



A tropical proof of the Brill–Noether Theorem

Filip Cools^a, Jan Draisma^b, Sam Payne^{c,*}, Elina Robeva^d

^a Department of Mathematics, K.U. Leuven, Celestijnenlaan 200B, B-3001 Leuven, Belgium

^b Department of Mathematics and Computer Science, TU/e, PO Box 513, 5600 MB Eindhoven, The Netherlands

^c Mathematics Department, Yale University, 10 Hillhouse Ave, New Haven, CT 06511, USA

^d Department of Mathematics, Stanford University, Bldg 380, Stanford, CA 94305, USA

Received 6 February 2011; accepted 20 February 2012

Communicated by Ravi Vakil

Abstract

We produce Brill–Noether general graphs in every genus, confirming a conjecture of Baker and giving a new proof of the Brill–Noether Theorem, due to Griffiths and Harris, over any algebraically closed field. © 2012 Elsevier Inc. All rights reserved.

Keywords: Tropical geometry; Brill–Noether theory

1. Introduction

Let X be a smooth projective curve of genus g . Brill–Noether theory studies the geometry of the subscheme $W_d^r(X)$ of $\text{Pic}_d(X)$ parametrizing linear equivalence classes of divisors of degree d that move in a linear system of dimension at least r , especially when the curve X is general. This subscheme can be realized as a degeneracy locus of a natural map of vector bundles, with naive expected dimension

$$\rho = g - (r + 1)(g - d + r).$$

* Corresponding author.

E-mail addresses: filip.cools@wis.kuleuven.be (F. Cools), j.draisma@tue.nl (J. Draisma), sam.payne@yale.edu (S. Payne), erobeva@stanford.edu (E. Robeva).

The Brill–Noether Theorem, due to Griffiths and Harris, says that the naive dimension count is essentially correct, for a general curve.

Brill–Noether Theorem. (See [14].) Suppose X is general.

- (1) If ρ is negative then $W_d^r(X)$ is empty.
- (2) If ρ is nonnegative then $W_d^r(X)$ has dimension $\min\{\rho, g\}$.

Curves that are easy to write down, such as complete intersections in projective spaces and Grassmannians, have many more special divisors than this dimension count predicts, so the generality hypothesis is crucial. The fact that W_d^r is nonempty and of dimension at least $\min\{\rho, g\}$ for an arbitrary curve, when ρ is nonnegative, is significantly easier, and was assumed by Griffiths and Harris in their proof of (2). General results on degeneracy loci say that W_d^r must support the expected cohomology class given by the Thom–Porteous determinantal formula in the Chern classes of the bundles. An explicit computation, due to Kempf, Kleiman, and Laksov, shows that this expected class is a nonzero multiple of a power of the theta divisor [19,20].

The nonexistence of special divisors when ρ is negative, and the upper bound on the dimension of W_d^r when ρ is nonnegative, for a general curve, is considered much deeper. The depth of this result is related to the difficulty of writing down a sufficiently general curve in high genus, or even a reasonable criterion for a curve to be sufficiently general. The original proof uses a degeneration to a g -nodal rational curve, followed by a very subtle transversality argument for certain Schubert varieties associated to osculating flags of a rational normal curve. Two subsequent proofs [10,21] also continue to be heavily cited after more than twenty years; the Eisenbud–Harris proof uses limits of linear series and a degeneration to a cuspidal rational curve. Lazarsfeld’s proof involves no degenerations; he shows that a general hyperplane section of a complex $K3$ surface of Picard number 1 is Brill–Noether general, by which we mean that W_d^r is empty when ρ is negative and of dimension $\min\{\rho, g\}$ otherwise. Here we give a novel “tropical” proof of the Brill–Noether Theorem, replacing the subtle transversality arguments in the original proof with the combinatorics of chip-firing on certain graphs. The graphs encode degenerations to semistable unions of rational curves without self-intersection, which may be realized as degenerations of the g -nodal rational curves used by Griffiths and Harris.

Our starting points are the theory of ranks of divisors on graphs, as developed by Baker and Norine in their groundbreaking paper [4], and Baker’s Specialization Lemma [2], which says that the dimension of the complete linear system of a divisor on a smooth curve over a discretely valued field is less than or equal to the rank of its specialization to the dual graph of the special fiber of a strongly semistable regular model. This allows one to translate geometric results about existence of special divisors between curves and metric graphs. For instance, when ρ is nonnegative, the nonemptiness of W_d^r for arbitrary X implies that every metric graph Γ of genus g with rational edge lengths has a divisor of degree d and rank at least r , and a rational approximation argument from [13] then shows that the same holds for metric graphs with arbitrary edge lengths.

For (1), the nonexistence part of the Brill–Noether Theorem, the natural implication goes in the other direction. Suppose X has a strongly semistable regular model whose special fiber has dual graph Γ . If Γ has no divisor of degree d and rank at least r , then neither does X . Similarly, if there is an effective divisor of degree $r + \rho + 1$ on Γ that is not contained in any effective divisor of degree d and rank at least r , then there is such a divisor on X , and it follows that the dimension of $W_d^r(X)$ is at most ρ . See Section 3 for further details.

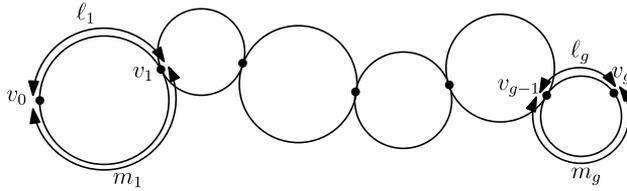


Fig. 1. The graph Γ is a chain of g loops with generic edge lengths.

The graph Γ that we consider is combinatorially a chain of g loops, as shown in Fig. 1, with generic edge lengths. Here, generic means that the tuple of lengths $(\ell_1, \dots, \ell_g, m_1, \dots, m_g)$ in $\mathbb{R}_{>0}^{2g}$ lies outside the union of a finite collection of hyperplanes. See Definition 4.1 for a precise statement.

Theorem 1.1. *Suppose Γ is a chain of g loops with generic edge lengths.*

- (1) *If ρ is negative then Γ has no effective divisors of degree d and rank at least r .*
- (2) *If ρ is nonnegative then Γ has no effective divisors of degree d and rank at least r that contain $(r + \rho + 1)v_0$.*

The existence of graphs with no special divisors when ρ is negative was conjectured by Baker. In particular, Theorem 1.1(1) confirms Conjectures 3.9(2), 3.10(2), and 3.15 of [2]. As a consequence of the theorem, we obtain the following criterion for a curve to be Brill–Noether general.

Corollary 1.2. *Let X be a curve of genus g over a discretely valued field with a regular, strongly semistable model whose special fiber has dual graph Γ . Then X is Brill–Noether general.*

Such curves exist over any complete, discretely valued field [2, Appendix B], and the existence of Brill–Noether general curves over an arbitrary algebraically closed field follows easily. See Section 3.

While our primary interest is the Brill–Noether Theorem, and the tropical criterion for a curve to be Brill–Noether general, we also prove tropical analogues of the dimension part of the Brill–Noether Theorem and an enumerative formula when ρ is zero. We write $W_d^r(\Gamma)$ for the subset of the real torus $\text{Pic}_d(\Gamma)$ parametrizing divisor classes of degree d and rank at least r .

