15.8. Consider the simplicial complex pictured in Figure 6 with facets $\bar{14}, \bar{24}, \bar{34}, \bar{123}$. It consists of a solid triangular base whose vertices are connected by edges to the vertex 4. The three triangular walls incident on the base are hollow.





Figure 6. Simplicial complex for Example 15.8.

What are the Betti numbers? The chain complex is:

$$0 \rightarrow \mathbb{Z} \xrightarrow{\partial_2} \mathbb{Z}^6 \xrightarrow{\partial_1} \mathbb{Z}^4 \xrightarrow{\partial_0} \mathbb{Z} \rightarrow 0.$$

$\begin{matrix} \overline{12} \\ \overline{13} \\ \overline{14} \\ \overline{23} \\ \overline{24} \\ \overline{34} \end{matrix} \begin{pmatrix} 1 \\ -1 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}$
 $\begin{matrix} \overline{1} \\ \overline{2} \\ \overline{3} \\ \overline{4} \end{matrix} \begin{pmatrix} -1 & -1 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & -1 & -1 & 0 \\ 0 & 1 & 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{pmatrix}$
 $\emptyset \begin{matrix} \overline{1} & \overline{2} & \overline{3} & \overline{4} \\ (1 & 1 & 1 & 1) \end{matrix}$

By inspection, $\text{rk}(\partial_2) = 1$ and $\text{rk}(\partial_1) = \text{rk}(\ker \partial_0) = 3$. Rank-nullity gives $\text{rk}(\ker \partial_1) = 6 - 3 = 3$. Therefore, $\tilde{\beta}_0 = \tilde{\beta}_2 = 0$ and $\tilde{\beta}_1 = 2$. It is not surprising that $\tilde{\beta}_0 = 0$, since Δ is connected. Also, the fact that $\tilde{\beta}_2 = 0$ is easy to see since $\overline{123}$ is the only face of dimension 2, and its boundary is not zero. Seeing that $\tilde{\beta}_1 = 2$ is a little harder. Given the cycles corresponding to the three hollow triangles incident on vertex 4, one might suppose $\tilde{\beta}_1 = 3$. However, as conveyed in Figure 7, those cycles are not independent: if properly oriented their sum is the boundary of the solid triangle, $\overline{123}$; hence, their sum is 0 in the first homology group.



$$(\overline{12} + \overline{24} - \overline{14}) + (\overline{23} + \overline{34} - \overline{24}) + (\overline{14} - \overline{34} - \overline{13}) = \underbrace{(\overline{12} + \overline{23} - \overline{13})}_{\partial_2(\overline{123})}$$

Figure 7. A tetrahedron with solid base and hollow walls. Cycles around the walls sum to the boundary of the base, illustrating a dependence among the cycles in the first homology group.

15.1.1. A quick aside on algebraic topology. Algebraic topology seeks an assignment of the form $X \mapsto \alpha(X)$ where X is a topological space and $\alpha(X)$ is some algebraic invariant (a group, ring, etc.). If $X \simeq Y$ as topological spaces, i.e., if X and Y are homeomorphic, then we should have $\alpha(X) \simeq \alpha(Y)$ as algebraic objects—this is what it means to be *invariant*. The simplicial homology we have developed provides the tool for creating one such invariant.

Let X be a 2-torus—the surface of a donut. Draw triangles on the surface so that neighboring triangles meet vertex-to-vertex or edge-to-edge. The triangulation is naturally interpreted as a simplicial complex Δ . An amazing fact, of fundamental importance, is that the associated homology groups do not depend on the choice

of triangulation! In this way, we get an assignment

$$X \mapsto \tilde{H}_i(X) := \tilde{H}_i(\Delta),$$

and, hence, also $X \mapsto \tilde{\beta}_i(X) := \tilde{\beta}_i(\Delta)$, for all i .

In a course on algebraic topology, one learns that these homology groups do not see certain aspects of a space. For instance, they do not change under certain contraction operations. A line segment can be continuously morphed into a single point, and the same goes for a solid triangle or tetrahedron. So these spaces all have the homology of a point—in other words: none at all (all homology groups are trivial). A tree is similarly contractible to a point, so the addition of a tree to a space has no effect on homology. Imagine the tent with missing walls depicted in Figure 6. Contracting the base to a point leaves two vertices connected by three line segments. Contracting one of these line segments produces two loops meeting at a single vertex. No further significant contraction is possible—we are not allowed to contract around “holes” (of any dimension). These two loops account for $\tilde{\beta}_1 = 2$ in our previous calculation. As another example, imagine a hollow tetrahedron. Contracting a facet yields a surface that is essentially a sphere with three longitudinal lines connecting its poles, thus dividing the sphere into 3 regions. Contracting two of these regions results in a sphere—a bubble—with a single vertex drawn on it. No further collapse is possible. This bubble accounts for the fact that $\tilde{\beta}_2 = 1$ is the only nonzero Betti number for the sphere, as calculated in Problem 15.5.

15.1.2. How to compute homology groups. As illustrated in the examples above, Betti numbers of simplicial complexes may be found by computing ranks of matrices. Homology groups are more subtle. They are finitely generated abelian groups and thus have—as we know from Theorem 2.23—a free part, which determines the rank, and a torsion part. For instance,

$$\mathbb{Z}_2 \times \mathbb{Z}_4 \times \mathbb{Z}^3$$

is a typical finitely generated abelian group of rank 3; it has free part \mathbb{Z}^3 and torsion part $\mathbb{Z}_2 \times \mathbb{Z}_4$.

The key to calculating homology groups is to compute the Smith normal forms of the matrices representing the boundary maps. Since the homology groups depend on comparing the kernel of one boundary mapping with the image of another, one might think that a complicated simultaneous reduction for successive boundary maps would be necessary. Happily, that is not the case.

Let Δ be a simplicial complex, and choose bases for each chain group C_i thus identifying the boundary mappings with matrices. Suppose the invariant factors—the diagonal entries of the Smith normal form—for ∂_{i+1} are $d_1, \dots, d_k, 0, \dots, 0$, with respect to bases e_1, \dots, e_s for C_{i+1} and e'_1, \dots, e'_t for C_i . We have

$$\partial_{i+1}(e_j) = \begin{cases} d_j e'_j & \text{for } 1 \leq j \leq k, \\ 0 & \text{for } k < j \leq s. \end{cases}$$

Therefore, $d_1 e'_1, \dots, d_k e'_k$ is a basis for $\text{im } \partial_{i+1}$, the group of i -boundaries of Δ .

Since $\partial^2 = 0$,

$$0 = \partial_i(\partial_{i+1}(e_j)) = d_j \partial_i(e'_j),$$

for $j = 1, \dots, k$. Therefore, e'_1, \dots, e'_k form part of a basis for $\ker \partial_i$, the group of i -cycles of Δ , which we complete to a basis of $\ker \partial_i$,

$$e'_1, \dots, e'_k, e''_{k+1}, \dots, e''_{k+r}$$

for some e''_j and r . It follows that

$$\begin{aligned} \tilde{H}_i(\Delta) &= \ker \partial_i / \text{im } \partial_{i+1} \\ &\simeq \mathbb{Z}\langle e'_1, \dots, e'_k, e''_{k+1}, \dots, e''_{k+r} \rangle / \mathbb{Z}\langle d_1 e'_1, \dots, d_k e'_k \rangle \\ &\simeq \mathbb{Z}_{d_1} \times \dots \times \mathbb{Z}_{d_k} \times \mathbb{Z}^r. \end{aligned}$$

In summary, the torsion part of $\tilde{H}_i(\Delta)$ is determined by the invariant factors of ∂_{i+1} . The rank r of $\tilde{H}_i(\Delta)$ is $\text{rk}(\ker \partial_i) - k$, where k is the number of nonzero invariant factors of ∂_{i+1} .

Example 15.9. The real projective plane, \mathbb{RP}^2 , is a topological space representing the set of 1-dimensional vector subspaces of \mathbb{R}^3 , i.e., the set of lines through the origin in \mathbb{R}^3 . To create it, let S^2 be a sphere centered at the origin in \mathbb{R}^3 . Each line meets S^2 in a pair of antipodal points. So to form \mathbb{RP}^2 , we glue each point of S^2 with its antipode. Proceed in two steps: first, glue every point strictly above the equator with its antipode to get a disc. Forgetting the labeled vertices, this disc is pictured in Figure 8. Next, we need to glue each point on the equator (the boundary of the disc) to its antipode. This cannot be done without crossings in 3 dimensions, but the vertex labels in Figure 8 indicate the gluing instructions. Thus, for instance, the edge $\overline{12}$ appearing in the lower-left should be glued to the edge $\overline{12}$ in the upper-right, matching up like vertices, and so on. Figure 8 may be thought of as an embedding of the complete graph K_6 in \mathbb{RP}^2 .




Figure 8. A triangulation of the real projective plane, \mathbb{RP}^2 .

Let Δ be the 2-dimensional simplicial complex pictured in Figure 8. By Problem 15.8, all the Betti numbers of Δ are 0. The Smith normal form for ∂_2 is the 15×10 matrix with diagonal elements $d_1 = \dots = d_9 = 1$ and $d_{10} = 2$. It follows that $\tilde{H}_1(\mathbb{RP}^2) = \mathbb{Z}_2$. All other homology groups are trivial (cf. Problem 15.8).

15.2. Higher-dimensional critical groups

Every simple connected graph $G = (V, E)$ determines a one-dimensional simplicial complex with E as its set of facets. Its chain complex is

$$0 \longrightarrow \mathbb{Z}E \xrightarrow{\partial_1} \mathbb{Z}V \xrightarrow{\partial_0} \mathbb{Z} \longrightarrow 0.$$

We have seen this sequence before, in Chapter 13. According to Theorem 13.10, it is exact at $\mathbb{Z}V$, and the kernel of ∂_1 —the first homology group of the simplicial complex—is the cycle space, \mathcal{C} , for G . Further, \mathcal{C} is a free abelian group of rank $\tilde{\beta}_1 = |E| - |V| + 1$. All other Betti numbers are 0.

Thinking of $\mathbb{Z}V$ as the group of divisors on G , the mapping ∂_0 is the degree mapping. Hence, $\ker \partial_0$ is the group of divisors of degree 0. In Exercise 13.1 we found $\partial_1 \partial_1^t$ is the Laplacian of G . Therefore,

$$\ker \partial_0 / \text{im } \partial_1 \partial_1^t = \text{Jac}(G) \simeq \mathcal{S}(G),$$

motivating the following definition.

Definition 15.10. Let Δ be a simplicial complex of dimension d . The i -dimensional Laplacian or i -Laplacian mapping for Δ is

$$L_i := \partial_{i+1} \partial_{i+1}^t : C_i(\Delta) \rightarrow C_i(\Delta).$$

The i -th critical group for Δ is

$$\mathcal{K}_i(\Delta) := \ker \partial_i / \text{im } \partial_{i+1} \partial_{i+1}^t.$$

The critical group of Δ is

$$\mathcal{K}(\Delta) := \bigoplus_{i=0}^d \mathcal{K}_i(\Delta).$$

The transpose ∂_{i+1}^t is the mapping

$$\begin{aligned} \text{Hom}(C_i(\Delta), \mathbb{Z}) &\rightarrow \text{Hom}(C_{i+1}(\Delta), \mathbb{Z}) \\ \phi &\mapsto \phi \circ \partial_{i+1}. \end{aligned}$$

Choosing bases for $C_i(\Delta)$ and $C_{i+1}(\Delta)$ identifies ∂_{i+1} with a matrix and ∂_{i+1}^t with the transpose of that matrix.

Remark 15.11. Since $\text{im } \partial_{i+1} \partial_{i+1}^t \subseteq \text{im } \partial_{i+1}$, the i -th homology group is a quotient of the i -th critical group:

$$\mathcal{K}_i(\Delta) \twoheadrightarrow \tilde{H}_i(\Delta).$$

Accordingly, the critical groups encode more information about Δ . Since the rank of a real matrix is equal to the rank of the product of the matrix with its transpose, $\mathcal{K}_i(\Delta)$ and $\tilde{H}_i(\Delta)$ have the same rank—these groups differ only in torsion.

If Δ is d -dimensional, then $\partial_{d+1} = 0$, hence,

$$\mathcal{K}_d(\Delta) = \tilde{H}_d(\Delta) = \ker \partial_d.$$

Example 15.12. The k -skeleton of a simplicial complex Δ is the simplicial complex $\text{Skel}_k(\Delta)$ formed by all faces of Δ of dimension at most k . The 1-skeleton is called the underlying graph of the simplicial complex. The 0-th critical group of Δ is the Jacobian group of its underlying graph:

$$\mathcal{K}_0(\Delta) = \text{Jac}(\text{Skel}_1(\Delta)) \simeq \mathcal{S}(\text{Skel}_1(\Delta)).$$

The same reasoning used to compute homology groups in Section 15.1.2 applies directly to computing critical groups, yielding the following result.

