

LIFTING RANK 2 TROPICAL DIVISORS

DUSTIN CARTWRIGHT

ABSTRACT. Tropical geometry gives a way to bound the ranks of divisors on curves in terms of the combinatorics of the dual graph of a degeneration. We show that for a family of examples, curves realizing can fail to exist or might only exist over certain characteristics or over certain fields of definition. In particular, we show that the problem of lifting a rank 2 divisor on a graph is as hard as the matroid realizability problem. Our examples also apply to the theory of metrized complexes and weighted graphs.

1. INTRODUCTION

The specialization inequality states that the rank of a divisor on a curve is at most the rank of its specialization to the dual graph [Bak08]. Many improvements have been made to sharpen this inequality by incorporating additional information about the components of the special fiber of the graph [AC13, AB]. In all cases, these bounds can be strict because there are many algebraic curves with the same specialization. Thus, the natural question is whether, for a given graph and divisor, there exists an algebraic curve and divisor of the same rank and specializing to the given ones. We call such a curve together with the divisor a *lifting* of the divisor on the graph. In this paper, we show that the existence of a lifting can depend on the field of definition:

Theorem 1.1. *Let X be a scheme of finite type over $\text{Spec } \mathbb{Z}$. Then there exists a graph Γ with a rank 2 divisor D such that, for any infinite field k , Γ and D are the specializations of a semistable curve over $k[[t]]$ and a rank 2 divisor on it if and only if X has a k -point.*

By taking X to be the spectrum of a product of finite fields or of $\mathbb{Z}[1/n]$, we see that the existence of a lift can depend on the characteristic:

Corollary 1.2. *Let P be any finite set of prime numbers. Then there exist graphs Γ and Γ' with rank 2 divisors D on Γ and D' on Γ' with the following property: For any infinite field k , Γ and D lift over $k[[t]]$ if and only if the characteristic of k is in P , and Γ' and D' lift over $k[[t]]$ if and only if the characteristic of k is not in P .*

Moreover, the existence of a lift doesn't depend only on the characteristic, as can be seen by taking X to be the spectrum of a ring of integers in Theorem 1.1:

Corollary 1.3. *Let k' be any number field. Then there exists a graph Γ with a rank 2 divisor D such that for any field k of characteristic 0, Γ and D lift over $k[[t]]$ if and only if k contains k' .*

Improvements to the specialization inequality take into account additional information about the special fiber and its divisors. For example, weighted graphs additionally record the genus of each of the components of the special fiber, and a sharper bound can be given when these genera are greater than 0 [AC13, Thm. 4.9]. However, the construction implicit in Theorem 1.1 uses rational curves and so it is also true for graph replaced by weighted graph.

Rather than just the genus of the curve, the metrized complexes introduced in [AB] also record the curve's isomorphism type and the points of intersection. Because, again, the metrized complex doesn't improve the inequality for rational components, any lift of Γ and D from Theorem 1.1 to a metrized complex with rational components will again have rank 2. In the same paper, the authors also introduce a notion of a limit g_d^r on metrized complexes. Specialization of the complete linear series of a divisor on an algebraic curve will yield a limit g_d^r , but we also have:

Proposition 1.4. *Let Γ and D be a graph and divisor constructed as in Theorem 1.1. Then for any lift of Γ to a metrized complex with rational components, there also exists a lift of D to a limit g_d^2 .*

The method of proof of Theorem 1.1 and its generalizations is through the theory of matroids and the universality theorem of Mnev [Mne88]. Matroids are combinatorial abstractions of vector configurations in linear algebra. However, not all matroids come from vector configurations and those that do are called realizable. Theorem 1.1 follows from Lafforgue's scheme-theoretic version of Mnev universality [Laf03, Thm. 1.14], together with the following:

Theorem 1.5. *Let M be any rank 3 matroid. Then there exists a graph Γ_M and a rank 2 divisor D_M such that if k is an infinite field, then Γ_M and D_M lift over $k[[t]]$ if and only if M is realizable over k .*

Matroids have been previously used to provide obstructions to lifting in tropical geometry. For any simple matroid, Ardila and Klivans defined a corresponding tropical linear space [AK06], which is the tropicalization of an algebraic variety if and only if the matroid is realizable [KP11, Cor. 1.5]. In fact, the graph Γ_M in Theorem 1.5 coincides with the link of the fine subdivision of the tropical linear space of M .