Theorem 1.3. *If ρ is nonnegative, then the dimension of $W_d^r(\Gamma)$ is $\min\{\rho, g\}$.*

This points toward a potentially interesting Brill–Noether theory entirely within tropical geometry, and it is natural to wonder whether these results can be extended to a larger class of graphs.

When ρ is zero and X is general, Griffiths and Harris show that W_d^r consists of finitely many reduced points, and the formula of Kempf, Kleiman, and Laksov says that the number of points is exactly

$$\lambda = g! \prod_{i=0}^r \frac{i!}{(g - d + r + i)!}.$$

This integer λ has many interpretations. It is the $(g - d + r)$ th Catalan number of dimension r [24, p. 133]. By the hook-length formula, it counts standard tableaux on the $(r + 1) \times (g - d + r)$ rectangle [11, Exercise 9, p. 54]. It is also the degree of the Grassmannian of r -planes in \mathbb{P}^{g-d+2r} in its Plücker embedding, and hence counts r -planes in \mathbb{P}^{g-d+2r} meeting g general $g - d + r - 1$ planes [16, Lecture 19]. We prove that it also counts divisor classes of degree d and rank r on Γ , with an explicit bijection to tableaux.

Theorem 1.4. *If ρ is zero, then there are exactly λ distinct divisor classes of degree d and rank r on Γ .*

This exact equality is somewhat surprising; enumerative formulas in tropical geometry often require counts with multiplicities. Here it seems that every divisor should be counted with multiplicity one.

Conjecture 1.5. *Let X be a smooth projective curve of genus g over a discretely valued field, for which the special fiber of a strongly semistable regular model has dual graph Γ . Then every divisor of degree d and rank r on Γ lifts to a divisor of degree d and rank r on X . Furthermore, if ρ is zero, then this lift is unique.*

Remark 1.6. Important refinements of the Brill–Noether Theorem include Gieseker’s proof of the Petri Theorem, which implies that $W_d^r(X)$ is smooth away from $W_d^{r+1}(X)$, for a general curve. Fulton and Lazarsfeld then applied a general connectedness theorem to prove that $W_d^r(X)$ is irreducible [12]. It should be interesting to see if tropical methods may be applicable to these properties of W_d^r as well. Furthermore, there is a close analogy between Brill–Noether loci and degeneracy loci studied by Farkas, where the natural map from $\text{Sym}^n(H^0(C, L))$ to $H^0(C, L^{\otimes n})$ drops rank, and the Strong Maximal Rank Conjecture predicts that these loci should have the naive expected dimension for a general curve [1, Conjecture 5.4]. Hyperplane sections of $K3$ surfaces are not sufficiently general for this conjecture, but it is tempting to hope that tropical methods may be useful instead.

Remark 1.7. When the genus g is small, the moduli space of curves is unirational over \mathbb{Q} . In these cases, rational points are dense in the moduli space and hence there exist Brill–Noether general curves defined over \mathbb{Q} . For large g , the moduli space is of general type, and Lang’s Conjectures predict that rational points should be sparse. In these cases, it is unclear whether there exists a Brill–Noether general curve defined over \mathbb{Q} . The rational numbers carry many discrete valuations, one for each prime, so the criterion given by Corollary 1.2 could potentially be used to produce such curves by explicit computational methods; see Section 6 of [6] for details on computational tests for faithful tropical representations of minimal skeletons. To the best of our knowledge, there are no known examples of Brill–Noether general curves defined over \mathbb{Q} when the moduli space is not unirational.

Remark 1.8. We briefly mention a few related results that have appeared since this paper was written. The special case of Conjecture 1.5 where $r = 1$ and $\rho = 0$ has been proved by Coppens and the first author [8, Theorem 2.3]. The two remaining open conjectures from [2], on the existence of special divisors on complete graphs when ρ is nonnegative, have now been proved by Caporaso [7, Theorem 6.3]. The idea of using Theorem 1.1(2) to prove the classical Brill–Noether Theorem has been developed into a theory of “ranks” for tropical Brill–Noether loci,

analogous to the Baker–Norine theory of ranks of divisors on graphs, by Lim, Potashnik, and the third author [22]. They have also given examples of open subsets of the moduli space of metric graphs where the equality in Theorem 1.3 does not hold.

2. Preliminaries

We briefly recall the theory of divisor classes and ranks of divisors on (metric) graphs, following Baker and Norine [4,2], to which we refer the reader for further details, references, and applications.

2.1. Divisors, equivalence, and chip-firing

Let G be a finite, connected, undirected graph, with a positive real number length assigned to each edge. For compatibility with [4], we allow G to have multiple edges but no loops. Let Γ be the associated metric graph, which is the compact connected metric space obtained by identifying the edges of G with segments of the assigned lengths. Such metric graphs are examples of *abstract tropical curves*, in the sense of [13]. Let g be the genus, or first Betti number, of Γ .

The group $\text{Div}(\Gamma)$ is the free abelian group on the points of Γ , and elements of $\text{Div}(\Gamma)$ are called *divisors* on Γ . The *degree* of a divisor

$$D = a_1 v_1 + \cdots + a_s v_s$$

is the sum of the coefficients $\text{deg}(D) = a_1 + \cdots + a_s$, and D is *effective* if each coefficient a_i is nonnegative.

The subgroup of principal divisors are given by corner loci of piecewise linear functions, as follows. Let ψ be a continuous function on Γ , and suppose that there is a finite subdivision of Γ such that ψ is given by a linear function with integer slope on each edge of the subdivision. Then, for each vertex v of this subdivision, the order $\text{ord}_v(\psi)$ is the sum of the incoming slopes of ψ along the edges containing v , and the divisor of ψ is

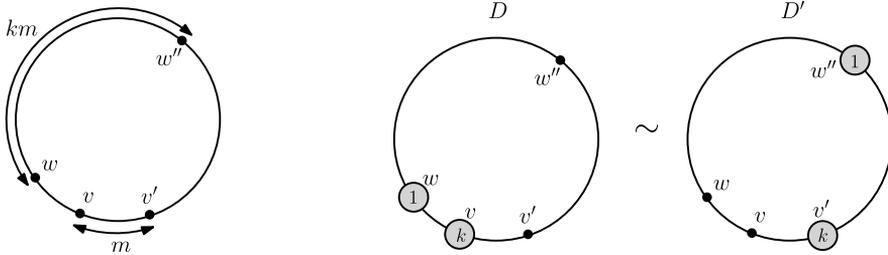
$$\text{div}(\psi) = \sum_v \text{ord}_v(\psi)v.$$

Two divisors D and D' are equivalent, and we write $D \sim D'$, if their difference $D - D'$ is equal to $\text{div}(\psi)$ for some piecewise linear function ψ .