Proposition 15.13. *Let Δ be a simplicial complex. Let d_1, \dots, d_k be the nonzero invariant factors of the i -Laplacian mapping, L_i , and let $r = \text{rk}(\ker \partial_i) - k$. Then*

$$\mathcal{K}_i(\Delta) \simeq \mathbb{Z}_{d_1} \times \cdots \times \mathbb{Z}_{d_k} \times \mathbb{Z}^r.$$

Example 15.14 (Tetrahedron). Let Δ be the simplicial complex whose facets are all 3-element subsets of $[4]$ —a triangulation of a hollow tetrahedron. We have

$$\partial_2: \mathbb{Z}^4 \xrightarrow{\begin{matrix} & \overline{123} & \overline{124} & \overline{134} & \overline{234} \\ \overline{12} & \begin{pmatrix} 1 & 1 & 0 & 0 \\ \overline{13} & \begin{pmatrix} -1 & 0 & 1 & 0 \\ \overline{14} & \begin{pmatrix} 0 & -1 & -1 & 0 \\ \overline{23} & \begin{pmatrix} 1 & 0 & 0 & 1 \\ \overline{24} & \begin{pmatrix} 0 & 1 & 0 & -1 \\ \overline{34} & \begin{pmatrix} 0 & 0 & 1 & 1 \end{pmatrix} \end{pmatrix} \end{pmatrix} \end{pmatrix} \end{pmatrix} \end{matrix}} \mathbb{Z}^6.$$

Thus, $\text{rk}(\partial_2) = 3$. Hence, $\text{rk}(\ker \partial_2) = 1$, and

$$\mathcal{K}_2(\Delta) = \tilde{H}_2(\Delta) = \ker \partial_2 \simeq \mathbb{Z}.$$

The nonzero invariant factors of

$$L_1 = \partial_2 \partial_2^t = \begin{pmatrix} 2 & -1 & -1 & 1 & 1 & 0 \\ -1 & 2 & -1 & -1 & 0 & 1 \\ -1 & -1 & 2 & 0 & -1 & -1 \\ 1 & -1 & 0 & 2 & -1 & 1 \\ 1 & 0 & -1 & -1 & 2 & -1 \\ 0 & 1 & -1 & 1 & -1 & 2 \end{pmatrix}$$

are $d_1 = d_2 = 1$ and $d_3 = 4$, and the kernel of ∂_1 has rank 3. Therefore, by Proposition 15.13,

$$\mathcal{K}_1(\Delta) \simeq \mathbb{Z}_4.$$

The underlying graph of Δ is the complete graph K_4 ; hence,

$$\mathcal{K}_0(\Delta) \simeq \mathcal{S}(K_4) \simeq \mathbb{Z}_4 \times \mathbb{Z}_4.$$

In general, it is shown in [41] that if Δ' is any simplicial complex drawn on a d -dimensional sphere, then $\mathcal{K}_{d-1}(\Delta') \simeq \mathbb{Z}_n$ where n is the number of facets of Δ' . On the other hand, $\tilde{H}_{d-1}(\Delta') = 0$. So the critical groups encode combinatorial structure not determined by the underlying topology of the sphere.

15.3. Simplicial spanning trees

Let Δ be a d -dimensional simplicial complex. A simplicial complex Δ' is a *subcomplex* of Δ , written $\Delta' \subseteq \Delta$, if the faces of Δ' are a subset of the faces of Δ , i.e., $F_i(\Delta') \subseteq F_i(\Delta)$ for all i . Recall: the i -skeleton $\text{Skel}_i(\Delta)$ is the subcomplex of Δ

formed by the collection of all faces of Δ of dimension at most i ; the i -th Betti number of Δ is $\tilde{\beta}_i(\Delta) := \text{rk } \tilde{H}_i(\Delta)$; and the f -vector (face-vector) for Δ is

$$f(\Delta) := (f_{-1}, \dots, f_d),$$

where $f_i := |F_i(\Delta)| = \text{rk } C_i(\Delta)$. (When we saw the f -vector in Chapter 14, it was notationally convenient to write it in terms of the *rank* ($= 1 + \dim \Delta$) rather than the dimension of Δ .)

Definition 15.15. Let Δ be a d -dimensional simplicial complex. A (*simplicial*) *spanning tree* of Δ is a d -dimensional subcomplex $\Upsilon \subseteq \Delta$ with $\text{Skel}_{d-1}(\Upsilon) = \text{Skel}_{d-1}(\Delta)$ and satisfying the three conditions:

- (1) $\tilde{H}_d(\Upsilon) = 0$;
- (2) $\tilde{\beta}_{d-1}(\Upsilon) = 0$;
- (3) $f_d(\Upsilon) = f_d(\Delta) - \tilde{\beta}_d(\Delta) + \tilde{\beta}_{d-1}(\Delta)$.

If $0 \leq k < d$, a k -dimensional (*simplicial*) *spanning tree* of Δ is a spanning tree of the k -skeleton $\text{Skel}_k(\Delta)$.

Exercise 15.16. What are the 0-dimensional spanning trees of a simplicial complex?

Example 15.17. Let Δ be a 1-dimensional simplicial complex, i.e., a graph, and suppose Υ is a simplicial spanning tree of Δ . The condition $\text{Skel}_0(\Upsilon) = \text{Skel}_0(\Delta)$ says that Υ contains all of the vertices of Δ . Condition (1) says Υ has no cycles. Condition (2) says Υ is connected. What about condition (3)? Since Υ contains all of the vertices and is connected, it follows that Δ is connected, and hence, $\tilde{\beta}_0(\Delta) = 0$. Since Δ is 1-dimensional, $\tilde{H}_1(\Delta) = \ker(\partial_1) = \mathcal{C}$, the cycle space of Δ . Therefore, $\tilde{\beta}_1(\Delta) = \text{rk } \mathcal{C}$. Then, by Theorem 13.10 (3), $\tilde{\beta}_1(\Delta) = f_1(\Delta) - f_0(\Delta) + 1$. Hence, condition (3) says the number of edges of Υ is one less than the number of vertices of Δ . So simplicial spanning trees are the same as ordinary spanning trees in the case of a graph.

Example 15.18. It turns out that if Δ is any triangulation of the real projective plane, \mathbb{RP}^2 , then Δ is its own spanning tree (cf. Example 15.9 and Problem 15.8). In this case, $\tilde{H}_1(\Delta) = \mathbb{Z}_2 \neq 0$, but condition (2) is still satisfied.

In fact, any two of the three conditions defining a spanning tree suffice:

Proposition 15.19. *Let Υ be a d -dimensional subcomplex of a d -dimensional simplicial complex Δ , and suppose that $\text{Skel}_{d-1}(\Upsilon) = \text{Skel}_{d-1}(\Delta)$. Then any two of the three conditions in Definition 15.15 imply the remaining one.*

Proof. Problem 15.9. □

Example 15.20. (Equatorial bipyramid.) The *equatorial bipyramid*, B , is pictured in Figure 9. It has three top facets, $\overline{124}$, $\overline{134}$, $\overline{234}$; three bottom facets, $\overline{125}$, $\overline{135}$, $\overline{235}$; and a middle facet, $\overline{123}$. By Problem 15.10, the only nontrivial homology group for B is $\tilde{H}_2(B) = \mathbb{Z}^2$. Its generating 2-cycles are the boundaries of the two missing 3-dimensional faces, $\overline{1234}$ and $\overline{1235}$. Condition (3) of the definition of a spanning tree requires any spanning tree of B to be formed by removing two facets




Figure 9. The equatorial bipyramid—a 2-dimensional simplicial complex with 7 facets.

from B . In order to satisfy condition (1), we must either remove a top face and a bottom face or remove the middle face and any other face. (And having satisfied conditions (1) and (3), condition (2) must hold as well.) Thus, in total, B has 15 spanning trees.

When does a simplicial complex have a spanning tree? For a graph, we have seen that possessing a spanning tree is equivalent to being connected, i.e., having 0-th Betti number 0. There is a simple generalization that characterizes the existence of spanning trees in higher dimensions.

Proposition 15.21. *Suppose that Δ is a d -dimensional simplicial complex. Then Δ possesses a simplicial spanning tree if and only if $\tilde{\beta}_{d-1}(\Delta) = 0$. We say that such complexes are acyclic in codimension 1.*

Proof. First suppose that Υ is a spanning tree of Δ . Then $C_d(\Upsilon) \subset C_d(\Delta)$, $C_{d-1}(\Upsilon) = C_{d-1}(\Delta)$, and $\text{im}(\partial_{\Upsilon,d}) \subseteq \text{im}(\partial_{\Delta,d})$. It follows that there is a surjection

$$\tilde{H}_{d-1}(\Upsilon) = C_{d-1}(\Upsilon)/\text{im}(\partial_{\Upsilon,d}) \twoheadrightarrow C_{d-1}(\Delta)/\text{im}(\partial_{\Delta,d}) = \tilde{H}_{d-1}(\Delta).$$

Since $\tilde{H}_{d-1}(\Upsilon)$ is finite, it follows that $\tilde{H}_{d-1}(\Delta)$ is also finite, hence $\tilde{\beta}_{d-1}(\Delta) = 0$.

Now suppose that Δ is acyclic in codimension 1. To construct a spanning tree, start with $\Upsilon = \Delta$. If $\tilde{H}_d(\Upsilon) = 0$, then Υ is a spanning tree, and we are done. If not, then there is an integer-linear combination of facets σ_i in the kernel of ∂_d :

$$a_1\sigma_1 + a_2\sigma_2 + \cdots + a_k\sigma_k,$$

where we assume that $a_1 \neq 0$. If we work over the rational numbers, then we may assume $a_1 = 1$. Still working over the rationals, we see that

$$\partial_d(\sigma_1) = -\sum_{i=2}^k a_i \partial_d(\sigma_i).$$

Hence, if we remove the facet σ_1 from Υ , we obtain a smaller subcomplex Υ' without changing the image of the rational boundary map: $\text{im}(\partial_{\Upsilon,d}) = \text{im}(\partial_{\Upsilon',d})$. It follows from rank-nullity that

$$\tilde{\beta}_d(\Upsilon') = f_d(\Upsilon') - \text{rk } \text{im}(\partial_{\Upsilon',d}) = f_d(\Upsilon) - 1 - \text{rk } \text{im}(\partial_{\Upsilon,d}) = \tilde{\beta}_d(\Upsilon) - 1,$$

and $\tilde{\beta}_{d-1}(\Upsilon') = \tilde{\beta}_{d-1}(\Upsilon) = 0$. Continuing to remove facets in this way, we eventually obtain a spanning tree of Δ . \square

Definition 15.22. A simplicial complex Δ is *acyclic in positive codimension (APC)* if $\tilde{\beta}_i(\Delta) = 0$ for all $i < \dim \Delta$.

Proposition 15.23. Let Δ be a d -dimensional simplicial complex. The following are equivalent:

- (1) Δ is APC.
- (2) Δ has a k -dimensional spanning tree for each $0 \leq k \leq d$.

Proof. Problem 15.11. \square

Geometrically, a prototypical APC complex is the triangulation of a d -sphere. For a slightly more complicated example, take several d -spheres and pick a point on each. Now glue these spheres together by identifying the chosen points. A triangulation of this *wedge* of d -spheres turns out to be APC. For example, consider the equatorial bipyramid. Contracting the middle face does not change the homology groups, and the resulting figure is the wedge of two 2-spheres.

15.3.1. Reduced Laplacians and the generalized matrix-tree theorem. In the case of a sandpile graph, the reduced Laplacian is obtained from the ordinary Laplacian by removing the row and column corresponding to the sink vertex. The cokernel of the reduced Laplacian is isomorphic to the sandpile group of the graph, and the determinant counts the number of spanning trees. Both of these results generalize to simplicial complexes of arbitrary dimension.

Let Δ be a d -dimensional simplicial complex. For each $0 \leq i \leq d$, define the i -th tree number:

$$\tau_i := \tau_i(\Delta) := \sum_{\Upsilon} |\tilde{H}_{i-1}(\Upsilon)|^2,$$

where the sum is over all i -dimensional spanning trees of Δ . Recall that for an i -dimensional spanning tree Υ we have $\tilde{\beta}_{i-1}(\Upsilon) = 0$; hence, $\tilde{H}_{i-1}(\Upsilon)$ consists only of torsion and is finite. Note that $\tau_i = 0$ if $\tilde{\beta}_{i-1}(\Delta) > 0$ (Proposition 15.21).

Suppose that $\tilde{\beta}_{i-1}(\Delta) = 0$, and fix an i -dimensional spanning tree Υ . Let

$$\tilde{F}_i := F_i(\Delta) \setminus F_i(\Upsilon),$$

the i -faces of Δ not contained in Υ , and let \tilde{L}_i be the i -Laplacian of Δ with the rows and columns corresponding to faces in $F_i(\Upsilon)$ removed.

Theorem 15.24 (Duval, Klivans, Martin [40, Theorem 1.3], [41, Theorem 3.4]).

(1) If $\tilde{H}_{i-1}(\Upsilon) = 0$, there is an isomorphism

$$\psi: \mathcal{K}_i(\Delta) \rightarrow \mathbb{Z}\tilde{F}_i / \text{im}(\tilde{L}_i),$$

defined by dropping i -faces of Υ : if $c = \sum_{f \in F_i(\Delta)} c_f \cdot f \in \ker \partial_i$ represents an element of $\mathcal{K}_i(\Delta)$, then

$$\psi(c) = \sum_{f \in \tilde{F}_i} c_f \cdot f \pmod{\text{im}(\tilde{L}_i)}.$$

(2) (*Simplicial matrix-tree theorem*)

$$\tau_{i+1} = \frac{|\tilde{H}_{i-1}(\Delta)|^2}{|\tilde{H}_{i-1}(\Upsilon)|^2} \det(\tilde{L}_i).$$

Example 15.25. Let G be the 1-skeleton of Δ , i.e., the underlying graph, and consider the $i = 0$ case of Theorem 15.24. In this case the spanning tree, Υ , is a single vertex. Part (1) says the Jacobian group of G is the cokernel of its ordinary reduced Laplacian. For part (2), note that $\tilde{H}_0(T) = 0$ for all spanning trees T of G , so τ_1 is the number of spanning trees of G , each counted with weight 1. Also, $\tilde{H}_{-1}(\Delta) = \tilde{H}_{-1}(\Upsilon) = 0$. Therefore, part (2) is the ordinary matrix-tree theorem applied to G .