2. MATROID DIVISORS

In this section, we construct the divisors and graphs that are used in Theorems 1.1 and 1.5. As in the latter theorem, the starting point is a rank 3

simple matroid. For a thorough explanation of matroid theory, see [Oxl92], but we include the following definition for convenience:

Definition 2.1. A rank 3 simple matroid M consists of a finite set E of elements and a collection F of subsets of E , called the flats of M , such that any pair of elements is contained in exactly one flat, and such that there are at least two flats.

Throughout this paper, we will write flat for the rank 2 flats of a rank 3 matroid. A basis of a matroid will be a triple of elements which are not all contained in a single flat.

A motivating example of a rank 3 simple matroid is the matroid of a configuration of distinct k -points in the projective plane \mathbb{P}_k^2 , not all contained in a line, where k is an arbitrary field. The elements of this matroid are the points of the configuration and the flats correspond to the lines in \mathbb{P}_k^2 which are spanned by two or more of the points. We identify the flats with the elements which are contained in them. The bases of this matroid then correspond to bases for the underlying 3-dimensional vector space. A matroid coming a point configuration is called *realizable over k* and in Section 3, we will use the fact that not all matroids are realizable.

Given a rank 3 simple matroid M with elements E and flats F , we let Γ_M be the bipartite graph with vertex set $E \amalg F$, and an edge between $e \in E$ and $f \in F$ when e is contained in f . We let D_M be the divisor on the graph Γ_M consisting of the sum of all vertices corresponding to elements of the ground set E .

Proposition 2.2. *The divisor D_M has rank 2.*

Proof. To prove the theorem, we first need to show that for any degree 2 effective divisor E , the difference $D_M - E$ is linearly equivalent to an effective divisor. We build up a “toolkit” of divisors linearly equivalent to D_M . First, for any flat f , we can reverse fire f . This moves the chip from each element contained in f to f itself. Thus, the result is an effective divisor whose multiplicity at f is the cardinality of f , which is at least 2. Our second chip-firing move is to reverse fire a vertex e as well as all flats containing e . The net effect will be no change at e but all of its neighbors will end with $\#f - 1 \geq 1$ chips. Third, we will use the second chip-firing move, after which all the flats which contain e have at least one chip, after which it is possible to reverse fire e again.

Now let E be any effective degree 2 divisor on Γ_M . Thus, E is the sum of two vertices of Γ_M . We consider the various combinations which are possible for these vertices. First, if $E = [e] + [e']$ for distinct elements e and e' , then $\Gamma_M - E$ is effective. Second, if $E = [e] + [f]$, then we have two subcases. If e is in f , then we reverse fire e and all flats containing it. If e is not in f , then we can reverse fire just f . Third, if $E = [f] + [f']$ for distinct flats f and f' , then there are again two subcases. If f and f' have no elements in common, then we can reverse fire f and f' . If f and f' have a common

element, say e , then we reverse fire e together with the flats which contain it. Fourth, if $E = 2[e]$, then we use the third chip-firing move, which will move one chip onto e for each flat containing e , of which there are at least 2. Fifth, if $E = 2[f]$, then we reverse fire f .

Finally, to show that the rank is at most 2, we give an effective degree 3 divisor E such that $D_M - E$ is not linearly equivalent to any effective divisor. For this, let e_1, e_2 , and e_3 form a basis for M and let f_{ij} be the unique flat containing e_i and e_j for $1 \leq i < j \leq 3$. We set $E = [f_{12}] + [f_{13}] + [f_{23}]$ and claim that $D_M - E$ is not linearly equivalent to any effective divisor. We reverse fire e_1 together with all flats containing it and to get the following divisor linearly equivalent to $D_M - E$:

$$(1) [e_1] + (\#f_{12} - 2)[f_{12}] + (\#f_{13} - 2)[f_{13}] - [f_{23}] + \sum_{\substack{f_k \in e_1 \\ f_k \neq f_{12}, f_{13}}} (\#f_k - 1)[f_k],$$

which is effective except at f_{23} .

