IBL CHIP FIRING NOTES

These notes are intended for use in a class taught in an inquiry-based learning format. They contain very few proofs. Instead, the reader is expected to provide proofs of the stated theorems, propositions, and exercises on their own. Everything is presented in an order that guides the reader through the process of discovery. There is a "teacher's guide" containing proofs, available upon request.

These notes draw from many different sources. None of the ideas presented herein are original to this work. For pedagogical reasons, line-by-line references are not included, but each chapter ends with a short section detailing references for that chapter.

Chip firing is a game played with poker chips on the vertices of a graph that has applications to many diverse fields of mathematics. In these notes, our primary focus is on applications to the geometry of algebraic curves. For this reason, we use terms like *divisors* and *Jacobians*, rather than *chip configurations*, *sandpiles*, or *critical groups*.

Outline. Section 1 develops the basic theory of chip firing games on graphs. Aside from the energy and monodromy pairings, this is standard material that can be found in almost every text on the subject. Section 2 is focused on reduced divisors and Dhar's Burning Algorithm. This is a very useful tool in applications, and while it is completely combinatorial, it powers many of the applications to algebraic geometry. Section 3 explores the connection between divisors on graphs and graph orientations. This material is necessary for the proof of Riemann-Roch, but is otherwise not used. Readers who want to skip the proof of Riemann-Roch should feel free to skip this section. Section 4 introduces an important invariant of a divisor, known as its *Baker-Norine rank*. This section includes the Riemann-Roch theorem for graphs, introduces tools for computing gonality of graphs, and covers the combinatorial steps in the tropical proof of the Brill-Noether theorem. My suggestion for an undergraduate course would be to cover Sections 1, 2, and 4, skipping the proof of Riemann-Roch. Section 5 covers metric graphs. Most of the literature relating chip firing to algebraic curves uses metric graphs, but we do not cover the material on nonarchimedean analytic geometry required to draw this connection. Section 6 is an introduction to algebraic geometry, and particularly to algebraic curves. Because it is not combinatorial in nature, this section has a distinctly different flavor from previous sections. Finally, in Section 7, we discuss specialization of divisors from curves to graphs, and provide sample applications of this theory to the geometry of algebraic curves.

1. Divisors on Graphs

1.1. **Basic Theory.** Throughout these notes, all graphs are assumed to be connected and loopless, though possibly with multi-edges. Given a graph G, we write $V(G)$ for the set of vertices and $E(G)$ for the set of edges. Our main object of study is divisors on graphs.

Definition 1.1. *A* divisor D *(or* chip configuration *or* sandpile*) on a graph G is a formal $\mathbb{Z}$-linear combination of vertices of G,*

$$D = \sum_{v \in V(G)} D(v) \cdot v$$

with $D(v) \in \mathbb{Z}$.

For example, Figure 4 depicts a divisor on the wedge of two triangles. Here and elsewhere, a vertex with coefficient zero is undecorated.

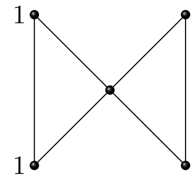

FIGURE 1. A divisor.

Divisors on graphs have been studied in combinatorics, computer science, and dynamics long before algebraic geometers got interested in them. In these disciplines it is more common to refer to divisors on graphs as *chip configurations* or *abelian sandpiles*. The term "chip configuration" comes from thinking of a divisor as a stack of poker chips on each vertex of the graph. Here we use the term *divisor* to emphasize the analogy with divisors on algebraic curves. Note that the divisors on a graph G form an abelian group, which we denote $\mathrm{Div}(G) = \mathbb{Z}^{V(G)}$.

The chip-firing game is a game played with divisors on graphs, in which there is only one move. Starting with a divisor, we may "fire" a vertex, which results in that vertex giving a chip to each of its neighbors. More concretely, we have the following definition.

Definition 1.2. *The* chip-firing move *at a vertex v takes a divisor D to D' where*

$$D'(w) = \begin{cases} D(v) - \mathrm{val}(v) & \text{if } w = v \\ D(v) + \#\text{ of edges between } w \text{ and } v & \text{if } w \neq v. \end{cases}$$

In our example, if we fire the top left vertex, we get the divisor pictured in Figure 2.

Exercise 1.3. Let D be a divisor on a graph G, and let $v, w \in V(G)$. Show that the divisor obtained from D by first firing v and then w is the same as the divisor obtained from D by first firing w and then v.

Exercise 1.4. Let D be a divisor on a graph G, and let $A \subseteq V(G)$. What is the effect of firing each of the vertices in A exactly once (in any order)?

Firing each vertex in the subset A is sometimes referred to as the *cluster-fire* of A.

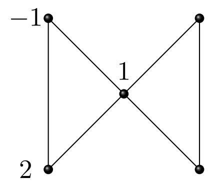

FIGURE 2. The result of firing a vertex.

Definition 1.5. *Two divisors D, D' are* linearly equivalent*, and we write $D \sim D'$, if D' can be obtained from D by a sequence of chip-firing moves.*

Lemma 1.6. *Linear equivalence of divisors is an equivalence relation.*

Lemma 1.7. *If $D_1 \sim D_2$ then, for any divisor E, we have $D_1 + E \sim D_2 + E$.*

Definition 1.8. *A divisor that is equivalent to 0 is called a* principal divisor*. We denote the set of principal divisors by* $\mathrm{Prin}(G)$*.*

The Picard group *of a graph G is the group of linear equivalence classes of divisors on G. That is,*

$$\mathrm{Pic}(G) = \mathrm{Div}(G)/\mathrm{Prin}(G).$$

1.2. **The Graph Laplacian.** To compute the Picard group of a graph G, it helps to have an algebraic description of the principal divisors. This is accomplished by way of the Laplacian matrix.

Definition 1.9. *The* graph Laplacian *of a graph G is the square matrix with rows and columns indexed by the vertices of G, and whose (i, j)th entry is*

$$\Delta_{i,j} = \begin{cases} \mathrm{val}(v_i) & \text{if } i = j \\ -\#\text{ of edges between } v_i \text{ and } v_j & \text{if } i \neq j. \end{cases}$$

That is, Δ is the difference of the valency matrix and the adjacency matrix.

Exercise 1.10. Compute the graph Laplacian of the graph pictured in Figure 3[1].

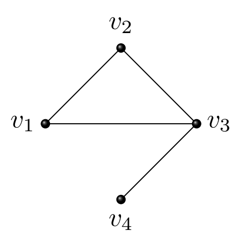

FIGURE 3. A simple graph.

[1]In the film "Good Will Hunting", the first of the two problems to appear on the blackboard is a four-parter, the first part of which is to compute the Laplacian of this graph.

The graph Laplacian defines a map:

$$\mathbb{Z}^{V(G)} \xrightarrow{\Delta} \mathbb{Z}^{V(G)} = \mathrm{Div}(G).$$

Theorem 1.11. *Let D be a divisor on a graph G. The following are equivalent:*

(1) *D is a principal divisor,*
(2) *$D \in \mathrm{Im}(\Delta)$, and*
(3) *There exists a function $f\colon V(G) \to \mathbb{Z}$ such that $D(v) = \sum_{e=vw}(f(w) - f(v))$ for all v, where the sum is over all edges with one endpoint v.*

Corollary 1.12. $\mathrm{Pic}(G) = \mathbb{Z}^{V(G)}/\,\mathrm{Im}(\Delta)$.

Exercise 1.13. Let G be the graph from Exercise 1.10.

(1) Compute the Smith normal form of the matrix Δ.
(2) Show that $\mathrm{Pic}(G) \cong \mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$.

1.3. **The degree of a divisor.** We now consider a fundamental invariant of divisors on graphs.

Definition 1.14. *The* degree *of a divisor $D = \sum_{v \in V(G)} D(v)v$ is the integer*

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

Exercise 1.15. The degree of a divisor is invariant under chip-firing.

Lemma 1.16. *The degree is a surjective group homomorphism from* $\mathrm{Pic}(G)$ *to* $\mathbb{Z}$.

Definition 1.17. *The* Jacobian $\mathrm{Jac}(G)$ *of a graph G is the group of linear equivalence classes of divisors of degree 0 on G.*

The Jacobian is also known as the *sandpile group* or *critical group*. Again, we use the term *Jacobian* to emphasize the connection with algebraic curves.

Proposition 1.18. *For any graph G, we have* $\mathrm{Pic}(G) \cong \mathbb{Z} \oplus \mathrm{Jac}(G)$.

For an integer d, let $\mathrm{Div}^d(G)$ denote the set of divisors on G of degree d, and let $\mathrm{Pic}^d(G)$ denote the set of equivalence classes of divisors on G of degree d.

Lemma 1.19. *The group* $\mathrm{Jac}(G)$ *acts freely and transitively on* $\mathrm{Pic}^d(G)$ *by addition. (In other words,* $\mathrm{Pic}^d(G)$ *is a* $\mathrm{Jac}(G)$*-torsor.)*

Exercise 1.20. Let T be a tree. Show that any two adjacent vertices of T are linearly equivalent. Conclude that *any* two vertices of T are linearly equivalent, and the Jacobian of T is trivial.

Exercise 1.21. Let G be a cycle with n vertices. Label the vertices of G counterclockwise with the elements of $\mathbb{Z}/n\mathbb{Z}$. Show that the map $\mathrm{Div}(G) \to \mathbb{Z}/n\mathbb{Z}$ given by

$$\sum_{i=0}^{n-1} a_i v_i \mapsto \sum_{i=0}^{n-1} a_i i \pmod{n}$$

is invariant under linear equivalence. Use this to prove that $\mathrm{Jac}(G) \cong \mathbb{Z}/n\mathbb{Z}$.

Exercise 1.22. Let G_1 and G_2 be graphs, and let G be the graph obtained by connecting a single vertex of G_1 to a single vertex of G_2 by an edge. Show that

$$\mathrm{Jac}(G) \cong \mathrm{Jac}(G_1) \times \mathrm{Jac}(G_2).$$

Corollary 1.23. *For any finite abelian group A, there is a graph G with $\mathrm{Jac}(G) \cong A$.*

1.4. **Finiteness of the graph Jacobian.** The following series of exercises shows that the Jacobian of a graph is a finite group. Later, in Corollary 2.23, we will compute the order of this group.

Lemma 1.24. *Let G be a graph, Δ its graph Laplacian, and $\mathbf{x} \in \mathbb{R}^{V(G)}$. Then*

$$\mathbf{x}^T \Delta \mathbf{x} = \sum_{i,j \in E(G)} (\mathbf{x}_i - \mathbf{x}_j)^2.$$

Corollary 1.25. *The graph Laplacian is a positive semi-definite matrix.*

Recall that we assume throughout that our graphs are connected.

Lemma 1.26. *Let G be a graph and Δ its Laplacian. The rank of Δ is $|V(G)| - 1$.*

Corollary 1.27. *The Jacobian of a graph is a finite group.*

The following corollary will be useful in the next subsection.

Corollary 1.28. *For all $D \in \mathrm{Div}^0(G)$, there exists $\mathbf{x} \in \mathbb{Q}^{V(G)}$ such that $\Delta \mathbf{x} = D$.*

1.5. **The Energy and Monodromy Pairings.** Lemma 1.26 shows that the graph Laplacian Δ is not invertible. However, every matrix has a generalized inverse.

Definition 1.29. *A* generalized inverse *of a matrix Δ is a matrix L such that*

$$\Delta L \Delta = \Delta.$$

The next lemma shows that a generalized inverse of the graph Laplacian exists.

Lemma 1.30. *Let G be a graph, Δ its graph Laplacian, and Δ_j the matrix obtained by deleting the jth row and jth column of Δ. The matrix Δ_j is invertible, and the matrix L_j obtained by adding a zero row and zero column to Δ_j^{-1} is a generalized inverse of Δ.*

Given a generalized inverse L of Δ, we define the *energy pairing*

$$\langle , \rangle \colon \mathrm{Div}^0(G) \times \mathrm{Div}^0(G) \to \mathbb{Q}$$

by

$$\langle D_1, D_2 \rangle = D_1^T L D_2.$$

Lemma 1.31. *The energy pairing is independent of the choice of generalized inverse.*

Lemma 1.32. *Let $D_1, D_2, E \in \mathrm{Div}^0(G)$. If $D_1 \sim D_2$, then*

$$\langle D_1, E \rangle = \langle D_2, E \rangle \pmod{\mathbb{Z}}.$$

By Lemma 1.32, the energy pairing descends to a pairing

$$\overline{\langle , \rangle} \colon \mathrm{Jac}(G) \times \mathrm{Jac}(G) \to \mathbb{Q}/\mathbb{Z}$$

called the *monodromy pairing*.

Lemma 1.33. *The monodromy pairing is nondegenerate. That is, if $\overline{\langle \cdot, D \rangle}$ is identically zero in $\mathbb{Q}/\mathbb{Z}$, then D is principal.*

1.6. **Further Reading.** Most of the material in this section is standard, and can be found in nearly any text on chip firing and graph Laplacians. We recommend [Kli19]. Section 1.5 borrows heavily from [Sho10], which introduces the energy and monodromy pairings.

2. Effective and Reduced Divisors

2.1. **Effective Divisors.** The main topic of this section is the theory of v-reduced divisors, which are canonical representatives of divisor classes on graphs, depending only on the choice of a base vertex v. We will prove the existence and uniqueness of v-reduced divisors, along with some of the fundamental properties that make them essential tools in this subject.

Definition 2.1. *A divisor $D = \sum_{v\in V(G)} D(v)v$ is* effective *if $D(v) \geq 0$ for all $v \in V(G)$.*

Not every divisor on a graph is equivalent to an effective divisor. For example, a divisor of negative degree cannot be effective. A divisor of degree 0 is effective if and only if it is the zero divisor. To determine whether a given divisor is equivalent to an effective divisor, we use the theory of v-reduced divisors. These are canonical representatives of divisor classes in $\mathrm{Pic}(G)$, depending only on the choice of a base vertex v.