Example 15.26. Let $K_{n,d}$ denote the d -skeleton of the $(n - 1)$ -dimensional simplex Δ_{n-1} (see Problem 15.4). Explicitly, the facets of $K_{n,d}$ consist of all subsets of $[n]$ of size $d + 1$. For $d = 1$, the complex $K_{n,1} = K_n$ is simply the complete graph on n vertices, and Cayley's formula 9.8 says that

$$\tau_1 = n^{n-2}.$$

In [60], Kalai established the following generalization of Cayley's result:

$$\tau_d = n^{\binom{n-d}{d}}.$$

Consider the particular case of $K_{6,2}$, the 2-skeleton of the 5-dimensional simplex with vertex set $[6]$. By Kalai's formula, the tree number τ_2 of this complex is $6^6 = 46656$. Some of the simplicial spanning trees for $K_{6,2}$ are contractible, hence have no homology, and thus contribute 1 to τ_2 . For instance, consider the tree T comprised of all 10 facets containing the vertex 1. On the other hand, $K_{6,2}$ also has spanning trees with nontrivial 1-dimensional homology. Indeed, one such tree is the triangulation of $\mathbb{R}P^2$ shown in Figure 8, having $\tilde{H}_1 = \mathbb{Z}_2$. This tree thus contributes $2^2 = 4$ to the quantity τ_2 .

Exercise 15.27. Let Δ be the simplicial complex whose facets are the faces of a tetrahedron (cf. Example 15.14). Verify the cases $i = 1$ and $i = 2$ of both parts of Theorem 15.24 for Δ (for a single fixed spanning tree in each case). In addition, by calculating Smith normal forms of reduced Laplacians, verify that $\mathcal{K}_1(\Delta) \simeq \mathbb{Z}_4$ and that $\mathcal{K}_2(\Delta) \simeq \mathbb{Z}$.

Remark 15.28. In Chapter 9 we proved a matrix-tree theorem for graphs in which spanning trees were allowed to have edges with arbitrary weights. For a generalization to simplicial complexes with weighted facets, see [40, Theorem 1.4].

15.4. Firing rules for faces

We know the columns of the ordinary Laplacian matrix of a graph encode vertex-firing rules (cf. Figure 1 from Chapter 2). In this section, we interpret the columns of the higher-dimensional Laplacians of a simplicial complex as firing rules for higher-dimensional faces. For this purpose, the group of i -dimensional chains, $C_i(\Delta)$, of a simplicial complex Δ are called i -flows and thought of as assignments of "flow" along the i -faces. Firing an i -face redirects flow around incident $(i + 1)$ -faces. An element of $\ker \partial_i$ is thought of as a *conservative* flow—one for which flow does

not build up along the $(i - 1)$ -dimensional boundary. Thus, the i -th critical group is interpreted as conservative i -flows modulo redirections of flow around incident $(i + 1)$ -faces.

Definition 15.29. Let $c \in C_i(\Delta)$ be an i -flow on a simplicial complex Δ and let $\sigma \in F_i(\Delta)$ be an i -dimensional face. Let L_i be the i -Laplacian of Δ . Firing σ from c produces the configuration $c - L_i\sigma$, and we use the notation:

$$c \xrightarrow{\sigma} c - L_i\sigma.$$

Example 15.30. Let Δ be the 2-dimensional simplicial complex with facets $\{1, 2, 3\}$ and $\{2, 3, 4\}$ pictured in Figure 10.




Figure 10. Simplicial complex for Example 15.30.

Figure 11 displays a 1-flow on Δ :

$$c = 2 \cdot \overline{12} - 2 \cdot \overline{13} + 3 \cdot \overline{23} - \overline{24} + \overline{34} \in C_1(\Delta).$$

Note that negative coefficients are interpreted visually as positive flow in the opposite direction: $-2 \cdot \overline{13}$ represents 2 units of flow directed from vertex 3 to vertex 1. The flow c is conservative since $\partial_1(c) = 0$, which means that for each vertex, the incoming flow balances the outgoing flow.

What happens when the edge $\overline{23}$ fires? In analogy with vertex firing, we subtract the $\overline{23}$ -column of the 1-dimensional Laplacian.




Figure 11. A conservative 1-flow on Δ . (Labels on directed edges represent units of flow; an unlabeled directed edge represents one unit of flow.)

The chain complex for Δ is

$$0 \rightarrow \mathbb{Z}^2 \xrightarrow{\partial_2} \mathbb{Z}^5 \xrightarrow{\partial_1} \mathbb{Z}^4 \xrightarrow{\partial_0} \mathbb{Z} \rightarrow 0,$$

$$\begin{matrix} \overline{12} & \overline{123} & \overline{234} \\ \overline{13} & \begin{pmatrix} 1 & 0 \\ -1 & 0 \end{pmatrix} \\ \overline{23} & \begin{pmatrix} 1 & 1 \\ 0 & -1 \end{pmatrix} \\ \overline{24} & \begin{pmatrix} 0 & -1 \\ 0 & 1 \end{pmatrix} \\ \overline{34} & \end{matrix} \quad \begin{matrix} \overline{1} & \overline{2} & \overline{3} & \overline{4} \\ \overline{1} & \begin{pmatrix} -1 & -1 & 0 & 0 & 0 \\ 1 & 0 & -1 & -1 & 0 \\ 0 & 1 & 1 & 0 & -1 \\ 0 & 0 & 0 & 1 & 1 \end{pmatrix} \\ \overline{2} & \\ \overline{3} & \\ \overline{4} & \end{matrix} \quad \begin{matrix} \overline{1} & \overline{2} & \overline{3} & \overline{4} \\ \partial_0 & \begin{pmatrix} 1 & 1 & 1 & 1 \end{pmatrix} \end{matrix}$$

and the 1-Laplacian is

$$L_1 = \partial_2 \partial_2^t = \begin{matrix} & \overline{12} & \overline{13} & \overline{23} & \overline{24} & \overline{34} \\ \overline{12} & \begin{pmatrix} 1 & -1 & 1 & 0 & 0 \\ -1 & 1 & -1 & 0 & 0 \\ 1 & -1 & 2 & -1 & 1 \\ 0 & 0 & -1 & 1 & -1 \\ 0 & 0 & 1 & -1 & 1 \end{pmatrix} \end{matrix}.$$

The firing rule determined by the $\overline{23}$ -column of L_1 is illustrated in Figure 12.




Figure 12. Firing the edge $\overline{23}$ diverts flow around the bordering 2-faces. The net flow into each vertex is unaffected by the firing.

Figure 13 shows the effect on c of firing the edge $\overline{23}$:

$$c \xrightarrow{\overline{23}} c - L_1(\overline{23}) = \overline{12} - \overline{13} + \overline{23}.$$




Figure 13. Firing the edge $\overline{23}$.

In Part 1 of this book, we interpreted graph vertex-firing rules as lending moves in the dollar game; in Part 2 we interpreted them as sandpile-topplings. We end this chapter by briefly indicating possibilities for corresponding interpretations of the face-firing rules on a simplicial complex.

15.4.1. The flow game. Let Δ be a d -dimensional simplicial complex, and consider a $(d-1)$ -flow $D = \sum_{f \in F_{d-1}(\Delta)} D(f) f \in C_{d-1}(\Delta)$. By analogy with the theory of divisors on graphs, we say that D is *effective* if $D(f) \geq 0$ for all $(d-1)$ -faces f . We think of effective flows as those which flow in the “correct” direction, as determined by the orientation coming from the vertex-ordering of Δ . Given two $(d-1)$ -flows D and D' , we say that D is *linearly equivalent* to D' if D may be obtained from D' via a sequence of face-firings/reverse-firings. With this setup, we have a natural generalization of the dollar game to d -dimensional simplicial complexes Δ :

The flow game: Is a given $(d-1)$ -flow D on Δ linearly equivalent to an effective $(d-1)$ -flow?

As of this writing, little is known about the flow game on simplicial complexes, although there are certainly many questions one could ask. Is there a greedy algorithm for winning the flow game? Is there a version of Dhar’s algorithm for flows? Most notably, is there a Riemann-Roch type theorem in this context?

15.4.2. Higher-dimensional chip firing. In order to interpret the face-firing rules in terms of a higher-dimensional version of the sandpile model, one would want notions of stability and recurrence for i -flows on a simplicial complex Δ . Unfortunately, the search for such notions is complicated by the fact that reduced i -Laplacians for $i \geq 1$ are not M -matrices (indeed, they are not even Z -matrices—see Section 12.3). However, in [53] Guzmán and Klivans propose a generalization of chip-firing in the context of an arbitrary invertible matrix \tilde{L} , at the price of choosing an auxiliary M -matrix M . In the case of graphs, where the invertible matrix \tilde{L} is the usual graph Laplacian, one may choose $M = \tilde{L}$ and obtain the usual sandpile theory. On the other hand, one can always choose $M = I$, the identity matrix. In the case where \tilde{L} is the reduced i -Laplacian of a simplicial complex, any choice of M leads to well-defined notions of *stable* flows, *superstable* flows, and *critical* flows (the latter being analogous to recurrent sandpiles). Moreover, just as the collections of superstable and recurrent sandpiles each form systems of representatives for the sandpile group, so do the superstable and critical i -flows form systems of representatives for the i -dimensional critical group.

However, many questions remain in this context as well. For instance, is there a good notion of duality between the superstable and critical i -flows, generalizing the duality between superstables and recurrents for sandpiles? Perhaps most importantly, is there a natural choice for the M -matrix M in the case where \tilde{L} is the reduced i -Laplacian of a simplicial complex?

Notes

Most of the results in this chapter are taken from the papers [40] and [41] by A. Duval, C. Klivans, and J. Martin. To avoid confusion, we note that in these original papers there are occasional minor errors with regard to precise acyclicity conditions, but the same authors corrected and generalized all of their results in the later book chapter [42].

Problems for Chapter 15

Note. For convenience, in all computations involving chain complexes in these problems, please use lexicographic ordering of the faces: $\overline{i_1 \cdots i_k}$ precedes $\overline{j_1 \cdots j_k}$ if the first nonzero component of $(i_1 - j_1, \dots, i_k - j_k)$ is negative.

15.1. (The boundary of a boundary is empty.) Show that for each i , the following relation holds for the boundary mappings in the chain complex of a simplicial complex: $\partial_{i-1} \circ \partial_i = 0$. This relation is often abbreviated as $\partial^2 = 0$. It justifies the use of the word *complex*.

15.2. Let Δ be the simplicial complex in Figure 1.

- Find the matrices representing the chain complex for Δ with respect to lexicographic ordering of the faces in each dimension.
- Find the Betti numbers by computing ranks.
- Find the homology groups by computing Smith normal forms. Find generators for each non-trivial group.

15.3. Let Δ be a simplicial complex. Show that $\tilde{\beta}_0(\Delta)$ is one less than the number of connected components of Δ .

15.4.

- Let Δ_n denote the n -simplex, i.e., the n -dimension simplicial complex consisting of all subsets of $[n + 1]$. Thus, Δ_0 is a point, Δ_1 is a line segment, Δ_2 is a solid triangle, Δ_3 is a solid tetrahedron, and so on. Show that $\tilde{H}_i(\Delta_n) = 0$ for all i .
- Let Δ_n° be the $(n - 1)$ -dimensional simplicial complex consisting of all subsets of $[n + 1]$ except $[n + 1]$, itself. Thus, Δ_1° is a pair of points, Δ_2° is a hollow triangle, Δ_3° is a hollow tetrahedron, and so on. Prove that

$$\tilde{H}_i(\Delta_n^\circ) \approx \begin{cases} \mathbb{Z} & \text{if } i = n - 1, \\ 0 & \text{otherwise.} \end{cases}$$

15.5. (Sphere.) The left side of Figure 14 shows a simplicial complex drawn on a sphere. On the right side is the same simplicial complex projected onto the face $\overline{134}$. It has eight two-dimensional facets, counting $\overline{134}$. Find the chain complex for this simplicial complex and compute the Betti numbers.




Figure 14. Simplicial complex on a sphere for Problem 15.5.

15.6. (Torus.) Gluing the same-numbered vertices in Figure 15 and their corresponding edges produces a torus. The figure represents an embedding of the complete graph K_7 on the torus. Find the chain complex and compute the Betti numbers for the corresponding simplicial complex (with 14 facets).




Figure 15. Simplicial complex on a torus for Problem 15.6.

15.7. (Klein bottle.) Gluing the same-numbered vertices in Figure 16 and their corresponding edges produces a Klein bottle. The graph determines a simplicial complex with 16 facets. Find the chain complex and compute the homology groups for the corresponding simplicial complex. Find generators for the homology. Is there any torsion?




Figure 16. Simplicial complex on a Klein bottle for Problem 15.7.

15.8. Let Δ be the simplicial complex triangulating the real projective plane in Example 15.9.

- Compute the critical groups and homology groups for Δ by computing Smith normal forms of Laplacians.
- Verify that Δ is a 2-dimensional spanning tree of itself.
- Verify both parts of Theorem 15.24 for the cases $i = 0, 1, 2$ (for a single spanning tree in each case). Verify your calculation of the critical groups by computing Smith normal forms of reduced Laplacians.