Example 2.1. Let Γ be a single loop, formed by two edges, of lengths ℓ and m respectively, joining vertices v and v' . We label the points of Γ by the interval $[0, \ell + m)$ according their distance from v in the counterclockwise direction. In particular, v is labeled 0 and v' is labeled m . Let $D = kv + w$, where w is either the zero divisor or the point of $\Gamma \setminus v$ labeled by $x \in (0, \ell + m)$, and assume km is not an integer multiple of $\ell + m$. Then D is linearly equivalent to

$$D' = \begin{cases} (k - 1)v' + w' & \text{if } w \text{ is zero,} \\ (k + 1)v' & \text{if } w \text{ is not zero and } x \equiv (k + 1)m \pmod{\ell + m}, \\ kv' + w'' & \text{otherwise.} \end{cases}$$

Here w' and w'' are the points of $\Gamma \setminus v'$ labeled by $x' \equiv -(k - 1)m$ and $x'' \equiv (x - km) \pmod{\ell + m}$, respectively. The equivalence can be seen easily from Dhar’s burning algorithm with base point v' , as presented in [23]. It is also possible to explicitly construct a piecewise linear function ψ such that $D - D' = \text{div}(\psi)$, depending on the combinatorial configuration of the points in question. For instance, in the last case, if x is greater than $(k + 1)m$, then w'' is at distance km from w on the segment of length ℓ , as shown.



A piecewise linear function ψ such that $D - D' = \text{div}(\psi)$ is then given by the constant functions k and zero on the segments $[v, w]$ and $[v', w'']$, respectively, and by linear functions of slopes $-k$ and -1 on the segments $[v, v']$ and $[w, w'']$, respectively.

The group of equivalence classes of divisors

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

is an extension of a real torus of dimension g by \mathbb{Z} [25,3]. The projection to \mathbb{Z} takes the class of a divisor D to its degree, which is well defined because the degree of the divisor of any piecewise linear function is zero. We write $\text{Pic}_d(\Gamma)$ for the space of divisor classes of degree d on Γ .

Remark 2.2. In combinatorics it is customary to refer to divisors on Γ (especially those supported on the vertices of G) as *chip configurations*. The equivalence relation \sim is generated by certain elementary equivalences called *chip-firing moves*. One imagines that $D = a_1 v_1 + \dots + a_s v_s$ is represented by a stack of a_i chips at each v_i (or $|a_i|$ antichips, if a_i is negative). If D is equivalent to D' then $D - D' = \text{div}(\psi)$ for some piecewise linear function ψ , and one imagines a path ψ_t from zero to ψ in the space of piecewise linear functions on Γ . Then

$$D_t = D - \text{div}(\psi_t)$$

is a path from D to D' in the space of divisors equivalent to D . The laws of the game allow chips and antichips to collide and annihilate each other, and pairs of chips and antichips may sometimes be created from the ether. However, if D and D' are both effective, Dhar’s algorithm [9] ensures that $\{\psi_t\}$ can be chosen so that D_t is effective for all t . In other words, two effective divisors are linearly equivalent if and only if one can continuously move the chips from one configuration to the other by a sequence of allowable chip-firing moves. The survey article [18] is an excellent entry point to the vast literature on the combinatorics of chip-firing on graphs, and [13] and [17] include helpful explanations on the generalization of chip-firing to tropical curves and metric graphs.

2.2. Ranks of divisors

Let D be a divisor on Γ . If D is not equivalent to an effective divisor then the *rank* of D , written $r(D)$, is defined to be -1 . Otherwise, $r(D)$ is the largest nonnegative integer r such $D - E$ is equivalent to an effective divisor for every effective divisor E of degree r on Γ .

Remark 2.3. This notion of rank is a natural analogue of the dimension of the complete linear system of a divisor on an algebraic curve. A divisor D on a smooth projective curve X moves in a linear series of dimension at least r if and only if $D - E$ is linearly equivalent to an effective divisor for every effective divisor E of degree r on X .

Remark 2.4. The set of effective divisors on Γ that are equivalent to D is naturally identified with the underlying set of a finite polyhedral complex [15]. The dimension of this complex is bounded below by the rank of D , but is often larger.

The canonical divisor K on Γ is defined as

$$K = \sum_v (\deg v - 2)v,$$

where the sum is over the vertices of Γ . The degree of K is $2g - 2$, as can be checked by a computation of the topological Euler characteristic of Γ , and the rank of K is $g - 1$. The latter is a special case of the following generalization to metric graphs of the Baker–Norine–Riemann–Roch Theorem for graphs.

Tropical Riemann–Roch Theorem. (See [13,25].) *Let D be a divisor on Γ . Then*

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

This formula has many beautiful and useful applications. For instance, any divisor D of degree greater than $2g - 2$ has rank exactly $\deg(D) - g$.

2.3. Reduced divisors and Luo’s Theorem

Two fundamental tools for computing ranks of divisors on graphs are the existence and uniqueness of v -reduced divisors and Luo’s Theorem on rank determining sets, which we now recall.

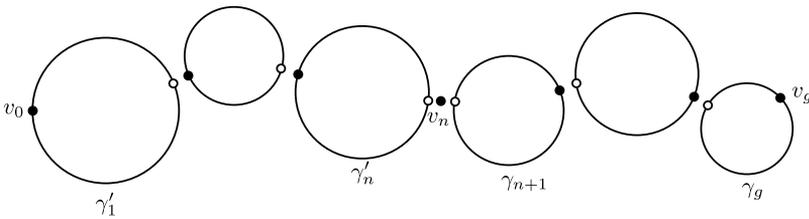
If we fix a basepoint v on Γ , then each divisor D on Γ is equivalent to a unique *v -reduced divisor*, denoted D_0 [17, Theorem 10]. This v -reduced divisor D_0 is characterized by two properties. First, it is effective away from v , so $D_0 + kv$ is effective for k sufficiently large. Second, the points in D_0 are, roughly speaking, as close to v as possible. More precisely, the multiset of distances to v of points in $D_0 + kv$ is lexicographically minimal among the multisets of distances to v for all effective divisors equivalent to $D + kv$. In particular, the coefficient of v in D_0 is maximal among all divisors equivalent to D and effective away from v , and hence D is linearly equivalent to an effective divisor if and only if D_0 is effective.

Remark 2.5. The existence and uniqueness of v -reduced divisors on graphs is a natural analogue of the following existence and uniqueness property for divisors on algebraic curves. If D is a divisor on a smooth projective curve X with a chosen basepoint x , then there is a unique divisor D_0 that is linearly equivalent to D and effective away from x , and whose coefficient of x is maximal among all such divisors. If D is effective, then D_0 is the zero locus of the section of $\mathcal{O}_X(D)$ with maximal order vanishing at x , which is unique up to scaling.

Example 2.6. Let Γ be the chain of loops shown in Fig. 1, so v_n is the point of intersection between the n th loop and the $(n + 1)$ th loop, for $1 \leq n \leq g - 1$, and v_0 and v_g are distinguished points on the first and last loops, respectively. We characterize v_n -reduced divisors on Γ , as follows. Let γ'_i be the i th loop minus v_i , and let γ_j be the j th loop minus v_{j-1} . Then Γ can be decomposed as a disjoint union

$$\Gamma = \left(\bigsqcup_{i=1}^n \gamma'_i \right) \sqcup v_n \sqcup \left(\bigsqcup_{j=n+1}^g \gamma_j \right),$$

as shown.



A divisor D on Γ is v_n -reduced if and only if it is effective away from v_n and each cell γ'_i for $1 \leq i \leq n$, and γ_j for $n + 1 \leq j \leq g$, contains at most one point of D . Indeed, if one of these cells contains more than one point of D then they can be moved closer to v_n by an equivalence similar to the one given in Example 2.1 (with ψ extended by a locally constant function on the complement of that cell), or by applying Dhar’s burning algorithm. Conversely, if each of these cells contains at most one point of D , then the fact that D is v_n -reduced can be checked by Dhar’s burning algorithm. See Section 2 of [23].

