Riemann Roch on Directed Graphs

A Thesis
Presented to
The Division of Mathematics and Natural Sciences
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Jeffrey Samuel Sebastian Weinmann Cruse

May 2008

Approved for the Division
(Mathematics)

David Perkinson

Acknowledgements

There are rather a lot of people who deserve thanks here.

First to David Perkinson, for turning me onto this subject in the first place, and for putting up with me (and sometimes for not putting up with me) as I put it all together. For believing that I would figure it out.

Also to Tom Weiting, who is almost singlehandedly responsible for the style of my adult understanding of mathematics, and incidentally physics, and perhaps the whole of human knowledge worthy of the name. Also to Rao Potluri, to Jamie Pommersheim, and to Ray Mayer, for most of the rest.

Next to my folks, whose advocacy and indignation on my behalf is probably more to blame for the fact that I am where I am, and that I do what I do, than the extent to which I'm good at it and enjoy it. Who never gave up, who kept sending me back.

And finally, to everyone else, to Evan Pierce and TC Proctor and Steve McClellan and Strom Borman and Josh Fieldstone and everyone else I've ever really collaborated with, everyone else who's followed along enough to laugh at the punchline at the end of my fifteen minutes of mathematical introduction (and everyone who's just smiled and nodded), everyone with whom I started this work of mine. Y'all've done more than y'all'd think.

Preface

"It is well known that a finite graph can be viewed, in many respects, as a discrete analogue of a Riemann surface." Matthew Baker and Sergei Norine [1]

Table of Contents

Abstract

Baker and Norine have recently published a number of papers proving results on weighted undirected graphs that were originally proven for Riemann surfaces. Intriguingly, this included a powerful theorem called Riemann Roch, which relates the linear systems of certain pairs of divisors on the surface, and whose undirected graph analogue built on the sandpile group. In this paper, we develop a few of the properties of the sandpile group on weighted directed graphs and the machinery of Riemann Roch, in pursuit of a proof of such in the directed case.

Baker and Norine are after a sort of category-theoretical correspondence between Riemann surfaces and weighted undirected graphs: it was our strong suspicion that such a correspondence can equally well be proven with directed graphs, and such a suspicion motivated our work here.

In the process of such we found a counterexample; we show that no such analogue of Riemann Roch exists on directed graphs.

Introduction

There has been a good deal of recent interest in the sandpile groups of finite directed and undirected graphs, as developed by a number of authors, and in particular in the preparation of a sort of correspondence between finite undirected graphs and Riemann surfaces, as developed by Baker and Norine. They have proven that a number of results—theorems on mappings, structural theorems on certain constructions, and so on—from the ancient lair of the Riemann surface could be applied, with most if not all wrinkles intact, in a new and unspoilt cavern of the finite undirected graph. And their proof of Riemann Roch, perhaps the most powerful such theorem, flows right through sandpile territory.

Well, that's all well and good. But sandpiles can—indeed, most naturally are—spoken about not in terms of finite undirected graphs, but finite directed graphs! Thus their tools, in this if not every case, are naturally fitted to a rather different and more dangerous manner of beast than that to which they have directed themselves thus far. And bolstered by their previous triumphs, fresh from the ruin of that smoten foe, their ultimate quarry became a general correspondence, of a sort so new that none of us can even write it down.

So one must wonder. Can what they're claiming for the finite undirected graph be won also for the finite directed graph? Are these structures linked in this same mysterious way? Could this be one step closer to whatever exactly we're after?

Our contention was that it was. This was to be a work in that direction.

But alas, it was not to be. We sharpened our swords and plotted our advance, sallied forth with heads held high and banners gleaming, but treacherous terrain and tumultuous tempest soon halted our advance, tried our resolve. And as our troops dug in after many months for the hard winter which was to come, we received a foul and damning omen.

The most foul and most damning omen, in fact. A counterexample.

Thus it is that you find me here, head bare and hands empty, with nothing to my name but a thousand thousand steps down a fruitless path, a manner of proof that it cannot be done. This is the story of my defeat, how a brave and mighty host was lost, and we were finished, spent and empty, on the slopes of RR2.

Chapter 1

Sandpiles and Sandcastles

1.1 Graphs

We spend rather a lot of time in this paper speaking about graphs, so we'll begin by specifying what exactly we mean by that. A graph is the simplest mathematical structure that speaks about a set of objects and relationships between them: the first thing we require to construct the graph is the set of objects in question, the vertices, $V(G)$. These can, of course, be anything we like—elements of a field, points in a topological space, vectors or whatnot—but for the purposes of graph theory we ignore any other structure they may have and consider $V(G)$ to be merely a set of objects.

Further, since we only have interesting things to say about finite graphs, we require that $V(G)$ be a finite set.

Now we need to speak about the connections between these objects. There are a number of ways that we could proceed, but we begin with the simplest: binary relationships. Either the two vertices are related in this way, or they're not. Thus we construct the set of edges on the graph, pairs of vertices, $E(G) \subset \{\{v, w\} : v \in V(G), v \neq w \in V(G)\}$. (To make things cleaner, we disqualify graphs that aren't connected; who have subsets of vertices that don't have any paths connecting them at all. We don't really lose anything by doing so—basically, if they weren't connected, they weren't a single graph to begin with.) And this is enough, so we say that a graph is the ordered pair of these two sets, the set of vertices and the set of edges. $G = (V(G), E(G))$. This is the basic structure, the finite simple undirected graph.

This is a rich abstraction, and a lot of powerful and interesting things can be said about it, but it's actually a bit too simple for our purposes. Binary relationships aren't quite general enough: we want our graph to contain some information about the strength of the relationship too. Not just whether or not the kindergartners under our charge fought, but how often they did so. Not just whether or not there are rivers flowing from one country to another, but how many. So we generalize a bit by introducing a weight function: $W(G) \colon E(G) \to \mathbb{Z}^+$, assigning a positive integer to each edge, and say that the new graph is the ordered triple of these three elements: $G := (V(G), E(G), W(G))$. This is our larger structure, the finite weighted graph.

Sometimes one is interested in weighing functions that don't map to the integers, but for the moment, we're not. Their sandcastles tend to be rather boring.

This is the structure that Baker and Norine consider in their paper on the subject, the structure that they are attempting to flesh out in a manner similar to Riemann surfaces, in pursuit of some sort of general correspondence between the structures. But it's still not quite general enough for our tastes. It often makes sense to consider not just binary relationships, but ordered binary relationships. Not just whether or not your kindergartners fought, but who started it; whether the river runs from Germany to France, or the other way around. So instead of considering pairs of vertices to be our edges, we consider ordered pairs: $E(G) \subset \{(v \to w) : v \in V(G), v \neq w \in V(G)\}$, and let the graph be the ordered triple of these new elements, $G := (V(G), E(G), W(G))$. This is the structure we're talking about, the finite weighted directed graph. Our goal is to show that all the results that Baker and Norine have proven about weighted graphs hold for weighted directed graphs too, and the correspondence they're after works on more objects than they at first considered.

Note that there's a natural injection from the undirected weighted graphs into the directed weighted graphs—for every edge $\{v, w\} \in E(G)$ in the undirected graph, put $(v \to w), (w \to v) \in E(G)$ in its image under the injection, and let $W(G)((v \to w)) = W(G)((w \to v)) = W(G)(\{v, w\})$. Most of our analysis takes place on directed graphs; when we refer to undirected graphs, we'll usually be referring to the properties under our analysis of the images of these graphs; the alternative is running two sets of analysis in parallel, which would be a lot of work.