Definition 2.2. *Let G be a graph and let $v \in V(G)$. A divisor $D = \sum_{w\in V(G)} D(w)w$ is* effective away from v *if $D(w) \geq 0$ for all $w \neq v$.*

Lemma 2.3. *Every divisor is equivalent to one that is effective away from v. (Hint: use Corollary 1.27.)*

2.2. **Reduced Divisors.** For a subset $A \subseteq V(G)$, let χ_A denote the characteristic function. Note that $-\Delta\chi_A$ is the divisor obtained from 0 by firing all vertices of A (see Exercise 1.4). The following is the key definition.

Definition 2.4. *A divisor D is v-*reduced *if it is effective away from v, and for any subset $A \subseteq V(G) \smallsetminus \{v\}$, $D - \Delta\chi_A$ is not effective away from v.*

The primary goal of this section is to prove that every divisor is equivalent to a unique v-reduced divisor. We break this into separate steps.

Lemma 2.5. *Let D and E be divisors on a graph G, both of degree zero. Let $A \subseteq V(G)$. If $E = D - \Delta\chi_A$, then*

$$\langle E, E\rangle = \langle D, D\rangle - \sum_{v\in A}(D + E)(v).$$

Lemma 2.6. *Let $v \in V(G)$, and let D be a divisor of degree zero that is effective away from v. If D is not v-reduced, then there exists $E \sim D$, also effective away from v, such that $\langle E, E\rangle < \langle D, D\rangle$.*

Proposition 2.7. *Let G be a graph, $v \in V(G)$. Every divisor is equivalent to a v-reduced divisor.*

Theorem 2.9 below shows that every divisor is equivalent to a *unique* v-reduced divisor. A useful tool for proving this is the following, which Baker and Shokrieh call the Maximum Principle.

Lemma 2.8. *Let $\mathbf{x} \in \mathbb{Z}^{V(G)}$, and let $A \subseteq V(G)$ be the set of vertices where $\mathbf{x}$ obtains its maximum. Then*

$$\Delta\mathbf{x}(v) \geq \mathrm{outdeg}_v(A)$$

for all $v \in A$.

Theorem 2.9. *Let G be a graph, $v \in V(G)$. Every divisor is equivalent to a* unique *v-reduced divisor.*

2.3. **Dhar's Burning Algorithm.** In the previous section, we saw that every divisor is equivalent to a unique v-reduced divisor. In this section, we describe a procedure for computing this v-reduced divisor. This is known as Dhar's Burning Algorithm.

Let G be a graph, $v \in V(G)$, and D a divisor on G that is effective away from v. Dhar's Burning Algorithm proceeds as follows:

(1) Start a fire at v.
(2) For all unburnt vertices $w \in V(G)$, if the number of burnt edges adjacent to w exceeds the number of chips $D(w)$ at w, burn w. If no vertex with this property exists, proceed to step (3).
(3) Burn every edge of the graph that is adjacent to a burnt vertex. If at least one edge burns at this step, return to step (2). Otherwise, proceed to step (4).
(4) Let A be the set of unburnt vertices. If A is non-empty, then $D + \Delta\chi_A$ is effective away from v. Replace D with $D - \Delta\chi_A$ and return to step (1). Otherwise, if A is empty, then D is v-reduced.

An example of Dhar's Burning Algorithm appears below.

Example 2.10. We run Dhar's Burning Algorithm to compute the v-reduced divisor equivalent to the divisor pictured in Figure 4, where v is the lower right vertex.

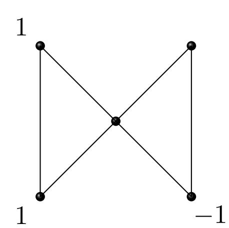

FIGURE 4. A divisor that is effective away from v

After burning vertices and edges until there are none left to burn, we see that the two vertices on the left remain unburnt. Therefore, the divisor pictured in Figure 5 is not v-reduced.

Firing the two vertices on the left, we obtain the divisor pictured in Figure 6.

Running a second iteration of Dhar's Burning algorithm, we see that the three left vertices are unburnt, as in Figure 7.

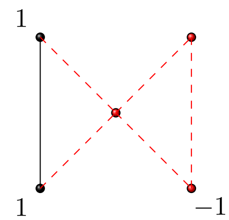

FIGURE 5. First iteration of Dhar's Burning Algorithm

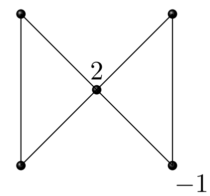

FIGURE 6. Result of firing the left two vertices

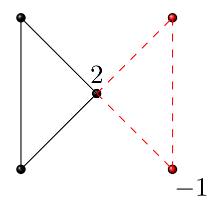

FIGURE 7. Second iteration of Dhar's Burning Algorithm

Firing these three vertices, we obtain the divisor pictured in Figure 8.

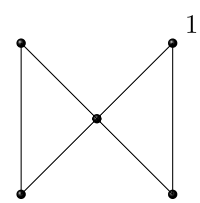

FIGURE 8. Result of firing the left three vertices

Running Dhar's Burning Algorithm a final time, we see that the whole graph burns. Therefore, the divisor depicted in Figure 8 is v-reduced.

Exercise 2.11. Let v be the top right vertex in Figure 9. Find the v-reduced divisor equivalent to the divisor depicted below.

Exercise 2.12. Prove that Dhar's Burning Algorithm terminates.

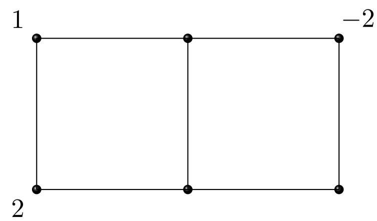

FIGURE 9

Exercise 2.13. Prove that Dhar's Burning Algorithm works. That is, in step (4), the divisor is v-reduced if and only if the set A of unburnt vertices is empty.

Proposition 2.14. *Let G be a graph and $v \in V(G)$. A divisor D is equivalent to an effective divisor if and only if its v-reduced representative is effective.*

Exercise 2.15. Is the divisor depicted in Figure 9 equivalent to an effective divisor? Why or why not?

2.4. **Cori-Le Borgne Algorithm.** In the previous section, we described Dhar's Burning Algorithm as burning multiple edges at once. In the Cori-Le Borgne Algorithm, we instead fix a total order on the set $E(G)$ of edges. Now we run Dhar's Burning Algorithm, but burn edges one at a time. Whenever a vertex is elligible to burn, it burns. Whenever there are multiple edges that are eligible to burn, we burn the smallest one first. Whenever a vertex w burns, we mark the edge along which the fire travled just prior to burning w.

Example 2.16. Consider the divisor pictured in Figure 10, with edges labeled from 1 to 6. If v is the bottom right vertex, then this divisor is v-reduced.

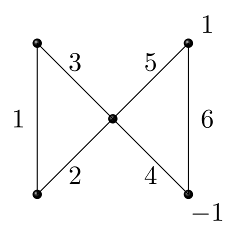

FIGURE 10. A v-reduced divisor of degree 0

Running the Cori-Le Borgne Algorithm, the edge 4 burns first, and then the central vertex, so we mark edge 4. The edge 2 burns next, and then the bottom left vertex, so we mark edge 2. The edge 1 burns next, and then the top left vertex, so we mark edge 1. The edge 3 burns, followed by edge 5, followed by edge 6. Note that the top right vertex does not burn until edge 6 does, so we mark edge 6. The resulting spanning tree is pictured in Figure 11.

Lemma 2.17. *Let D be a v-reduced divisor. The set of marked edges under the Cori-Le Borgne Algorithm is a spanning tree.*

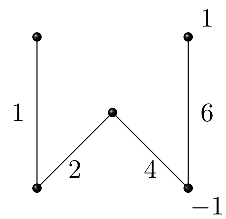

Figure 11. The corresponding spanning tree

The key insight of Cori and Le Borgne is that this procedure produces a bijection between v-reduced divisors of degree zero and spanning trees. To see this, we describe an inverse algorithm, which takes as input a spanning tree T, and returns a v-reduced divisor of degree zero. Cori and Le Borgne also describe this as a burning algorithm, but we find it helpful in our arguments to distinguish the edges and vertices that are burn in the Cori-Le Borgne algorithm from those that are burnt in the inverse algorithm. For this reason, we use the metaphor of "freezing" instead.

(1) Freeze the vertex v.
(2) At the moment that an unfrozen vertex w is adjacent to a frozen edge in T, set $D(w)$ to be one less than the number of frozen edges adjacent to w, and then freeze w. If all of the vertices are frozen, proceed to step (4).
(3) If no unfrozen vertex is adjacent to a frozen edge in T, freeze the smallest unfrozen edge that is adjacent to a frozen vertex.
(4) Once all the vertices are burnt, set $D(v) = -\sum_{w \neq v} D(w)$.

Example 2.18. Consider the spanning tree in Figure 11 from Example 2.16. Running the inverse Cori-LeBorgne Algorithm, the edges freeze in the order 4, 2, 1, 3, 5, 6. All vertices w but the top right are adjacent to only 1 frozen edge at the moment they freeze, so we set $D(w) = 0$. The top right vertex x does not freeze until edge 6 freezes, at which point it is adjacent to 2 frozen edges, so we set $D(x) = 1$. Finally, we set $D(v) = 0$, so that the total degree is zero.

Lemma 2.19. *Let T be a spanning tree in the graph G, and let D be the divisor produced by the inverse Cori-Le Borgne algorithm, starting with T. If we run the inverse Cori-Le Borgne algorithm starting with T, then the order that the vertices and edges of G are frozen is the same as the order in which they are burnt if we run the Cori-Le Borgne algorithm, starting with D.*

Lemma 2.20. *The output of the inverse Cori-Le Borgne Algorithm is a v-reduced divisor.*

Lemma 2.21. *Let D be a v-reduced divisor, and let T be the spanning tree produced by the Cori-Le Borgne algorithm, starting with D. If we run the inverse Cori-Le Borgne algorithm starting with T, then the order that the vertices and edges of G are frozen is the same as the order in which they are burnt if we run the Cori-Le Borgne algorithm, starting with D.*

Proposition 2.22. *The two algorithms defined by Cori and Le Borgne are inverses.*

Corollary 2.23. *For any graph G, $|\operatorname{Jac}(G)|$ is equal to the number of spanning trees in G.*

Exercise 2.24. Use Corollary 2.23 to give short solutions to Exercises 1.20 and 1.21.

2.5. **Further Reading.** Reduced divisors and Dhar's burning algorithm were introduced in [Dha90], but our treatment follows that of Baker and Shokrieh in [BS13]. The Cori-Le Borgne bijection can be found in [CLB03]. The "standard" proof of Corollary 2.23, that $|\operatorname{Jac}(G)|$ is equal to the number of spanning trees in G, uses the matrix tree theorem, first proved by Kirchhoff in [Kir47], though we do not take this approach here.

3. Orientable Divisors and Break Divisors

In this section, we introduce divisors corresponding to graph orientations. This was an essential ingredient in the Baker-Norine proof of the Riemann Roch theorem for graphs. Readers who wish to skip the proof of Riemann-Roch may choose to read Subsection 3.1 and then skip the rest of this section.

3.1. **The Canonical Divisor.**

Definition 3.1. *The* genus *of a graph G is*

$$g = |E(G)| - |V(G)| + 1.$$

The genus of a graph is its first Betti number. In other words, it is the rank of $H^1(G, \mathbb{Z})$. We use the term genus to emphasize the analogy between divisors on graphs and divisors on algebraic curves. This should not be confused with the other common graph invariant known as the genus, which is the minimal genus of a surface in which the graph can be embedded without crossings.

Another invariant of a graph is its canonical divisor, defined as follows.

Definition 3.2. *The canonical divisor of a graph G is the divisor*

$$K_G = \sum_{v \in V(G)} (\operatorname{val}(v) - 2)v.$$

The degree of the canonical divisor is computed in Corollary 3.6 below. While this computation follows directly from the theory of orientable divisors, this theory is not necessary to prove Corollary 3.6.

3.2. **Orientable Divisors.** Much of the theory we have developed concerning divisors on graphs can also be formulated in terms of graph orientations. The connection between divisors and graph orientations begins with orientable divisors.

Definition 3.3. *Let G be a graph and $\mathcal{O}$ an orientation of G. The corresponding* orientable divisor *is*

$$D_{\mathcal{O}} := \sum_{v \in V(G)} (\operatorname{indeg}_{\mathcal{O}}(v) - 1)v.$$

Lemma 3.4. *If $\mathcal{O}$ is an orientation and $\overline{\mathcal{O}}$ is the reverse orientation, then*

$$D_{\mathcal{O}} + D_{\overline{\mathcal{O}}} = K_G.$$

Another simple fact about orientable divisors is that they all have the same degree.

Lemma 3.5. *Let G be a graph of genus g. Every orientable divisor on G has degree $g-1$.*

Corollary 3.6. *If G is a graph of genus g, then $\deg(K_G) = 2g-2$.*

A key connection between chip firing and graph orientations is the following observation.

Proposition 3.7. *Let G be a graph and $\mathcal{O}$ an orientation of G. Let $A \subseteq V(G)$ be a subset of the vertices with the property that all edges in the cut (A, A^c) are directed toward A, and let $\mathcal{O}'$ be the orientation obtained from $\mathcal{O}$ by reversing this directed cut. Then*

$$D_{\mathcal{O}'} = D_{\mathcal{O}} - \Delta\chi_A.$$

In particular, $D_{\mathcal{O}'}$ is equivalent to $D_{\mathcal{O}}$.