The existence and uniqueness of v -reduced divisors facilitate checking whether any given divisor is equivalent to an effective divisor. However, to check if a divisor D has rank at least r , in principle we must check whether $D - E$ is equivalent to an effective divisor for all divisors E of rank r , of which there are uncountably many if r is positive. Luo has recently improved this situation by showing that there is a small, finite set of points in Γ with the property that, for any D and any r , it is enough to check for divisors E whose support is contained in A . Here, the support of an effective divisor is the set of points that appear in it with nonzero coefficient.

Luo’s Theorem. (See [23].) Let A be a finite subset of Γ such that the closure in Γ of each connected component of $\Gamma \setminus A$ is contractible. Let D be a divisor on Γ and suppose that, for every effective divisor E of degree r whose support is contained in A , the difference $D - E$ is effective. Then the rank of D is at least r .

Remark 2.7. Such a subset A can always be chosen with size at most $g + 1$. Luo’s Theorem is then a natural analogue of the following fact about divisors on curves. Let $A = \{x_0, \dots, x_g\}$ be a set of distinct points on a smooth projective curve of genus g . Then a divisor D moves in a linear series of dimension at least r if and only if $D - E$ is linearly equivalent to an effective divisor for every effective divisor E with support in A . For a proof, due to Varley, see Remark 3.13 of [23].

Example 2.8. Let Γ be the chain of loops shown in Fig. 1, and let $A = \{v_0, \dots, v_g\}$. Then the closure of each connected component of $\Gamma \setminus A$ is either the top half or the bottom half of one of the loops, and hence is contractible. In this case, Luo’s Theorem says that a divisor D on Γ has rank at least r if and only if $D - E$ is equivalent to an effective divisor for any effective divisor $E = r_0v_0 + \dots + r_gv_g$ of degree r .

2.4. Specialization

We conclude this preliminary section with a review of Baker’s Specialization Lemma, which relates dimensions of complete linear series on certain curves over discretely valued fields to ranks of divisors on graphs.

Let K be a discretely valued field, with valuation ring R and residue field k , and let X be a smooth projective curve over K . A *strongly semistable regular model* of X is a regular scheme \mathfrak{X} over $\text{Spec } R$ whose general fiber \mathfrak{X}_K is isomorphic to X and whose special fiber \mathfrak{X}_k is a reduced union of geometrically irreducible smooth curves X_0, \dots, X_s that meet only at simple nodes defined over k . The dual graph G of the special fiber has vertices v_0, \dots, v_s corresponding to the irreducible components of \mathfrak{X}_k and one edge joining v_i to v_j for each point of intersection in $X_i \cap X_j$. Let Γ be the associated metric graph, where each edge is assigned length 1.

Each point in $X(K)$ specializes to a smooth point in the special fiber. We write $\tau : X(K) \rightarrow \Gamma$ for the induced map which takes a point x to the vertex v_i corresponding to the irreducible component of \mathfrak{X}_k that contains the specialization of x . This map τ is compatible with finite field extensions, as follows. If K' is a finite extension of K then there is a unique relatively minimal strongly semistable regular model \mathfrak{X}' of $X \times_K K'$ that dominates $\mathfrak{X} \times_R R'$. Let G' be the dual graph of the special fiber of \mathfrak{X}' , and let Γ' be the associated metric graph in which each edge is assigned length $1/e$, where e is the ramification index of K'/K . Then Γ' is naturally isomorphic to Γ , and the induced specialization maps $X(K') \rightarrow \Gamma$ for all finite extensions K'/K together give a well-defined geometric specialization map

$$\tau : X(\overline{K}) \rightarrow \Gamma.$$

Furthermore, the induced map on free abelian groups $\tau_* : \text{Div}(X_{\overline{K}}) \rightarrow \text{Div}(\Gamma)$ respects linear equivalence, and hence descends to a degree preserving group homomorphism

$$\tau_* : \text{Pic}(X_{\overline{K}}) \rightarrow \text{Pic}(\Gamma).$$

See Section 2 of [2] for details, further references, and a proof of the following.

Specialization Lemma. *Let D be a divisor on $X_{\overline{K}}$. Then*

$$r(\tau_*(D)) \geq r(D),$$

where $r(D)$ is the dimension of the complete linear system $|D|$ on $X_{\overline{K}}$.

In particular, if Γ has no divisors of degree d and rank at least r , then the Brill–Noether locus W_d^r in $\text{Pic}_d(X)$ is empty.

3. From tropical to classical Brill–Noether theory

Here we explain how the classical Brill–Noether theorem, over \mathbb{C} or an arbitrary algebraically closed field, follows from Theorem 1.1. First, we deduce Corollary 1.2, which establishes the existence of Brill–Noether general curves over an arbitrary complete, discretely valued field.

Proof of Corollary 1.2. Suppose X is a curve of genus g over a discretely valued field K with a regular semistable model whose special fiber has dual graph Γ . By Theorem 1.1(1), if ρ is negative then Γ has no divisors of degree d and rank at least r . Therefore, by Baker’s Specialization Lemma, the Brill–Noether locus $W_d^r(X)$ is empty. It remains to show that if ρ is nonnegative then $\dim W_d^r(X)$ is at most ρ .

If X has a divisor D of degree d with $r(D) \geq r$, then every effective divisor of degree r on X is contained in an effective divisor of degree d and rank r that is equivalent to D . If the class of D moves in a 1-dimensional algebraic family in $W_d^r(X)$, then every effective divisor of degree $r + 1$ on X is contained in an effective divisor whose class lies in that family, and a straightforward induction on dimension shows that every effective divisor of degree $r + \dim W_d^r(X)$ is contained in an effective divisor whose class lies in $W_d^r(X)$. Therefore, to prove that $\dim W_d^r(X)$ is at most ρ , it will suffice to produce an effective divisor of degree $r + \rho + 1$ that is not contained in any effective divisor of degree d such that $r(D) \geq r$.

Now, choose a point x in $X_{\bar{K}}$ specializing to v_0 . Then $(r + \rho + 1)x$ specializes to $(r + \rho + 1)v_0$, which is not contained in any effective divisor of degree d and rank at least r , by Theorem 1.1(2). By Baker’s Specialization Lemma, it follows that $(r + \rho + 1)x$ is not contained in any effective divisor D such that $r(D) \geq r$. \square

Remark 3.1. The above proof uses the fact that v_0 is a vertex of Γ , and hence is in the image of $X_{\bar{K}}$. These arguments can be extended to show that if X' is a curve over a discretely valued field with a regular semistable model whose special fiber has dual graph Γ' then every effective divisor of degree $r + \dim W_d^r(X')$ on Γ' is contained in an effective divisor of degree d and rank at least r . See [22] for details.