We wish to show the divisor in (1) is not linearly equivalent to any effective divisor, which we will do by showing that it is f_{23} reduced using Dhar's burning algorithm [Dha90]. We claim that the burning procedure leads to $\#f_{1i} - 1$ independent "fires" arriving at $\#f_{1i}$ for $i = 2, 3$. Without loss of generality, it suffices to prove this for f_{12} . Let e be an element in $f_{12} \setminus e_1$. Then e_1, e_3 , and e form a basis in M . Thus, the unique flat f containing both e_3 and e does not contain e_1 . The path in Γ from f_{23} to e_3 to f to e to f_{12} does not encounter any chips until f_{12} and thus "burns" one of these chips. The last such path then continues on to e_1 , whose sole chip is burned. The other path, via f_{13} does similarly, but then reaches all flats containing e_1 . Therefore, the divisor is f_{23} -reduced and thus not linearly equivalent to an effective divisor. \square

Proposition 2.2 also shows that if Γ_M is made into a weighted graph by giving all vertices genus 0, then D_M has rank 2 on the weighted graph. The rank is, again, unchanged for any lifting of the weighted graph to a metrized complex. To show that D_M is also a crude limit g_d^2 , we also need to choose 3-dimensional vector spaces of rational functions on the variety attached to each vertex.

Proof of Proposition 1.4. We recall from [AB] that a lift of Γ_M to a metrized complex means associated a \mathbb{P}_k^1 for each vertex v of the graph, which we denote C_v , and a point on C_v for each edge incident to v . A lift of the divisor D_M is a choice of a point on C_e for each element e of M .

The data of a limit g_d^2 is a 3-dimensional vector space H_v of rational functions on each C_v [AB, Sec. 5], which we choose as follows. For each flat f , we arbitrarily choose two distinguished elements from it and let $p_{f,1}$ and $p_{f,2}$ be the corresponding two points on C_f . Our vector space H_f consists of the rational functions which have at worst simple poles at $p_{f,1}$ and $p_{f,2}$. For each element e , we choose an arbitrary flat containing e and let q_e be

the corresponding point on C_e . Our vector space H_e consists of the rational functions which have at worst poles at q_e and at the point of the lift of D_M .

Now to check that these vector spaces form a limit g_d^2 , we need to show that the refined rank is 2. For this, we use the same “toolkit” functions as in the proof of Proposition 2.2, but we augment them with rational functions from the prescribed vector spaces on the algebraic curves. The first item from our toolkit was reverse firing a flat f to produce at least two points on C_f . We can use rational functions with poles at $p_{f,1}$ and $p_{f,2}$ to produce any degree two effective divisor on C_f . Note that we also need to use a rational function with a pole at the point of the divisor for each C_e such that e is an element of f .

The second item we needed in our toolkit was reverse firing an element e together with all of the flats which contain it. Here, for each element e' other than e , we use the rational function with a pole at the divisor and a zero at the point corresponding to the edge to the unique flat containing both e' and e . At each flat e containing e , we can use any function with a pole at $p_{e,i}$, where $i \in \{1, 2\}$ can be chosen to not be the edge leading to e . This produces a divisor at an arbitrary point of C_e .

The third and final function we used was the previous item followed by a reverse firing of e . Here, we use the same rational functions as before, but we can choose any rational function on C_e which has poles at the point of the divisor and q_e , thus giving us two arbitrary points on C_e . Thus, rational functions can be found from the prescribed vector spaces to induce a linear equivalence between the lift of D_M and any two points on the metrized complex. \square

In the case of rank 1 divisors, lifts can be constructed using the theory of harmonic maps of metrized complexes, which gives a complete theory for tamely ramified maps to \mathbb{P}^1 [ABBR]. A sufficient condition for lifting a rank 1 divisor is for it to be the underlying graph of a metrized complex which has a finite étale harmonic morphism to a tree (see [ABBR, Sec. 2] for precise definitions). These definitions are limited to the rank 1 case, but for rank 2 divisors we can subtract points to obtain a divisor of rank at least 1. In particular, if D_M lifts, then $D_M - [e]$ will be the specialization of a rank 1 effective divisor for any element e . However, for all matroid divisors, the lifting criterion of [ABBR] is satisfied.