But in order to say anything interesting on these directed graphs, we need a bit more structure. Let a path from one vertex v to another w be a nonempty subset of $E(G)$ such that the edges link up end to end, from v to w. Further, let a cycle be a path from a vertex back to itself.

Definition 1.1.1. *Let k be an integer greater than 0, and let u_0 and u_k be, not necessarily distinct, elements of $V(G)$ for a weighted undirected graph G. Let $\{u_1, \dots, u_{k-1}\}$ be an ordered subset of $V(G)$. If $\{u_0, \dots, u_k\}$ is such that $u_i \neq u_j$ for any i, j not $0, k$, and $\{u_i, u_{i+1}\} \neq \{u_j, u_{j+1}\}$ for any i, j, then we say that the ordered set $p = \{\{u_0, u_1\}, \dots, \{u_{k+1}, u_k\}\} \subseteq E(G)$ is a path of length k from u_0 to u_k.*

Similarly, let k be an integer greater than 0, and let u_0 and u_k be not necessarily distinct elements of $V(G)$ for a weighted directed graph G, and $\{u_1, \dots, u_{k-1}\}$ is an ordered subset of $V(G)$. If $\{u_0, \dots, u_k\}$ is such that $u_i \neq u_j$ for any i, j not $0, k$, and $(u_i \to u_{i+1}) \neq (u_j \to u_{j+1})$ for any i, j, then we say that the ordered set $p = \{(u_0 \to u_1), \dots, (u_{k+1} \to u_k)\} \subseteq E(G)$ is a path of length k from u_0 to u_k.

For us to be able to define a sandcastle on a graph, we need there to be at least one vertex such that there exists a path from every other vertex to that first vertex. This is automatic in the undirected case: since any edge goes in both directions, the only way the existence of a path can fail is if the graph isn't connected. But if your edges are directed, connected graphs can lose this property in all sorts of ways. So we need to explicitly require that our graphs don't.

We'll want to be able to speak a bit more formally about this. A subgraph of a graph G is a graph G' such that $V(G') \subseteq V(G)$, $E(G') \subseteq E(G)$, and $W(G')(e) \leq$

$W(G)(e)$ for all $e \in E(G')$; that is, a graph with vertices, edges, and weights taken from the parent graph, and no extras. A spanning tree of a weighted undirected graph G is a subgraph such that there is one path between any two points, and no extra edges; a watershed draining into a vertex v on an undirected graph G is a subgraph such that there exists one path from each vertex into the vertex v, and no extra edges.

Definition 1.1.2. *A spanning tree of a weighted undirected graph is a subgraph G' such that $V(G') = V(G)$, $W(G')(e) = W(G)(e)$ for all $e \in E(G')$, and $E(G')$ is such that there exists a path between any two distinct vertices, but no paths from a vertex to itself. Further, let the weight of the spanning tree be*

$$\prod_{e \in E(G')} W(G')(e).$$

Further, let a watershed draining into v be a subgraph G' of a subgraph G for a $v \in V(G)$, such that $V(G') = V(G)$, $W(G')(e) = W(G)(e)$ for all $e \in E(G')$, and $E(G')$ is such that there exists a unique path from each $w \in V(G) \setminus v$ to v, but no paths from a vertex to itself. Further, let the weight of the watershed be

$$\prod_{e \in E(G')} W(G')(e).$$

Note that spanning trees always exist on connected graphs, but watersheds don't necessarily; for reasons that will become clear, we will require that our directed graphs contain a watershed into at least one vertex. Note also that a spanning tree is symmetric among all vertices on the graph—it doesn't preference a specific sink—but a watershed necessarily does, and this breaks the symmetry of our analysis all the way down the line.

So these are the important structures: connected undirected weighted graphs, on which Riemann Roch has already been proven by Baker and Norine, and connected directed weighted graphs containing at least one watershed, which we are concerned with now.

1.2 Sand

The sandpile group is usually framed in terms of quantities of sand (and antisand) sitting on the vertices of the graph, which can be shuffled around in a specific fashion according to the structure of the graph itself. The basic unit of such shuffling is the Vertex Firing—one selects a vertex, and for each edge coming out of the vertex, one moves a number of grains of sand equal to the weight of the edge to the vertex at the other end.

There's a natural group structure to these distributions of sand—we just add them componentwise, combining the sand at each vertex. We focus for the moment on the subgroup such that the grains of sand on all vertices sum to 0.

Definition 1.2.1. *Let* $\mathrm{Con}(G)$ *be the set of functions $D \colon V(G) \to \mathbb{Z}^n$, and let* $\mathrm{Con}_0(G)$ *be the subset such that $\sum_{v \in V(G)} D(v) = 0$.*

Note that this is a group under componentwise addition: $(D + D')(u) := D(u) + D'(u)$, $u \in V(G)$, and $\mathrm{Con}_0(G)$ is a subgroup. To make things easier on ourselves, we index the vertices of G: $V(G) = \{v_1, v_2, \ldots, v_n\}$ This yields an isomorphism $\mathrm{Con}(G) \to \mathbb{Z}^n$, which maps $D(v_i)$ to the ith component of the corresponding element of $\mathbb{Z}^n$—in true mathematical fashion, now that we have the isomorphism, we'll ignore the distinction.

Now, the sandpile group is basically the set of all configurations on the graph modulo vertex firing, so we develop this in a more formal way.

We define the laplacian L of G, which is a matrix that encodes in a natural way what happens to a configuration on the graph when you fire vertices.

Definition 1.2.2. *Let the* laplacian L *of a weighed directed graph* G *be a matrix such that* $L_{ij} = -W(G)((v_j \to v_i))$ *for all* $i \neq j$*, and* $L_{ii} = \sum_j W(G)((v_i \to v_j))$

The laplacian is defined such that firing a vertex v_i changes the configuration by $-L\, e_i$. Note that this means anything in the image has zero total sand lying around, so the sum of the entries in any column of the laplacian is 0, and since this is true of all columns the sum of the rows of the laplacian gives you the zero vector. This is a property we'll make use of later.

We invoke the Matrix Tree Theorem, which states that the determinant of the laplacian of a graph whose ith row and column have been deleted is equal to the sum of the weights of the watersheds—being the products of the weights of all edges in the watershed—draining into v_i. Since we've assumed that our connected weighted directed graph has a watershed draining into at least one vertex, then there exists one row and column we may delete to obtain a matrix of maximal rank, so L has rank $n - 1$. Further, we will show that its kernel is the vector formed from the weights of the watersheds draining into each vertex.

(Note that we're interested in the kernel of $L\colon \mathbb{Z}^n \to \mathbb{Z}^n$ rather than as an honest linear mapping like $\mathbb{R}^n \to \mathbb{R}^n$, which means we have to be more careful when writing it down. For L as a mapping on $\mathbb{R}$, $\ker_{\mathbb{R}} L = a\mathbb{R}$ for any nonzero $a \in \ker_{\mathbb{R}}(L)$; the kernel is insensitive to our choice of a. But when L is a mapping on $\mathbb{Z}$, $\ker_{\mathbb{Z}} = b\mathbb{Z}$ only when $b \neq z \cdot c$ for some $z \in \mathbb{Z}$, $c \in \mathbb{Z}^n$, so we have to be careful to pick the 'smallest.')