Brill–Noether Theorem over \mathbb{C} . The existence of Brill–Noether general curves over \mathbb{C} follows easily from Corollary 1.2. The field of Laurent series $\mathbb{Q}((t))$ is complete and discretely valued. Therefore, Corollary 1.2 and the existence theorem from Appendix B of [2] show that for each genus g there is a Brill–Noether general curve X of genus g over $\mathbb{Q}((t))$. The algebraic closure of $\mathbb{Q}((t))$ is isomorphic to \mathbb{C} as an abstract field, because any two uncountable algebraically closed fields of the same cardinality and characteristic are isomorphic. Therefore, there is a (noncontinuous) embedding of $\mathbb{Q}((t))$ as a subfield of \mathbb{C} . Fix such an embedding. Then

$$X_{\mathbb{C}} = X \times_{\text{Spec } \mathbb{Q}((t))} \text{Spec } \mathbb{C}$$

is a Brill–Noether general curve over \mathbb{C} . \square

The argument above is insufficient to prove the Brill–Noether Theorem over fields such as $\bar{\mathbb{F}}_p$ that have no subfields with nontrivial valuations. Nevertheless, the existence of Brill–Noether

general curves over $\overline{\mathbb{F}}_p$, or an arbitrary algebraically closed field, follows from Corollary 1.2 and the existence of the moduli space of curves.

Brill–Noether Theorem over an algebraically closed field. The coarse moduli space M_g of smooth projective curves of genus g is a scheme of finite type over $\text{Spec } \mathbb{Z}$, and the locus U of Brill–Noether general curves is Zariski open and hence is also a scheme of finite type over $\text{Spec } \mathbb{Z}$. By Corollary 1.2, the scheme U has points over $\mathbb{Q}((t))$ and $\mathbb{F}_p((t))$ for all primes p , and hence surjects onto $\text{Spec } \mathbb{Z}$. Therefore, the fiber of U over any prime field is a nonempty scheme of finite type, and hence U has points over any algebraically closed field. \square

4. Chip-firing on a generic chain of loops

Let Γ be a chain of g loops, as pictured in Fig. 1. The top and bottom segments of the i th loop connect the vertex v_{i-1} to v_i and have length ℓ_i and m_i , respectively. In particular, the total length of the i th loop is $\ell_i + m_i$.

Definition 4.1. The graph Γ is *generic* if none of the ratios ℓ_i/m_i is equal to the ratio of two positive integers whose sum is less than or equal to $2g - 2$.

Some genericity condition on these lengths is necessary for nonexistence of special divisors. For instance, if $\ell_i = m_i$ for all i then Γ is hyperelliptic [5], meaning that it has a divisor of degree two and rank one. Our genericity condition is easily achieved with integer edge lengths. For instance, one may take $\ell_i = 2g - 2$ and $m_i = 1$ for all i .

Notation 4.2. Throughout the remainder of the paper, we assume that Γ is generic, in the sense of Definition 4.1. Our main results, stated in the introduction, are trivial when g is zero or one and follow from the Tropical Riemann–Roch Theorem when d is greater than $2g - 2$, so we also assume that g is at least two and d is at most $2g - 2$.

Our main combinatorial tool in the proof of Theorems 1.1, 1.3, and 1.4 is a *lingering lattice path* associated to each v_0 -reduced divisor of degree d on Γ , defined as follows.

Definition 4.3. A lingering lattice path P in \mathbb{Z}^r is a sequence p_0, \dots, p_g of points in \mathbb{Z}^r such that each successive difference $p_i - p_{i-1}$ is either $(-1, \dots, -1)$, a standard basis vector, or zero.

When $p_i - p_{i-1}$ is zero, we say that the lattice path *lingers* at the i th step.

Recall that a divisor D on Γ is v_0 -reduced if and only if it is effective away from v_0 and each cell γ_i in the decomposition given in Example 2.6 contains at most one point of D . We label such a point by its distance from v_{i-1} in the counterclockwise direction, so v_i is labeled by m_i . This leads to a natural bijection

$$\{v_0\text{-reduced divisors on } \Gamma\} \longleftrightarrow \mathbb{Z} \times \mathbb{R}/(\ell_1 + m_1) \times \dots \times \mathbb{R}/(\ell_g + m_g),$$

taking a v_0 -reduced divisor D to the data $(d_0; x_1, \dots, x_g)$, where d_0 is the coefficient of v_0 in D , and x_i is the location of the unique point in D on γ_i , if there is one, and zero otherwise. The lingering lattice path associated to D is defined in terms of this data as follows. We label the

coordinates of \mathbb{Z}^r and \mathbb{R}^r from zero to $r - 1$, and write $p_i(j)$ for the j th coordinate of p_i . We write \mathcal{C} for the open Weyl chamber

$$\mathcal{C} = \{y \in \mathbb{R}^r \mid y(0) > \dots > y(r - 1) > 0\}.$$

Definition 4.4. Let D be the v_0 -reduced divisor of degree d corresponding to $(d_0; x_1, \dots, x_g)$. Then the associated lingering lattice path P in \mathbb{Z}^r starts at $(d_0, d_0 - 1, \dots, d_0 - r + 1)$ with steps given by

$$p_i - p_{i-1} = \begin{cases} (-1, \dots, -1) & \text{if } x_i = 0; \\ e_j & \text{if } x_i \equiv (p_{i-1}(j) + 1)m_i \pmod{\ell_i + m_i}, \text{ and} \\ & p_{i-1} \text{ and } p_{i-1} + e_j \text{ are in } \mathcal{C}; \\ 0 & \text{otherwise,} \end{cases}$$

where e_0, \dots, e_{r-1} are the standard basis vectors in \mathbb{Z}^r .

To see that the i th step in the lingering lattice path is well defined when p_{i-1} is in \mathcal{C} , one uses the genericity condition on Γ , as follows. Note first that the coordinates of each p_i satisfy the inequalities $p_i(0) > \dots > p_i(r - 1)$. This is because p_0 satisfies these inequalities, and the inequalities are preserved by each step in the lingering lattice path. Furthermore, since only $d - d_0$ of the x_i are nonzero and the coordinates of d_0 are bounded above by d_0 , all of the coordinates of each p_i are bounded above by d , and hence by $2g - 2$. In particular, if p_{i-1} is in \mathcal{C} then its coordinates are distinct integers between 1 and $2g - 2$. The genericity condition then ensures that $x_i \equiv (p_{i-1}(j) + 1)m_i \pmod{\ell_i + m_i}$ holds for at most one j , and hence the i th step is well defined, as required.

As noted above, the coordinates of each point in the lingering lattice path are strictly decreasing, so p_i is in the chamber \mathcal{C} if and only if its last coordinate is positive. Note also that if p_{i-1} is in \mathcal{C} then $p_{i-1} + e_j$ is not in \mathcal{C} exactly when $p_{i-1}(j - 1)$ is only one more than $p_{i-1}(j)$. Since $p_0(j - 1)$ is only one more than $p_0(j)$, the condition on $p_{i-1} + e_j$ guarantees that, among the first i steps of P , the number of steps in direction e_j is always less than or equal to the number in direction e_{j-1} , for all $i, j \geq 1$.

Remark 4.5. The three cases for $p_i - p_{i-1}$ are parallel to the three cases for D' in Example 2.1 and the three cases in the proof of Theorem 4.6, below. For an interpretation in terms of a chip-firing game on Γ , see Remark 4.9.

Let D be a v_0 -reduced divisor of degree d on Γ , and let P be the associated lingering lattice path in \mathbb{Z}^r . Our main technical result is then the following.

Theorem 4.6. *The divisor D has rank at least r if and only if the associated lingering lattice path P lies entirely in the open Weyl chamber \mathcal{C} .*

Before giving the proof of Theorem 4.6, we explain how it implies Theorems 1.1, 1.3, and 1.4, and give examples illustrating the bijection between divisors and tableaux when ρ is zero.