15.9. Let Δ be a d -dimensional simplicial complex. The *reduced Euler characteristic of Δ* is the alternating sum of the reduced Betti numbers of Δ :

$$\tilde{\chi}(\Delta) := \sum_{i=0}^d (-1)^i \tilde{\beta}_i(\Delta) = \tilde{\beta}_0(\Delta) - \tilde{\beta}_1(\Delta) + \cdots + (-1)^d \tilde{\beta}_d(\Delta).$$

Now consider the chain complex for Δ :

$$0 \longrightarrow C_d(\Delta) \xrightarrow{\partial_d} \cdots \xrightarrow{\partial_2} C_1(\Delta) \xrightarrow{\partial_1} C_0(\Delta) \xrightarrow{\partial_0} C_{-1}(\Delta) \longrightarrow 0.$$

Recall that the f -vector $(f_{-1}, f_0, f_1, \dots, f_d)$ is defined by $f_i = \text{rk}(C_i(\Delta))$. Show that the reduced Euler characteristic of Δ may be computed as the alternating sum of the numbers f_i :

$$\tilde{\chi}(\Delta) = \sum_{i=-1}^d (-1)^i f_i = -f_{-1} + f_0 - f_1 + \cdots + (-1)^d f_d.$$

Now suppose that Υ is a simplicial spanning tree of Δ . Show that $f_i(\Upsilon) = f_i(\Delta)$ for $-1 \leq i \leq d-1$ and $\tilde{\beta}_i(\Upsilon) = \tilde{\beta}_i(\Delta)$ for $0 \leq i \leq d-2$. Use these observations to prove Proposition 15.19 by computing $\tilde{\chi}(\Upsilon) - \tilde{\chi}(\Delta)$ in two different ways.

15.10. Compute the homology groups for the equatorial bipyramid (cf. Example 15.20).

15.11. Prove Proposition 15.23.

15.12. Calculate the 1-dimensional Laplacian for the simplicial complex in Figure 1, and draw a picture like that in Figure 12 showing the effect of firing the edge $\overline{23}$.

15.13. Consider the *solid bipyramid*, obtained from the equatorial bipyramid in Figure 9 by adding the 3-dimensional facets $\overline{1234}$ and $\overline{1235}$. Calculate the 2-dimensional Laplacian, and draw a picture analogous to that in Figure 12 showing the effect of firing the face $\overline{123}$.

Appendix

In this appendix, we briefly recall some basic notions from graph theory.

A.1. Undirected multigraphs

Definition A.1. A *multiset* is a set A with a *multiplicity function* $m: A \rightarrow \mathbb{N}_{\geq 1}$ from A to the positive integers. Informally, we think of a multiset as a set A where each element $a \in A$ appears $m(a) \geq 1$ times. A *submultiset* of A is a subset $B \subseteq A$ together with a multiplicity function m' satisfying $m'(b) \leq m(b)$ for all $b \in B$.

Example A.2. The multiset $\{a, a, a, b, b, c\}$ has underlying set $A = \{a, b, c\}$ and multiplicity function m given by $m(a) = 3$, $m(b) = 2$, and $m(c) = 1$. The multiset $\{a, a, b, b\}$ is a submultiset of $\{a, a, a, b, b, c\}$.

Definition A.3. An (*undirected*) *multigraph* $G = (V, E)$ is a pair consisting of a set of *vertices* V and a multiset of *edges* E comprised of unordered pairs $\{v, w\}$ of vertices. We generally write vw or wv for the edge $\{v, w\}$. A *loop* is a singleton $\{v\}$, corresponding to an edge connecting the vertex v to itself. A multigraph G is *finite* if both V and E are finite; G is a *simple graph* if E is a set, i.e., if each edge occurs with multiplicity 1.

Definition A.4. Two vertices v and w are *adjacent* if vw is an edge of G , and the edge $e = vw$ is said to be *incident* to v and w . The *degree* of a vertex v is its number of incident edges, denoted $\deg_G(v)$.

Definition A.5. A *submultigraph* of a multigraph $G = (V, E)$ is a pair (V', E') with $V' \subseteq V$ and E' a submultiset of E such that every vertex appearing in an edge from E' is an element of V' . If $W \subseteq V$, then the submultigraph *induced by* W is the pair (W, E') where E' consists of every edge in E whose endpoints are both in W .

Definition A.6. A *path* in a multigraph G is an alternating sequence of vertices and edges $v_1, e_1, v_2, e_2, \dots, v_n$ such that $e_i = v_i v_{i+1}$ for all $i = 1, \dots, n - 1$. A finite multigraph is *connected* if for all vertices $v, w \in V$ there is a path in G with $v_1 = v$ and $v_n = w$. The maximal connected submultigraphs of G are its *connected components*.

In Part 1 of this text, we use the generic term *graph* to mean a finite, connected, undirected multigraph without loop edges.

Example A.7. Here are several infinite families of graphs, indexed by positive integers $n \geq 1$:

- (1) The *path graph* P_n consists of n vertices connected in a line; it has *length* $n - 1$ (Figure 1).



Figure 1. The path graph P_n of length $n - 1$.

- (2) The *cycle graph* C_n consists of n vertices connected in a circle (Figure 2). Note that C_2 has 2 edges connecting its two vertices.




Figure 2. The cycle graph C_9 on 9 vertices.

- (3) The *complete graph* K_n consists of n vertices and all possible edges, each with multiplicity 1 (Figure 3).




Figure 3. The complete graph K_7 on 7 vertices.

- (4) The *banana graph* B_n consists of 2 vertices connected by n edges (Figure 4). Note that $B_2 = C_2$.




Figure 4. The banana graph B_6 with 6 edges.

Definition A.8. A *cycle* in a multigraph G is a path $C = v_1, e_1, v_2, e_2, \dots, v_n$ such that $v_1 = v_n$, the first $n - 1$ vertices are distinct, and no edge is repeated. We consider two cycles to be the same if they differ by a cyclic shift—i.e., cycles do not have distinguished starting points. Note that a loop is a cycle of length 1. We treat multiple edges between the same two vertices as distinct, so that the graph C_2 is a cycle. A multigraph with no cycles is called *acyclic*. The path graph P_n in Figure 1 is acyclic, while the banana graph in Figure 4 has $\binom{n}{2}$ cycles of length 2.

Definition A.9. A *forest* is an acyclic graph, and a *tree* is a connected forest. Note that forests are simple graphs, since if G has two edges e_1, e_2 of the form vw , then v, e_1, w, e_2, v is a cycle in G .

Proposition A.10. G is a tree if and only if there is a unique path in G between any two vertices.

Proof. The existence of a path between any two vertices is connectivity, and the uniqueness is acyclicity. \square

Proposition A.11. Suppose that T is a finite tree with at least two vertices. Then T has at least two vertices of degree 1, called leaf vertices.

Proof. By induction on n , the base case $n = 2$ being the path graph P_2 . So suppose that every tree on $2 \leq k < n$ vertices has at least 2 leaves, and that T is a tree on n vertices. Consider any edge $e = vw$ in T , which is the unique path between v and w in T . Removing e thus yields a disconnected graph $T_1 \sqcup T_2$, where $v \in T_1, w \in T_2$, and T_1 and T_2 are trees with fewer than n vertices.

First suppose that T_1 and T_2 each have at least 2 vertices. Then by the induction hypothesis, T_1 and T_2 each have at least 2 leaves, so that the disjoint union has at least 4 leaves. Hence, at least 2 leaf vertices remain in T after replacing the edge $e = vw$.

Now consider the case where T_1 is the single vertex v , while T_2 has at least 2 vertices. Then v is a leaf of the original tree T . Since T_2 has at least 2 leaves by the induction hypothesis, at least one remains in T after replacing $e = vw$, and this leaf is distinct from v . \square

Proposition A.12. Suppose that G is a multigraph on n vertices. The following are equivalent:

- (1) G is a tree;
- (2) G is minimal connected: G is connected and removing any edge from G yields a disconnected multigraph;

- (3) G is maximal acyclic: G is acyclic and adding any edge between vertices of G produces a cycle;
- (4) G is connected and has $n - 1$ edges;
- (5) G is acyclic and has $n - 1$ edges.

Proof. (1) \implies (2): Consider the removal of an edge $e = vw$ from G . If after removal there is still a path P from v to w , then P and v, e, w are two different paths in G from v to w , contradicting uniqueness.

(2) \implies (3): If G contained a cycle, then removing any edge contained in the cycle would not disconnect the graph, so G would not be minimal connected. Hence G is acyclic. Consider any two vertices v, w of G . Since G is connected, there exists a path P from v to w in G —choose P to be of minimal length, so that no vertex or edge is repeated. But then adding a new edge of the form $e = vw$ to G would yield a cycle P, e, v .

(3) \implies (1): Consider the graph H obtained by adding an edge $e = vw$ to G . By assumption, H has a cycle of the form $C = v, e, w, P$. Then P must be a path from w to v in G . Acyclicity implies that P is in fact the unique path from w to v in G .

(1) \implies (4) and (5): We prove that G has $n - 1$ edges by induction on n , the base case $n = 1$ being clear. So suppose that every tree on $n - 1$ vertices has $n - 2$ edges, and that G is a tree on n vertices. Choose a leaf vertex v , and let G' be the multigraph obtained from G by removing v and the unique edge incident to v . Then G' is a tree on $n - 1$ vertices, hence has $n - 2$ edges by the induction hypothesis. It follows that G has $n - 1$ edges as required.

(4) \implies (5) and (1): Suppose that G is connected with $n - 1$ edges. To get a contradiction, suppose that G has a cycle, and choose an edge e contained in the cycle. Removing e does not disconnect the graph G , so G is not minimal connected. Let G' be any minimal connected subgraph of G containing all n vertices. By (2) \implies (1), G' is a tree with fewer than $n - 1$ edges, contradicting the implication (1) \implies (4). Thus G is acyclic, and hence a tree.

(5) \implies (4): Finally, suppose that G is acyclic with $n - 1$ edges. Again we proceed by contradiction: suppose that G is not connected, and choose two vertices v, w in different connected components. Then adding the edge $e = vw$ does not produce a cycle, so G is not maximal acyclic. Let M be a maximal acyclic multigraph on the same n vertices as G and containing G as a subgraph. By (3) \implies (1), M is a tree with more than $n - 1$ edges, contradicting the implication (1) \implies (5). \square

Definition A.13. A *spanning forest* of a multigraph G is a maximal (with respect to inclusion of edge sets) acyclic subgraph F that contains all of the vertices of G . A *spanning tree* is a connected spanning forest. Note that G has a spanning tree if and only if G is connected, since by part 2 of Proposition A.12, a spanning tree is simply a minimal connected subgraph containing all of the vertices of G . Every multigraph has a spanning forest, obtained by choosing a spanning tree for each connected component and forming the disjoint union.

Definition A.14. A multigraph G is *k-edge connected* if it has at least 2 vertices and if every submultigraph obtained from G by removing fewer than k edges is

connected. The *edge-connectivity* of G is the largest k for which G is k -edge connected. Referring to Example A.7, the path graph P_n has edge-connectivity 1, the cycle graph C_n has edge-connectivity 2, the complete graph K_n has edge-connectivity $n - 1$, and the banana graph B_n has edge-connectivity n .

Definition A.15. A multigraph G is *k-connected* if it has more than k vertices and every submultigraph obtained from G by removing fewer than k vertices (together with all incident edges) is connected. The *connectivity* of G is the largest k for which G is k -connected. Referring to Example A.7, the path graph P_n has connectivity 1 for $n \geq 2$, the cycle graph C_n has connectivity 2 for $n \geq 3$, the complete graph K_n has connectivity $n - 1$ for $n \geq 1$ (since it only has n vertices), and the banana graph B_n has connectivity 1 for $n \geq 1$ (since it only has 2 vertices).

A.2. Directed multigraphs

Definition A.16. A *directed multigraph* or *multidigraph* $G = (V, E)$ is a pair consisting of a set of *vertices* V and a multiset of *directed edges* E comprised of ordered pairs (v, w) of vertices. In this case, we write vw for the directed edge (v, w) , which is distinct from the directed edge (w, v) , denoted wv .

Definition A.17. A *submultidigraph* of a multidigraph $G = (V, E)$ is a pair (V', E') with $V' \subseteq V$ and E' a submultiset of E such that every vertex appearing in an edge from E' is an element of V' . If $W \subseteq V$, then the submultidigraph *induced by* W is the pair (W, E') where E' consists of every edge in E whose endpoints are both in W .

Definition A.18. We say that an edge $e = vw$ in a multidigraph G *emanates* from v . The vertex $e^- = v$ is the *tail* and $e^+ = w$ is the *head* of e . The *outdegree* of a vertex v is the number of edges emanating from v , denoted $\text{outdeg}_G(v)$. Equivalently, $\text{outdeg}_G(v)$ is the number of edges with tail v . The *indegree* of v , denoted $\text{indeg}_G(v)$, is the number of edges with head v . A vertex v is a *source* if $\text{indeg}_G(v) = 0$, and v is a *sink* if $\text{outdeg}_G(v) = 0$. If $W \subseteq V$ and $v \in V$, then the outdegree and indegree *with respect to* W are, respectively,

$$\begin{aligned}\text{outdeg}_W(v) &:= |\{w \in W : (v, w) \in E\}| \\ \text{indeg}_W(v) &:= |\{w \in W : (w, v) \in E\}|.\end{aligned}$$

Thus, if $v \in W$, then these are the outdegree and indegree for v as a vertex in the the subgraph of G induced by W . In particular, $\text{indeg}_G(v) = \text{indeg}_W(v)$. If G is clear from context, we write $\text{outdeg}(v)$ and $\text{indeg}(v)$ for $\text{outdeg}_G(v)$ and $\text{indeg}_G(v)$, respectively.

Definition A.19. Every directed multigraph G has an *underlying undirected multigraph* G^{ud} obtained by replacing each directed edge (v, w) by the undirected edge $\{v, w\}$. It also has an *underlying simple undirected graph* formed from G^{ud} by setting the multiplicity of each of the edges of G^{ud} equal to 1.

Definition A.20. If $G = (V, E)$ is an undirected multigraph, then an *orientation* \mathcal{O} on G is the choice of a directed edge (v, w) or (w, v) for each copy of the undirected edge $\{v, w\}$ appearing in the multiset E . An orientation \mathcal{O} on G yields a directed multigraph (G, \mathcal{O}) with underlying undirected multigraph G .

Definition A.21. A multidigraph G is *connected* (sometimes called *weakly connected*) if its underlying undirected multidigraph G^{ud} is connected.