Definition 1.2.3. *Let* $\kappa(i)$ *for all* $v_i \in V(G)$ *be the sum of the weights of the watersheds draining into* v_i*, let* $\mu = \gcd_i \kappa(i)$ *and let* $\lambda(i) = \kappa_i/\mu$*.*

Then

Theorem 1.2.1. $\ker_{\mathbb{Z}} L = \lambda\mathbb{Z}$*.*

Proof. First, we show that the cofactors of a matrix in which the rows sum to the zero vector are constant for all cofactors on a given column. Let c_k be the kth row of the laplacian, and let $c_k^{(i)}$ be the kth row of the laplacian with the ith entry deleted.

Consider the (k, i)th cofactor. Since the rows sum to zero, the kth row of L is just minus the sum of the matrix's other rows: $-\sum_{j \neq k} c_j$. So to obtain the $(k + 1, i)$ cofactor submatrix, one switches out the $k + 1$th row of L for the row above, with the ith entry deleted: $-\sum_{j\neq k} c_j^{(i)}$. But by linear algebra, the $-c_j^{(i)}$ terms are invisible to the determinant for all $j \neq k+1$, and $j = k+1$ is minus the row that was just replaced! Thus the determinants of the (k,i) and $(k+1,i)$ cofactor submatrices differ by a factor of -1, and their cofactors are equal.

Now consider the cofactor expansion of the laplacian along the kth row:

$$\begin{aligned} 0 &= \det L \\ &= \sum_i (L)^{ki} L_{ki}. \end{aligned}$$

Since the cofactors are constant on each column,

$$= \sum_i (L)^{ii} L_{ki},$$

and by the Matrix Tree Theorem, $(L)^{ii}$ is the number of watersheds draining into v_i, κ_i, times $(-1)^{i+i} = 1$, so

$$\begin{aligned} &= \sum_i \kappa_i L_{ki} \\ &= (L\kappa)_k, \end{aligned}$$

so $\kappa \in K(L)$. But we want the smallest integer vector in the real subspace: $\kappa / \gcd_i \kappa(i) = \lambda$, and $\ker_{\mathbb{Z}} L = \lambda\mathbb{Z}$. □

1.3 The Sandpile and Other Groups

We now address ourselves to the definitions of, and relationships between, the sandpile group and related groups on a weighted directed graph, containing at least one watershed.

Note! *Thus the graph G now always refers to a* weighted directed graph *containing at least one watershed, and v_i is always a* vertex into which a watershed drains, *and the official sink for the purposes of the sandpile group.*

Now for the purposes of the sandpile group we want to restrict ourselves to a certain subset of vertex firing—that is, firing any vertex but the sink.

Definition 1.3.1. *Let $\alpha(j)$ for a given i be such that $\alpha(j) = j$ for $j < i$, $\alpha(j) = j+1$ for $j \geq i$. Let the* winnowed laplacian *L^i be the matrix of dimension $n \times (n-1)$ such that $L^i_{jk} = L_{j\alpha(k)}$. Let the* winnowed principals *be $\mathrm{Prin}^i(G) = \mathrm{img}_{\mathbb{Z}} L$. And let the* sandpile group *with respect to sink v_i be $\mathcal{S}^i = \mathrm{Con}_0(G)/\mathrm{Prin}^i(G)$.*

There are a number of definitions of the sandpile group in the literature; we'll devote some time to proving that they're equivalent. Holroyd and Levine *et.al.*[2], for example, do basically the same thing by removing the sink entirely.

Definition 1.3.2. *Let the* deleted configurations $\operatorname{Con}^i(G)$ *be the set of functions* $D\colon V(G)\setminus\{v_i\}\to\mathbb{Z}^{n-1}$, *let* $\alpha(j)$ *be as above, and let the* deleted laplacian L^{ii} *be the matrix of dimension* $(n-1)\times(n-1)$ *such that* $L^{ii}_{jk}=L_{\alpha(j)\alpha(k)}$. *Let the* deleted principals $\operatorname{Prin}^{ii}(G)=\operatorname{img}_{\mathbb{Z}} L^{ii}$, *and let the* Holroyd group *be* $H^i=\operatorname{Con}^i(G)/\operatorname{Prin}^{ii}(G)$.

Proof of the following contentions is by a thimbleful of Category Theory—basically, rather than requiring that our mappings be isomorphisms (it's not always possible to say the important things nicely with isomorphisms and quotient structures) we require that they preserve mathematical structure, and then construct sequences such that the image of the one is the kernel of the next: that is, *exact* sequences. That is, we construct sequences of mappings that preserve as much information as possible subject to the constraint that two steps along the sequence kills anything. This turns out to be a powerful idea with a number of theorems attached, and we make use of one of the first—the Snake Lemma—to do all our housecleaning.

The statement of the theorem is as follows: assume a diagram that looks like so:

$$\begin{array}{ccccccccc}
 & & 4 & & 5 & & 6 & & \\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & D & \longrightarrow & E & \longrightarrow & F & \longrightarrow & 0\\
 & & \uparrow & & \uparrow & & \uparrow & & \\
0 & \longrightarrow & A & \longrightarrow & B & \longrightarrow & C & \longrightarrow & 0\\
 & & \uparrow & & \uparrow & & \uparrow & & \\
 & & 1 & & 2 & & 3 & &
\end{array}$$

such that the horizontal and vertical sequences are exact, and the diagram commutes—that is, you may follow whichever path you like from one letter to another, without it affecting the image. Then there exists an exact sequence

$$0\to 1\to 2\to 3\to 4\to 5\to 6\to 0.$$

So, getting started.

Theorem 1.3.1. $\mathcal{S}^i \approx H^i$.

Proof. We consider the diagram

$$\begin{tikzcd}
 & 0 & 0 & \boxed{0} & \\
0 \arrow[r] & \mathbb{Z}^{n-1} \arrow[u] \arrow[r, "L^{ii}"] & \operatorname{Con}^i(G) \arrow[u] \arrow[r, "+\operatorname{Prin}^{ii}(G)"] & H^i \arrow[u] \arrow[r] & 0 \\
0 \arrow[r] & \mathbb{Z}^{n-1} \arrow[u, equal] \arrow[r, "L^{i}"] & \operatorname{Con}_0(G) \arrow[u, "F^i"] \arrow[r, "+\operatorname{Prin}^{i}(G)"] & \mathcal{S}^i \arrow[u] \arrow[r] & 0 \\
 & 0 \arrow[u] & 0 \arrow[u] & \boxed{0} \arrow[u] &
\end{tikzcd}$$

Since a watershed drains into v_i, L^{ii} has maximal rank by the Matrix Tree theorem. Further, since H^i is the group theoretic quotient of $\operatorname{Con}^i(G)$ and the image of $\mathbb{Z}^{n-1}$ under L^{ii}, $\operatorname{Prin}^{ii}(G)$, then the natural map to the quotient structure makes the top row exact, as advertised.

Further, L^i is of maximal rank if L^{ii} is, and $\mathcal{S}^i$ admits the same sort of map and is also the group theoretic quotient of just the right spaces. So the second row is also exact.

Now, the $\mathbb{Z}^{n-1}$s are the same spaces, and behave well mapped to one another. The natural map $F^i\colon \operatorname{Con}_0(G) \to \operatorname{Con}^i(G)$ of forgetting the ith entry commutes with the identification of the $\mathbb{Z}^{n-1}$s. Thus, by squinting at our diagram and invoking the Snake Lemma, we can fill in the boxed zeros, implying that

$$0 \to \mathcal{S}^i \to H^i \to 0$$

is exact—that is, the two are isomorphic, as desired. □

Exploiting this, we can prove that the order of the sandpile group is $\kappa(i)$, the sum of the weights of the watersheds draining into v_i.