Proposition 2.3. *Let M be any rank 3 matroid and e any element of M . Then the divisor $D_M - [e]$ is the fiber of a finite, generically étale morphism from a modification of Γ_M to a tree.*

Proof. We construct a modification $\tilde{\Gamma}_M$ of Γ_M together with a generically étale morphism from $\tilde{\Gamma}_M$ to a tree T . The tree T will be a star tree with a central vertex w , together with an unbounded edge, denoted r_f , for each flat f which does not contain e , and a single unbounded edge r_e corresponding to e . Our modification of Γ_M consists of adding the following unbounded

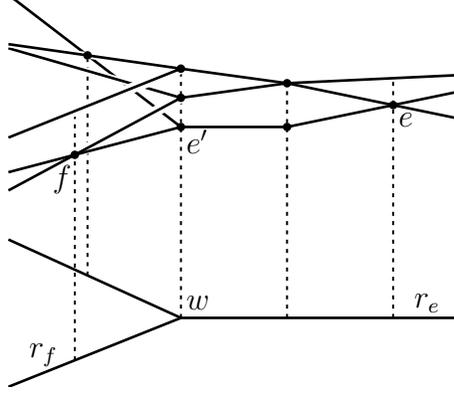


FIGURE 1. Generically étale morphism from Γ_M to a tree, such that the fiber over the central vertex of the tree is $D_M - [e]$. In this figure, f is a flat which doesn't contain e and e' is an arbitrary element of f .

edges: At e , we add one unbounded edge $s_{e,f}$ for each flat f containing e . At each element e' other than e , we add one unbounded edge $s_{e',f}$ for each flat f which contains neither e nor e' . At a flat f , we add unbounded edges $s_{f,i}$ where i ranges from 1 to $|f|$ if $e \notin f$ and from 1 to $|f| - 2$ if $e \in f$.

We now construct a generically étale map ϕ from $\tilde{\Gamma}_M$ to T . Each element other than e maps to the central vertex w of T and thus the fiber of w will be $D_M - [e]$, as desired. Each flat f not containing e maps to a point one unit of distance along the corresponding ray r_f of T . Then the rays $s_{e',f}$ and $s_{f,i}$ also map to the ray r_f , starting at w and $\phi(w)$ respectively.

We map the vertex e to its unbounded ray r_e , at a distance of 2 from w , which leaves all of the flats containing e along the same ray at a distance of 1. The rays $s_{e,f}$ and $s_{f,i}$, for flats f containing e also map to r_e , starting distances of 2 and 1 from w respectively.

To check that ϕ is harmonic, we need to verify that locally, around each vertex v of $\tilde{\Gamma}_M$, the same number of edges map to each of the edges incident to $\phi(v)$ [BN09, Sec. 2]. We do this for the case when v is a flat f and the case of elements is similar. If f does not contain e , then there are $|f|$ rays mapping to the unbounded side of r_f and the same number of edges mapping to the bounded side, connecting f to the elements it contains. If f does contain e , then there are $|f| - 2$ rays mapping to the unbounded side of r_e together with the edge connecting f to e . On the bounded side of r_e , there are also $|f| - 1$ edges, connecting f to the elements $f \setminus \{e\}$.

Recall that ϕ is generically étale if the ramification divisor R is supported on the infinite vertices of $\tilde{\Gamma}_M$ [ABBR, Def. 2.17]. In our case, at a flat f or the element e , the degrees at the tangent directions of $\phi(f)$ or $\phi(e)$ are (d, d) , where d is either $|f|$, $|f| - 1$, or the number of flats containing e . At an

element $e' \neq e$, the degrees at the tangent directions of $\phi(e')$ are $(1, 1, \dots, 1)$. In either case, (2.14.1) from [ABBR] evaluates to 0.

Finally, we can lift ϕ to a harmonic morphism of totally degenerate metrized complexes by attaching to a vertex v of $\tilde{\Gamma}_M$ either the map $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ defined by $z \mapsto z^d$, with d as above, or the identity map $\mathbb{P}^1 \rightarrow \mathbb{P}^1$ for elements $e' \neq e$. \square

3. LIFTING MATROID DIVISORS

In this section, we characterize the existence of lifts of matroid divisors. Our characterization is in terms of realizability of matroids. Recall from Section 2 that a matroid is realizable over k if there exists a configuration of points in \mathbb{P}_k^2 giving that matroid. For our lifting condition, we will need a notion which is slightly weaker when the field k is not algebraically closed.