In Part 2 of this text, we use the generic term *graph* to mean a finite, connected multidigraph, with loop edges allowed.

Definition A.22. A (*directed*) *path* in a multidigraph G is a sequence of directed edges e_1, e_2, \dots, e_n such that $e_i^+ = e_{i+1}^-$ for $1 \leq i \leq n-1$; such a path has *length* n . A finite multidigraph is *strongly connected* if for all vertices $v, w \in V$ there is a directed path in G with $e_1^- = v$ and $e_n^+ = w$.

Definition A.23. A (*directed*) *cycle* in a multidigraph G is a directed path e_1, e_2, \dots, e_n such that $e_n^+ = e_1^-$, no other vertex appears more than once, and no directed edge is repeated. We consider two cycles to be the same if they differ by a cyclic shift. A multidigraph with no directed cycles is called *acyclic*.

Proposition A.24. *Suppose that G is a finite acyclic multidigraph. Then G has at least one sink and at least one source.*

Proof. Start at any vertex v and begin walking along directed edges—the walk will only stop when a sink is encountered. Since G is acyclic, the resulting directed path will never close up to produce a directed cycle. By finiteness, this means the walk must end at a sink after finitely many steps. Applying this argument to a “wrong-way” walk yields the existence of a source. \square

Definition A.25. Let G be a multidigraph and $s \in V$ a chosen *root* vertex. A *directed spanning tree of G rooted at s* is a subdigraph T with the property that for every $v \in V$, there exists a unique directed path in T from v to s .

Proposition A.26. *Suppose that G is a multidigraph, $s \in V$ is fixed, and T is a subdigraph. Then T is a directed spanning tree of G rooted at s if and only if*

- (1) T contains all vertices of G ;
- (2) T is acyclic;
- (3) All non-root vertices $v \neq s$ have $\text{outdeg}_T(v) = 1$, and $\text{outdeg}_T(s) = 0$.

Proof. Suppose that T is a directed spanning tree rooted at s . Then T clearly contains all vertices of G . In addition, T must be acyclic, because the existence of a directed cycle would violate the uniqueness of directed paths to s . Similarly, two distinct edges emanating from a non-root vertex would also violate uniqueness. Finally, any edge emanating from the root s would yield a cycle through s .

For the other direction, suppose that T satisfies all three stated conditions. Let $v \neq s$ be an arbitrary non-root vertex. Since each non-root vertex has outdegree 1 and there are no cycles in G , following the directed edges yields a unique directed path from v to s . \square

Proposition A.27. *Suppose that T is a directed spanning tree rooted at s in the multidigraph G . Then T^{ud} is a spanning tree of the underlying undirected multidigraph G^{ud} . In particular, if G has n vertices, then T has $n-1$ edges.*

Proof. The underlying undirected multigraph T^{ud} is connected and contains all vertices of G^{ud} . We must show that it is also acyclic as an undirected graph. So suppose that T^{ud} contains a cycle C , and let \mathcal{O} denote the orientation on C determined by the original directed tree T . Then (C, \mathcal{O}) is a subdigraph of T , hence acyclic. By Proposition A.24, it follows that (C, \mathcal{O}) contains a source, i.e., a vertex v with $\text{outdeg}_{(C, \mathcal{O})}(v) = 2$. But then $\text{outdeg}_T(v) \geq 2$, which contradicts Proposition A.26. \square

Definition A.28. An *Eulerian cycle* in a multidigraph G is an edge-disjoint union of cycles that contains all of the edges of G . A multidigraph is *Eulerian* if it contains an Eulerian cycle.

Proposition A.29. *Suppose that G is a connected multidigraph with at least one edge. Then G is Eulerian if and only if $\text{indeg}_G(v) = \text{outdeg}_G(v)$ for all vertices v .*

Proof. First suppose that C is an Eulerian cycle in G . Each subcycle of C passing through a vertex v contributes 1 to the indegree of v and 1 to the outdegree. Since each edge occurs exactly once in C , we see that $\text{indeg}_G(v) = \text{outdeg}_G(v)$.

Now suppose that $\text{indeg}_G(v) = \text{outdeg}_G(v)$ for all vertices v . In particular, G has no sources or sinks, so that G must contain a directed cycle. Let C be an edge-disjoint union of cycles with the maximal number of edges. If C does not contain all edges of G , we will derive a contradiction. Consider the subgraph H of G formed by the edges of G not contained in C . The subgraph H still has the property that $\text{indeg}_H(v) = \text{outdeg}_H(v)$ for all vertices v in H . As before, H must contain a directed cycle. But this cycle can then be combined with C to produce a larger edge-disjoint union, contradicting the maximality of C . \square

Definition A.30. Let G be an undirected multigraph. Then the *associated directed multigraph* G^{dir} is the multidigraph obtained from G by replacing each copy of the undirected edge $\{v, w\}$ by the pair of directed edges (v, w) and (w, v) . By the previous proposition, if G is connected then G^{dir} is Eulerian.

In this appendix, we collect some definitions, terminology, and basic results from algebra.

B.1. Monoids, groups, rings, and fields

Definition B.1. A *monoid* is a set S together with an associative binary operation $\star: S \times S \rightarrow S$ possessing an identity element e , meaning that every element $s \in S$ satisfies $e \star s = s \star e = s$. If the operation is commutative, then S is called a *commutative monoid*; in this case the operation is generally denoted by $+$ and the identity denoted by 0 .

As an example of a commutative monoid, consider the set \mathbb{N}^m of m -tuples of natural numbers with componentwise addition; the identity element is the zero vector $(0, 0, \dots, 0)$.

Definition B.2. A *group* is a monoid (G, \star) for which every element $g \in G$ has an inverse, denoted g^{-1} and satisfying $g \star g^{-1} = g^{-1} \star g = e$. An abelian group $(A, +)$ is a group in which the operation is commutative; in this case the inverse of an element $a \in A$ is denoted by $-a$ and satisfies $a + (-a) = 0$.

As an example of an abelian group, consider \mathbb{Z}^m with componentwise addition; this is called the *free abelian group of rank m* . For another example, consider $\mathbb{Z}_m := \{0, 1, 2, \dots, m-1\}$ with the operation of addition modulo m ; this is called the *cyclic group of order m* .

A *subgroup* B of an abelian group A is a subset $B \subseteq A$ that is itself a group under the addition law of A . For any such subgroup B , the *quotient group* A/B is defined to be the collection of distinct *cosets* $\{a + B : a \in A\}$ under the operation $(a + B) + (a' + B) := (a + a') + B$; the identity element is the coset $0 + B = B$. As an example of the quotient construction, consider the subgroup $m\mathbb{Z} \subset \mathbb{Z}$ consisting of multiples of a fixed integer $m > 0$. Then $\mathbb{Z}/m\mathbb{Z} = \mathbb{Z}_m$, where we identify the coset $k + m\mathbb{Z}$ with the element $k \in \mathbb{Z}_m$.

If G and H are groups, then a *group homomorphism* $f: G \rightarrow H$ is a mapping that preserves the operations: $f(g_1 \star g_2) = f(g_1) \star f(g_2)$ for all $g_1, g_2 \in G$. The

kernel of f is the subgroup of G given by $\ker(f) := f^{-1}(e)$ where e is the identity element of H . Continuing the previous example, the mapping $\mathbb{Z} \rightarrow \mathbb{Z}/m\mathbb{Z}$ defined by $k \mapsto k + m\mathbb{Z}$ is a surjective group homomorphism with kernel $m\mathbb{Z}$.

Definition B.3. Suppose that (G, \star) is a group and S is a set. A *group action of G on S* is a mapping $G \times S \rightarrow S$ (denoted by juxtaposition) satisfying

- $es = s$ for all $s \in S$;
- $(g \star h)s = g(hs)$ for all $g, h \in G$ and $s \in S$.

Given such a group action, the *stabilizer subgroup of $s \in S$* is the subgroup of G defined as $G_s := \{g \in G : gs = s\}$. The *orbit of $s \in S$* is the subset of S defined as $\mathcal{O}_G(s) := \{gs : g \in G\}$. In the case where G and S are both finite, the *orbit-stabilizer theorem* says that $|G| = |G_s| |\mathcal{O}_G(s)|$.

Definition B.4. A *commutative ring (with unity)* is an abelian group $(R, +)$ together with a second associative and commutative operation $R \times R \rightarrow R$ (called multiplication and denoted by juxtaposition) satisfying:

- $r(s + t) = rs + rt$ for all $r, s, t \in R$ (distributive law);
- there exists $1 \in R$ such that $1r = r$ for all $r \in R$ (unity).

As examples, consider the ring of integers \mathbb{Z} or the ring of integer-coefficient polynomials $\mathbb{Z}[x]$ with the usual addition and multiplication of numbers and polynomials. Also, the group \mathbb{Z}_m becomes a ring when multiplication is performed modulo m .

An *ideal I* of a commutative ring R is a subgroup that is closed under outside multiplication: for all $r \in R$ and $i \in I$ we required that $ri \in I$. In this case, the *quotient ring R/I* is the group of distinct cosets $r + I$ for $r \in R$ with multiplication given by $(r + I)(r' + I) := rr' + I$; the unity is given by the coset $1 + I$.

If R and S are commutative rings with unity, then a *ring homomorphism $f: R \rightarrow S$* is a group homomorphism that also preserves unity and multiplication: $f(1) = 1$ and $f(r_1 r_2) = f(r_1) f(r_2)$ for all $r_1, r_2 \in R$. The *kernel of f* is the ideal of R given by $\ker(f) := f^{-1}(0)$. The mapping $\mathbb{Z} \rightarrow \mathbb{Z}/m\mathbb{Z}$ defined by $k \mapsto k + m\mathbb{Z}$ is a surjective ring homomorphism with kernel $m\mathbb{Z}$.

Definition B.5. A *field* is a commutative ring \mathbb{F} with unity $1 \neq 0$ such that every non-zero element $a \in \mathbb{F}$ has a multiplicative inverse, denoted a^{-1} and satisfying $aa^{-1} = 1$.

As an example, consider the field of rational numbers \mathbb{Q} , with the usual addition and multiplication. Also, if p is any prime number, then the ring \mathbb{Z}_p with addition and multiplication modulo p is a field.

B.2. Modules

Let R be a commutative ring with unity $1 \neq 0$. An *R -module* is an abelian group M together with a mapping $R \times M \rightarrow M$ called *scalar multiplication* satisfying, for all $r, s \in R$ and $m, n \in M$:

$$(1) \quad (r + s)m = rm + sm$$

- (2) $r(m + n) = rm + rn$
- (3) $(rs)m = r(sm)$
- (4) $1m = m$.

A subgroup N of an R -module M is a *submodule* if $rn \in N$ for all $r \in R$ and $n \in N$. In that case, the *quotient module* M/N is the group of cosets $m + N$ for $m \in M$, with scalar multiplication given by $r(m + N) := rm + N$ for all $r \in R$ and $m \in M$.

A *homomorphism of R -modules* M and N is a group homomorphism $f: M \rightarrow N$ satisfying $f(rm) = rf(m)$ for all $r \in R$ and $m \in M$. The set of all R -module homomorphisms from M to N is denoted $\text{Hom}_R(M, N)$. It is, itself, naturally an R -module letting $(rf)(m) := r(f(m))$ for all $r \in R$ and $f \in \text{Hom}_R(M, N)$. Two R -modules M, N are *isomorphic* if there are R -module homomorphisms $f: M \rightarrow N$ and $g: N \rightarrow M$ that are inverses of each other: $g \circ f = \text{id}_M$ and $f \circ g = \text{id}_N$. If M and N are isomorphic, we write $M \simeq N$.

Examples

- (1) A \mathbb{Z} -module is just an abelian group.
- (2) If R is a field, an R -module is exactly a vector space over R .
- (3) For any ring R , every ideal of R is an R -submodule.
- (4) If $f: M \rightarrow N$ is a homomorphism of R -modules, then
 - (a) the *kernel* of f is the R -submodule of M defined by $\ker(f) := f^{-1}(0)$;
 - (b) the *image* of f is the R -submodule of N defined by $\text{im}(f) := f(M)$;
 - (c) the *cokernel* of f is the quotient R -module $N/\text{im}(f)$;
 - (d) the mapping f is injective iff $\ker(f) = 0$ and is surjective iff $\text{coker}(f) = 0$;
 - it is an isomorphism iff it is injective and surjective.
- (5) For every set S and R -module N , let N^S be the collection of all set-mappings $S \rightarrow N$. Then N^S is an R -module with the usual addition of functions and scalar multiplication defined by $(rf)(s) := r(f(s))$ for all $r \in R$, $s \in S$, and $f \in N^S$. If M is another R -module, then $N^M := \text{Hom}_R(M, N)$.

B.2.1. Direct products and sums. Suppose M_i is an R -module for each $i \in I$, where I is some index set. The *direct product* of the M_i is the cartesian product $\prod_{i \in I} M_i$ with R -module structure defined component-wise: $(rm + m')_i := r m_i + m'_i$ for all $r \in R$ and $m, m' \in \prod_{i \in I} M_i$. The *direct sum* of the M_i , denoted $\bigoplus_{i \in I} M_i$, is the submodule of $\prod_{i \in I} M_i$ consisting of those elements with only finitely many nonzero components. If I is finite, then $\bigoplus_{i \in I} M_i = \prod_{i \in I} M_i$. If n is a positive integer, then $M^n := \prod_{i=1}^n M = \bigoplus_{i=1}^n M$.

B.2.2. Free modules. An R -module M has a *basis* $E \subset M$ if

- (1) E *generates* M : for each $m \in M$, we have $m = \sum_{i=1}^k r_i e_i$ for some $k \in \mathbb{N}$, $r_i \in R$, and $e_i \in E$.
- (2) E is *linearly independent*: if $\sum_{i=1}^k r_i e_i = 0$ for distinct $e_1, \dots, e_k \in E$, then $r_i = 0$ for all i .