Theorem 1.3.2. *$|\mathcal{S}^i| = \kappa(i)$.*

Proof.

$$\begin{aligned}
|\mathcal{S}^i| &= |\operatorname{Con}_0(G)/\operatorname{Prin}^i(G)| \\
&= |\operatorname{Con}^i(G)/\operatorname{Prin}^{ii}(G)| \\
&= \det L^{ii} \\
&= \kappa_i.
\end{aligned}$$

□

While we're at it, Baker and Norine [1] prove their results about the Picard Group: our sandpile group, except that they disqualify no vertex firings. Though this is isomorphic to $\mathcal{S}^i$ in the undirected case, in general it's a smaller.

Definition 1.3.3. *Let the* Picard group *be* $P = \mathrm{Con}_0(G)/\mathrm{Prin}(G)$.

Theorem 1.3.3. $|P| = \mu$, *(cf. Definition 1.2.3) and*

$$0 \to \mathbb{Z}_{\lambda(i)} \to \mathcal{S}^i \to P \to 0,$$

where $\mathbb{Z}_{\lambda(i)} = \mathbb{Z} \bmod \lambda(i)$.

Proof. We consider the diagram

$$\begin{tikzcd}
 & \mathbb{Z}_{\lambda(i)} & 0 & \boxed{0} & \\
0 \arrow[r] & \mathbb{Z}^n/\lambda\mathbb{Z} \arrow[r, "L"] \arrow[u] & \mathrm{Con}_0(G) \arrow[r, "+\mathrm{Prin}(G)"] \arrow[u] & P \arrow[r] \arrow[u] & 0 \\
0 \arrow[r] & \mathbb{Z}^{n-1} \arrow[r, "L^i"] \arrow[u, "\iota"] & \mathrm{Con}_0(G) \arrow[r, "+\mathrm{Prin}^i(G)"] \arrow[u, equal] & \mathcal{S}^i \arrow[r] \arrow[u] & 0 \\
 & 0 \arrow[u] & 0 \arrow[u] & \boxed{\mathbb{Z}_{\lambda(i)}} \arrow[u] &
\end{tikzcd}$$

Since L is not of maximal rank, modding out by $\ker_{\mathbb{Z}}(L) = \lambda\mathbb{Z}$ is necessary to make the top row exact; again P is the quotient of the relevant spaces. Again the $\mathrm{Con}_0(G)$ spaces can be identified. This time, however, the map from $\mathbb{Z}^{n-1} \to \mathbb{Z}^n/\ker(L)$ is not surjective. The natural map ι, that commutes with the previous identification, takes an element of $\mathbb{Z}^{n-1}$ and inserts a 0 in the ith position to get an element of $\mathbb{Z}^n/\lambda\mathbb{Z}$. But this doesn't quite surject—though we can get whatever we want in any position but the ith, ι gives only representatives with some multiple of $\lambda(i)$ in the ith entry. So $\mathbb{Z}^n/\mathrm{img}(\iota) = \mathbb{Z}_{\lambda(i)}$, the exact sequence has to be as above, and filling in the boxed terms via the Snake Lemma, get that the exact sequence has to be

$$0 \to \mathbb{Z}_{\lambda(i)} \to \mathcal{S}^i \to P \to 0,$$

as desired.

However, it's an easy consequence of the exactitude of the above sequence that $|P||\mathbb{Z}_\lambda(i)| = |\mathcal{S}^i|$. But then $|P| = \kappa(i)/\lambda(i) = \mu$, as desired. □

Finally, we consider the structure most important to Riemann Roch, the sandcastle at i: the group of all configurations, modulo the winnowed principals at i. For the first time we're working with all configurations, not just the ones whose values on all vertices sum to 0. It turns out that the sandcastle is just an infinite column of copies of the sandpile group, one of each degree.

Definition 1.3.4. *Let the* sandcastle $\mathfrak{S}^i = \mathrm{Con}(G)/\mathrm{Prin}^i(G)$, *and let the* degree *of a configuration* $\deg\colon \mathrm{Con}(G) \to \mathbb{Z}$ *be such that* $\deg(D) = \sum_j D(v_j)$.

Theorem 1.3.4. *$\mathfrak{S}^i \approx \mathcal{S}^i \oplus \mathbb{Z}$.*

Proof.

$$\begin{tikzcd}
 & 0 & \mathbb{Z} & \boxed{\mathbb{Z}} & \\
0 \arrow[r] & \mathbb{Z}^{n-1} \arrow[u] \arrow[r, "L^i"] & \mathrm{Con}(G) \arrow[u, "\deg"'] \arrow[r, "+\mathrm{Prin}^i(G)"] & \mathfrak{S}^i \arrow[u, "\deg"'] \arrow[r] & 0 \\
0 \arrow[r] & \mathbb{Z}^{n-1} \arrow[u, equal] \arrow[r, "L^i"] & \mathrm{Con}_0(G) \arrow[u, "\iota"'] \arrow[r, "+\mathrm{Prin}^i(G)"] & \mathcal{S}^i \arrow[u, "\iota"'] \arrow[r] & 0 \\
 & 0 \arrow[u] & 0 \arrow[u] & \boxed{0} \arrow[u] &
\end{tikzcd}$$

The top row is exact for all the same reasons, and we've seen the bottom row twice already. The $\mathbb{Z}^{n-1}$s are the same space, but although $\mathrm{Con}(G)$ contains a copy of $\mathrm{Con}_0(G)$ and our mapping upwards is just the natural injection, it's bigger by a factor of $\mathbb{Z}$, and deg picks out exactly that structure. So by using the same mappings and invoking the Snake Lemma, we get two more boxed terms, and that

$$0 \longrightarrow \mathcal{S}^i \xrightarrow{\ \iota\ } \mathfrak{S}^i \xrightarrow{\ \deg\ } \mathbb{Z} \longrightarrow 0$$

is exact.

In general this isn't enough to prove our contention. But all we require in addition is a homomorphism $\beta\colon \mathbb{Z} \to \mathfrak{S}^i$ such that composed with the mapping in the sequence, we have the identity: $\deg \circ \beta = I_{\mathbb{Z}}$. After all, since the sequence is exact, the image of ι is the kernel of deg, and since the sequence is exact having a homomorphism β back would allow us to express each element as the sum of an element in the image of ι and in the image of β, which is exactly what we need. Finally, it's easy to convince oneself that $\beta\colon d \mapsto (v_i \mapsto d)$ is indeed such a homomorphism, since $(\deg \circ \beta)\, d = d$. Thus, $\mathfrak{S}^i \approx \mathcal{S}^i \oplus \mathbb{Z}$. □

Chapter 2

Riemann Roch on Sandcastles

2.1 Further Structure

Now we have the sandcastle, the group of all configurations modulo nonsink vertex firings, the group direct sum of the sandpile group and the integers. This is where the magic happens, this is the structure in which contact is made with the theory of Riemann surfaces—the elements of the sandcastle are analogous to the *divisor classes* of Riemann surfaces, and to emphasize the correspondence (after the work of Baker and Norine, [1]) we name them such. But in order to define the relationship, we need a bit of additional structure.

Note first that there's a natural partial order to the sandcastle, defined componentwise on representatives.