Definition 3.1. Let k be a field. We say that a matroid M has a *Galois-invariant realization over an extension of k* if there exists a finite scheme in \mathbb{P}_k^2 which becomes a union of distinct points over \bar{k} , and these points realize M .

Equivalently, a Galois-invariant realization is a realization over a finite Galois extension k' of k such that the Galois group $\text{Gal}(k'/k)$ permutes the points of the realization. Thus, the distinction between a realization and a Galois-invariant realization is only relevant for matroids which have non-trivial symmetries. Moreover, any matroid can be extended to one with no symmetries, without affecting realizability over infinite fields, by Lemma 3.10.

Example 3.2. Let M be the matroid determined by all 21 points of $\mathbb{P}_{\mathbb{F}_4}^2$. Then M is not realizable over $\mathbb{P}_{\mathbb{F}_2}^2$ because it contains more than 7 elements, and there are only 7 points in $\mathbb{P}_{\mathbb{F}_2}^2$. However, M is clearly realizable over \mathbb{F}_4 and the Galois group $\text{Gal}(\mathbb{F}_4/\mathbb{F}_2) \cong \mathbb{Z}/2$ acts on these points by swapping pairs. Thus, M has a Galois-invariant realization over an extension of \mathbb{F}_2 .

Example 3.3. Let M be the Hesse matroid of 9 elements and 12 flats. Then M is not realizable over \mathbb{R} by the Sylvester-Gallai theorem. However, the flex points of any elliptic curve are a realization of M over the complex numbers. If the elliptic curve is defined over \mathbb{R} , then the set of all flexes points is also defined over \mathbb{R} , so M has a Galois-invariant realization over an extension of \mathbb{R} .

Theorem 3.4. Let Γ_M and D_M be the graph and divisor obtained from a rank 3 simple matroid M as in Section 2. Also, let R be any discrete valuation ring with residue field k . If D_M lifts over a discrete valuation ring R , then the matroid M has a Galois-invariant realization over an extension of k .

Proof. Let \mathfrak{X} be the semistable family over R and D a rank 2 divisor on the general fiber of \mathfrak{X} specializing to D_M . First, we make the simplifying assumption that the components of the special fiber are geometrically

irreducible. Let \overline{D} denote the closure in \mathfrak{X} of the divisor on the general fiber. By assumption, $H^0(\mathfrak{X}, \mathcal{O}(\overline{D}))$ is isomorphic to the free R -module R^3 . By restricting a basis of these sections to the special fiber \mathfrak{X}_0 , we have a rank 2 linear series on the reducible curve \mathfrak{X}_0 . Since \overline{D} doesn't intersect the components of \mathfrak{X}_0 corresponding to the flats of M and Γ_M is bipartite, \overline{D} doesn't intersect any of the nodes of \mathfrak{X}_0 . Thus, the base locus of our linear series consists of a finite number of smooth points of \mathfrak{X}_0 . Since the base locus consists of smooth points, we can subtract the base points to get a regular, non-degenerate morphism $\phi: \mathfrak{X}_0 \rightarrow \mathbb{P}_k^2$.

By the assumption that D specializes to D_M , we have an upper bound on the degree of ϕ restricted to each component of \mathfrak{X}_0 . For a flat f of M , the corresponding component C_f has degree 0 under ϕ , so $\phi(C_f)$ consists of a single point. For an element e , the corresponding component C_e has either degree 1 or degree 0 depending on whether or not the intersection of \overline{D} with C_e is contained in the base locus. If it is in the base locus, then C_e again maps to a point, and if not, C_e must map isomorphically to a line in \mathbb{P}_k^2 . Thus, the image $\phi(\mathfrak{X}_0)$ is a union of lines in \mathbb{P}_k^2 , which we will show to be a realization of the matroid M . Let f be a flat of M . Since the component of \mathfrak{X}_0 corresponding to f maps to a point, the images of the components corresponding to the elements in f all have a common point of intersection.