Definition 3.8. *Let G be a graph and $v \in V(G)$. An orientation $\mathcal{O}$ of G is called* v-connected *if, for every vertex w in G, there is a directed path in $\mathcal{O}$ from v to w.*

Corollary 3.9. *Let G be a graph and $\mathcal{O}$ an orientation of G. There exists a v-connected orientation $\mathcal{O}'$ such that $D_{\mathcal{O}} \sim D_{\mathcal{O}'}$. (Hint: let $A \subseteq V(G)$ be the set of vertices that can be reached from v by a directed path in $\mathcal{O}$, and use Proposition 3.7.)*

Our next goal is to show that every divisor of degree $g-1$ on a graph of genus g is equivalent to an orientable divisor. To see this, we present an algorithm that starts with a divisor D of degree $g-1$ and an orientation $\mathcal{O}$. At each step, it either modifies the orientation $\mathcal{O}$ or replaces D with an equivalent divisor. When the algorithm terminates, we will have $D = D_{\mathcal{O}}$. The algorithm proceeds as follows:

(1) Define

$$A^+ = \{v \in V(G) \mid D(v) > D_{\mathcal{O}}(v)\}$$
$$A^- = \{v \in V(G) \mid D(v) < D_{\mathcal{O}}(v)\}$$
$$A = \{v \in V(G) \mid \text{ there exists a directed path from a vertex in } A^+ \text{ to } v\}.$$

(2) If $A^+ = \emptyset$, then $D = D_{\mathcal{O}}$ and the algorithm terminates.
(3) Otherwise, if $A \cap A^- \neq \emptyset$, then there is an oriented path in $\mathcal{O}$ from a vertex $v \in A^+$ to a vertex $w \in A^-$. Let $\mathcal{O}'$ be the orientation obtained from $\mathcal{O}$ by reversing this directed path, and replace $\mathcal{O}$ with $\mathcal{O}'$. Return to step (1).
(4) Otherwise, if $A \cap A^- = \emptyset$, then by definition, the cut (A, A^c) is directed away from A. Let $\mathcal{O}'$ be the orientation obtained by reversing this directed cut. Replace $\mathcal{O}$ with $\mathcal{O}'$ and D with $D - \Delta\chi_A$. Return to step (1).

Theorem 3.10. *This algorithm terminates. As a consequence, if G is a graph of genus g, then every divisor on G of degree $g-1$ is equivalent to an orientable divisor.*

Theorem 3.10 has the following interesting consequence.

Corollary 3.11. *Let G be a graph of genus g. Every divisor of degree at least g on G is equivalent to an effective divisor.*

The bound in Corollary 3.11 is sharp. In other words, on any graph of genus g, there exist divisors of degree $g-1$ that are not equivalent to an effective divisor.

Lemma 3.12. *Let $\mathcal{O}$ be an acyclic orientation of a graph G. Then $D_{\mathcal{O}}$ is not equivalent to an effective divisor. (Hint: use Lemma 2.8.)*

Corollary 3.13. *Let G be a graph of genus g. There exists a divisor of degree $g-1$ on G that is not equivalent to an effective divisor.*

In fact, acyclic orientations can be used to distinguish between divisors that are equivalent to effective divisors and those that are not.

Lemma 3.14. *Let G be a graph, $v \in V(G)$, and let D be a v-reduced divisor. Define an orientation $\mathcal{O}$ on G by running Dhar's Burning Algorithm and orienting each edge from the first endpoint that burns to the second endpoint that burns. Then $D(w) \leq D_{\mathcal{O}}(w)$ for all $w \neq v$.*

Corollary 3.15. *For any divisor D on a graph G, either D is equivalent to an effective divisor or there is an acyclic orientation $\mathcal{O}$ such that $D_{\mathcal{O}} - D$ is equivalent to an effective divisor (but not both).*

Corollary 3.16. *Let G be a graph of genus g. A divisor D on G of degree $g-1$ is equivalent to an effective divisor if and only if $K_G - D$ is equivalent to an effective divisor.*

By Theorem 3.10, every divisor of degree $g-1$ on a graph of genus g is equivalent to a v-connected orientable divisor. We will now show that this v-connected orientable divisor is unique.

Theorem 3.17. *Let $\mathcal{O}$ and $\mathcal{O}'$ be orientations of a graph G, and let v be a vertex of G. If $D_{\mathcal{O}}$ and $D_{\mathcal{O}'}$ are equivalent but not equal, then at most one of them is v-connected. (Hint: let $D_{\mathcal{O}'} = D_{\mathcal{O}} + \Delta\mathbf{x}$, let A be the set where $\mathbf{x}$ achieves its maximum, and consider $\deg(D_{\mathcal{O}'}|_A)$.)*

3.3. Break Divisors.

In previous sections, we have studied canonical representatives of divisor classes on graphs. More specifically, given a vertex v on a graph of genus g, every divisor class contains a unique v-reduced representative, and every divisor class of degree $g-1$ contains a unique v-connected orientable representative. A drawback of these theories is that these representatives depend on the choice of vertex v. In this lecture, we discuss canonical representatives of degree g divisors that do not depend on any choices.

Definition 3.18. *Let G be a graph, and let T be a spanning tree of G. For each edge e not in T, let v_e be one of its endpoints. A divisor of the form*

$$D = \sum_{e \notin T} v_e$$

is called a break divisor.

By definition, a break divisor is effective of degree g.

Exercise 3.19. Find all break divisors in the graph pictured in Figure 12.

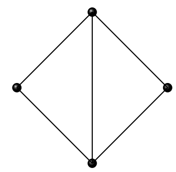

Figure 12. A graph of genus 2

There is a direct connection between break divisors and orientable divisors.

Proposition 3.20. *Let G be a graph and $v \in V(G)$. A divisor D on G is a break divisor if and only if $D - v$ is a v-connected orientable divisor.*

Corollary 3.21. *Let G be a graph of genus g. Every divisor of degree g on G is equivalent to a unique break divisor.*

3.4. **Further Reading.** Most of the material on orientable divisors can be found in Baker and Norine's proof of the Riemann-Roch theorem for graphs [BN07]. A much more extensive treatment of this theory can be found in [Bac17]. The theory of break divisors is introduced in [MZ08] and developed further in [ABKS14].

4. The Rank of a Divisor

4.1. **Basics on Rank.** A key invariant of a divisor is its (Baker-Norine) *rank*. It is this invariant that powers the connection between chip firing and algebraic geometry.

Definition 4.1. *Let D be a divisor on a graph. If D is not equivalent to an effective divisor, we say that D has* rank -1*. Otherwise, we define the* rank *of D to be the largest integer r such that, for all effective divisors E of degree r, $D - E$ is equivalent to an effective divisor.*

Computing the rank of a divisor can be thought of as a game, in which our opponent is allowed to "steal" r chips from wherever they like, and our task is to perform a sequence of chip firing moves that eliminates the debt created by our opponent. If we can win this game regardless of which r chips our opponent chooses to steal, then the divisor has rank at least r.

Exercise 4.2. A divisor has nonnegative rank if and only if it is equivalent to an effective divisor.

Exercise 4.3. Compute the ranks of the two divisors of degree 2 depicted in Figure 13.

Exercise 4.4. Let T be a tree, and let D be a divisor on T of nonnegative degree. Show that $\mathrm{rk}(D) = \deg(D)$.

Exercise 4.5. Let G be a cycle with n vertices, and let D be a divisor on G of positive degree. Show that $\mathrm{rk}(D) = \deg(D) - 1$.

We record a few other observations about ranks of divisors.

Lemma 4.6. *Let D be a divisor on a graph of genus g. Then $\mathrm{rk}(D) \geq \deg(D) - g$.*

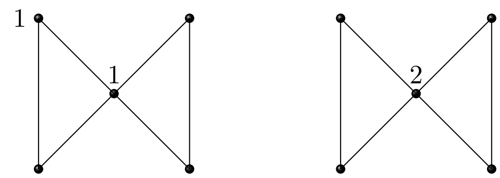

Figure 13. Two divisors of the same degree on a graph of genus 2

Lemma 4.7. *Let D_1, D_2 be divisors of nonnegative rank on a graph G. Then*

$$\mathrm{rk}(D_1 + D_2) \geq \mathrm{rk}(D_1) + \mathrm{rk}(D_2).$$

Exercise 4.8. Let G be a simple bipartite graph, and let D be the sum of the vertices of a single color. Show that $\mathrm{rk}(D) \geq 1$.

Lemma 4.9. *A divisor D has rank at least 1 if and only if, for all $v \in V(G)$, the v-reduced divisor D_v equivalent to D satisfies $D_v(v) \geq 1$.*

4.2. **Riemann-Roch.** In this section, we prove possibly the most important result about ranks of divisors on graphs, the Riemann-Roch Theorem. While we have developed all the necessary background to prove this theorem, the results in this section are nevertheless still quite difficult. We begin with the following definition.

Definition 4.10. *Let D be a divisor on a graph G. We define*

$$\deg^+(D) = \sum_{v \in V(G), D(v) > 0} D(v).$$

Baker and Norine give an alternate characterization of the rank.

Proposition 4.11. *Let D be a divisor on a graph G. Then*

$$\mathrm{rk}(D) = \min_{\substack{D' \sim D \\ \mathcal{O} \text{ acyclic}}} \{\deg^+(D' - D_{\mathcal{O}})\} - 1.$$

(Hint: prove two inequalities, one using Lemma 3.12 and the other using Lemma 3.14.)

Theorem 4.12 (Riemann-Roch Theorem)**.** *Let D be a divisor on a graph G of genus g. Then*

$$\mathrm{rk}(D) - \mathrm{rk}(K_G - D) = \deg(D) - g + 1.$$

4.3. **Consequences of Riemann-Roch.** We now explore some consequences of the Riemann-Roch theorem.

Corollary 4.13. *Let D be a divisor on a graph of genus g. If $\deg(D) > 2g - 2$, then $\mathrm{rk}(D) = \deg(D) - g$.*

By Corollary 4.13, if a divisor has large degree, then its rank is completely determined by its degree. Similarly, if a divisor has negative degree, then it has rank -1. It follows that, on a given graph G, there are only finitely many divisors whose rank is not determined by their degree. In the edge cases, when the degree of a divisor is 0 or $2g - 2$, there are two possibilities.

Corollary 4.14. *Let D be a divisor on a graph of genus g. If $\deg(D) = 2g-2$, then*

$$\mathrm{rk}(D) = \begin{cases} g-1 & \text{if } D \sim K_G \\ g-2 & \text{otherwise.} \end{cases}$$

Our next consequence of Riemann-Roch is usually referred to as the Clifford bound.

Theorem 4.15. *Let D be a divisor on a graph G, and suppose that both D and $K_G - D$ have nonnegative rank. Then*

$$\mathrm{rk}(D) \leq \frac{1}{2}\deg(D).$$

Exercise 4.16. Let G be a graph of genus g, and let D be a divisor on G of degree 2 and rank 1. Compute the rank of mD for all integers m. Conclude that $K_G \sim (g-1)D$.

As mentioned above, there are only finitely many divisors on a given graph G whose rank is not determined by their degree. For each such divisor, we see that

$$\max\{-1, \deg(D) - g\} \leq \mathrm{rk}(D) \leq \frac{1}{2}\deg(D).$$

There are therefore only finitely many possibilities for the rank of such a divisor. The possible pairs (d, r), where d and r are the degree and rank, respectively, of a divisor, are illustrated in Figure 14.

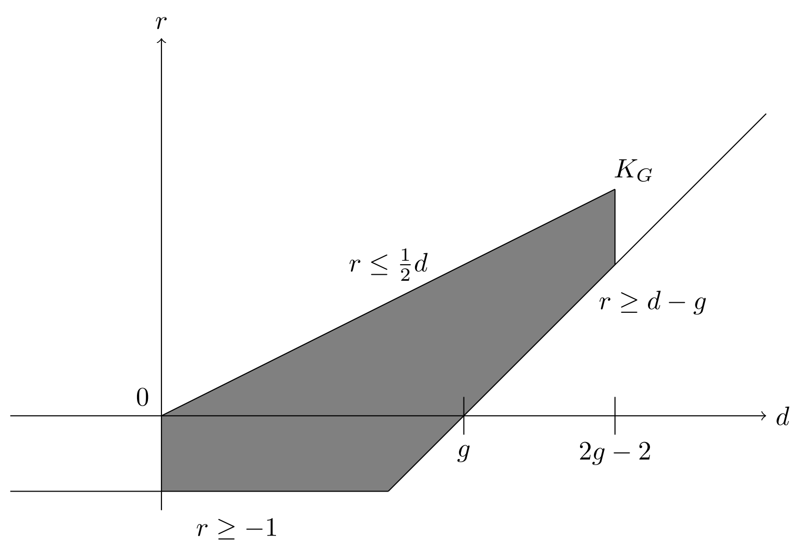

FIGURE 14. Possibilities for the degree and rank of a divisor.

Of particular interest are the divisors that have larger than expected rank. These divisors correspond to lattice points in Figure 14 that lie above the lower bound.

Definition 4.17. *A divisor D on a graph of genus g is called* special *if*

$$\mathrm{rk}(D) > \max\{-1, \deg(D) - g\}.$$

Exercise 4.18. Use Riemann-Roch to give short proofs of Exercises 4.4 and 4.5.

4.4. **Gonality and Scrambles.** One of the most studied graph invariants in this area is the *gonality* of a graph.

Definition 4.19. *Then* gonality *of a graph G is the minimum integer d such that there exists a divisor on G of degree d and positive rank.*

Lemma 4.20. *A graph G has gonality 1 if and only if it is a tree.*

Exercise 4.21. Find the gonality of every graph of genus 0, 1, and 2.

For graphs of genus greater than 2, the gonality is not determined by the genus.

Exercise 4.22. Find two graphs, each of genus 3, one of which has gonality 2 and the other of which does not.

Exercise 4.23. The complete graph K_n has gonality $n-1$.

Exercise 4.24. The complete bipartite graph $K_{m,n}$ has gonality $\min\{m,n\}$.