Definition 2.1.1. *For any two S, $S' \in \mathfrak{S}^i$, we say that $S \geq S'$ if there exist representatives $s \in S$ and $s' \in S'$ such that $s(v) \geq s'(v)$ for all $v \in V(G)$.*

Since vertex firing respects degree, divisor classes have a well-defined degree, and $S \geq S'$ implies $\deg(S) \geq \deg(S')$, and $\deg(S) = \deg(S')$ only if $S = S'$; thus it's not hard to convince oneself that it's actually a partial order.

Further, since both this and the group operation are defined componentwise, the partial order respects the group operation.

Now, we're also interested in divisor classes S which have a certain property—that is, that there are no representatives $s \in S$ such that $s(v) \geq 0$ for all S, or such that $S \not\geq 0$. Naturally this is true of all divisors of degree less that 0, and true of all divisors other than 0 of degree 0. For a sufficiently complicated graph, there will be a number of divisor classes of degree greater than 0 for which this is also true, and the maximum degree of those divisors is a measure of the complicatedness of the graph. We will have theorems about this later.

2.2 $r\colon \mathfrak{S}^i \to \mathbb{Z}_{\geq -1}$

Now, Riemann Roch is a theorem about the dimensions of the linear systems of divisor classes on Riemann surfaces. Thus, again after [1], we're motivated to define a similar function on the sandcastle, r.

Definition 2.2.1. *Let $r\colon \mathfrak{S}^i \to \mathbb{Z}_{\geq -1}$ be such that $r(S)$ is the maximum k such that $S \geq E$ for all $E \geq 0$ of degree k,* i.e., *$S - E \geq 0$ for all $E \geq 0$ such that* $\deg(E) = k$, *if such a k exists. If there exists no such k—if $S \ngeq 0$—then we define $r(S) = -1$.*

Thus r is a measure of how far removed a divisor class S is from the divisors $\ngeq 0$, or how hard it is to 'break' a divisor class S by removing sand from it. It's the number of levels in the pyramid beneath S with respect to the partial order that one can go down before one first meets a divisor class $\ngeq 0$. And it's somehow analogous to the dimension of the linear system of a divisor.

We claim for it a few nifty properties:

Theorem 2.2.1. *$S \geq S' \Rightarrow r(S) \geq r(S')$ and for $S, S' \geq 0$, $r(S+S') \geq r(S)+r(S')$.*

Proof. Since $S \geq S' \geq E \geq 0$, for all $E \geq 0$ such that $\deg(E) = r(S')$, the first contention follows.

Further, since $S \geq E \geq 0$ and $S' \geq E' \geq 0$ for all $E \geq 0$ and $E' \geq 0$ of degrees $r(S)$ and $r(S')$ respectively, we have that $S + S' \geq E + E' \geq 0$. And since all $E'' \geq 0$ of degree $r(S)+r(S')$ can be written as the sum of $E \geq 0$ and $E' \geq 0$ of the necessary degrees, the second contention follows. □

Now, we consider what the sandcastle looks like under r.

Naturally for all S of degree ≤ 0, $S \ngeq 0$ and $r(S) = -1$. So up until degree 0, the sandcastle under r is rather boring. At degree 0 itself, naturally $0 \geq 0$ so $r(0) = 0$; for $\deg(S) = 0$, $S \neq 0$, $r(S) = -1$. From there, the divisor classes $\ngeq 0$ rise in a number of pinnacles of equal height.

That is,

Definition 2.2.2. *Let the* foundation *be $\mathcal{F} = \{S \in \mathfrak{S}^i : S \ngeq 0\}$, let $f = max_{F\in\mathcal{F}} \deg(F)$, and let the* nonspecial divisors *be $\mathcal{N} = \{F \in \mathcal{F} : \deg(F) = f\}$.*

Then the behavior that we were attempting to demonstrate on the sandcastles are

Condition 2.1 (RR1). *For all $S \in \mathfrak{S}^i$, $S \ngeq 0 \Leftrightarrow S \leq \nu$ for some $\nu \in \mathcal{N}$*

From here, there's a good deal of behavior consistent with Theorem 2.2.1 that we won't remark upon, but things get interesting again at $\deg(S) = 2\deg(\mathcal{N})$. At this point, there's exactly one divisor class with $r(S) \geq \deg(\mathcal{N})$, and all the rest $\mapsto \deg(\mathcal{N}) - 1$. That is,

Condition 2.2 (RR2). *There exists a divisor class K such that* $\deg(K) = 2\deg(\mathcal{N})$ *and $r(K) \geq \deg(\mathcal{N})$.*

Finally, it turns out that under our assumptions, these contentions are equivalent to Riemann Roch.

Condition 2.3 (RR). $r(S) - r(K - S) = \deg(S) - \deg(\mathcal{N})$

2.3 Riemann Roch

Note that this links the r values of an element S, and $K - S$ of degree $2\deg(\mathcal{N}) - \deg(S)$. That is, every degree below $\mathcal{N}$ is mapped in a nonobvious way to the corresponding level above, and the difference of the r values of the paired divisor classes is always $\deg(S) - \deg(\mathcal{N})$. You can picture this as a theorem about the 'safe pyramids' of divisor classes $\leq S$ and $\leq K - S$, how far down you can go before you hit something $\not\geq 0$—the deeper one goes while staying safe, the deeper the other one must.

We now proceed with the proof of equivalence. (Which, along with just about everything else in this subsection, is taken straight from BN07.)

But first, a small lemma.

Theorem 2.3.1. *RR2 is equivalent to* $\nu \in \mathcal{N} \Leftrightarrow K - \nu \in \mathcal{N}$.

Proof. If RR2 holds, then $K \geq E$ for all $E \geq 0$ of degree $\deg(\mathcal{N})$. But then $K - E \geq 0$—that is, subtracting from K maps a divisor class ≥ 0 to another ≥ 0, of degree $\deg(\mathcal{N})$. And since $K - (K - S) = S$, subtraction from K has to map a divisor class $\not\geq 0$ to another $\not\geq 0$, and $\nu \in \mathcal{N} \Leftrightarrow K - \nu \in \mathcal{N}$.

Conversely, if $\nu \in \mathcal{N} \Leftrightarrow K - \nu \in \mathcal{N}$, then since $K - (K - S) = S$, subtraction from K must map a divisor class of degree $\deg(\mathcal{N})$ and ≥ 0 to another ≥ 0, and $r(K) \geq \deg(\mathcal{N})$. □

Now, a simple observation—if the difference of two functions, bounded below on the same domain, is constant, then they must both obtain any minima on the same domain element, and the difference of the minima returns that constant value, which is the minimum of the difference.

Finally,

Definition 2.3.1. *We define* $\deg^+ \colon \operatorname{Con}(G) \to \mathbb{Z}_{\geq 0}$ *by*

$$\deg^+(s) = \sum_{\substack{v \in V(G) \\ s(v) \geq 0}} s(v).$$

(Note that $\deg^+$ is defined on configurations, not divisor classes.)

Theorem 2.3.2. *If RR1 holds, then*

$$r(S) = r'(S) := \left(\min_{\substack{s \in S \\ n \in \nu \in \mathcal{N}}} \deg^+(s - n) \right) - 1.$$

Proof. If $r(S) < r'(S)$, then there exists $E \geq 0$ of degree $r'(S)$ such that $S \not\geq E$, and by RR1 $S - E \leq \nu$ or $0 \leq \nu - S + E$. But then that means there exists some $E' \geq 0$ such that $E' = \nu - S + E$, or $S - \nu = E - E'$, and there exist representatives such that $s - n = e - e'$. But then $\deg^+(s - n) - 1 \leq \deg(E) - 1 = r'(S) - 1$, contradicting the definition of r'. Thus $r(S) \geq r'(S)$.