Now let e_1 be an element of M and suppose that the component C_{e_1} maps to a point $\phi(C_{e_1})$. Since every other element e' is in a flat with e_1 , that means that $\phi(C_{e'})$, the image of the corresponding component must contain the point $\phi(C_{e_1})$. Since ϕ is non-degenerate, there must be at least one component C_{e_2} which maps to a line. Let e_3 be an element of M which completes $\{e_1, e_2\}$ to a basis. Thus, the flat containing e_2 and e_1 is distinct from the flat containing e_2 and e_3 . Since ϕ maps C_{e_2} isomorphically onto its image, this means that $\phi(C_{e_3})$ must meet $\phi(C_{e_2})$ at a point distinct from the point $\phi(C_{e_1})$. Thus, $\phi(C_{e_3})$ must equal $\phi(C_{e_2})$. Any other element e'' in M forms a basis with e_1 and either e_2 or e_3 (or both). In either case, the same argument again shows that $C_{e''}$ must map to the same line as C_{e_2} and C_{e_3} . Again, since ϕ is non-degenerate, this is impossible. Thus, we conclude that ϕ maps each component C_e corresponding to an element e isomorphically onto a line in \mathbb{P}_k^2 . We've already shown that for any set of elements in a flat, the corresponding lines intersect at the same point. Moreover, because each component C_e maps isomorphically onto its image, distinct flats must correspond to distinct points in \mathbb{P}_k^2 . Thus, $\phi(\mathfrak{X}_0)$ is a realization of the matroid M .

If the components of the special fiber are not geometrically irreducible, then we can find a finite étale extension R' of R over which they are. In our construction of a realization of M over the residue field of R' , we can assume that we've chosen a basis of $H^0(\mathfrak{X} \times_R R', \mathcal{O})$ that's defined over R . Then, the matroid realization will be the base extension of a map of k -schemes $\mathfrak{X}_0 \rightarrow \mathbb{P}_k^2$. We let k' be the Galois closure of the residue field of R' . Then $\text{Gal}(k'/k)$ acts on the realization of M over k' , but the total collection of hyperplanes is

defined over k , and thus invariant. Thus, M has a Galois-invariant realization over an extension of k as desired. \square

For the converse of Theorem 3.4, we need to consider realizations of matroids over discrete valuations ring R . We say that a rank 3 matroid M is *realizable over R* if there exist R -points in \mathbb{P}^2 whose images in both the residue field and the fraction field realize M . For example, if R contains a field over which M is realizable, then M is realizable over R . We say that M has a *Galois-invariant realization over an extension of R* if there exists a finite, flat scheme in \mathbb{P}^2 whose special and general fiber are Galois-invariant realizations of M over extensions of the residue field and fraction field of R , respectively.

For a rank 3 matroid, a *complete flag* consists of the pair of an element e and a flat f such that e is contained in f .

Theorem 3.5. *Let R be the discrete valuation ring with residue field k . Let M be a simple rank 3 matroid with a Galois-invariant realization over an extension of R . Assume that $|k| > m - 2n + 1$, where n is the number of elements of M and m is the number of complete flags. Then there exists a regular semistable model \mathfrak{X} over R whose dual graph is Γ_M and a rank 2 divisor D on the general fiber of \mathfrak{X} such that D specializes to D_M .*

Note that Theorem 3.5 puts no assumptions on the DVR other than the cardinality of the residue field. In contrast, a semistable model is only known to exist for an arbitrary graph when the valuation ring is complete.

We construct the semistable family in Theorem 3.5 using a blow-up of projective space. We begin with a computation of the Euler for this blow-up.

Lemma 3.6. *Let S be the blow-up of \mathbb{P}_K^2 at the points of intersection of an arrangement of n lines. If A is the union of the strict transforms of the lines and the exceptional divisors, then the dimension of $H^0(S, \mathcal{O}(A))$ is at least $2n + 1$.*

Proof. We first use Riemann-Roch to compute that $\chi(\mathcal{O}(A))$ is $2n + 1$. Let m be the number of complete flags of M . We let H denote the pullback of the class of a line on \mathbb{P}^2 and C_f to denote the exceptional lines. Then, we have the following linear equivalences

$$A \sim nH - \sum_f (\#f - 1)C_f$$

$$K_S \sim -3H + \sum_f C_f$$

Now, Riemann-Roch for surfaces tells us that

$$\begin{aligned} \chi(\mathcal{O}(A)) &= \frac{A^2 - A \cdot K_S}{2} + 1 = \frac{n^2 - \sum_f (\#f - 1)^2 + 3n - \sum_f (\#f - 1)}{2} + 1 \\ (2) \qquad \qquad &= \frac{n^2 + 3n - \sum_f (\#f)^2 + \sum_f (\#f)}{2} + 1. \end{aligned}$$

We can think of the summation $\sum_f (\#f)^2$ as an enumeration of all triples of a flat and two elements of the flat. Since two distinct elements uniquely determine a flat, we have the identity that $\sum_f (\#f)^2 = n(n-1) + m$. Moreover, $\sum_f (\#f) = m$, so (2) simplifies to $\chi(\mathcal{O}(A)) = 2n + 1$.