To bound the gonality from above, it suffices to construct a single divisor of positive rank. In practice, it is typically much more difficult to bound the gonality from below. For this reason, much of the literature on gonality is focused on computing lower bounds. Here, we present one such lower bound, known as the *scramble number*.

A *scramble* on a graph G is a collection $\mathscr{S} = \{E_1, \dots, E_n\}$ of connected subsets of $V(G)$, called *eggs*. A set $C \subseteq V(G)$ is a *hitting set* for the scramble $\mathscr{S}$ if $C \cap E \neq \emptyset$ for all eggs $E \in \mathscr{S}$. We define the *hitting number* $h(\mathscr{S})$ to be the minimum size of a hitting set for $\mathscr{S}$.

A set $F \subseteq E(G)$ is an *egg cut* for the scramble $\mathscr{S}$ if, when the edges in F are deleted, the graph G is disconnected into two connected components, each containing an egg. We define the *egg cut number* $e(\mathscr{S})$ to be the minum size of an egg cut for $\mathscr{S}$, if an egg cut exists, and to be infinity otherwise. The *order* of a scramble $\mathscr{S}$ is defined to be $\|\mathscr{S}\| = \min\{h(\mathscr{S}), e(\mathscr{S})\}$.

Exercise 4.25. On the complete graph K_n, consider the scramble where every vertex is an egg. What is the order of this scramble?

Exercise 4.26. On the complete bipartite graph $K_{m,n}$, consider the scramble where every vertex is an egg. What is the order of this scramble?

Exercise 4.27. Let G be the grid graph with m rows and n columns, pictured in Figure 15, and consider the scramble where every column in an egg and every row is an egg. What is the order of this scramble?

The *scramble number* of a graph G is the maximum order of a scramble on G. To see that the scramble number is a lower bound on the gonality, we first need the following result.

Proposition 4.28. *Let D, D_0 be equivalent effective divisors on a graph G. Then there exists a chain of subsets $A_1 \subseteq A_2 \subseteq \cdots \subseteq A_k$ such that $D_i := D_{i-1} - \Delta\chi_{A_i}$ is effective for all i, and $D_k = D$. (Hint: by definition, there exists an $x \in \mathbb{Z}^{V(G)}$ such that $D = D_0 - \Delta x$. Consider the level sets of x.)*

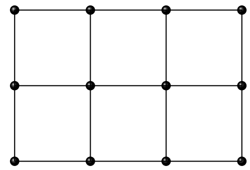

Figure 15. The grid graph with 3 rows and 4 columns.

Theorem 4.29. *The gonality of a graph G is bounded below by the scramble number.*

Exercise 4.30. Find the gonality of the grid graph with m rows and n columns.

Exercise 4.31. Find the gonality of each of the two graphs depicted in Figure 16.

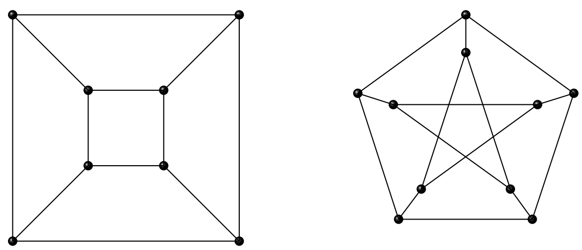

Figure 16. The cube graph and the Petersen graph

If $e \in E(G)$ is an edge with endpoints v and w, the *edge refinement* of G at e is obtained by introducing a new vertex in the middle of e. More precisely, the vertex set of the edge refinement is $V(G) \cup \{u\}$, and the edge set is $E(G) \cup \{uv, uw\} \smallsetminus \{e\}$. A *refinement* of a graph G is a graph that can be obtained from G by finitely many edge refinements.

Exercise 4.32. Every refinement of the complete graph K_n has gonality $n-1$.

In fact, the scramble number of any graph is invariant under refinement, but we will not prove this here.

4.5. **Ramification.**

Definition 4.33. *Let D be a divisor of rank r on a graph G, and let v be a vertex of G. The sequence*

$$a_0 < a_1 < \cdots < a_r$$

defined by

$$a_i := \max\{m | \operatorname{rk}(D - mv) \geq r - i\}$$

is called the ramification sequence *of D at v. We say that v is a* ramification point *of D if the ramification sequence of D at v is anything other than $0 < 1 < \cdots < r$.*

Exercise 4.34. Show that a divisor on a tree has no ramification points.

Exercise 4.35. Let D be a divisor of degree $d > 0$ on a cycle, and let v be a vertex of the cycle. Show that v is a ramification point of D if and only if $D \sim dv$.

Exercise 4.36. In Figure 17, which of the vertices v and w is a ramification point of the canonical divisor? Ramification points of the canonical divisor are typically referred to as *Weierstrass points.*

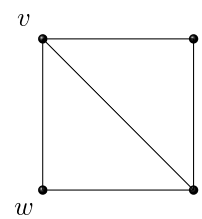

FIGURE 17. A graph of genus 2.

It is traditional to express the ramification sequence using partitions. In what follows, we identify the boxes in the Ferrers diagram of a partition with lattice points in $\mathbb{Z}^2_{>0}$.

Definition 4.37. *Let D be a divisor on a graph G of genus g, and let v be a vertex of G. We define the* Weierstrass partition *of D at v to be the partition*

$$\lambda_{G,v}(D) := \{(r+1, g-d+r) | \operatorname{rk}(D - (\deg(D) - d)v) \geq r\}.$$

Lemma 4.38. *The Weierstrass partition is a partition. That is, if $(x, y) \in \lambda_{G,v}(D)$, then:*

(1) *either $x = 1$ or $(x-1, y) \in \lambda_{G,v}(D)$, and*
(2) *either $y = 1$ or $(x, y-1) \in \lambda_{G,v}(D)$.*

We now record several other simple facts about Weierstrass partitions.

Lemma 4.39. *Let G be a graph of genus g and let v be any vertex of G. A divisor D on G has rank at least r if and only if*

$$(r+1, g - \deg(D) + r) \in \lambda_{G,v}(D).$$

One nice aspect of this definition is that the Weierstrass partition is invariant under addition of the vertex v.

Lemma 4.40. *Let G be a graph and let D be a divisor on G. For any vertex v of G, we have $\lambda_{G,v}(D) = \lambda_{G,v}(D+v)$.*

The term $g - d + r$ in the definition of the Weierstrass partition may appear mysterious, but it is motivated by Riemann-Roch. There is a natural involution on the set of partitions given by the transpose. There is also a natural involution on the set of divisors given by mapping a divisor D to $K_G - D$. The Weierstrass partition is defined so that these two involutions agree.

Proposition 4.41. *Let G be a graph and v a vertex of G. For any divisor D on G, we have $\lambda_{G,v}(K_G - D) = \lambda^T_{G,v}(D)$.*

Exercise 4.42. Let G be the graph in Figure 17. Compute the Weierstrass partitions $\lambda_{G,w}(K_G)$ and $\lambda_{G,v}(K_G)$.

4.6. **Extended Example: A Chain of Loops.** In this section, we compute Weierstrass partitions (and therefore ranks) of divisors on a certain family of graphs. To begin, we consider the graph G pictured in Figure 18. Specifically, we let G' be a graph of genus $g-1$, and v a vertex of G. We let C be a cycle with m vertices, labeled counterclockwise by $v_0, \ldots, v_{m-1}$. We let G be the graph obtained by connecting the vertex v of G' to the vertex v_0 of C. Our goal is to compute the ranks of divisors on this graph G.

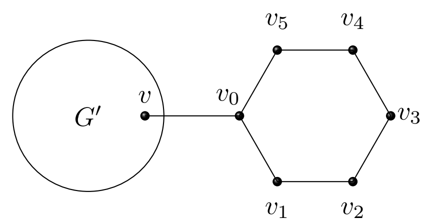

FIGURE 18. A graph with an attached cycle.

Lemma 4.43. *Let D be a divisor on G'. Then $\mathrm{rk}_G(D+v_i) \geq r$ if and only if:*

(1) $\mathrm{rk}_{G'}(D) \geq r$ *when* $i \neq 0$, *or*
(2) $\mathrm{rk}_{G'}(D+v) \geq r$ *and* $\mathrm{rk}_{G'}(D-v) \geq r-1$, *when* $i=0$.

We now translate Lemma 4.43 into the language of Weierstrass partitions.

Proposition 4.44. *Let D be a divisor on G'. Then $\lambda_{G',v}(D) \subseteq \lambda_{G,v_j}(D+v_i)$. Moreover, a box $(x,y) \notin \lambda_{G',v}(D)$ is contained in the Weierstrass partition $\lambda_{G,v_j}(D+v_i)$ if and only if the following conditions hold:*

(1) *either* $x=1$ *or* $(x-1,y) \in \lambda_{G',v}(D)$,
(2) *either* $y=1$ *or* $(x,y-1) \in \lambda_{G',v}(D)$, *and*
(3) $i \equiv (\deg(D) - g - x + y + 2)j \pmod{m}$.

We now use Proposition 4.44 to compute the ranks of divisors on the graph pictured in Figure 19. These graphs were studied heavily by Cools, Draisma, Payne and Robeva in their tropical proof of the Brill-Noether theorem. They were later studied in more depth by Pflueger, and our analysis follows his closely. As a warmup, we have the following exercises.

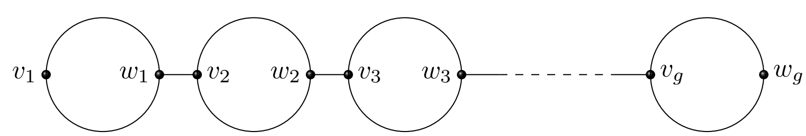

FIGURE 19. A chain of loops.

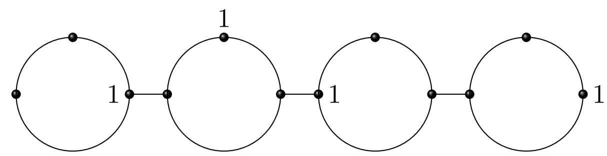

FIGURE 20. A divisor of degree 4 on a chain of 4 loops.

Exercise 4.45. Use Proposition 4.44 to compute the Weierstrass partition of the divisor pictured in Figure 20.

Exercise 4.46. Compute the rank of the divisor pictured in Figure 21.

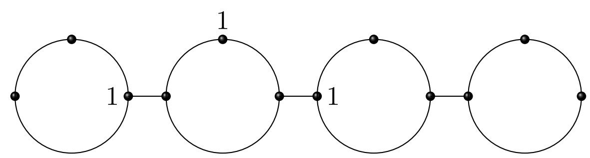

FIGURE 21. A divisor of degree 3 on a chain of 4 loops.

Returning to the general case, we assume that the bottom part of each cycle is a single edge, while the top part of the kth cycle consists of $m_k - 1$ edges. (So the total number of edges in the kth cycle is m_k.) We define $\vec{m} = (m_1, \ldots, m_g)$, and we refer to this graph as the chain of g loops with torsion profile $\vec{m}$.

Let D be a divisor on this graph. By Lemma 4.40, we have

$$\lambda_{G,w_g}(D) = \lambda_{G,w_g}(D + (g - \deg(D))w_g),$$

so we may reduce to the case that D has degree g. By Corollary 3.21, every divisor of degree g is equivalent to a unique break divisor, so we may reduce to the case that D is a break divisor. In other words, D has exactly 1 "chip" on each cycle of G. That is, the restriction of D to any individual cycle in G has degree 1.

Let G_k be the union of the first k cycles of G, and for ease of notation, let

$$\lambda_k = \lambda_{G_k,w_k}(D|_{G_k}).$$

By Proposition 4.44, we have

$$\emptyset = \lambda_0 \subseteq \lambda_1 \subseteq \lambda_2 \subseteq \cdots \subseteq \lambda_g.$$

Moreover, a box $(x, y) \notin \lambda_{k-1}$ is contained in λ_k if and only if:

(1) either $x = 1$ or $(x - 1, y) \in \lambda_{k-1}$,
(2) either $y = 1$ or $(x, y - 1) \in \lambda_{k-1}$, and
(3) the distance from w_k to the chip of D on γ_k, in the counterclockwise direction, is equivalent to $y - x \pmod{m_k}$.

This sequence of partitions defines a tableau t on the partition λ_g, defined by

$$t(x, y) = k \text{ if } (x, y) \in \lambda_k \smallsetminus \lambda_{k-1}.$$

This tableau has the property that, if $t(x,y) = t(x',y') = k$, then $y - x \equiv y' - x' \pmod{m_k}$. Equivalently, the lattice distance between the boxes (x,y) and (x',y') is divisible by m_k. We say that a tableau with this property is an *$\vec{m}$-displacement tableau*.

Conversely, given an $\vec{m}$-displacement tableau t, we define

$$\lambda_k = \{(x,y)|t(x,y) \leq k\}.$$

We may then construct a break divisor D such that $\lambda_k \subseteq \lambda_{G_k,w_k}(D|_{G_k})$ for all k, as follows. If $t(x,y) = k$, then we place a chip on the kth loop, at a distance of $y - x \pmod{m_k}$ from w_k, in the counterclockwise direction. This is well-defined by the definition of $\vec{m}$-displacement tableaux. If the symbol k does not appear in the tableau t, then we place a chip at any vertex of the kth loop.

Lemma 4.47. *There exists a divisor of degree d and rank at least r on the chain of g loops with torsion profile $\vec{m}$ if and only if there exists an $\vec{m}$-displacement tableau with alphabet $\{1,\ldots,g\}$ on the rectangular partition with $r+1$ columns and $g-d+r$ rows.*

Corollary 4.48. *Suppose that $m_k > g$ for all k. Then there exists a divisor of degree d and rank r on the chain of g loops with torsion profile $\vec{m}$ if and only if*

$$g - (r+1)(g-d+r) \geq 0.$$

The number $\rho(g,r,d) \colon = g - (r+1)(g-d+r)$ in Corollary 4.48 is known as the Brill-Noether number. The fact that these graphs possess a divisor of degree d and rank r if and only if the Brill-Noether number is nonnegative is a key step in the tropical proof of the Brill-Noether theorem.