Now, let's pick $s \in S$ and $n \in \nu \in \mathcal{N}$ achieving the minimum in $r'(S)$. Then $\deg^+(s-n) = r'(S)+1$, and thus there exist $E \geq 0$ such that $\deg(E) = r'(S)-1$ and $E' \geq 0$ such that $E - E' = S - \nu$, or $S - E = \nu - E' \leq \nu$. So $S - E \not\geq 0$, and $r(S) \leq r'(S)$, which proves the theorem. □

Finally, we use this to prove Riemann Roch.

Theorem 2.3.3. *RR1 & RR2* $\Leftrightarrow$ *RR.*

Proof. First we assume RR1 and RR2. Select a divisor class S and a representative s, an $n \in \nu \in \mathcal{N}$, and a $k \in K$. Since $K - \nu \in \mathcal{N}$, $\overline{n} = k - n \in \overline{\nu} \in \mathcal{N}$, and

$$\begin{aligned}\deg^+(s-n) - \deg^+((K-s)-\overline{n}) &= \deg^+(s-n) - \deg^+(n-s)\\ &= \deg(S-\nu)\\ &= \deg(S) - \deg(\mathcal{N}).\end{aligned}$$

Now, since the two values on the left are bounded below at 0, they achieve their minima somewhere on the domain, and since their difference is a constant it must be at the same domain element. But as s and $n \in \nu \in \mathcal{N}$ run over all possible values, so do s and $\overline{n} \in \overline{\nu} \in \mathcal{N}$, and we have

$$r(S) - r(K-S) = \deg(S) - \deg(\mathcal{N}).$$

Now, assume Riemann Roch. Since $r(K-\nu) = \deg(\mathcal{N}) - \deg(\nu) + r(\nu) = -1$, we have RR2.

For RR1, by the additivity of $r(S)$, we can't have both $r(S) \geq 0$ and $r(\nu - S) \geq 0$, for some $\nu \in \mathcal{N}$, or we'd have $r(\nu) \geq 0$. RR1 states that either $S \not\geq 0$ or $\nu - S \not\geq 0$ for some $\nu \in \mathcal{N}$, so it remains to show that they can't both be. So, pick a divisor class S of degree less than or equal to $\deg(\mathcal{N})$, and add sand until you get there. By Riemann Roch, if there is no way to do that such that the divisor you end up with is in $\mathcal{N}$, then your divisor class is effective. Thus, if you start out with a divisor class such that $S \not\geq 0$, there is a way to add sand to S to wind up in $\mathcal{N}$. But that means $S \leq \nu \in \mathcal{N}$, or $\nu - S \geq 0$, which is RR1.

That is, pick $E \geq 0$, $\deg(E) = \deg(\mathcal{N}) - \deg(S)$. If $S + E \geq 0$, then by Riemann Roch,

$$\begin{aligned}0 &\leq r(S+E)\\ &= \deg(S+E) - \deg(\mathcal{N}) + r(K-S-E)\\ &= r(K-S-E),\end{aligned}$$

and

$$0 \leq K - S - E$$

or

$$E \leq K - S.$$

If this is true for all $E \geq 0$, $\deg(E) = \deg(\mathcal{N}) - \deg(S)$, then

$$\begin{aligned} \deg(\mathcal{N}) - \deg(S) &\leq r(K - S) \\ &= \deg(K - S) - \deg(\mathcal{N}) + r(S) \\ &= \deg(\mathcal{N}) - \deg(S) + r(S). \end{aligned}$$

so $r(S) \geq 0$.

Thus, if $S \ngeq 0$, then there exists an $E \geq 0$ with $\deg(E) = \deg(\mathcal{N}) - S$ such that $S + E = \nu \in \mathcal{N}$, or $S \leq \nu \in \mathcal{N}$, which proves the theorem. $\square$

Chapter 3

Nuts and Bolts

This is all exciting stuff, but we have still given no constructive method for determining, from a connected weighted directed graph with a watershed draining into at least one vertex and a choice of sink, $\mathcal{N}$ and K, nor have we proven that they have the properties we require in RR1 and RR2. That is the task we address ourselves to here. RR1 is unproven on the directed case, and RR2 is false in general for the directed case, as we'll see in Chapter 4, but there is some intermediate work of intermediate interest.

3.1 Washouts

We begin with a specific set of strict partial orders on the vertices of the graph. A strict partial order is a binary relation that satisfies irreflexivity $\neg(u \prec u)$, asymmetry $u \prec v \Rightarrow \neg(v \prec u)$ and transitivity, $u \prec v$ and $v \prec w$ implies $u \prec w$. One can show that there is a one-to-one correspondence between strict partial orders and partial orders; we use strict partial orders here.

In constructing the washouts, we consider the set of all pairs of adjacent vertices. On each of these, we select either $u \prec v$ or $v \prec u$ subject to the constraint that there exists a path from every vertex into the sink that can be followed in descending order. (Since we require our graphs to contain at least one watershed, and since we require our sinks to have at least one watershed draining into them, we can always pick at least one set of such relationships with this property.) Finally, since this doesn't yet satisfy transitivity, we pick the smallest, 'most partial' partial ordering consistent with our preferences among adjacent vertices.

Definition 3.1.1. *We call two vertices $u, v \in E(G)$* adjacent *if either $(u \to v) \in E(G)$ or $(v \to u) \in E(G)$. Select a sink v_i such that a watershed drains into v_i, and consider the set of binary relations on adjacent vertices such that there exists a path p from each vertex in $V(G)$ to v_i such that for each edge $(u_i \to u_{i+1}) \in p$, $u_i \succ u_{i+1}$. For each such binary relation, let the intersection of all partial orders containing that binary relation be the corresponding washout, and let the set of all* washouts *into v_i be denoted d.*

We can think of a washout as being a determination of relative height among the vertices of the graph, and $u \succ v$ as telling you whether or not you can get from u to v by going exclusively downhill. We require that everything flow downhill into the sink, along the edges that exist within the graph, but for sufficiently complicated graphs there are, in general, multiple ways to do this, multiple ways to preference the vertices such that the downhill edges give you a path into the sink. The set of washouts is the set of such preferences.

Now, a few important terms.

Definition 3.1.2. *Let $\delta \in d$ be a washout into v_i, and let the* allowed *edges* $E_\delta = \{(u \to v) \in E(G) : u \succ_\delta v\}$. *Let the* damage *of the washout* $\operatorname{dam}(\delta) = \sum_{e \in E_\delta} W(G)(e)$. *Let* $h = \max_{\delta \in d} \operatorname{dam}(\delta)$, *and let the* disasters *be* $\mathcal{D} = \{\delta \in d : \operatorname{dam}(\delta) = h\}$.

3.2 Hills and $\sigma\colon d \to m$

Now we develop a closely related object. Remember that a subgraph of a graph G is a graph G' such that $V(G') \subseteq V(G)$, $E(G') \subseteq E(G)$, and $W(G')(e) \leq W(G)(e)$ for all $e \in E(G')$; that is, a graph with vertices, edges, and weights taken from the parent graph, and no extras. A hill on G is a subgraph satisfying a few additional properties.

Definition 3.2.1. *Let the set of all subgraphs G' of G such that G' contains a watershed draining into the sink, and no paths p from a vertex $u \in V(G')$ to u, be a* hill, *and let the set of all hills be denoted m.*

We can think of a hill as being the set of ways that water can flow downhill, such that it all winds up in the sink. Now, just as we did with washouts, we define a manner of 'weight,' and the objects maximal with respect to such.