A module with a basis is called a *free module*.

Bases do not always exist. For example, the \mathbb{Z} -module \mathbb{Z}_2 has no basis since $2a = 0$ for all $a \in \mathbb{Z}_2$ even though $2 \neq 0$ in \mathbb{Z} . On the other hand, if R is a field, then an R -module is vector space over R and thus has a basis.

It turns out that any two bases of a free module have the same cardinality. This cardinality is called the *rank* of the module. The *rank* of an arbitrary module M , denoted $\text{rk}(M)$, is the maximal rank of a free submodule of M . Then M is a free module of rank $n \in \mathbb{N}$ if and only if $M \simeq R^n$.

B.2.3. Finite generation and Noetherian modules. An R -module M is *finitely generated* if there exist finitely many *generators* $m_1, m_2, \dots, m_n \in M$ such that every element $m \in M$ may be written as $m = \sum_{i=1}^n r_i m_i$ for some $r_i \in R$.

A commutative ring R is *Noetherian* if every ideal $I \subseteq R$ is finitely generated as an R -module. More generally, an R -module M is Noetherian if every submodule $N \subseteq M$ is finitely generated. A fundamental result states that if R is a Noetherian ring and M is a finitely generated R -module, then M is a Noetherian module—this may be proved by induction on the number of generators for M .

B.2.4. Principle ideal domains. A *principle ideal domain (PID)* is a commutative ring R such that every ideal I is generated by a single element $I = (a) := Ra$ for some $a \in R$. The extended Euclidean algorithm for finding the greatest common divisor of two integers shows that \mathbb{Z} is a PID.

Clearly, if R is a PID, then R is a Noetherian ring. It follows that the finitely generated free module R^m is a Noetherian R -module. Hence, every submodule of R^m is finitely generated. This fact plays a key role in the proof of the following structure theorem for finitely generated modules over a principle ideal domain:

Theorem B.6 (Structure theorem for f.g. modules over a PID). *Let R be a PID and M a finitely generated R -module. Then there exists an integer $r \geq 0$ and a unique chain of $k \geq 0$ principle ideals $0 \neq (a_k) \subsetneq (a_{k-1}) \subsetneq \dots \subsetneq (a_1) \neq R$ such that*

$$M \simeq R^r \oplus R/(a_1) \oplus R/(a_2) \oplus \dots \oplus R/(a_k).$$

We provide a complete proof of this structure theorem in the case $R = \mathbb{Z}$ (i.e., abelian groups) in Section 2.4.

B.2.5. Exact sequences. A sequence of R -module homomorphisms

$$M' \xrightarrow{f} M \xrightarrow{g} M''$$

is *exact* (or *exact at M*) if $\text{im}(f) = \ker(g)$. A *short exact sequence* of R -modules is a sequence of R -module homomorphisms

$$0 \longrightarrow M' \xrightarrow{f} M \xrightarrow{g} M'' \longrightarrow 0$$

exact at M' , M , and M'' .

Exercise B.7. For a short exact sequence of R -modules as above,

- (1) f is injective;
- (2) g is surjective;

- (3) M'' is isomorphic to $\text{coker}(f)$;
- (4) If R is a field, so that M' , M , and M'' are vector spaces, then

$$\dim M = \dim M' + \dim M''.$$

In general, a sequence of R -module mappings

$$\cdots \rightarrow M_i \rightarrow M_{i+1} \rightarrow \cdots$$

is *exact* if it is exact at each M_i (except the first and last, if they exist).

Consider a commutative diagram of R -modules with exact rows

$$\begin{array}{ccccccc} M' & \xrightarrow{f} & M & \xrightarrow{g} & M'' & \longrightarrow & 0 \\ \downarrow \phi' & & \downarrow \phi & & \downarrow \phi'' & & \\ 0 & \longrightarrow & N' & \xrightarrow{h} & N & \xrightarrow{k} & N'' \end{array} .$$

(By *commutative*, we mean $\phi \circ f = h \circ \phi'$ and $\phi'' \circ g = k \circ \phi$.)

The *snake lemma* says there is an exact sequence

$$\ker \phi' \rightarrow \ker \phi \rightarrow \ker \phi'' \rightarrow \text{coker } \phi' \rightarrow \text{coker } \phi \rightarrow \text{coker } \phi''.$$

If f is injective, then so is $\ker \phi' \rightarrow \ker \phi$, and if k is surjective, so is $\text{coker } \phi \rightarrow \text{coker } \phi''$.

Exercise B.8. Prove the snake lemma.

Exercise B.9. Show that the following are equivalent for a short exact sequence

$$0 \rightarrow M' \xrightarrow{f} M \xrightarrow{g} M'' \rightarrow 0$$

- (1) There exists a homomorphism $f': M \rightarrow M'$ such that $f' \circ f = \text{id}_{M'}$.
- (2) There exists a homomorphism $g': M'' \rightarrow M$ such that $g \circ g' = \text{id}_{M''}$.
- (3) There exists an isomorphism $h: M \rightarrow M' \oplus M''$ such that the following diagram commutes:

$$\begin{array}{ccccccc} 0 & \longrightarrow & M' & \xrightarrow{f} & M & \xrightarrow{g} & M'' \longrightarrow 0 \\ & & \swarrow \pi_1 & & \downarrow h & & \searrow \pi_2 \\ & & & & M' \oplus M'' & & \end{array}$$

where π_1 and π_2 are the projection mappings onto the first and second components, respectively.

Definition B.10. A short exact sequence satisfying any of the conditions in Exercise B.9 is called *split exact*.

Glossary of Symbols

$\text{Div}(G)$	divisor group 9
$\text{deg}(D)$	degree of a divisor 10
$\text{Div}^k(G), \text{Div}_+(G)$	degree k divisors, nonnegative divisors 10
$D \xrightarrow{v} D', D \xrightarrow{W} D'$	vertex and set firings/lendings for divisors 10
$D \sim D'$	linear equivalence of divisors 11
$[D]$	divisor class 12
$\text{Jac}(G), \text{Pic}(G)$	Jacobian and Picard groups 14
$ D $	complete linear system 15
$\mathcal{M}(G)$	group of firing scripts 20
$\text{div}(\sigma)$	divisor of a firing script 20
$\text{Prin}(G)$	principal divisors 20
L	Laplacian 21 , 102
\tilde{V}	non-sink vertices 25 , 97
$\text{Config}(G)$	group of configurations 25 , 97
$\text{deg}(c)$	degree of a configuration 25
\tilde{L}	reduced Laplacian 26 , 101
$\text{Pic}^d(G)$	degree d part of the Picard group 55
S_q	Abel-Jacobi map 56
\mathcal{O}	orientation 62

$\text{indeg}_{\mathcal{O}}, \text{outdeg}_{\mathcal{O}}$	orientation indegree, outdeg 62
$\mathbf{D}(\mathcal{O}), \mathbf{c}(\mathcal{O})$	divisor, configuration corresponding to an orientation 63
g	$ E - V + 1$, genus (cycle rank) 65
$r(D)$	rank of divisor D 69
\mathcal{O}^{rev}	reversed orientation 71
c°	stabilization of c 93
\rightsquigarrow	sequence of legal vertex firings 93
e^{-}, e^{+}	tail and head of a directed edge 96, 307
$\text{outdeg}(v), \text{indeg}(v)$	outdegree and indegree of a vertex 96, 307
$c \xrightarrow{v} c', c \xrightarrow{\sigma} c'$	vertex and script firings for configurations 98
$\mathcal{L}, \tilde{\mathcal{L}}$	Laplacian and reduced Laplacian lattices 103
c_{\max}	maximal stable configuration 104
$a \otimes b$	stable addition 105
$\mathcal{S}(G)$	sandpile group of G 107
$\text{supp}(c)$	support of a configuration 120
ζ_{τ}	threshold density 144
$\mathcal{B}(G, s)$	basic alive divisors 144
$\beta_v(c)$	burst size 146
ζ_{st}	stationary density 147
$\phi_*(D)$	push-forward of D by ϕ 198
$\phi^*(D)$	pull-back of D by ϕ 201
g_d^r	complete linear system for rank r , degree d divisor 208
$\mu(G)$	minimal number of generators for $\mathcal{S}(G)$ 232
\mathcal{C}	integral cycle space 246
\mathcal{C}^*	integral cut space 247
$M(G)$	cycle matroid of G 258
$T(G; x, y)$	Tutte polynomial of G 259

$\tilde{\chi}(\Delta)$	reduced Euler characteristic of a simplicial complex 272
$C_i(\Delta)$	i -chains of Δ 280
∂_i	i -th boundary mapping 280
$\tilde{H}_i(\Delta), \tilde{\beta}_i(\Delta)$	i -th reduced homology group and Betti number 282
$\mathcal{K}(\Delta)$	critical group of Δ 287
$\deg_G(v)$	degree of vertex 303
P_n, C_n, K_n, B_n	path, cycle, complete, and banana graphs 304
$\text{outdeg}_W(v), \text{indeg}_W(v)$	outdegree and indegree with respect to a subset 307
\mathbb{Z}_m	cyclic group of order m 311

Bibliography

1. L. Babai and E. Toumpakari, *A structure theory of the sandpile monoid for directed graphs*, [arXiv:1401.3309](#), to appear in Journal of Combinatorics.
2. R. Bacher, P. de la Harpe, and T. Nagnibeda, *The lattice of integral flows and the lattice of integral cuts on a finite graph*, Bull. Soc. Math. France **125** (1997), no. 2, 167–198.
3. S. Backman, *Riemann-Roch theory for graph orientations*, Adv. Math. **309** (2017), 655–691.
4. P. Bak, C. Tang, and K. Wiesenfeld, *Self-organized criticality: an explanation of $1/f$ noise*, Phys. Rev. Lett. **59** (1987), no. 4, 381–384.
5. M. Baker, *Specialization of linear systems from curves to graphs*, Algebra Number Theory **2** (2008), no. 6, 613–653, With an appendix by Brian Conrad.
6. M. Baker and S. Norine, *Riemann-Roch and Abel-Jacobi Theory on a Finite Graph*, Adv. Math. **215** (2007), 766–788.
7. ———, *Harmonic morphisms and hyperelliptic graphs*, Int. Math. Res. Not. IMRN (2009), no. 15, 2914–2955.
8. M. Baker and F. Shokrieh, *Chip Firing Games, Potential Theory on Graphs, and Spanning Trees*, J. Combinatorial Theory Series A **120** (2013), no. 1, 164–182.
9. M. Baker and Y. Wang, *The Bernardi Process and Torsor Structures on Spanning Trees*, International Mathematics Research Notices (2017).
10. M. Beck and S. Robins, *Computing the Continuous Discretely*, 2 ed., Undergraduate Texts in Mathematics, Springer-Verlag, New York, 2015, Integer-Point Enumeration in Polyhedra.
11. B. Benson, D. Chakrabarty, and P. Tetali, *G -parking functions, acyclic orientations and spanning trees*, Discrete Math. **310** (2010), no. 8, 1340–1353.
12. S. Bhupatiraju, J. Hanson, and A. A. Járai, *Inequalities for critical exponents in d -dimensional sandpiles*, [arXiv:1602.06475](#), 2016.
13. N. Biggs, *Algebraic potential theory on graphs*, Bull. London Math. Soc. **29** (1997), no. 6, 641–682.
14. ———, *Chip-firing and the critical group of a graph*, J. Algebraic Combin. **9** (1999), no. 1, 25–45.

15. ———, *The Tutte polynomial as a growth function*, J. Algebraic Combin. **10** (1999), no. 2, 115–133. MR 1719148
16. A. Björner, *The homology and shellability of matroids and geometric lattices*, Matroid applications, Encyclopedia Math. Appl., vol. 40, Cambridge Univ. Press, Cambridge, 1992, pp. 226–283. MR 1165544 (94a:52030)
17. A. Björner and L. Lovász, *Chip-firing games on directed graphs*, J. Algebraic Combin. **1** (1992), no. 4, 305–328.
18. A. Björner, L. Lovász, and P. W. Shor, *Chip-firing games on graphs*, European J. Combin. **12** (1991), no. 4, 283–291. MR 1120415
19. M. Boij, J. C. Migliore, R. M. Miró-Roig, U. Nagel, and F. Zanello, *On the shape of a pure O -sequence*, Mem. Amer. Math. Soc. **218** (2012), no. 1024, viii+78.
20. B. Bollobás, *Modern Graph Theory*, Graduate Texts in Mathematics, vol. 184, Springer, 1998.
21. B. Bond and L. Levine, *Abelian networks I. Foundations and examples*, SIAM J. Discrete Math. **30** (2016), no. 2, 856–874.
22. M. Brandt, *An alternative proof of a theorem by Gessel*, <https://math.berkeley.edu/~brandtm/research/divisors.pdf>, 2015.
23. M. Chan, T. Church, and J. A. Grochow, *Rotor-routing and spanning trees on planar graphs*, Int. Math. Res. Not. IMRN (2015), no. 11, 3225–3244.
24. M. Chan, D. Glass, M. Macauley, D. Perkinson, C. Werner, and Q. Yang, *Sandpiles, spanning trees, and plane duality*, SIAM J. Discrete Math. **29** (2015), no. 1, 461–471.
25. S. H. Chan, *Abelian sandpile model and Biggs-Merino for directed graphs*, [arXiv:1412.4837](https://arxiv.org/abs/1412.4837), 2015.
26. R. Cori and Y. Le Borgne, *The sand-pile model and Tutte polynomials*, Adv. in Appl. Math. **30** (2003), no. 1-2, 44–52, Formal power series and algebraic combinatorics (Scottsdale, AZ, 2001).
27. ———, *On computation of Baker and Norine’s rank on complete graphs*, Electron. J. Combin. **23** (2016), no. 1, Paper 1.31, 47.
28. R. Cori and D. Rossin, *On the sandpile group of dual graphs*, European J. Combin. **21** (2000), no. 4, 447–459.
29. Robert Cori, Dominique Rossin, and Bruno Salvy, *Polynomial ideals for sandpiles and their Gröbner bases*, Theoret. Comput. Sci. **276** (2002), no. 1-2, 1–15.
30. S. Corry, *Genus bounds for harmonic group actions on finite graphs*, Int. Math. Res. Notices **2011** (2011), no. 19, 4515–4533.
31. ———, *Harmonic Galois theory for finite graphs*, Galois-Teichmüller Theory and Arithmetic Geometry (H. Nakamura, F. Pop, L. Schneps, and A. Tamagawa, eds.), vol. 63, Advanced Studies in Pure Mathematics, no. 19, Math. Soc. of Japan, 2012, pp. 121–140.
32. ———, *Maximal harmonic group actions on finite graphs*, Discrete Mathematics **338** (2015), no. 5, 784–792.
33. The Sage Developers, *Sagemath, the Sage Mathematics Software System*, 2018, <http://www.sagemath.org>.
34. D. Dhar, *Self-organized critical state of sandpile automaton models*, Phys. Rev. Lett. **64** (1990), no. 14, 1613–1616. MR 1044086
35. ———, *The abelian sandpile and related models*, Physica A **263** (1999), no. 4, 4–25.
36. ———, *Studying self-organized criticality with exactly solved models*, [arXiv:cond-mat/9909009](https://arxiv.org/abs/cond-mat/9909009), 1999.