Proof of Theorem 1.1. Suppose D is a divisor of degree d and rank at least r on Γ , and let $(d_0; x_1, \dots, x_g)$ be the data associated to the v_0 -reduced divisor equivalent to D . We must show that $\rho = g - (r + 1)(g - d + r)$ is nonnegative and d_0 is less than or equal to $r + \rho$.

Exactly $d - d_0$ of the x_i are nonzero, so the lingering lattice path P includes $g - d + d_0$ steps in the direction $(-1, \dots, -1)$. In particular, the last coordinate of p_g is

$$p_g(r - 1) = d - g - r + 1 + \#\{\text{steps of } P \text{ in direction } e_{r-1}\}.$$

By Theorem 4.6, the lattice path P lies in the open Weyl chamber \mathcal{C} , so $p_g(r - 1)$ is strictly positive. Therefore P includes at least $g - d + r$ steps in direction e_{r-1} . By construction, the number of steps of P in the e_i direction is at least the number of steps in the e_{i+1} direction, for all i , so P must include at least $g - d + r$ steps in each of the r coordinate directions. Therefore, the total number steps is

$$g \geq r(g - d + r) + (g - d + d_0).$$

Rearranging terms then shows that $d_0 \leq r + \rho$, which proves part (2) of the theorem. It remains to show that ρ is nonnegative.

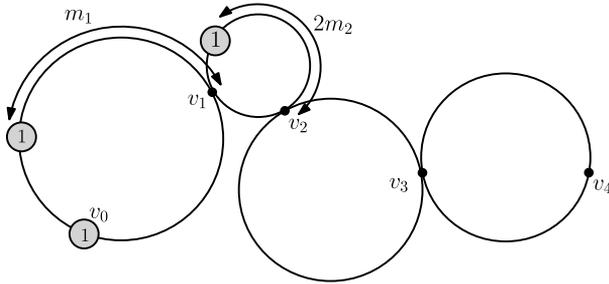
Since D is v_0 -reduced and of rank at least r , $D - rv_0$ must be effective. Hence d_0 is at least r , and $\rho \geq 0$, as required. \square

Proof of Theorem 1.3. Each divisor class of degree d and rank r on Γ is associated to some lattice path in \mathcal{C} , and the proof of Theorem 1.1 above shows that this path must include at least $g - d + r$ steps in each of the r coordinate directions, as well as at least $g - d + r$ steps in the direction $(-1, \dots, -1)$. Therefore, the lattice path has at most $\rho = g - (r + 1)(g - d + r)$ lingering steps. Given a lingering lattice path p_0, \dots, p_g with $\min\{\rho, g\}$ steps, the v_0 -reduced realizations of this lattice path are chip configurations with $p_0(0)$ chips at v_0 , a chip in location $x_i \equiv (p_{i-1}(j) + 1)m_i \pmod{\ell_i + m_i}$ in the cell γ_i if $p_i - p_{i-1} = e_j$, a chip anywhere else in γ_i if the i th step is lingering, and no chips on γ_i otherwise. Since there are finitely many such paths, and the dimension of the space of realizations is equal to the number of lingering steps, it follows that the dimension of $W_d^r(\Gamma)$ is $\min\{\rho, g\}$, as required. \square

Proof of Theorem 1.4. Suppose ρ is zero. Then each v_0 -reduced divisor D of degree d and rank r has $d_0 = r$ and the associated lattice path P has exactly $g - d + r$ steps in each of the coordinate directions, $g - d + r$ steps in the direction $(-1, \dots, -1)$, and no lingering steps. Moreover, each such lattice path corresponds to a unique divisor of degree d and rank r . Therefore, there is a natural bijection between divisors of degree d and rank r on Γ and g -step lattice paths from $(r, \dots, 1)$ to itself in the open Weyl chamber \mathcal{C} . These lattice paths are in natural bijection with standard tableaux on the rectangular shape $(r + 1) \times (g - d + r)$, as follows. We label the columns of the tableau from zero to r . Then the number i appears in the column j column of the tableau corresponding to P for $0 \leq j \leq r - 1$ if the i th step of P is in the j th coordinate direction, and in column r if the i th step of P is in the direction $(-1, \dots, -1)$. \square

Example 4.7. Consider the case where $(g, r, d) = (4, 1, 3)$. Then ρ is zero and λ is two, corresponding to the classical fact that there are exactly two lines meeting four general lines in \mathbb{P}^3 . Theorem 1.4 says that there are exactly two v_0 -reduced divisors of degree three and rank one on Γ , corresponding to the lattice paths $1, 2, 3, 2, 1$ and $1, 2, 1, 2, 1$ in \mathbb{Z} , respectively.

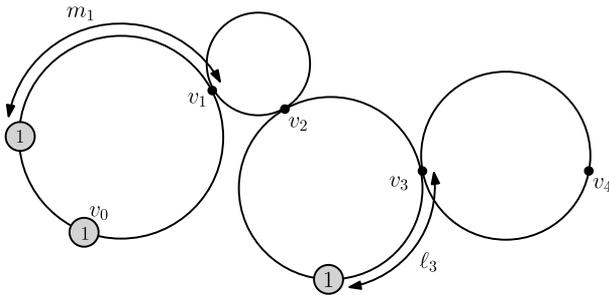
The first path is associated to the divisor



which has points on the first two loops, since the first two steps of the lattice path are in the positive direction. This corresponds to the tableau

1	3
2	4

in which one and two appear in the first column. The second path is associated to the divisor



which has points on the first and third loops. This corresponds to the tableau

1	2
3	4

in which one and three appear in the first column. It is a pleasant exercise to use the chip-firing moves in Example 2.1 to show that each of these divisors has rank one.

Example 4.8. Consider the case $(g, r, d) = (12, 3, 12)$. Then ρ is zero and λ is 462. The v_0 -reduced divisors of degree 12 and rank 3 correspond to the 462 standard tableaux on a 4×3 rectangle. For instance, the tableau

1	3	4	6
2	5	7	9
8	10	11	12

corresponds to the lattice path

$$(3, 2, 1), (4, 2, 1), (5, 2, 1), (5, 3, 1), (5, 3, 2), (5, 4, 2), (4, 3, 1), \\ (4, 3, 2), (5, 3, 2), (4, 2, 1), (4, 3, 1), (4, 3, 2), (3, 2, 1).$$

The associated divisor has multiplicity three at v_0 and one carefully chosen point on each cell γ_i for $i \neq 6, 9, 12$. The locations of these points can be read off from the tableau and lattice path, as follows. The columns of the tableau are labeled from zero to three and 10 appears in the column labeled one, so the point in γ_{10} is at distance

$$x_{10} \equiv 3m_{10} \pmod{\ell_{10} + m_{10}}$$

from v_9 , in the counterclockwise direction, with the coefficient of m_{10} given by the formula $p_9(1) + 1 = 3$.

Remark 4.9. One may think of a divisor of degree d and rank at least r on Γ as a winning strategy in the following game, which we call the *Brill–Noether game*. First, Brill chooses a divisor D of degree d on Γ , and shows it to Noether. Then, Noether chooses an effective divisor E of degree r on Γ and shows it to Brill. Finally, Brill performs chip-firing operations on $D - E$, and wins if he reaches an effective divisor. Otherwise Noether wins. One can imagine a similar game with divisors on an algebraic curve.