Exercise 4.49. Let G be a chain of 4 loops, and suppose that $m_k \neq 2$ for all k. Show that there are exactly two divisor classes of degree 3 and rank 1 on G.

Exercise 4.50. Let G be a chain of 4 loops. Show that G has a divisor of degree 2 and rank 1 if and only if $m_2 = m_3 = 2$.

Exercise 4.51. Let G be a chain of 6 loops, and suppose that $m_k > 3$ for all k. Show that there are exactly five divisor classes of degree 4 and rank 1 on G.

Exercise 4.52. Let G be a chain of $2d-2$ loops, and suppose that $m_k \geq d$ for all k. Show that the number of divisor classes of degree d and rank 1 on G is equal to the Catalan number C_{d-1}.

4.7. **Further Reading.** The rank of a divisor is defined in [Bak08], which also introduces the gonality of a graph, and the Riemann-Roch theorem for graphs is proved in [BN07]. Scrambles were first introduced in [HJJS22], where it is proved that the scramble number is a lower bound on gonality. The argument in that paper is based heavily on [vDdBG20], which shows that the bramble number is a lower bound on graph gonality. (Indeed, one of the two reasons for the naming convention is that *scramble* sounds like *bramble*.) Weierstrass partitions are introduced in [Pfl17]. Divisors on chains of loops were first studied in [CDPR12], where they were used to give a tropical proof of the Brill-Noether theorem. We follow the approach of [Pfl17]. In the years that followed, chains of loops have been used in several applications of chip firing to algebraic geometry, too numerous to list here.

5. Metric Graphs

5.1. **Divisors on Metric Graphs.** We now turn from discrete graphs to metric graphs.

Definition 5.1. *A* metric graph *is a compact, connected metric space* Γ *obtained by identifying the edges of a graph* G *with line segments of fixed positive real length. The graph* G *is called a* model *for* Γ

Example 5.2. If we assign lengths to the edges of a cycle, we obtain a circle. Thus, the circle is a metric graph.

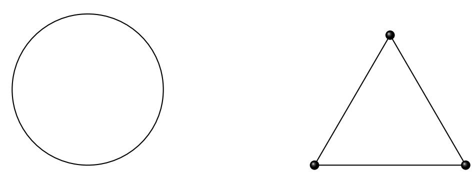

Figure 22. A metric graph and one of its models

A metric graph Γ does not have a unique model. Two graphs are models for the same metric graph if and only if they admit a common refinement.

Definition 5.3. *The* divisor group $\mathrm{Div}(\Gamma)$ *of a metric graph* Γ *is the free abelian group on points of the metric space* Γ*.*

Note that a divisor on a metric graph Γ is supported on *points* of Γ, not *vertices* of a model. In particular, a divisor can have a chip at any point on the interior of an edge. Many properties of divisors can be defined in a way that is completely analogous to the discrete graph case.

Definition 5.4. *A divisor* $D = \sum a_i v_i$ *on a metric graph is* effective *if* $a_i \geq 0$ *for all* i*. Its* degree *is defined to be*

$$\deg(D) := \sum a_i.$$

As in the case of discrete graphs, we want to talk about equivalence of divisors. For this, we need a notion of rational functions on metric graphs.

Definition 5.5. *A* rational function *on a metric graph* Γ *is a continuous, piecewise linear function* $\varphi : \Gamma \to \mathbb{R}$ *with integer slopes. We write* $\mathrm{PL}(\Gamma)$ *for the group of rational functions on* Γ*.*

Let $v, w \in \Gamma$ be contained in an edge of a model for Γ, and let $\ell(v, w)$ denote the length of the interval from v to w. If $\varphi \in \mathrm{PL}(\Gamma)$ is linear on the interval from v to w, then its slope along this interval is $(\varphi(w) - \varphi(v))/\ell(v, w)$. By definition, the slope of φ in the opposite direction – that is, on the interval from w to v – is the negative of its slope on the interval from v to w. Thus, when discussing the slope of a function in $\mathrm{PL}(\Gamma)$, it is necessary to identify the direction in which the slope is measured.

Example 5.6. Figure 23 indicates the domains of linearity and slopes of a rational function φ on a circle. It therefore determines the rational function up to translation. Note that, in order for the function to be continuous, the two regions on which the function has slope 1 must be of equal length.

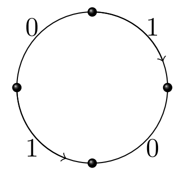

FIGURE 23. The domains of linearity and slopes of φ

Definition 5.7. *Given $\varphi \in \mathrm{PL}(\Gamma)$ and $v \in \Gamma$, we define the* order of vanishing *of φ at v, $\mathrm{ord}_v(\varphi)$, to be the sum of the incoming slopes of φ at v. Note that $\mathrm{ord}_v(\varphi)$ is nonzero for only finitely many points $v \in \Gamma$. We define the divisor associated to φ to be*

$$\mathrm{div}(\varphi) = \sum_{v \in \Gamma} \mathrm{ord}_v(\varphi) \cdot v.$$

Divisors of the form $\mathrm{div}(\varphi)$ are called principal.

Exercise 5.8. Let φ be the rational function of Example 5.6. Find $\mathrm{div}(\varphi)$.

Note that $\mathrm{div}(\varphi)$ is equal to $\mathrm{div}(\varphi + c)$ for any real number c. This is analogous to the fact that, on an algebraic curve, the divisor associated to a rational function is invariant under scaling the function by a non-zero constant. Indeed, we have the following.

Lemma 5.9. *Let Γ be a metric graph and let $\varphi, \psi \in \mathrm{PL}(\Gamma)$. We have $\mathrm{div}(\varphi) = \mathrm{div}(\psi)$ if and only if there exists a constant c such that $\varphi = \psi + c$.*

Example 5.10. On any metric graph Γ, given a point $v \in \Gamma$, let $\epsilon \in \mathbb{R}$ be sufficiently small so that the open ball $B_\epsilon(v)$ contains no points of valence greater than 2 other than possibly v. Let χ be the rational function that takes the value ϵ on $\Gamma \smallsetminus B_\epsilon(v)$, the value 0 at v, and has slope 1 on the edges in $B_\epsilon(v)$ emanating from v. Then χ has order of vanishing $-\mathrm{val}(v)$ at v and 1 at each of the boundary points of $B_\epsilon(v)$. In this way, we can view addition of $\mathrm{div}(\chi)$ as a continuous version of chip firing, where we specify not only the vertex v that we fire from, but also the distance ϵ that we fire the chips.

We note the following.

Lemma 5.11. *The degree of a principal divisor is zero.*

Lemma 5.12. *The map $\mathrm{div}\colon \mathrm{PL}(\Gamma) \to \mathrm{Div}(\Gamma)$ is a group homomorphism.*

Now that we have a notion of principal divisors on metric graphs, we can use it to define equivalence of divisors.

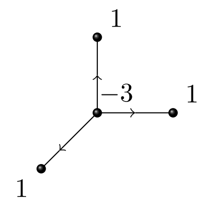

FIGURE 24. A rational function that is constant outside a local neighborhood, and its associated divisor

Definition 5.13. *We say that two divisors D and D' on a metric graph Γ are* equivalent *if $D - D'$ is principal. We define the* Picard group *of Γ to be the group of equivalence classes of divisors on Γ. That is,*

$$\mathrm{Pic}(\Gamma) = \mathrm{Div}(\Gamma)/\operatorname{div}(\mathrm{PL}(\Gamma)).$$

The Jacobian $\mathrm{Jac}(\Gamma)$ *of Γ is the group of equivalence classes of divisors of degree zero.*

The Jacobian of a metric graph is related to the Jacobian of a discrete graph.

Lemma 5.14. *Let G be a graph, let Γ be the associated metric graph with all edge lengths 1, and let $D \in \mathrm{Div}(G)$. Then D is principal on G if and only if it is principal on Γ.*

Corollary 5.15. *Let G be a graph, let Γ be the associated metric graph with all edge lengths 1, and let $D_1, D_2 \in \mathrm{Div}(G)$. Then D_1 and D_2 are equivalent divisors on G if and only if they are equivalent divisors on Γ.*

Corollary 5.16. *Let G be a graph and let Γ be the associated metric graph with all edge lengths 1. Then $\mathrm{Jac}(G)$ is a subgroup of $\mathrm{Jac}(\Gamma)$.*

5.2. **Reduced Divisors on Metric Graphs.** We now develop the theory of reduced divisors on metric graphs, analogous to the corresponding theory for discrete graphs. We begin with the definition.

Definition 5.17. *Let Γ be a metric graph and $v \in \Gamma$. A divisor D on Γ is* v-reduced *if*

1. *D is effective away from v and*
2. *every closed subset $A \subseteq \Gamma \smallsetminus \{v\}$ contains a point x with $D(x) < \mathrm{outdeg}_A(x)$.*

This is reminiscent of the definition for discrete graphs, but with a set of vertices replaced with a connected closed set. As in the discrete case, we will show that every divisor on a metric graph is equivalent to a unique v-reduced divisor. In this way, v-reduced divisors give a natural choice of representatives for divisors classes on metric graphs. We begin by showing that v-reduced divisors, if they exist, are unique.

Theorem 5.18. *If $D \sim D'$ are v-reduced, then $D = D'$. (Hint: mimic the proof in the discrete case.)*

To prove the existence of v-reduced divisors, we first need the following lemma.

Lemma 5.19. *Let $D \in \mathrm{Div}(\Gamma)$, $v \in \Gamma$, and let G be a model for Γ with vertex set containing $\{v\} \cup \mathrm{supp}(D)$. Then D is v-reduced if and only if the corresponding divisor D on G is v-reduced.*

Remark 5.20. Note that Lemma 5.19 does not imply that the v-reduced divisor on Γ corresponds to the v-reduced divisor on G. In particular, since our choice of model G depends on the divisor D, replacing D with an equivalent divisor may force us to change the model.

We prove the existence of a v-reduced divisor equivalent to a given divisor D in two parts, starting with the case where D is effective away from v. Let D be an effective divisor of degree k. Let v_1 be the point in the support of D that is closest to v. Recursively define v_i to be the point in the support of $D - \sum_{i=1}^{i-1} v_i$ that is closest to v. Let d_i denote the distance from v to v_i, and let $d_v(D) \in \mathbb{R}^k$ be the vector $(d_1, \ldots, d_k)$. We write $d_v(D) < d_v(D')$ if the first is former preceeds the latter in the lexicographic order on $\mathbb{R}^k$. In other words, $d_v(D) < d_v(D')$ if there exists an $\ell \leq k$ such that $d_i = d'_i$ for all $i < \ell$ and $d_\ell < d'_\ell$.

Lemma 5.21. *Let D be a divisor that is effective away from v. If D is not v-reduced, then there is a divisor D' equivalent to D, also effective away from v, such that $d_v(D') < d_v(D)$.*

Given $D \in \mathrm{Div}(\Gamma)$, the *complete linear series* of D is

$$|D| := \{D' \sim D \mid D' \geq 0\}.$$

Lemma 5.21 does not, on its own, imply the existence of v-reduced divisors. This is because, unlike the discrete case, the set of possible vectors $d_v(D)$ is inifinite. For this reason, we need the following lemma.

Lemma 5.22. *The set $|D|$ is a closed subset of* $\mathrm{Sym}^d \Gamma$.

Theorem 5.23. *If D is effective away from v, then D is equivalent to a v-reduced divisor.*

To handle the case where D is not effective, we will introduce the rank of divisors on metric graphs. This is defined in exactly the same way as on a discrete graph.

Definition 5.24. *The* rank *of D is the largest integer r such that $|D - E| \neq \emptyset$ for all effective divisors E of degree r.*

Lemma 5.25. *Any effective divisor of degree at least $g + r$ on a metric graph Γ of genus g has rank at least r. (Hint: induct on r, and us Lemma 5.19.)*

We now complete the proof that every divisor on a metric graph is equivalent to a unique v-reduced divisor. Recall that, in the discrete case, we used the fact that the Jacobian is finite, and hence every element of the Jacobian is torsion, to complete the proof. This argument will not work in the case of metric graphs, where the Jacobian has non-torsion elements.

Corollary 5.26. *Every divisor on Γ is equivalent to a divisor that is effective away from v.*

Corollary 5.27. *Let Γ be a metric graph, $v \in \Gamma$. Every divisor on Γ is equivalent to a unique v-reduced divisor.*

5.3. **Riemann-Roch for Metric Graphs.** In this section, we prove the Riemann-Roch theorem for metric graphs. In their seminal paper, Baker and Norine describe a general strategy for proving theorems of Riemann-Roch type. Specifically, let X be a non-empty set and let $\mathrm{Div}(X)$ be the free abelian group on X. As usual, we define the degree of a divisor to be the sum of its coefficients, and we say that a divisor is effective if all of its coefficients are nonnegative. Let $\sim$ be an equivalence relation on $\mathrm{Div}(X)$ satisfying:

(1) if $D \sim D'$, then $\deg(D) = \deg(D')$, and
(2) if $D_1 \sim D_1'$ and $D_2 \sim D_2'$, then $D_1 + D_2 \sim D_1' + D_2'$.

We can define a rank function on $\mathrm{Div}(X)$ by declaring $\mathrm{rk}(D)$ to be -1 if D is not equivalent to an effective divisor, and otherwise declaring $\mathrm{rk}(D)$ to be the largest integer r such that $D - E$ is equivalent to an effective divisor for all effective divisors E of degree r. Finally, let g be an integer, and let $K \in \mathrm{Div}(X)$ be a divisor of degree $2g - 2$. Baker and Norine prove the following.