37. ———, *Theoretical studies of self-organized criticality*, Phys. A **369** (2006), no. 1, 29–70.
38. D. Dhar and S. N. Majumdar, *Equivalence between the abelian sandpile model and the $q \rightarrow 0$ limit of the potts model*, Physica A **185** (1992), 129–145.
39. R. Diestel, *Graph Theory*, fourth ed., Graduate Texts in Mathematics, vol. 173, Springer, Heidelberg, 2010. MR 2744811 (2011m:05002)
40. A. M. Duval, C. J. Klivans, and J. L. Martin, *Simplicial matrix-tree theorems*, Trans. Amer. Math. Soc. **361** (2009), no. 11, 6073–6114.
41. ———, *Critical groups of simplicial complexes*, Ann. Comb. **17** (2013), no. 1, 53–70.
42. ———, *Simplicial and cellular trees*, vol. Recent Trends in Combinatorics, IMA Volumes in Mathematics and its Applications, no. 159, pp. 713–752, Springer, 2016.
43. J. Ellenberg, *What is the sandpile torsor?*, <https://mathoverflow.net/questions/83552>, 2011.
44. C. Fish, A. Grant, et al., *Web sandpiles*, http://www.reed.edu/~davidp/web_sandpiles, 2017.
45. P. Flajolet, P. J. Grabner, P. Kirschenhofer, and H. Prodinger, *On Ramanujan’s Q -function*, J. Comput. Appl. Math. **58** (1995), no. 1, 103–116.
46. A. Gabrielov, *Asymmetric abelian avalanches and sandpile*, preprint 93-65, MSI, Cornell University, 1993.
47. I. M. Gessel, *Enumerative applications of a decomposition for graphs and digraphs*, Discrete Mathematics **139** (1995), no. 1-3, 257–271.
48. E. Goles and M. Margenstern, *Universality of the chip-firing game*, Theoret. Comput. Sci. **172** (1997), no. 1-2, 121–134.
49. E. Goles and E. Prisner, *Source reversal and chip firing on graphs*, Theoret. Comput. Sci. **233** (2000), no. 1-2, 287–295.
50. C. Greene and T. Zaslavsky, *On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs*, Trans. Amer. Math. Soc. **280** (1983), no. 1, 97–126.
51. B. Gutenberg and C. F. Richter, *Frequency of earthquakes in California*, Bull. Seismol. Soc. Amer. **34** (1944), 185–188.
52. J. Guzmàn and C. Klivans, *Chip-firing and energy minimization on m -matrices*, [arXiv:1403.1635](https://arxiv.org/abs/1403.1635), 2014.
53. ———, *Chip-firing on general invertible matrices*, [arXiv:1508.04262](https://arxiv.org/abs/1508.04262), 2015.
54. Kestent H., *Renewal theory for functionals of a Markov chain with general state space*, Ann. Probability **2** (1974), 355–386.
55. B. Head, *Sandpile App*, <http://www.reed.edu/~davidp/sand/program/program.html>, 2009.
56. A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson, *Chip-firing and rotor-routing on directed graphs*, In and Out of Equilibrium II (V. Sidovicius and M. E. Vares, eds.), Progress in Probability, vol. 60, Birkhauser, 2008, pp. 331–364.
57. S. Hopkins and D. Perkinson, *Biographical arrangements*, Trans. Amer. Math. Soc. **368** (2016), no. 1, 709–725.
58. T. Hoppenfeld, *The Structure of the Jacobian Group of a Graph*, <http://www.reed.edu/~davidp/homepage/students/hoppenfeld.pdf>, 2014.

59. S. Janson, D. E. Knuth, T. Łuczak, and B. Pittel, *The birth of the giant component*, Random Structures Algorithms **4** (1993), no. 3, 231–358, With an introduction by the editors.
60. G. Kalai, *Enumeration of \mathbf{Q} -acyclic simplicial complexes.*, Israel J. Math. **45** (1983), no. 4, 337–351.
61. N. Kalinin and M. Shkolnikov, *Tropical curves in sandpiles*, C. R. Math. Acad. Sci. Paris **354** (2016), no. 2, 125–130.
62. V. Kiss and L. Tóthmérész, *Chip-firing games on Eulerian digraphs and NP-hardness of computing the rank of a divisor on a graph*, [arXiv:1407.6958](https://arxiv.org/abs/1407.6958), 2014.
63. D. E. Knuth, *The art of computer programming. Vol. 1*, Addison-Wesley, Reading, MA, 1997, Fundamental algorithms, Third edition [of MR0286317].
64. G. Kreweras, *Une famille de polynômes ayant plusieurs propriétés énumératives*, Period. Math. Hungar. **11** (1980), no. 4, 309–320.
65. D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009, With a chapter by James G. Propp and David B. Wilson.
66. L. Levine, *Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle*, Comm. Math. Phys. **335** (2015), no. 2, 1003–1017.
67. L. Levine, W. Pegden, and C. K. Smart, *Apollonian structure in the Abelian sandpile*, Geom. Funct. Anal. **26** (2016), no. 1, 306–336.
68. ———, *The Apollonian structure of integer superharmonic matrices*, Ann. of Math. (2) **186** (2017), no. 1, 1–67.
69. D. J. Lorenzini, *Arithmetical graphs.*, Math. Ann. **285** (1989), no. 3, 481–501.
70. ———, *A finite group attached to the Laplacian of a graph*, Discrete Math. **91** (1991), no. 3, 277–282.
71. ———, *Arithmetical properties of Laplacians of graphs*, Linear and Multilinear Algebra **47** (2000), no. 4, 281–306.
72. ———, *Smith normal form and Laplacians*, J. Combin. Theory Ser. B **98** (2008), no. 6, 1271–1300.
73. C. Merino López, *Chip firing and the Tutte polynomial*, Ann. Comb. **1** (1997), no. 3, 253–259. MR 1630779 (99k:90232)
74. James Oxley, *Matroid Theory*, Oxford Graduate Texts in Mathematics, no. 3, Oxford UP, 2006.
75. I. Pak, *Computation of Tutte polynomials of complete graphs*, http://www.math.ucla.edu/~pak/papers/Pak_Computation_Tutte_polynomial_complete_graphs.pdf, 1993.
76. W. Pegden and C. K. Smart, *Convergence of the Abelian sandpile*, Duke Math. J. **162** (2013), no. 4, 627–642.
77. D. Perkinson, *Sage thematic tutorial: The abelian sandpile model*, 2018, http://doc.sagemath.org/html/en/thematic_tutorials/sandpile.html.
78. D. Perkinson, Q. Yang, and K. Yu, *G-parking functions and tree inversions*, Combinatorica **37** (2017), no. 2, 269–282.
79. K. Perrot and T. V. Pham, *Chip-firing game and a partial Tutte polynomial for Eulerian digraphs*, Electron. J. Combin. **23** (2016), no. 1, Paper 1.57, 34.
80. R. J. Plemmons, *M-matrix characterizations. I. Nonsingular M-matrices*, Linear Algebra and Appl. **18** (1977), no. 2, 175–188.

81. V. Reiner, *Lectures on matroids and oriented matroids*, <http://www-users.math.umn.edu/~reiner/Talks/Vienna05/index.html>, 2005.
82. S. Ross, *A first course in probability*, second ed., Macmillan Co., New York; Collier Macmillan Ltd., London, 1984.
83. J. P. Sethna, *Statistical Mechanics: Entropy, Order Parameters and Complexity*, Oxford Master Series in Physics, Oxford UP, 2006.
84. E. R. Speer, *Asymmetric abelian sandpile models*, J. Statist. Phys. **71** (1993), no. 1-2, 61–74.
85. R. P. Stanley, *Cohen-Macaulay complexes*, Higher combinatorics (Proc. NATO Advanced Study Inst., Berlin, 1976), Reidel, Dordrecht, 1977, pp. 51–62. NATO Adv. Study Inst. Ser., Ser. C: Math. and Phys. Sci., 31.
86. ———, *Enumerative combinatorics. Vol. 2*, Cambridge Studies in Advanced Mathematics, vol. 62, Cambridge University Press, Cambridge, 1999, With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.
87. ———, *Algebraic combinatorics*, Undergraduate Texts in Mathematics, Springer, New York, 2013, Walks, trees, tableaux, and more.
88. H. N. V. Temperley, *On the mutual cancellation of cluster integrals in Mayer's fugacity series*, Proc. Phys. Soc. **83** (1964), 3–16.
89. W. T. Tutte, *The dissection of equilateral triangles into equilateral triangles*, Proc. Cambridge Philos. Soc. **44** (1948), 463–482.
90. D. Wagner, *The critical group of a directed graph*, [arXiv:math/0010241](https://arxiv.org/abs/math/0010241), 2000.
91. M. E. Walter, *Earthquakes and weatherquakes: mathematics and climate change*, Notices of the AMS **57** (2010), no. 10, 1278–1284.
92. N. W. Watkins, G. Pruessner, S. C. Chapman, N. B. Crosby, and H. J. Jensen, *25 years of self-organized criticality: concepts and controversies*, Space Sci. Rev. **198** (2016), 3–44.
93. H. S. Wilf, *generatingfunctionology*, third ed., A K Peters, Ltd., Wellesley, MA, 2006.
94. C. H. Yuen, *Geometric bijections between spanning trees and break divisors*, J. Combin. Theory Ser. A **152** (2017), 159–189.
95. D. Zeilberger, *A combinatorial approach to matrix algebra*, Discrete Math. **56** (1985), no. 1, 61–72.

Index

- 2-isomorphism, 262
 - Whitney's theorem, 263
- Abel-Jacobi map, 54–58
 - definition, 56
 - injectivity and edge-connectivity, 57
 - is universal for harmonic functions, 57
 - surjectivity and genus, 66
- abelian property
 - for dollar game, 10
 - for sandpiles, 98–99
- abelian sandpile, *see* Markov chain,
 - abelian sandpile model
- accessible vertex, 97

- banana graph, 304
- basic alive divisor, 144
- burning sandpile, 120–123
 - and test for recurrent sandpiles, 120
 - definition, 120
 - existence and uniqueness, 121
- burning script, 120
 - algorithm, 123
 - for undirected graphs, 123
- burst size, 146
 - upper bound for, 160

- canonical divisor, 71
- Cayley's formula, 172, 191, 218
 - higher-dimensional, 292
- Chinese remainder theorem, 32
- Clifford's theorem, 73
- complete graph, 304
 - and parking functions, 217

- Cori-Le Borne algorithm for
 - computing rank of a divisor on, 218
 - divisors on, 215–226
 - Jacobian/sandpile group for, 40
 - size of sandpile group of, 218
 - stationary density of, 275
 - Tutte polynomial for, 266
- complete linear system
 - and convex polytopes, 28–31
 - and edge-connectivity, 49
 - definition, 15
 - harmonic morphisms and trees, 208
 - of divisors on C_n , 41
- configuration, 24–28, 97
 - definition, 25, 97
 - degree of, 25, 97
 - relation to Jacobian group, 26
 - script-firing, 25
 - stabilization of, 98, 114
 - existence, 101
 - uniqueness, 100
 - superstabilization of, 50
 - superstable, 50–54
 - and acyclic orientations, 64
 - definition, 50
 - maximal, 61, 65
 - relation to q -reduced divisors, 50
 - vertex and set-firing, 25
 - legal, 50
- Cori-Le Borgne algorithm
 - for computing ranks of divisors
 - on K_n , 218
- cuts, *see* cycles and cuts