Our main result says that Brill wins the game if and only if ρ is nonnegative, and our proof is a combinatorial classification of Brill’s winning strategies. In terms of this game, the theory of v_0 -reduced divisors implies that Brill has an optimal strategy in which he chooses an effective v_0 -reduced divisor with $d_0 \geq r$, and Luo’s Theorem implies that Noether has an optimal strategy in which he chooses a divisor supported in $\{v_0, \dots, v_g\}$.

When the game is played, Brill typically starts with a large pile of chips at v_0 and moves this pile to the right using chip-firing moves, as in Example 2.1. The pile grows when it picks up an additional chip that Brill placed at x_i in the i th loop, if x_i is carefully chosen. The pile shrinks when it encounters Noether’s antichips in $-E$, or when it crosses over an empty loop and one chip is left behind in the process of moving the pile to the next vertex. The dynamics of this growing and shrinking pile of chips are encoded in the coordinates of the points in the lattice path P . Roughly speaking, the j th coordinate $p_i(j)$ is the size of the pile when it reaches v_i , if Noether has distributed j antichips over the vertices v_0, \dots, v_i , assuming optimal play. For instance, if Noether places j antichips at v_0 , then Brill is left with $p_0(j) = d_0 - j$ chips at v_0 .

The preceding remarks are made precise in the following proposition, which will be used in the proof of Theorem 4.6. As before, let γ_i be the i th loop minus v_{i-1} , which is an open cell in the decomposition of $\Gamma \setminus v_0$, described in Example 2.6. The restriction of a divisor $D = a_1 w_1 + \dots + a_s w_s$ to γ_j is defined as

$$D|_{\gamma_j} = \sum_{w_i \in \gamma_j} a_i w_i.$$

Proposition 4.10. *Suppose p_0, \dots, p_{n-1} are in \mathcal{C} . Let E_n be an effective divisor of degree $j < r$ with support contained in $\{v_0, \dots, v_n\}$, and let D_n be the v_n -reduced divisor equivalent to $D - E_n$. Then*

- (1) the coefficient of v_n in D_n is at least $p_n(j)$; and
- (2) for $i > n$, the restriction $D_n|_{\gamma_i}$ is equal to $D|_{\gamma_i}$.

Furthermore, for each $j < r$ there exists an effective divisor E_n of degree j with support in $\{v_0, \dots, v_n\}$ such that equality holds in (1).

Proof. Write $E = r_0v_0 + \dots + r_nv_n$. If $n = 0$, then D_0 is $D - r_0v_0$, and the proposition is clear. We proceed to prove (1) and (2) by induction on n . Let D_{n-1} be the v_{n-1} -reduced divisor equivalent to $D - r_0v_0 - \dots - r_{n-1}v_{n-1}$. By the induction hypothesis, we may assume that the coefficient k of v_{n-1} in D_{n-1} is at least $p_{n-1}(j - r_n)$ and the restriction of D_{n-1} to γ_i is equal to $D|_{\gamma_i}$ for $i \geq n$.

Let D'_{n-1} be the divisor equivalent to D_{n-1} obtained by moving the chips at v_{n-1} to v_n , as in Example 2.1. Then D'_{n-1} is effective and each cell of $\Gamma \setminus v_n$ contains at most one point of D'_{n-1} , so D'_{n-1} is v_n -reduced. It follows that $D_n = D'_{n-1} - r_nv_n$ and the restriction of D_n to γ_i is equal to $D|_{\gamma_i}$ for $i > n$, as required. We now show that the coefficient of v_n in D_n is at least $p_n(j)$ by considering three cases according to the position of the point of D , if any, in γ_n .

Case 1. There is no point of D in γ_n . In this case, the n th step of P is in the direction $(-1, \dots, -1)$, so $p_n(j)$ is equal to $p_{n-1}(j) - 1$. In terms of the Brill–Noether game, the pile of chips at v_{n-1} shrinks by one as it moves from v_{n-1} to v_n , since one chip must be left behind in γ'_n . In other words, the equivalence $D'_{n-1} \sim D_{n-1}$ is given by the first case in Example 2.1, and hence the coefficient of v_n in D'_{n-1} is $k - 1$, and that of D_n is $k - 1 - r_n$. Now k is at least $p_{n-1}(j - r_n)$, which is at least $p_{n-1}(j) + r_n$, since the coordinates of p_{n-1} are strictly decreasing integers. Therefore, the coefficient of v_n in D_n is greater than or equal to $p_{n-1}(j) - 1$, as required. For later use note that equality holds if $k = p_{n-1}(j - r_n)$ and moreover the entries of p_i (and hence of p_{i-1}) at positions $j - r_n, \dots, j$ are consecutive integers.

Case 2. The point of γ_n in D is at $x_n \equiv (p_{n-1}(j - r_n) + 1)m_n \pmod{\ell_n + m_n}$. By hypothesis k is at least $p_{n-1}(j - r_n)$. If it is equal to $p_{n-1}(j - r_n)$ then the pile of chips picks up one extra as it moves from v_{n-1} to v_n . In other words, the equivalence $D_{n-1} \sim D'_{n-1}$ is given by the second case in Example 2.1, and hence the coefficient of v_n in D_n is $p_{n-1}(j - r_n) + 1 - r_n \geq p_n(j - r_n) - r_n \geq p_n(j)$, as required. For later use note that equality holds if $p_{n-1}(j - r_n) + 1$ does not occur among the entries of p_{n-1} and moreover the entries of p_n at positions $j - r_n, \dots, j$ are consecutive integers.

On the other hand, if k is greater than $p_{n-1}(j - r_n)$ then the equivalence $D_{n-1} \sim D'_{n-1}$ is given by the third case in Example 2.1, and the coefficient of v_n in D_n is $k - r_n$, which is again greater than or equal to $p_n(j)$.

Case 3. There is a point of γ_n in D , but not at x_n . In this case, the pile does not shrink as it moves from v_{n-1} to v_n , and $p_n(j - r_n)$ is equal to $p_{n-1}(j - r_n)$. So the coefficient of v_n in D_n is at least $k - r_n \geq p_n(j)$, as required. It is worth noting that if k is greater than $p_{n-1}(j - r_n)$ and the point of D is at $(k + 1)m_n$ then the pile will grow as it moves from v_{n-1} to v_n , as in Case 2, but the lattice path still lingers. However, if $k = p_{n-1}(j - r_n)$ then the pile does not grow, and moreover the entries of p_n at positions $j - r_n, \dots, j$ are consecutive integers.

It remains to show that E_n can be chosen so that equality holds in (1). Again, we proceed by induction on n . For $n = 0$ we can take $E_n = jv_0$. For the induction step from $n - 1$ to n let r_n be maximal such that the entries of p_n at positions $j - r_n, \dots, j$ are consecutive integers, and let E_{n-1} be an effective divisor of degree $j - r_n$ with support in $\{v_0, \dots, v_{n-1}\}$ such that the

coefficient k of v_{n-1} of the v_{n-1} -reduced divisor equivalent to $D - E_{n-1}$ equals $p_{n-1}(j - r_n)$. To prove that $E_n := E_{n-1} + r_n v_n$ has the required property we go through the Cases 1–3 above. In Cases 1 and 3 the equality $k = p_{n-1}(j - r_n)$ and the fact that the entries of p_n at positions $j - r_n, \dots, j$ are consecutive suffice to conclude that the coefficient of v_n in E_n equals $p_{n-1}(j)$. In Case 2 the only thing that could go wrong is that the pile of $k = p_{n-1}(j - r_n)$ chips at v_{n-1} picks up an additional chip when moved to v_n but p lingers at the n th step since the entries of $p_{n-1} + e_{j-r_n}$ are not all distinct. Then we have $j - r_n > 0$ and the entry of p_{n-1} at position $j - r_n - 1$ equals $p_{n-1}(j - r_n) + 1$, which is then also the entry of p_n at position $j - r_n - 1$. But then the entries of p_n at positions $j - r_n - 1, \dots, j$ are consecutive integers, contradicting the maximality of r_n . We conclude that E_n has the required property. \square

We conclude by applying the proposition to prove Theorem 4.6.