Theorem 5.28. *The Riemann-Roch formula*

$$r(D) - r(K_\Gamma - D) = \deg(D) - g + 1$$

holds for every $D \in \mathrm{Div}(X)$ if and only if the following 2 conditions hold:

(1) *For every $D \in \mathrm{Div}(X)$, either $\mathrm{rk}(D) \geq 0$, or there exists a divisor D' of degree $g-1$ and rank -1 such that $\mathrm{rk}(D' - D) \geq 0$.*
(2) *For every $D \in \mathrm{Div}(X)$ of degree $g-1$, if $\mathrm{rk}(D) = -1$, then $\mathrm{rk}(K - D) = -1$.*

We now show that the 2 conditions hold for divisors on metric graphs. For the first condition, we must identify a suitably large collection of divisors of degree $g-1$ and rank -1. In the case of discrete graphs, these were the orientable divisors.

Definition 5.29. *An* orientation *on a metric graph Γ is an orientation of some model for Γ. As in the discrete case, given an orientation $\mathcal{O}$, define*

$$D_{\mathcal{O}} = \sum_{v \in \Gamma} (\mathrm{indeg}_{\mathcal{O}}(v) - 1)v.$$

Lemma 5.30. *If $\mathcal{O}$ is an acyclic orientation, then $D_{\mathcal{O}}$ is not equivalent to an effective divisor.*

We now show that the first condition holds.

Lemma 5.31. *For any $D \in \mathrm{Div}(\Gamma)$, either D is equivalent to an effective divisor, or there exists an acyclic orientation $\mathcal{O}$ such that $D_{\mathcal{O}} - D$ is equivalent to an effective divisor. Moreover, for any $v \in \Gamma$, $\mathcal{O}$ can be taken to have unique source v. (Hint: choose an appropriate model for Γ, and then use Corollary 3.15.)*

The second condition follows from the previous two lemmas.

Lemma 5.32. *Let $D \in \mathrm{Div}(\Gamma)$ be a divisor of degree $g-1$. If D is not equivalent to an effective divisor, then $K_\Gamma - D$ is not equivalent to an effective divisor.*

Combining Lemmas 5.31 and 5.32 with Theorem 5.28, we obtain the Riemann-Roch theorem for metric graphs.

Theorem 5.33 (Riemann-Roch for Metric Graphs)**.** *Let Γ be a metric graph of genus g. For any $D \in \mathrm{Div}(\Gamma)$,*

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

5.4. **Rank Determining Sets.** On a discrete graph G, the rank of a divisor D can be computed as follows. Choose a vertex v of G. For each effective divisor E of degree r, run Dhar's Burning Algorithm to compute the v-reduced divisor equivalent to $D - E$. The divisor D has rank at least r if and only if this v-reduced is effective for all E.

On a metric graph, however, this procedure is impossible to implement because for $r > 0$ there are infinitely many effective divisors of degree r. The goal of this section is to show that there exists a finite set of "test" divisors E such that, if $|D - E|$ is nonempty for all E in this finite set, then the divisor D has rank at least r. This will make it feasible to compute the ranks of divisors on metric graphs. This idea is made precise by the notion of rank determining sets. We first make the following definition.

Definition 5.34. *Let D be a divisor on a metric graph Γ. We define the* support *of the complete linear series $|D|$ to be*

$$\mathrm{supp}(|D|) := \{v \in \Gamma \mid D'(v) > 0 \textit{ for some } D' \in |D|\}.$$

We say that a divisor D has support in A if $\mathrm{supp}(D)$ is contained in A.

Definition 5.35. *Let Γ be a graph, and let A be a subset of Γ.*

(1) *The A-rank $r_A(D)$ of a divisor D is the largest integer r such that $|D - E|$ is nonempty for all effective divisors E of degree r with support in A.*
(2) *The set A is* rank determining *if $r_A(D) = r(D)$ for all $D \in \mathrm{Div}(G)$.*

Lemma 5.36. *For any subset $A \subseteq \Gamma$ and any divisor D, we have $r_A(D) \geq r(D)$.*

Definition 5.37. *Let $A \subseteq \Gamma$ be a subset. We define $\mathcal{L}(A)$ to be*

$$\mathcal{L}(A) = \bigcap_{\mathrm{supp}|D| \supseteq A} \mathrm{supp}|D|.$$

Proposition 5.38. *Let A be a nonempty subset of Γ. The following are equivalent:*

(1) *$\mathcal{L}(A) = \Gamma$.*
(2) *If D is a divisor with $r_A(D) \geq 1$, then $r(D) \geq 1$.*
(3) *A is a rank-determining set.*

(Hint: to show that (2) implies (3), use induction on $r_A(D)$.)

We now provide a topological condition to determine when $\mathcal{L}(A) = \Gamma$. To do this we define a YL set. These sets are named after Ye Luo, whose work this section is based upon.

Definition 5.39. *Let Γ be a metric graph and $U \subseteq \Gamma$ a connected open subset. We call U a* YL set *if every connected component X of the complement $\Gamma \smallsetminus U$ contains a point v such that* $\mathrm{outdeg}_X(v) > 1$.

We can characterize YL sets in terms of divisor theory.

Lemma 5.40. *Let $U \subseteq \Gamma$ be a nonempty connected open subset. Then U is a YL set if and only if $D = \sum_{v \in \partial U} v$ is w-reduced for any $w \in U$.*

Given a divisor D on Γ, we may use Lemma 5.40 to find YL sets that are disjoint from the support of $|D|$.

Lemma 5.41. *For $v \in \Gamma$, let D be a v-reduced divisor, and let U be the set of vertices that can be reached from v by a path that does not pass through $\mathrm{supp}(D) \smallsetminus \{v\}$. Then U is a YL set. Moreover, if $v \notin \mathrm{supp}(D)$, then U is disjoint from $\mathrm{supp}|D|$.*

The following consequence of Lemma 5.41 is not necessary for our other results on rank-determining sets, but may be of independent interest.

Corollary 5.42. *Let D be a divisor on Γ. Then $(\mathrm{supp}|D|)^c$ is a disjoint union of YL sets.*

We now turn to the main theorem of this lecture, which gives a sufficient condition for subsets of the vertices to be rank-determining.

Theorem 5.43. *Let $A \subseteq \Gamma$ be a nonempty subset. Then*

$$\mathcal{L}(A) \supseteq \bigcap_{\substack{U \text{ is } YL \\ A \cap U = \emptyset}} U^c.$$

Moreover, if all YL sets intersect A, then A is a rank determining set.

Remark 5.44. Luo proves the stronger result that the containment of Theorem 5.43 is in fact an equality, from which he derives that this sufficient condition for subsets to be rank-determining is also necessary. For our purposes, we will only need the fact that this condition is sufficient.

We note the following interesting property of YL sets.

Lemma 5.45. *If Γ is a metric graph of genus g and U is a YL set in Γ, then the closure $\overline{U}$ has genus at least 1. As a consequence, a collection of disjoint YL sets in Γ can contain at most g elements.*

The condition for rank determining sets provided in Theorem 5.43 is useful for many reasons. An important consequence of this result is that every metric graph contains a finite rank determining set.

Theorem 5.46. *Let Γ be a metric graph of genus g, let G be a model for Γ, let T be a spanning tree in G, and let $e_1, \ldots, e_g$ be the edges of G not in T. Choose a point $v_0 \in T$, and a point v_i in the interior of e_i for each i. Then $A = \{v_0, v_1, \ldots, v_g\}$ is a rank determining set. In particular, there exists a rank-determining set of cardinality $g+1$.*

A metric graph of genus g may have a rank determining set of cardinality less than $g+1$.

Exercise 5.47. Let Γ be a metric graph with model the complete graph K_4. Then Γ has a rank determining set of cardinality 3.

Exercise 5.48. Let Γ be a metric graph with model the complete graph K_4, where the edges have arbitrary lengths. Let $v_1, v_2, v_3 \in V(K_4)$ be distinct vertices. Show that $D = v_1 + v_2 + v_3$ has rank 1.

Theorem 5.49. *Let G be a graph, let Γ be the associated metric graph with all edge lengths 1, and let $D \in \mathrm{Div}(G)$. Then $\mathrm{rk}_G(D) = \mathrm{rk}_\Gamma(D)$.*

5.5. **The Tropical Abel-Jacobi Map.** A *1-form* on a graph G is an element of the real vector space generated by the formal symbols de, as e ranges over all directed edges of G, subject to the relations that if $e, \overline{e}$ represent the same edge with opposite orientations, then $d\overline{e} = -de$. A 1-form $\omega = \sum \omega_e de$ is called *harmonic* if, for all vertices v, the sum $\sum_{e=vw} \omega_e$ over the outgoing edges at v is equal to 0. Denote by $\Omega(G)$ the set of all harmonic 1-forms on G.

Lemma 5.50. *If G has genus g, then $\Omega(G)$ is a real vector space of dimension g.*

Given a spanning tree T in G, let $e_1, \ldots, e_g$ be the edges of G that are not contained in T. Let v_i and w_i be the endpoints of e_i. Since T is a tree, there is a unique path $f_{i,1}, \ldots, f_{i,k}$ in T that starts at w_i and ends at v_i. Define the 1-form $\omega_i = de_i + df_{i,1} + \cdots + df_{i,k}$.

Lemma 5.51. *The 1-forms $\omega_1, \ldots, \omega_g$ form a basis for $\Omega(G)$.*

Lemma 5.52. *If G and G' are models for the same metric graph Γ, then $\Omega(G')$ is canonically isomorphic to $\Omega(G)$.*

If Γ is a metric graph, we define the space $\Omega(\Gamma)$ to be $\Omega(G)$ for any model G of Γ. By Lemma 5.52, this is well defined. Given an isometric path $\gamma \colon [a, b] \to \Gamma$, any harmonic 1-form ω on Γ pulls back to a classical 1-form on the interval, and we can thus define the integral $\int_\gamma \omega$.

Proposition 5.53. *Fix a point $v_0 \in \Gamma$. For any point $v \in \Gamma$, let γ be a path from v_0 to v. The map*

$$AJ_{v_0} \colon \Gamma \to \Omega(\Gamma)^* / H_1(\Gamma, \mathbb{Z})$$

given by $AJ_{v_0}(v) = \int_\gamma$ is well-defined. That is, it does not depend on the choice of path γ.

Extend the map AJ_{v_0} linearly to $\mathrm{Div}(\Gamma)$ and then restrict to $\mathrm{Div}^0(\Gamma)$ to obtain the *Abel-Jacobi map*:

$$AJ \colon \mathrm{Div}^0(\Gamma) \to \Omega(\Gamma)^* / H_1(\Gamma, \mathbb{Z}).$$

Lemma 5.54. *The map AJ does not depend on the choice of point v_0.*

Lemma 5.55. *The map AJ is surjective. (Hint: use Lemma 5.51.)*

Lemma 5.56. *The kernel of the map AJ is $\mathrm{div}(\mathrm{PL}(\Gamma))$.*

As a corollary of Lemmas 5.55 and 5.56, we obtain the following.

Theorem 5.57 (Tropical Abel-Jacobi Theorem)**.** *For any metric graph Γ, the Jacobian $\mathrm{Jac}(\Gamma)$ is isomorphic to the torus $\Omega(\Gamma)^* / H_1(\Gamma, \mathbb{Z})$.*

Exercise 5.58. Let Γ be the circle with circumference ℓ, pictured in Figure 25. Compute $\Omega(\Gamma)$ and then use the tropical Abel-Jacobi Theorem to compute $\mathrm{Jac}(\Gamma)$.

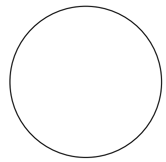

Figure 25. A circle with circumference ℓ

Exercise 5.59. Let Θ be the graph pictured in Figure 26, and let ℓ_1, ℓ_2, ℓ_3 be the lengths of the 3 edges.

(1) Find a basis for $\Omega(\Theta)$.
(2) Describe the lattice $H_1(\Gamma, \mathbb{Z})$ in $\Omega(\Theta)^*$.
(3) Draw a picture of a fundamental domain for the torus $\Omega(\Theta)^*/H_1(\Gamma, \mathbb{Z})$.

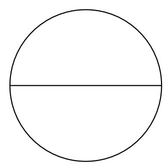

Figure 26. The graph Θ

5.6. Further Reading.

The divisor theory of metric graphs is also introduced in [Bak08]. The Riemann-Roch theorem for metric graphs was proved independently by [MZ08] and [HKN13]. These notes do not follow the approach of either of these papers. Rank determining sets were first defined by Luo in [Luo11]; Section 5.4 borrows heavily from this paper. Similarly, Section 5.5 follows [BF11], which proves the Abel-Jacobi theorem for metric graphs.

6. Introduction to Algebraic Curves

As mentioned in the first section, we use the term *divisors* to emphasize the analogy with divisors on algebraic curves. In this section, we cover some introductory material in algebraic geometry. We have tried to keep this material to a minimum, exploring only those topics that are necessary to understand the connection to chip firing.

6.1. Algebraic Curves.

Throughout this section, we let k be an algebraically closed field and let $S = k[x_0, \ldots, k_r]$ be the polynomial ring in $r+1$ variables. The *projective space* $\mathbb{P}^r_k$ of dimension r over k is the set of equivalence classes of nonzero vectors in k^{r+1}, where $(a_0, \ldots, a_r)$ is equivalent to $(b_0, \ldots, b_r)$ if there exists $\lambda \in k^*$ such that $b_i = \lambda a_i$ for all i. For ease of notation, we sometimes write $\mathbb{P}^r$ instead of $\mathbb{P}^r_k$.

Recall that a polynomial $f \in S$ is *homogeneous of degree d* if it is a linear combination of monomials of degree d. Polynomial functions are not well-defined on projective space, but the zeros of a homogenous polynomial are well-defined. If $T \subset S$ is a set of homogeneous elements, we define

$$Z(T) := \{p \in \mathbb{P}^r_k \mid f(p) = 0 \text{ for all } f \in T\}.$$

Definition 6.1. *A subset $Y \subseteq \mathbb{P}^r_k$ is a* projective variety *if there exists a set of homogeneous elements $T \subset S$ such that $Y = Z(T)$.*

Exercise 6.2. Show that $\mathbb{P}^r$ is a projective variety.