- cycle graph, C_n , 304
- cycle rank, cyclomatic number, 65
- cycles and cuts, 245–254
 - bond, 247
 - cut, directed, 246
 - cycle and cut spaces, 246, 247
 - bases and dimensions, 248
 - duality for plane graphs, 251
 - orthogonality, 248
 - relation to sandpile group, 249
 - vertex cut, 247
- degree
 - of configuration, 25, 97
 - of divisor, 10
 - of harmonic morphism, 200
 - of sandpile, 97
- deletion-contraction
 - and the Tutte polynomial, 259
 - for matroids, 259
 - proof of matrix-tree theorem, 170
 - universal property, 261
- determinantal divisor, 35
- Dhar's algorithm, 50–54
 - and acyclic orientations, 63–67
 - and tree bijections, 175–190
 - efficient implementation of, 54
 - for recurrences, 120, 127
 - pseudocode for superstables, 51
- divisor, 9
 - alive, 82–84, 141–142
 - avalanche, 141
 - basic alive divisor, 144
 - basic decomposition, 145
 - canonical, 71
 - class of, 12
 - complete linear system of, 15
 - corresponding to an orientation, 63
 - degree of, 10
 - effective, 12
 - and edge-connectivity, 49
 - legal firing/toppling of vertex, 141
 - lending or borrowing move from, 10
 - linear equivalence of, 11
 - maximal stable, 83
 - principal, 20
 - q -reduced, 46–49
 - algorithm for, 46, 54
 - and edge-connectivity, 49
 - existence and uniqueness of, 48
 - relation to dollar game, 48
 - relation to superstables, 50
- rank, 69–71
 - and degree, 70, 71, 73
 - and edge-connectivity, 70, 208
 - definition, 69
 - for divisors on complete graphs, 218
 - formula for, 86
 - lower bound for, 70
 - NP-hard to compute, 70
- script-firing, 20
- stability, 83
- stable/stabilizable, 141
- vertex and set-firing, 10
 - legal, 47
- Weierstrass points of, 87
- winnable, unwinnable, 13
 - maximal unwinnable, 61, 65
- dollar game
 - decided by q -reducing, 48
 - first example, 7–9
 - greedy algorithm for, 43–45
 - higher-dimensional, 295
 - household-solutions for, 206
 - restatement of, 13, 15, 22, 54–56
- duality between recurrences and superstables, 124, 240
- Dyck path, 220
- effective divisor, *see* divisor, effective
- elementary divisors, 32
- Eulerian digraph, 309
 - equivalent condition for, 309
- exact sequence, 314
 - short, 314
 - split exact, 315
- externally active edge, 180, 266
- field, 312
- finitely generated abelian group, 32–33
 - definition, 32
 - elementary divisors, invariant factors, primary factors of, 32
 - rank of, 33
 - structure theorem for, 32
 - torsion element of, 33
- firing script, 20, 98
 - divisor of, 20
 - uniqueness, 103
- forbidden subconfiguration, 125–127
- forest, 305
 - spanning, 306

- genus, 65
- globally accessible vertex, 97
- graph morphism, 195, 196
 - push-forward mapping for divisors, 196
- graph, directed, 307–309
 - connected
 - strongly, 308
 - weakly, 308
 - cycle in, 308
 - directed spanning tree for, 308
 - Eulerian cycle of, 309
 - path in, 308
 - subdigraph, 307
 - induced, 307
 - underlying simple undirected graph, 307
 - underlying undirected graph, 307
 - see also* graph, undirected, *see also* tree
- graph, undirected, 303–307
 - acyclic, 305, 308
 - associated digraph, 309
 - connected, 303
 - connected components, 303
 - cycle in, 305
 - cycle matroid for, 258
 - degree of vertex, 303
 - edge-connectivity, 306
 - hyperelliptic, 212
 - loop, 303
 - morphism, 195
 - path in, 303
 - plane duality, 250
 - quotient, 212
 - simple, 303
 - spanning forest of, 306
 - spanning tree of, 306
 - subgraph, 303
 - induced, 303
 - see also* graph, directed, *see also* tree
- greedy algorithm
 - and injectivity of pull-back, 202
 - and matroids, 276
 - and the dollar game, 43
 - and uniqueness of stabilization, 84, 114
 - better than benevolence, 59
 - for burning script, 122
 - uniqueness of script for, 45
- group, 311
 - action, 312
 - homomorphism, 311
- h -vector conjecture, 269–273
 - proof for cographic matroids, 273
- harmonic function, 24, 56
 - universal property, 57
- harmonic morphism, 195–211
 - and hyperelliptic graphs, 212
 - and quotient graphs, 212
 - compositions of, 212
 - definition, 196
 - degenerate, 200
 - degree, 200
 - horizontal multiplicity of, 196
 - pull-back mapping for divisors, 201
 - injectivity of, 202
 - preserves linear equivalence, 201
 - push-forward mapping for divisors, 198
 - comparing ranks, 207
 - preserves linear equivalence, 196
 - surjectivity of, 200
 - surjectivity of, 200
 - trees and complete linear systems, 208
 - vertical multiplicity of, 201
- indegree, 96
 - of a vertex, 307
 - with respect to subset, 307
- internally active edge, 266
- invariant factors
 - of a f.g abelian group, 32
 - of a matrix, 35
- Jacobian group, 14–15
 - definition, 14
 - of complete graph, 40
 - of cycle graph, 40
 - of one-point join, 40
 - relation to sandpile group, 230
 - size of, 38
 - structure of, 31
- Kreweras' formula, 184
- Kruskal's algorithm, 276
- Laplacian, 19–39
 - and vertex firing, 103
 - definition, 19, 102
 - higher-dimensional, 287
 - kernel of, 23, 174

- Laplacian matrix, 21
- lattice, 103
- reduced Laplacian, 26–28, 101–103
 - cokernel is isomorphic to sandpile group, 105
 - definition, 26, 101
 - higher-dimensional, 291
 - kernel of, 27
 - relation to the Jacobian group, 26, 38
 - relation to continuous Laplacian, 23
 - relation to the Picard group, 22
- least action principle, 99–100
 - for configurations, 99
 - for divisors, 142
 - for M -matrices, 239
- legal firing sequence, 98
- lending and borrowing moves, 10, 25
- linear equivalence of divisors, 11–14
 - definition, 11
- local confluence property, 99
- M -matrix, 238
 - equivalent conditions for, 239
 - recurrents and superstables for, 239
 - duality, 240
 - existence and uniqueness, 240
- Markov chain, 131–159
 - abelian sandpile model, 133
 - recurrent states are recurrent sandpiles, 136
 - stationary density of, 147
 - stationary distribution is uniform, 140
 - toppling probabilities encoded in the Laplacian, 140
 - accessible state, 134
 - aperiodic, 139
 - associated directed graph, 132
 - definition, 132
 - essential/inessential state, 135
 - fixed-energy sandpile, 141–151
 - comparison with abelian sandpile, 144
 - definition, 142
 - epicenter, 144
 - threshold, 143
 - threshold density, 144
 - threshold state, 144
 - irreducible, 134
 - random walk on a group, 137
 - recurrent state, 135
 - renewal theorem, 155
 - stationary distribution, 136
 - existence and uniqueness, 139
 - transient state, 135
 - transition matrix for, 132
- matrix-tree theorem, 166
 - edge-weighted version, 172
 - higher-dimensional, 291
 - proof by deletion-contraction, 170
 - proof by sign-reversing transposition, 166
- matroid, 257–273
 - graphic/cycle matroid, 258
 - and greedy algorithms, 276
 - basis for, 258
 - bridge, 259
 - cographic, 273
 - definition, 257
 - deletion-contraction, 259
 - determines sandpile group, 262
 - dual, 258
 - independent sets of, 258
 - linear, 258
 - loop, 259
 - rank of, 258
 - uniform, 258
 - universal property, 261
- matroid complex, 273
- Merino’s theorem, 263–266
 - and Gessel’s formula, 268
 - for complete graphs, 266
 - implies h -vector conjecture for cographic matroids, 273
- module, 312
 - direct products and sums, 313
 - homomorphism, 313
 - quotient by submodule, 313
 - snake lemma, 315
- modules
 - exact sequences of, 314
 - finitely generated, 314
 - free, 313
 - Noetherian, 314
- monoid, 311
- multicomplex, 270
- O -sequence, 271
- odometer function, 146
- orientation, 62, 307
 - acyclic, 62
 - and maximal superstables, 64
 - corresponding configuration, 63

- corresponding divisor, 63
 - indegree sequence for, 62
 - reversed, 71
 - sink or source vertex for, 62
- outdegree, 96, 307
 - with respect to subset, 307
- parking functions, 215–218
 - and superstables on K_n , 217
 - characterization of, 216
 - count, 217, 227
 - definition, 215
- path graph, P_n , 304
- Picard group, 14–15
 - definition, 14
 - structure of, 31
- plane duality, 250–254
 - and the sandpile group, 252
- poset (partially ordered set), 270
 - order ideal in, 270
- q -reduced, *see* divisor, q -reduced
- quotient graph, 212
- Ramanujan Q -function, 275
- rank, *see* divisor, rank
- reduced Laplacian, *see* Laplacian, reduced Laplacian
- Riemann surfaces, 74–82
 - branched coverings of, 204–206
 - divisors and mappings to projective space, 81
 - Riemann-Hurwitz formula, 205
 - Riemann-Roch theorem, 75
- Riemann-Hurwitz formula
 - and quotient graphs, 212
 - for graphs, 203
 - for Riemann surfaces, 205
- Riemann-Roch theorem for graphs, 69–85
 - statement and proof, 71–72
- ring, 312
 - ideal in, 312
 - quotient by ideal, 312
- rotor-router model, 186–188
 - and a free transitive action of $\mathcal{S}(G)$ on spanning trees, 188
- sandpile
 - and tropical curves, 109
 - building a Turing machine, 115
 - burst size for recurrent, 146
 - definition, 97
 - degree of, 97
 - duality between recurrents and superstables, 124
 - images of, 108–109
 - maximal stable, 104
 - monoid, 105, 115
 - recurrent, 103–108
 - accessible from maximal stable sandpile, 104
 - definition, 104
 - test for, 120, 126, 127
 - scaling limit, 109
 - stabilization of, 98, 114
 - existence, 101
 - uniqueness, 100
 - stable addition, 105
 - superstable
 - definition, 124
- sandpile graph, 97
- sandpile group
 - and plane duality, 252
 - as a principal ideal, 115
 - bijections with spanning trees, 174–190
 - calculating the identity, 107
 - cyclic, 236
 - definition, 107
 - dependence on sink, 229–232
 - effect of removing a bridge or loop, 129
 - effect of removing an edge, 235
 - introduction to, 91–95
 - isomorphic to Jacobian iff graph is Eulerian, 232
 - isomorphic to $\text{cok}(\tilde{L})$, 105
 - minimal number of generators for, 232–238
 - number of superstables of each degree, 148, 265
 - of complete graph, 40
 - of cycle graph, 40
 - of one-point join, 40
 - orders of elements, 114
 - relation to cycles and cuts, 249, 255
 - relation to Jacobian group, 230
 - size equals number of rooted spanning trees, 172
 - size is $\det(\tilde{L})$, 107
- script-firing
 - for configurations, 98, 123
 - for divisors, 20

- legal, 124
- self-organized criticality (SOC), 240–242
- selfish vertex, 120
- set-firing
 - for configurations, 123
 - for divisors, 10, 47
 - legal, 47, 50, 124
- simplicial complex, 271
 - f -vector, 272
 - h -vector, 272
 - k -skeleton of, 287
 - acyclic in positive codimension (APC), 291
 - Betti numbers of, 282
 - boundary mapping, 245, 280
 - chain complex of, 280
 - chains, group of, 280
 - critical groups of, 287–295
 - definition, 287
 - relation to reduced Laplacian, 291
 - structure as an abelian group, 288
 - cycles and boundaries of, 282
 - dimension, 272
 - Euler characteristic, 272, 299
 - examples
 - n -simplex, 297
 - equatorial bipyramid, 289
 - Klein bottle, 298
 - real projective plane, 286
 - sphere, 297
 - torus, 298
 - face, 271
 - facet, 272
 - firing faces of, 292–295
 - flow game for, 295
 - homology groups of, 282
 - how to compute, 285–286
 - Laplacians for, 287
 - matrix-tree theorem for, 291
 - matroid complex, 273
 - pure, 272
 - spanning trees of, 288–292
 - and acyclicity, 290, 291
- sink, 97
 - in digraph, 307
- Smith normal form, 33–36
 - definition, 35
 - existence and uniqueness of, 35
- snake lemma, 315
- source in digraph, 307
- spanning tree
 - and depth-first search, 182
 - and eigenvalues of the Laplacian, 173
 - and kernel of the Laplacian, 174
 - another counting formula for, 173
 - bijections with $\mathcal{S}(G)$, 174–190
 - count for complete bipartite graph, 192
 - Dhar’s tree bijection algorithm, 175
 - and external activity, 179–182
 - and tree inversions, 182–186
 - directed, rooted, 166
 - inversion number for, 184
 - κ -inversion number, 185
 - matrix-tree theorem, 166
 - of a simplicial complex, 288
 - random sampling, 188–190
 - see also* tree
 - stable addition, 105
 - stable/unstable vertex, 98
 - stationary density, 147
 - and unicycles, 274
 - of a tree, 160
 - of banana graph, 149, 161
 - of complete graph, 275
 - of cycle graph, 149
 - superstable configuration, *see* configuration, superstable
- threshold density, 144
 - of a tree, 160
 - of banana graph, 149, 161
 - of cycle graph, 149
- threshold density theorem, 151–159
 - statement and proof, 152
- torsor, 55, 188
- tree, 305
 - equivalent conditions for, 305
 - internal and external activity, 266
 - of a simplicial complex, 292
 - rooted spanning, 308
 - spanning, 306
 - stationary and threshold density, 160
 - see also* spanning tree
- tree ordering, 47
- tropical curves, 109
- Turing machine, 115
- Tutte polynomial, 257–273
 - and internal/external activity, 267
 - definition, 259
 - determined by deletion-contraction, 259

-
- for the complete graph, [268](#)
 - special values of, [261](#)
 - unicycle, [274](#)
 - vertex firing, [10](#), [98](#)
 - legal, [98](#)
 - reverse-firing, [98](#)
 - Weierstrass points, [87](#)
 - Whitney twist, [262](#)
 - and 2-isomorphisms, [263](#)
 - winnable, unwinnable divisor, [13](#)
 - see also* dollar game
 - Zelda, [7](#)