Definition 3.2.2. *Let $\mu \in m$ be a hill, and let* $\operatorname{slp}(\mu) = \sum_{e \in E(\mu)} W(e)$. *Let* $k = \max_{\mu \in m} \operatorname{slp} \mu$, *and let* $\mathcal{M} = \{\mu \in m : \operatorname{slp}(\mu) = k\}$.

Naturally, there's a nice way to get from washouts to hills: we define $\sigma\colon d \to m$ to be the map that takes a washout $\delta \in d$, and picks the subgraph defined by the allowed edges.

Definition 3.2.3. *Let $\sigma\colon d \to m$ be the map such that* $V(\sigma(\delta)) = V(G)$, $E(\sigma(\delta)) = \{(u \to v) \in E(G) : u \succ_\delta v\}$, *and* $W(\sigma(\delta))(e) = W(G)$ *for all* $e \in E(\sigma(\delta))$.

It's easy to convince oneself that this is indeed a hill—the allowed edges cannot contain a cycle on pain of violating irreflexivity, and the requirement that there exists a descending path from each vertex into the sink translates to the requirement that there exist a watershed. Note that though σ is injective—two distinct washouts will allow different sets of edges, since they must disagree on some adjacent vertices—it's not surjective. For example, the graph with two vertices, and one edge going from one to the other of weight two, has one washout, but two hills—the complete graph, and the subgraph where the edge has weight one.

Further, we define a quantity analogous to the cyclomatic number on an undirected graph, and analogous to the genus of a Riemann surface.

Definition 3.2.4. *Let the genus of a weighted directed graph G that contains a watershed draining into sink $v_i \in V(G)$ be $g^i = \sum_{e \in E(\mathcal{M})} W(\mathcal{M})(e) - |V(\mathcal{M})| + 1$*

Since the set of mountains is maximal with respect to the sums of weights of edges, this is constant among all mountains, and the quantity g^i is well-defined. Note that, unlike the directed case, this depends on one's choice of sink.

Finally, we note something nifty.

Theorem 3.2.1. *σ is a bijection between $\mathcal{D}$ and $\mathcal{M}$.*

Proof. Since two distinct washouts lead under σ to distinct hills, injectivity of disasters is clear. Since σ of all disasters are of equal slope, to show that σ of the set of all disasters surjects onto the set of all mountains, it suffices to show that all mountains have a preimage under σ.

So consider a specific mountain. By maximality, adding any additional edges of G must produce a cycle. Thus for every pair of adjacent vertices, either an edge connecting them is in the mountain, or the edge connecting them cannot be added to the mountain without producing a cycle, which means there's a path through the mountain connecting the second to the first. Thus, for each pair of adjacent vertices, there is a path going from one to the other. Now, for each pair of adjacent vertices, if a path exists from a vertex u to a vertex w, say that $u \succ w$, and pick the smallest partial ordering that contains these relations. This is a disaster, and σ of this disaster is the mountain.

Since we constructed the partial ordering by choosing its behavior on adjacent vertices, it leads easily to our original mountain. Since the mountain has a watershed draining into the sink, by the method in which we chose the relations on adjacent vertices the partial ordering must allow a descending path into the sink from every vertex, and so is a washout. And if the washout were not a disaster, then σ of a disaster would have a larger sum of weights of edges than our mountain, contrary to definition. Thus this is a disaster, which proves the theorem. □

3.3 $\mathcal{N}$, $\tau\colon m \to \mathfrak{S}^i$ and RR1

Finally, we consider the following map from the set of mountains into the sandcastle.

Definition 3.3.1. *For all $\mu \in m$, let*

$$\tau(\mu)(v) = \sum_{\substack{u \in V(\mu) \\ (u \to v) \in E(\mu)}} W((u \to v)) - 1.$$

That is, take all the edges in the mountain, add up all the weights of the edges heading into each vertex, and subtract one. This gives you a representative of the divisor class $\tau(\mu)$, a map whose properties dovetail beautifully with RR1: proof of the following conditions on the map τ would yield RR1.

I'm not aware of any examples in which RR1 fails, but I've no proof.

Condition 3.3.1. *$\tau(\mu) \ngeq 0$ for all $\mu \in m$, and for all $S \ngeq 0$, $S \leq \tau(\mu)$ for some $\mu \in \mathcal{M}$.*

Condition 3.3.2. *τ is a bijection between $\mathcal{M}$ and $\mathcal{N}$, and $\deg(\mathcal{N}) = g^i - 1$.*

3.4 K and RR2

Now we have a constructive method for finding $\mathcal{N}$, given a directed graph and a choice of sink. So we move on to K.

Theorem 3.4.1. *For finite $\mathcal{N}$, RR2 implies*

$$K = \frac{2}{|\mathcal{N}|} \sum_{\nu \in \mathcal{N}} \nu.$$

Proof. Since RR2 states that for all $\nu \in \mathcal{N}$, $K - \nu \in \mathcal{N}$

$$\begin{aligned} 2 \sum_{\nu \in \mathcal{N}} \nu &= \sum_{\nu \in \mathcal{N}} \nu + \sum_{\nu \in \mathcal{N}} K - \nu \\ &= \sum_{\nu \in \mathcal{N}} K \\ &= |\mathcal{N}| K. \end{aligned}$$

□

So RR2 guarantees the existence of a divisor K such that $|\mathcal{N}|K = 2\sum_{\nu \in \mathcal{N}} \nu$; one nevertheless has to be careful about representatives when constructing it. Note that RR2 is false on directed graphs, so this guarantee isn't worth much, but I'm not aware of any example in which K as defined is not in the sandcastle; rather, it merely lacks the properties we require of it.

And finally, RR2 can be rephrased one more way.

Theorem 3.4.2. *RR2 $\Leftrightarrow K \ngeq \nu$ for all $\nu \in \mathcal{N}$.*

Proof. Since $\deg(K) = 2\deg(\mathcal{N})$, $K - \nu \in \mathcal{N} \Leftrightarrow K - \nu \ngeq 0 \Leftrightarrow K \ngeq \nu$. □

Chapter 4

The Counterexample

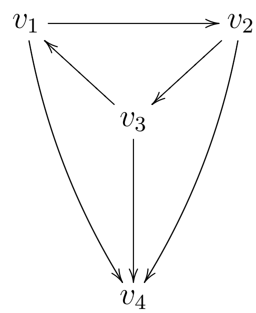

We examine the sandpile group and sandcastle for the graph above, and find that it's a counterexample to RR2.

Let v_4 be the sink, as it's the only vertex into which a watershed drains. We note that there are seven watersheds draining into v_4,

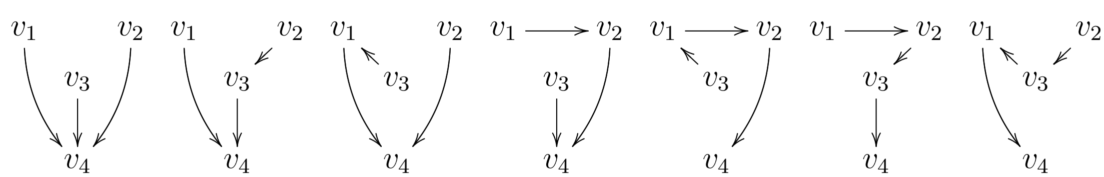

each of weight one, so the order of the sandpile group is seven, and it remains to find seven pairwise nonequivalent representatives.