Exercise 6.3. Show that the empty set is a projective variety.

Lemma 6.4. *If $T_1, T_2 \subset S$ are sets of homogeneous elements, let T_1T_2 denote the set of all products of an element of T_1 by an element of T_2. Then*

$$Z(T_1T_2) = Z(T_1) \cup Z(T_2).$$

Lemma 6.5. *Let $T_\alpha \subset S$ be a set of homogeneous elements, as α ranges over some index set A. Then*

$$Z(\bigcup_{\alpha \in A} T_\alpha) = \bigcap_{\alpha \in A} Z(T_\alpha).$$

Corollary 6.6. *There is a topology on $\mathbb{P}^r$ whose closed sets are the projective varieties.*

This toplogy is called the *Zariski topology* on $\mathbb{P}^r$. Any projective variety $Y \subset \mathbb{P}^r$ inherits the subspace topology, which is called the Zariski topology on Y.

Lemma 6.7. *Let $T \subset S$ be a set of homogeneous elements, and let T' be the set of homogeneous elements in the ideal generated by T. Then $Z(T) = Z(T')$.*

Definition 6.8. *A nonempty topological space Y is called* irreducible *if, for any two closed subsets $Y_1, Y_2 \subseteq Y$ such that $Y_1 \cup Y_2 = Y$, we have either $Y_1 = Y$ or $Y_2 = Y$.*

Exercise 6.9. The space $\mathbb{P}^n$, with the Zariski topology, is irreducible.

Lemma 6.10. *Any two nonempty open susbets of an irreducible space have nonempty intersection.*

Lemma 6.11. *Any nonempty open subset of an irreducible space is dense.*

Definition 6.12. *Let Y be a topological space. We define the* dimension *of Y to be the supremum of all integers n such that there exists a chain $Y_0 \subsetneq Y_1 \subsetneq \cdots \subsetneq Y_n$ of irreducible closed subsets of Y.*

A curve *is an irreducible projective variety of dimension 1.*

Exercise 6.13. The projective variety $\mathbb{P}^1$ is a curve.

Lemma 6.14. *The Zariski topology on a curve is the cofinite topology. That is, if C is a curve, then the only closed subsets of C are the empty set, the set C itself, and all finite subsets of C.*

6.2. **Valuations and Regular Functions.** Before doing more geometry, we require a digression about discrete valuations.

Definition 6.15. *Let k be a field. A* discrete valuation *on k is a function $\nu\colon k \to \mathbb{Z} \cup \{\infty\}$ such that:*

(1) *$\nu(a) = \infty$ if and only if $a = 0$,*
(2) *$\nu(ab) = \nu(a) + \nu(b)$ for all $a, b \in k$, and*
(3) *$\nu(a + b) \geq \min\{\nu(a), \nu(b)\}$ for all $a, b \in k$.*

Exercise 6.16. Let k be a field. The map $\nu\colon k(t) \to \mathbb{Z} \cup \{\infty\}$ given by $\nu(\frac{f}{g}) = \deg(g) - \deg(f)$ is a discrete valuation.

Exercise 6.17. Let k be a field. For $f \in k[t]$, define $\nu(f)$ to be the maximum power of t that divides f. Then the map $\nu\colon k(t) \to \mathbb{Z} \cup \{\infty\}$ given by $\nu(\frac{f}{g}) = \nu(f) - \nu(g)$ is a discrete valuation.

Lemma 6.18. *Let k be a field and $\nu\colon k \to \mathbb{Z} \cup \{\infty\}$ a discrete valuation. The set $R = \{a \in k \mid \nu(a) \geq 0\}$ is a ring.*

The ring R defined in Lemma 6.18 is called the *(discrete) valuation ring*, or DVR, for the valuation ν.

Lemma 6.19. *Let k be a field, ν a discrete valuation, and R the valuation ring. The set $m = \{a \in k \mid \nu(a) > 0\}$ is the unique maximal ideal in R.*

In commutative algebra, a ring with a unique maximal ideal is called a *local ring*. The reason for this terminology comes from geometry, as we will see at the end of the section. Because m is a maximal ideal, the ring R/m is a field, called the *residue field* of R.

Lemma 6.20. *Let k be a field, ν a discrete valuation, R the valuation ring, and m the unique maximal ideal in R. Then m is principal.*

Lemma 6.21. *Let k be a field, ν a discrete valuation, R the valuation ring, m the unique maximal ideal in R, and let $a \in R$ such that $\nu(a) = 1$. For $b \in R$, define $\varphi(b)$ to be the maximum power of a that divides b. Then $\nu(b) = \varphi(b)$ for all $b \in R$.*

Proposition 6.22. *Let R be a PID with a unique maximal ideal. Then R is a discrete valuation ring.*

Returning to geometry, let $Y \subset \mathbb{P}^r$ be a projective variety and let $U \subseteq Y$ be a nonempty open set. A function $f\colon U \to k$ is *regular* on U if there exist homogeneous polynomials $g, h \in S$ of the same degree, with h nowhere zero on U, such that $f = \frac{g}{h}|_U$.

Definition 6.23. *Let $Y \subseteq \mathbb{P}^r$ be an irreducible projective variety. The* function field *$K(Y)$ of Y is the set of equivalence classes of pairs (U, f), where $U \subseteq Y$ is a nonempty open set and $f\colon U \to k$ is regular on U, where (U, f) is equivalent to (V, g) if the restriction of f to $U \cap V$ is equal to that of g.*

If g and h are homogeneous polynomials of the same degree, then $\frac{g}{h}$ is regular on the open set $Y \smallsetminus Z(h)$. In practice, when describing elements of $K(Y)$, we typically suppress the open set U and just write f for the equivalence class of the regular function f.

Exercise 6.24. The function field of $\mathbb{P}^1$ is isomorphic to $k(t)$.

Exercise 6.25. Let $C = Z(y^2z - x^3) \subset \mathbb{P}^2$. The function field of C is isomorphic to $k(t)$. (Hint: consider the map $\varphi\colon \mathbb{P}^1 \to C$ given by $\varphi(x, y) = (x^2y, x^3, y^3)$.)

Definition 6.26. *Let $Y \subseteq \mathbb{P}^r$ be an irreducible projective variety and let $p \in Y$. The* local ring $\mathcal{O}_{Y,p}$ *of Y at p is the set of equivalence classes of pairs (U, f) as in Definition 6.23, where U contains p.*

Lemma 6.27. *Let $Y \subseteq \mathbb{P}^r$ be an irreducible projective variety and let $p \in Y$. Then $K(Y)$ is the field of fractions of $\mathcal{O}_{Y,p}$.*

Exercise 6.28. Let $p = [0 : 1] \in \mathbb{P}^1$. Then

$$\mathcal{O}_{\mathbb{P}^1,p} \cong \{\frac{f}{g} \in k(t) \mid g(0) \neq 0\}.$$

Exercise 6.29. Let $C = Z(y^2z - x^3) \subset \mathbb{P}^2$, and let $p = [0 : 0 : 1] \in C$. Then

$$\mathcal{O}_{C,p} \cong \{\frac{f}{g} \mid f \in k[t^2, t^3], g \in k[t], g(0) \neq 0\}.$$

Definition 6.30. *A curve C is* smooth *at a point p if the local ring $\mathcal{O}_{C,p}$ is a discrete valuation ring. We say that C is* smooth *if it is smooth at all points $p \in C$.*

Exercise 6.31. The curve $\mathbb{P}^1$ is smooth, and the curve $C = Z(y^2z - x^3) \subset \mathbb{P}^2$ is not.

By Lemma 6.21, if C is smooth at the point p, then the discrete valuation on $\mathcal{O}_{C,p}$ is unique. This discrete valuation is called the *order of vanishing* at p and denoted ord_p.

Lemma 6.32. *If the curve C is smooth at the point p, then $\mathrm{ord}_p(f) \geq 1$ if and only if $f(p) = 0$.*

6.3. **Divisors on Algebraic Curves.** We are primarily interested in divisors on algebraic curves.

Definition 6.33. *A* divisor *D on a curve C is a formal $\mathbb{Z}$-linear combination of points of C,*

$$D = \sum_{p \in C} D(p) \cdot p$$

with $D(p) \in \mathbb{Z}$.

As in the case of graphs, the *degree* of a divisor $D = \sum_{p\in C} D(p) \cdot p$ is the integer $\deg(D) = \sum_{p\in C} D(p)$. The divisor D is *effective* if $D(p) \geq 0$ for all $p \in C$. Given a nonzero regular function $f \in K(C)$, we define the corresponding *principal divisor* to be

$$\mathrm{div}(f) = \sum_{p \in C} \mathrm{ord}_p(f) \cdot p.$$

In order to see that this is well-defined, we need the following fact.

Lemma 6.34. *Let $f \in K(C)$ be nonzero. Then $\mathrm{ord}_p(f) = 0$ for all but finitely many points $p \in C$.*

We say that two divisors D, D' on a curve C are *linearly equivalent* if their difference is principal.

Lemma 6.35. *Linear equivalence of divisors is an equivalence relation.*

The degree of a divisor is invariant under linear equivalence. We have chosen not to include this among the exercises, as proving it would require a longer foray into commutative algebra than we would like. The *Picard group* $\mathrm{Pic}(C)$ of a curve C is the set of linear equivalence classes of divisors on C, and the *Jacobian* $\mathrm{Jac}(C)$ is the set of linear equivalence classes of degree zero.

Definition 6.36. *Let D be a divisor on a curve C. The* complete linear series $\mathcal{L}(D)$ *is defined to be*

$$\mathcal{L}(D) = \{f \in K(C)^* \mid \mathrm{div}(f) + D \textit{ is effective}\} \cup \{0\}.$$

Lemma 6.37. *Let D be a divisor on a curve C. Then $\mathcal{L}(D)$ is a k-vector space.*

Definition 6.38. *Let D be a divisor on a curve C. An $(r+1)$-dimensional subspace $V \subseteq \mathcal{L}(D)$ is called a* linear series of rank r. *The rank of D is defined to be the rank of $\mathcal{L}(D)$.*

Proposition 6.39. *Let C be a curve over k, and let $K \subseteq k$ be a subfield such that C has infinitely many K-points. Let D be a K-divisor on C and $V \subseteq \mathcal{L}(D)$ a linear series. The rank of D is equal to the maximum integer r such that, for all effective K-divisors E of degree r, there exists a nonzero $f \in V$ such that* $\mathrm{div}(f) + D - E$ *is effective.*

Corollary 6.40. *Let C be a curve, and let D be a divisor on C. Then $\mathcal{L}(D)$ is finite dimensional.*

Algebraic geometers are mainly interested in linear series because of the correspondence between linear series and maps to projective space. Let $V \subseteq \mathcal{L}(D)$ be a linear series of rank r, and let $f_0, \dots, f_r \in V$ be a basis. Then there exists a map $F \colon C \to \mathbb{P}^r$ given by $F(p) = [f_0(p) : \cdots : f_r(p)]$. Conversely, given a map $F \colon C \to \mathbb{P}^r$, let D be the "divisor at infinity", $D = F^*(Z(x_r))$. Let $f_i = F^* \frac{x_i}{x_r}$. Then the functions f_i are in $\mathcal{L}(D)$, and they generate a linear series.

6.4. Further Reading.

The material in this section is standard in algebraic geometry, and can be found in any introductory text, such as [Har77] or [Ful89].

7. Specialization of Divisors from Curves to Graphs

In this section, we discuss a tool for connecting the divisor theory of algebraic curves to the divisor theory of graphs.

7.1. Graph Curves.

We begin with the definition of a graph curve.

Definition 7.1. *A* graph curve *is a connected, 1-dimensional projective variety C such that:*

1. *each component of C is isomorphic to $\mathbb{P}^1$,*
2. *any point of C is contained in at most 2 components, and*

(3) *at any such point where 2 components intersect, they intersect transversally.*

Some authors include a fourth condition in the definition of a graph curve, that each component meets the rest of the curve in at most 3 points. We will not require this here. The interesting geometric properties of a graph curve are encoded by the combinatorics of which components meet each other.

Definition 7.2. *The* dual graph *of a graph curve C is the graph whose vertices are in bijection with the components of C, and where there is and edge between vertices v_i and v_j for every point where the corresponding components C_i and C_j intersect.*

Exercise 7.3. The dual graph of the union of n general lines in $\mathbb{P}^2$ is the complete graph K_n.

Exercise 7.4. The dual graph of the union of lines in $\mathbb{P}^1 \times \mathbb{P}^1$, n of which are in one ruling and m of which are in the other, is the complete bipartite graph $K_{n,m}$.

Exercise 7.5. The dual graph of the union of all -1 curves on a del Pezzo surface of degree 5 is the Petersen graph.

7.2. **Specialization.** In this section, we introduce the main construction for relating smooth curves to graph curves. This requires some concepts from algebraic geometry that we have not introduced, but we will explain the ideas that are necessary to complete the exercises below. We begin by fixing some notation. Throughout, R will denote a complete discrete valuation ring with field of fractions K and algebraically closed residue field k. For example, K could be the field $k(t)$ with the valuation given as in Exercise 6.17. Let C be a smooth curve over K.