It now suffices to prove that $H^2(S, \mathcal{O}(A))$ is zero, for which it will be sufficient, by Serre duality, to show that $K_S - A$ is not linearly equivalent to an effective divisor. The push-forward of $K_S - A$ to \mathbb{P}^2 is $-(n+3)H$, which is not effective, so $K_S - A$ is not linearly equivalent to any effective divisor, and so $H^2(S, \mathcal{O}(A))$ must be zero. Thus,

$$\chi(\mathcal{O}(A)) = H^0(S, \mathcal{O}(A)) - H^1(S, \mathcal{O}(A)) \leq H^0(S, \mathcal{O}(A)),$$

which together with the computation above yields the desired inequality. \square

Proof of Theorem 3.5. We first assume that M is realizable over R , and then at the end, we'll handle Galois-invariant realizations over extensions. Thus, we can fix a realization of M as a set of hyperplanes in \mathbb{P}_R^2 , and let S be the blow-up of \mathbb{P}_R^2 at all the points of intersections of the hyperplanes. We let the divisor $A \subset S$ be the sum of the strict transforms of the hyperplanes and the exceptional divisors. Note that A is a simple normal crossing divisor whose dual complex is Γ_M . As in the proof of Theorem 3.4, we denote the components of A as C_f and C_e corresponding to a flat f and an element e of M respectively.

We claim that A is a base point free divisor on S . Any two hyperplanes of the matroid configuration are linearly equivalent in \mathbb{P}_R^2 . The preimage of a linear equivalence between hyperplanes corresponding to elements e and e' is the divisor:

$$[C_{e'}] - [C_e] + \sum_{f: e' \in f, e \notin f} [C_f] - \sum_{f: e \in f, e' \notin f} [C_f].$$

Thus, we have a linear equivalence between A and a divisor which doesn't contain C_e , nor C_f for any of the flats containing e but not e' . By varying e and e' , we get linearly equivalent divisors whose common intersection is empty. Therefore, the complete linear series of A defines a regular morphism $\phi: S \rightarrow \mathbb{P}_R^N$ for some N .

We now look for a function $g \in H^0(S, \mathcal{O}(A))$ which does not vanish at the nodes of A . For each of the m nodes, the condition of containing that node amounts to one linear condition on the reduction $H^0(S, \mathcal{O}(A)) \otimes_R k$. Since A is base point free, this is a non-trivial linear condition, defining a hyperplane. Moreover, because of the degrees of the intersection of A with its components, the only functions vanishing on all of the nodes are multiples of the defining equation of A . By Lemmas 3.6 and 3.7 and our assumption on the cardinality of the residue field, we can find an element of $H^0(S, \mathcal{O}(A)) \otimes_R k$ and thus of $H^0(S, \mathcal{O}(A))$ avoiding all these hyperplanes and we will call this g . Now we set \mathfrak{X} to be the scheme defined by $h + \pi g$, where h is the defining equation of A and π is a uniformizer of R . It is clear that \mathfrak{X} is a flat family of curves over R whose special fiber is A and thus

has dual graph Γ_M . It remains to check that \mathfrak{X} is regular and for this it is sufficient to check the nodes of A_k . In the local ring of a node, h is in the square of the maximal ideal, but by construction πg is not, and thus, at this point \mathfrak{X} is regular.