Definition 4.0.1. *Let* $(0,0,0,0) \in \mathbf{0}$*,* $(0,0,1,-1) \in \mathbf{1}$*,* $(0,1,0,-1) \in \mathbf{2}$*,* $(1,0,0,0) \in \mathbf{3}$*,* $(1,1,0,-2) \in \mathbf{4}$*,* $(1,0,1,-2) \in \mathbf{5}$*,* $(0,1,1,-2) \in \mathbf{6}$*.*

Theorem 4.0.3. $\mathbf{n} \neq \mathbf{m}$ *for* $n \neq m$*, and* $\mathcal{S}^4 = \{\mathbf{n}\}$*.*

Proof. Since the graph is symmetric, $\mathbf{1} \neq \mathbf{0}$ else we'd have a trivial sandpile group. Noting further that adding $\mathbf{1}$ takes

$$\mathbf{0} \to \mathbf{1} \to \mathbf{3} \to \mathbf{5} \to \mathbf{2} \to \mathbf{6} \to \mathbf{4} \to \mathbf{0}$$

suffices. $\square$

These additionally give us a complete characterization of the divisor classes in the sandcastle at 4—linear equivalence naturally respects the group operation, so for the divisor classes of $\deg S = k$, we need merely add k to the v_4 term. Let the divisor class S of degree k such that $S = \mathbf{n} + k\, e_4$ be $\mathbf{n}^k$; then since $\mathfrak{S}^4 \approx \mathcal{S}^4 \oplus \mathbb{Z}$, $\mathfrak{S}^4 = \{\mathbf{n}^k : k \in \mathbb{Z}\}$.

Now, we characterize the partial order on the sandcastle at 4. Since the graph has four vertices, each divisor class is $\geq$ exactly four divisor classes of the previous degree—$S \geq S - e_j$ for all $v_j \in V(G)$. By firing a few vertices, it can be shown that

$$\begin{aligned}
\mathbf{0}^{k+1} &\geq \mathbf{0}^k, \mathbf{4}^k, \mathbf{5}^k, \mathbf{6}^k \\
\mathbf{1}^{k+1} &\geq \mathbf{0}^k, \mathbf{1}^k, \mathbf{2}^k, \mathbf{4}^k \\
\mathbf{2}^{k+1} &\geq \mathbf{0}^k, \mathbf{2}^k, \mathbf{3}^k, \mathbf{5}^k \\
\mathbf{3}^{k+1} &\geq \mathbf{0}^k, \mathbf{3}^k, \mathbf{1}^k, \mathbf{6}^k \\
\mathbf{4}^{k+1} &\geq \mathbf{1}^k, \mathbf{2}^k, \mathbf{4}^k, \mathbf{5}^k \\
\mathbf{5}^{k+1} &\geq \mathbf{2}^k, \mathbf{3}^k, \mathbf{5}^k, \mathbf{6}^k \\
\mathbf{6}^{k+1} &\geq \mathbf{3}^k, \mathbf{1}^k, \mathbf{6}^k, \mathbf{4}^k.
\end{aligned}$$

Thus, considering $k = 0$, we have that $\mathbf{4}^1, \mathbf{5}^1$ and $\mathbf{6}^1 \not\geq 0$, and since $\mathbf{4}^2 \geq \mathbf{1}^1 \geq 0$ and so on, $\mathcal{N} = \{\mathbf{4}^1, \mathbf{5}^1, \mathbf{6}^1\}$, consistent with RR1 and the machinery of Chapter 3. (The representatives given by τ are $(-1, 0, 0, 2), (0, -1, 0, 2), (0, 0, -1, 2)$, but by inverse firing vertices $1, 2, 3$ we get the representatives used above.)

However, for $k = 2 = 2\deg(\mathcal{N})$, we have that for all $\mathbf{n}$, $\mathbf{n}^2 \geq \nu$ for some $\nu \in \mathcal{N}$, contradicting RR2. Constructing K regardless, according to our formula (cf Theorem 3.4.1) yields

$$\begin{aligned}
K &= \frac{2}{|\mathcal{N}|} \sum_{\nu \in \mathcal{N}} \nu \\
&= \frac{2}{3} \left(\mathbf{4}^1 + \mathbf{5}^1 + \mathbf{6}^1\right) \\
&= \frac{2}{3} \langle (1, 1, 0, -1) + (1, 0, 1, -1) + (0, 1, 1, -1) \rangle \\
&= \frac{2}{3} \langle (1, -3, 2, 1) + (-3, 2, 1, 1) + (2, 1, -3, 1) \rangle \\
&= \langle (0, 0, 0, 2) \rangle \\
&\geq \mathbf{4}^1, \mathbf{5}^1, \mathbf{6}^1 \\
&= \mathcal{N}.
\end{aligned}$$

That is, $K \geq \nu$ for *all* $\nu \in \mathcal{N}$, again contradicting RR2. (cf. Theorem 3.4.2) Note also that since $\lambda(4) = 1$ here, $P = \mathcal{S}^4$, and this is also a counterexample for RR2 on the Picard group.

Thus, since by Theorem 2.3.3 RR1 and RR2 are equivalent to RR, we have that RR is false in general on weighted directed graphs containing at least one watershed—that is, that there exists no choice of K such that $r(S) - r(K - S) = \deg(S) - \deg(\mathcal{N})$ holds.

We can, in fact, go slightly farther; there exists no integer g^i such that $r(S) - r(K-S) \deg(S) - g^i + 1$. The sandcastle of our counterexample has only two degrees on which $r(S)$ is nonconstant; 0, for which $r(\mathbf{0}^0) = 0$, $r(\mathbf{n}^0) = -1$ for $n > 0$, and 1, for which $r(\mathbf{n}^1) = 0$ for $n \leq 4$, $r(\mathbf{n}^1) = -1$ else. An equation like Riemann Roch would have to match up the behavior with respect to r in a one-to-one way on the degrees in which r is nonconstant; that's just not possible here.

Nor is there any room to redefine r or deg to make things work. r falls straight out of the componentwise partial order on the sandcastle; any redefinition of r would lose us that correspondence, and the behavior described by Theorem 2.2.1. And abandoning deg would change the theory unrecognizably.

Baker and Norine proved their Riemann Roch Criterion for a general free abelian group under a degree-preserving equivalence relation; we've examined this for configurations on the vertices modulo the equivalence relation that gives the sandpile groups, vertex firing on vertices sufficient to grant a finite group. The above is a counterexample for all such—to prove any such Riemann Roch on weighted directed graphs, we would need a new equivalence relation, either vertex firing on vertices insufficient to grant a finite group, or something else entirely.

We suspect it's possible that Riemann Roch could still be proven on some subset of weighed directed graphs with a watershed draining into at least one vertex; that is, graphs that lack 'disastrous cycles,' cycles such that for each edge in the cycle, a disaster exists that allows every edge in the cycle but that one. The author suspects, however, that this is not a fruitful direction for further inquiry.

References

[1] Matthew Baker and Serguei Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. http://front.math.ucdavis.edu/0608.5360. E-print arXiv:0608.5360.

[2] Alexander E. Holroyd, Lionel Levine, Karola Meszaros, Yuval Peres, James Propp, and David B. Wilson. Chip-firing and rotor-routing on directed graphs. http://www.citebase.org/abstract?id=oai:arXiv.org:0801.3306, 2008. E-print arXiv:0801.3306.