One can think of a curve over R as a curve whose defining equations have coefficients in the ring R. The ring R admits two natural maps to fields: the inclusion $R \hookrightarrow K$ and the quotient $R \to k$. By specializing the coefficients via these two maps, one obtains a curve over K, called the *generic fiber* and a curve over k, called the *special fiber.*

Definition 7.6. *A curve $\mathfrak{C}$ over R is called a* strongly semistable model *for C over R if it satisfies the following:*

(1) *$\mathfrak{C}$ is proper and flat over R,*
(2) *$\mathfrak{C}$ is regular,*
(3) *the generic fiber $\mathfrak{C}_K$ is C, and*
(4) *The special fiber $\mathfrak{C}_k$ is a graph curve.*

Exercise 7.7. Let $K = k(t)$ with the valuation as in Exercise 6.17. Let $F(x, y, z)$ be a general homogeneous polynomial of degree 4, and let

$$\mathfrak{C} = \{(x, y, z) \in \mathbb{P}^2_R \mid tF(x, y, z) + xyz(x + y + z) = 0\}.$$

What is the special fiber of $\mathfrak{C}$? What is its dual graph?

Exercise 7.8. Let $K = k(t)$ with the valuation as in Exercise 6.17. Let Q be a general homogeneous polynomial of degree 2, let F be a general homogeneous polynomial of degree 3, and let $L_1, \cdots, L_5$ be general homogeneous polynomials of degree 1. Finally, let

$$\mathfrak{C} = \{(x, y, z, w) \in \mathbb{P}^3_R \mid tQ + L_1L_2 = tF + L_3L_4L_5 = 0\}.$$

What is the special fiber of $\mathfrak{C}$? What is its dual graph?

Let G be the dual graph of $\mathfrak{C}_k$, and let p be a point on C. Since $\mathfrak{C}$ is regular, p specializes to a smooth point of $\mathfrak{C}_k$. Its specialization therefore lies on a unique component of $\mathfrak{C}_k$, which corresponds to a vertex $v(p)$ of the dual graph G. Extending linearly, we define a map $\rho\colon \operatorname{Div}(C) \to \operatorname{Div}(G)$ by

$$\rho(\sum_{p \in C} D(p) \cdot p) := \sum_{p \in C} D(p) \cdot v(p).$$

Exercise 7.9. Let $\mathfrak{C}$ be as in Exercise 7.7, and let D be the intersection of $C = \mathfrak{C}_K$ with a general line in $\mathbb{P}^2_K$. Compute the divisor $\rho(D)$.

Exercise 7.10. Let $\mathfrak{C}$ be as in Exercise 7.7, and let $D = \operatorname{div}(\frac{x}{y})$ on $C = \mathfrak{C}_K$. Compute the divisor $\rho(D)$.

Proposition 7.11. *If D is a principal divisor on C, then $\rho(D)$ is a principal divisor on G.*

Corollary 7.12. *The map ρ induces a map* $\operatorname{Trop}\colon \operatorname{Pic}(C) \to \operatorname{Pic}(G)$.

Let $q \in \mathfrak{C}_k$ be a smooth point of the special fiber. In order to prove Baker's Specialization Lemma below, we will need the following fact: there exists a K-point $p \in C$ such that p specializes to q in $\mathfrak{C}_k$. Proving this fact is beyond the scope of these notes. It uses Hensel's Lemma, and requires our assumptions that $\mathfrak{C}$ is regular and that the residue field k is algebraically closed. The following is an immediate consequence.

Lemma 7.13. *For all $v \in V(G)$, there exists $p \in C$ such that $\rho(p) = v$. Moreover, C has infinitely many K-points.*

The last result in this section – possibly the most important in this field – is known as Baker's Specialization Lemma.

Theorem 7.14. *For all divisors D on C, we have $r(\operatorname{Trop}(D)) \geq r(D)$.*

7.3. **Refinements and Compatibility with Base Change.** Let G be a graph. For a positive integer e, we define the *eth refinement* of G, denoted $\sigma_e(G)$, to be the graph obtained from G by replacing each edge with a path of length e. Before returning to algebraic geometry, we first explore several properties of the eth refinement.

Lemma 7.15. *Let G be a graph of genus g. Then $|\operatorname{Jac}(\sigma_e(G))| = e^g |\operatorname{Jac}(G)|$.*

Note that each vertex of G corresponds to a vertex in $\sigma_e(G)$. Extending linearly, we obtain a map from the set of divisors on G to the set of divisors on $\sigma_e(G)$.

Lemma 7.16. *Let D be a divisor on G. The divisor D is principal on G if and only if it is principal on $\sigma_e(G)$.*

Proposition 7.17. *Let G be a graph of genus g. There is a short exact sequence*

$$0 \to \operatorname{Jac}(G) \to \operatorname{Jac}(\sigma_e(G)) \to \Big(\mathbb{Z}/e\mathbb{Z}\Big)^g \to 0.$$

Corollary 7.18. *Let G be a graph of genus g, and let $n_1, \ldots, n_k$ be nonnegative integers, such that* $\mathrm{Jac}(G) \cong \oplus \mathbb{Z}/n_i\mathbb{Z}$. *Then*

$$\mathrm{Jac}(\sigma_e(G)) \cong \Big(\mathbb{Z}/e\mathbb{Z}\Big)^{g-k} \oplus \Big(\bigoplus_{i=1}^{k} \mathbb{Z}/en_i\mathbb{Z}\Big).$$

Lemma 7.19. *Let D be a divisor on a graph G. Then $r_G(D) = r_{\sigma_e(G)}(D)$.*

Returning to specialization from curves to graphs, one issue with the theory we have developed so far is that algebraic geometers typically work over algebraically closed fields. In our setup, however, the field K is never algebraically closed.

Lemma 7.20. *Let K be a field with a nontrivial discrete valuation. Then K is not algebraically closed.*

For this reason, it makes sense to consider finite extensions of the discretely valued field K. We first consider extensions of discrete valuation rings. We say that $A \subseteq B$ is an *extension of discrete valuation rings* if A and B are DVRs with maximal ideals m_A, m_B, respectively, and $m_A B \subseteq m_B$.

Lemma 7.21. *Let $A \subseteq B$ be an extension of discrete valuation rings with maximal ideals m_A, m_B. Then there exists a positive integer esuch that $m_A B = m_B^e$.*

The integer e is called the *ramification index* of B over A.

Lemma 7.22. *Let $A \subseteq B$ be an extension of discrete valuation rings with fields of fractions K and L. If the extension L/K is finite, then $e \leq [L : K]$.*

A fact that we will not prove here is that, if K is a discretely valued field with valuation ring A, and L/K is a finite, separable extension, then there exists an extension of discrete valuation rings $A \subseteq B$ such that L is the field of fractions of B. In particular, there is a discrete valuation on L such that the valuation ring is B. We call this a finite extension of discretely valued fields, and the ramification index of the extension $(L, \nu_L)/(K, \nu_K)$ is defined to be the ramification index of the extension $A \subseteq B$.

Exercise 7.23. Let $K = k(t)$ with the valuation ν_K as in Exercise 6.17, and let $L = k(t^{1/e})$. Show that there is a unique discrete valuation ν_L on L such that $(L, \nu_L)/(K, \nu_K)$ is an extension of discretely valued fields. What is the ramification index of this extension?

Let K be a discretely valued field with valuation ring A, let L/K be a finite extension of discretely valued fields, and let B be the valuation ring of L. Let C be a curve over K and let $C_L = C \times_K L$. If $\mathfrak{C}$ is a regular semistable model for C over A, then $\mathfrak{C}_B := \mathfrak{C} \times_A B$ is not necessarily regular. However, there is a unique relatively minimal regular semistable model $\mathfrak{C}'$ for C_L over B that dominates $\mathfrak{C}$. If G is the dual graph of the special fiber of $\mathfrak{C}$, then $\sigma_e(G)$ is the dual graph of the special fiber of $\mathfrak{C}'$, where e is the ramification index of the extension.

One way to understand the importance of metric graphs to this subject is via refinements. Let G be a graph, and let Γ be the corresponding metric graph in which all edges have length 1. Then the vertices of $\sigma_e(G)$ can be thought of as all the points on Γ whose distance from a vertex in G is an integer multiple of $\frac{1}{e}$. The vertices of $\lim_e \sigma_e(G)$ is the set of points of rational distance from the vertices of G, and Γ is the metric completion of this set of points.

7.4. **Applications.** In this section, we discuss some applications of chip firing to algebraic geometry. Although many of the results of this section can be proven without specializing to graphs, we encourage the reader to think about proving them with the earlier material.

Proposition 7.24. *Let K be a discretely valued field, and let C be a smooth plane curve of degree n over K. Then any K-divisor on C of positive rank has degree at least $n-1$.*

Theorem 7.25. *Let K be a discretely valued field, let C be a smooth plane curve of degree n over K, and let L/K be a finite extension of discretely valued fields. Then any L-divisor on C of positive rank has degree at least $n-1$.*

Lemma 7.26. *Let C be the complete intersection of a quadric and a del Pezzo surface of degree 5, over the algebraic closure of a discretely valued field. Then C has gonality 4.*

Theorem 7.27. *Let C be a curve of genus g over a discretely valued field K, and suppose that $\mathfrak{C}$ is a regular semistable model such that the dual graph of the special fiber is a chain of g loops with $m_k > g$ for all k. If D is a divisor on C of degree d and rank r, then*

$$g-(r+1)(g-d+r) \geq 0.$$

For a given r and d, the *Brill-Noether variety* of a curve C is defined to be

$$W_d^r(C) = \{D \in \mathrm{Pic}^d(C) \mid r(D) \geq r\}.$$

Theorem 7.27 shows that there exists a curve C of genus g such that $W_d^r(C) = \emptyset$ if $g-(r+1)(g-d+r) < 0$. With a little work, we can improve on this, to bound the dimension of $W_d^r(C)$.

For a given r and d, the *Brill-Noether rank* w_d^r of a graph, or a metric graph, or a curve, is defined to be the largest integer w such that, for all effective divisors E of degree $r+w$, there exists a divisor D of degree d and rank at least r such that $D-E$ is effective.

Proposition 7.28. *Let C be a curve. Then $w_d^r(C) = \dim W_d^r(C)$.*

Corollary 7.29. *Let C be a curve over a discretely valued field K. Let $\mathfrak{C}$ be a regular semistable model for C over K and let G be the dual graph of the special fiber. Then $\dim W_d^r(C) \leq w_d^r(G)$.*

The following is known as the Brill-Noether Theorem.

Theorem 7.30. *Let k be the algebraic closure of a discretely valued field. Then there exists a curve C over k such that, for all r and d,*

$$\dim W_d^r(C) \leq g-(r+1)(g-d+r).$$

7.5. **Further Reading.** All of the material in this section comes from [Bak08]. We also recommend the survey article [BJ16]. The tropical proof of the Brill-Noether theorem comes from [CDPR12].

References

[ABKS14] Y. An, M. Baker, G. Kuperberg, and F. Shokrieh. Canonical representatives for divisor classes on tropical curves and the matrix-tree theorem. *Forum Math. Sigma*, 2:e24, 25, 2014.

[Bac17] S. Backman. Riemann-Roch theory for graph orientations. *Adv. Math.*, 309:655–691, 2017.

[Bak08] M. Baker. Specialization of linear systems from curves to graphs. *Algebra Number Theory*, 2(6):613–653, 2008.

[BF11] M. Baker and X. Faber. Metric properties of the tropical Abel-Jacobi map. *J. Algebraic Combin.*, 33(3):349–381, 2011.

[BJ16] Matthew Baker and David Jensen. Degeneration of linear series from the tropical point of view and applications. In *Nonarchimedean and tropical geometry*, Simons Symp., pages 365–433. Springer, [Cham], 2016.

[BN07] M. Baker and S. Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215(2):766–788, 2007.

[BS13] M. Baker and F. Shokrieh. Chip-firing games, potential theory on graphs, and spanning trees. *J. Combin. Theory Ser. A*, 120(1):164–182, 2013.

[CDPR12] F. Cools, J. Draisma, S. Payne, and E. Robeva. A tropical proof of the Brill-Noether theorem. *Adv. Math.*, 230(2):759–776, 2012.

[CLB03] Robert Cori and Yvan Le Borgne. The sand-pile model and Tutte polynomials. volume 30, pages 44–52. 2003. Formal power series and algebraic combinatorics (Scottsdale, AZ, 2001).

[Dha90] D. Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990.

[Ful89] William Fulton. *Algebraic curves*. Advanced Book Classics. Addison-Wesley Publishing Company, Advanced Book Program, Redwood City, CA, 1989. An introduction to algebraic geometry, Notes written with the collaboration of Richard Weiss, Reprint of 1969 original.

[Har77] R. Hartshorne. *Algebraic geometry*. Springer-Verlag, New York, 1977. Graduate Texts in Mathematics, No. 52.

[HJJS22] Michael Harp, Elijah Jackson, David Jensen, and Noah Speeter. A new lower bound on graph gonality. *Discrete Appl. Math.*, 309:172–179, 2022.

[HKN13] J. Hladký, D. Král', and S. Norine. Rank of divisors on tropical curves. *J. Combin. Theory Ser. A*, 120(7):1521–1538, 2013.

[Kir47] Gustav Kirchhoff. Ueber die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Vertheilung Galvanischer Ströme geführt wird. *Annalen der Physik und Chemie*, 72:497–508, 1847.

[Kli19] Caroline J. Klivans. *The mathematics of chip-firing*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton, FL, 2019.

[Luo11] Y. Luo. Rank-determining sets of metric graphs. *J. Combin. Theory Ser. A*, 118:1775–1793, 2011.

[MZ08] G. Mikhalkin and I. Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and abelian varieties*, volume 465 of *Contemp. Math.*, pages 203–230. Amer. Math. Soc., Providence, RI, 2008.

[Pfl17] N. Pflueger. Special divisors on marked chains of cycles. *J. Combin. Theory Ser. A*, 150:182–207, 2017.

[Sho10] F. Shokrieh. The monodromy pairing and discrete logarithm on the Jacobian of finite graphs. *J. Math. Cryptol.*, 4(1):43–56, 2010.

[vDdBG20] Josse van Dobben de Bruyn and Dion Gijswijt. Treewidth is a lower bound on graph gonality. *Algebr. Comb.*, 3(4):941–953, 2020.