Proof of Theorem 4.6. Suppose the lattice path P lies in the open Weyl chamber \mathcal{C} . Let $E = r_0 v_0 + \dots + r_g v_g$ be an effective divisor of degree r , and let n be the largest index such that r_n is strictly positive. Let $E_n = E - v_n$. By Proposition 4.10, the difference $D - E_n$ is equivalent to an effective divisor D'_n in which the coefficient of D'_n is at least $p_n(r - 1)$, which is strictly positive since P is in \mathcal{C} . Therefore $D - E$ is equivalent to the effective divisor $D'_n - v_n$. Since E is an arbitrary effective divisor of degree r with support in $\{v_0, \dots, v_g\}$, it follows by Luo’s Theorem that D has rank at least r .

For the converse, suppose the lattice path P does not lie in \mathcal{C} , and let n be the smallest index such that p_n is not in \mathcal{C} . By the construction of P , all coordinates of p_i are nonnegative for $i \leq n$, and $p_n(r - 1) = 0$. By Proposition 4.10, there exists an effective divisor $E_n = r_0 v_0 + \dots + r_n v_n$ of degree $r - 1$ such that the coefficient of v_n in the v_n -reduced divisor D'_n equivalent to $D - E_n$ is zero. Then $E = E_n + v_n$ is an effective divisor of degree r , and the v_n -reduced divisor equivalent to $D - E$ is $D'_n - v_n$, which is not effective. Therefore, $D - E$ is not equivalent to any effective divisor, and hence D has rank less than r , as required. \square

Acknowledgments

This research was done in part during the special semester on Tropical Geometry at MSRI in Berkeley, and we are grateful for the hospitality and ideal working environment provided by this program. We thank the participants of the chip-firing seminar at MSRI, including E. Brugallé, E. Cotterill, C. Haase, E. Katz, D. Maclagan, G. Musiker, J. Yu, and I. Zharkov, for many enlightening discussions, and M. Baker, F. Schroeter, and the referee for helpful comments on an earlier draft.

The work of F.C. is supported by a postdoctoral fellowship from the Research Foundation – Flanders. The work of J.D. was partially supported by MSRI and by a Vidi grant of the Netherlands Organization for Scientific Research (NWO). The work of S.P. was partially supported by the Clay Mathematics Institute and NSF grant DMS 1068689. The work of E.R. was supported by a research grant from VPUE of Stanford University.

References

[1] M. Aprodu, G. Farkas, Koszul cohomology and applications to moduli, in: D. Ellwood, E. Previato (Eds.), Grassmannians, Moduli Spaces, and Vector Bundles, CMI/AMS, 2011.
 [2] M. Baker, Specialization of linear systems from curves to graphs, Algebra Number Theory 2 (6) (2008) 613–653.
 [3] M. Baker, X. Faber, Metric properties of the tropical Abel–Jacobi map, J. Algebraic Combin. 33 (3) (2011) 349–381.

- [4] M. Baker, S. Norine, Riemann–Roch and Abel–Jacobi theory on a finite graph, *Adv. Math.* 215 (2) (2007) 766–788.
- [5] M. Baker, S. Norine, Harmonic morphisms and hyperelliptic graphs, *Int. Math. Res. Not.* 2009 (15) (2009) 2914–2955.
- [6] M. Baker, S. Payne, J. Rabinoff, Nonarchimedean geometry, tropicalization, and metrics on curves, preprint, arXiv:1104.0320v1, 2011.
- [7] L. Caporaso, Algebraic and combinatorial Brill–Noether theory, preprint, arXiv:1106.1140v1, 2011.
- [8] F. Cools, M. Coppens, Linear pencils on graphs and real curves, arXiv:1006.1770v1, 2010.
- [9] D. Dhar, Self-organized critical state of sandpile automaton models, *Phys. Rev. Lett.* 64 (14) (1990) 1613–1616.
- [10] D. Eisenbud, J. Harris, A simpler proof of the Gieseker–Petri theorem on special divisors, *Invent. Math.* 74 (2) (1983) 269–280.
- [11] W. Fulton, *Young Tableaux*, London Math. Soc. Stud. Texts, vol. 35, Cambridge University Press, Cambridge, 1997.
- [12] W. Fulton, R. Lazarsfeld, On the connectedness of degeneracy loci and special divisors, *Acta Math.* 146 (3–4) (1981) 271–283.
- [13] A. Gathmann, M. Kerber, A Riemann–Roch theorem in tropical geometry, *Math. Z.* 259 (1) (2008) 217–230.
- [14] P. Griffiths, J. Harris, On the variety of special linear systems on a general algebraic curve, *Duke Math. J.* 47 (1) (1980) 233–272.
- [15] C. Haase, G. Musiker, J. Yu, Linear systems on tropical curves, *Math. Z.* (2009), <http://dx.doi.org/10.1007/s00209-011-0844-4>, in press, arXiv:0909.3685.
- [16] J. Harris, *Algebraic Geometry: A First Course*, Grad. Texts in Math., vol. 133, Springer-Verlag, New York, 1992.
- [17] J. Hladký, D. Král', S. Norine, Rank of divisors on tropical curves, preprint, arXiv:0709.4485v2, 2008.
- [18] A. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, D. Wilson, Chip-firing and rotor-routing on directed graphs, in: *In and Out of Equilibrium 2*, in: *Progr. Probab.*, vol. 60, Birkhäuser, Basel, 2008, pp. 331–364.
- [19] G. Kempf, *Schubert Methods with an Application to Algebraic Curves*, Publ. Math. Centrum, 1971.
- [20] S. Kleiman, D. Laksov, On the existence of special divisors, *Amer. J. Math.* 94 (1972) 431–436.
- [21] R. Lazarsfeld, Brill–Noether–Petri without degenerations, *J. Differential Geom.* 23 (3) (1986) 299–307.
- [22] C.M. Lim, S. Payne, N. Potashnik, A note on Brill–Noether theory and rank-determining sets for metric graphs, *Int. Math. Res. Not.* (2012), <http://dx.doi.org/10.1093/imrn/rnr233>, in press.
- [23] Y. Luo, Rank-determining sets of metric graphs, *J. Combin. Theory Ser. A* 118 (2011) 1775–1793.
- [24] P. MacMahon, *Combinatory Analysis*, Chelsea Publishing Co., New York, 1960.
- [25] G. Mikhalkin, I. Zharkov, Tropical curves, their Jacobians and theta functions, in: *Curves and Abelian Varieties*, in: *Contemp. Math.*, vol. 465, Amer. Math. Soc., Providence, RI, 2008, pp. 203–230.