Finally, we can take D to be the preimage of any hyperplane in \mathbb{P}_R^2 which misses the points of intersection. Again, by Lemma 3.7, it is sufficient that $|k| > \ell - 2$, where ℓ is the number of flats. We claim that $\ell - 2 \leq m - 2n + 1$, which we've assumed to be less $|k|$. This claim can be proved using the induction in Theorem 4.1, but it also follows from Riemann-Roch for graphs [BN07, Thm. 1.12]. Since Γ_M has genus $m - \ell - n + 1$, then the Riemann-Roch inequality tells us that

$$2 = r(D_M) \geq n - (m - \ell - n + 1) = \ell - m + 2n - 1,$$

which is equivalent to the desired inequality.

Now, we assume that M may only have a Galois-invariant realization over an extension of R . We can construct the blow-up S in the same way, since the singular locus of the hyperplane configuration is defined over R . Again, the divisor A is base point free, because we've already checked that it is base point free after passing to an extension where the hyperplanes are defined. Finally, we need to choose the function g and the hyperplane which pulls back to D by avoiding certain linear conditions defined over an extension of k . However, when restricted to k , these remain linear conditions, possibly of higher codimension, so we can again avoid them under our hypothesis on $|k|$. \square

Lemma 3.7. *Let k be the field of cardinality q and let H_1, \dots, H_m be hyperplanes in the vector space k^N . Let c be the codimension of the intersection $H_1 \cap \dots \cap H_m$. If $q > m - c + 1$, then there exists a point in k^N not contained in any hyperplane.*

Proof. We first quotient out by the intersection $H_1 \cap \dots \cap H_m$, so we're working in a vector space of dimension c and we know that no non-zero vector is contained in all hyperplanes. This means that the vectors defining the hyperplanes span the dual vector space, so we can choose a subset as a basis. Thus, we assume that the first c hyperplanes are the coordinate hyperplanes. The complement of these consists of all vectors with non-zero coordinates, of which there are $(q - 1)^c$. Each of the remaining $m - c$ hyperplanes contains at most $(q - 1)^{c-1}$ of these. Our assumption is that $q - 1 > m - c$, so there must be at least one point not contained in any of the hyperplanes. \square

Example 3.8. Let M be the Fano matroid, which whose realization in $\mathbb{P}_{\mathbb{F}_2}^2$ consists of all 7 \mathbb{F}_2 -points. Then M is realizable over a field if and only if the field has equicharacteristic 2. Thus, by Theorem 3.4, a necessary condition for Γ_M and D_M to lift over a valuation ring R is that the residue field of R has characteristic 2. On the other hand, M has 7 elements and 21 complete flags, so Theorem 3.5 says that if R has equicharacteristic 2 and the residue

field of R has more than 8 elements, then Γ_M and D_M lift over R . We do not know if there exists a lift of Γ_M and D_M over any valuation ring of mixed characteristic 2.

Example 3.9. Conversely, let M be the non-Fano matroid, which is realizable over k if and only if k has characteristic not equal to 2. Moreover, M is realizable over any valuation ring R in which 2 is invertible. Thus, if Γ_M and D_M lift over a valuation ring R , then the residue field of R must have characteristic different than 2 by Theorem 3.4. The converse is true, so long as the residue field has more than 11 elements by Theorem 3.5.

Since Theorems 3.4 and 3.5 refer to Galois-invariant realizations, we will need the following lemma, which relates these to the more standard notion of matroid realizations.

Lemma 3.10. *Let M be a matroid. Then there exists a matroid M' such that for any infinite field k , the following are equivalent:*

- (1) M has a realization over k .
- (2) M' has a realization over k .
- (3) M' has a Galois-invariant realization over an extension of k .

Proof. We use the following construction of an extension of a matroid. Suppose that M is a rank 2 matroid and f is a flat of M . Then we construct a matroid M'' which contains the elements of M , together with an additional element x . The flats of M'' are those of M , except that f is replaced by $f \cup \{x\}$, and two-element flats for x and every element not in f . By repeated applications of this construction, we can construct a matroid M' such that every flat of M has a different number of elements in M' .

Now we prove that the conditions in the lemma statement are equivalent for this choice of M' . First, assume that M has a realization over an infinite field k . We can inductively extend this to a realization of M' . At each step, when adding an element x as above, it is sufficient to place x at a point along the line corresponding to F such that it doesn't coincide with any of the other points, and it is not contained in any of the lines spanned by two points not in F . We can choose such a point for x since k is infinite. Second, if M' has a realization over k , then by definition, it has a Galois-invariant realization over an extension of k .

Finally, we suppose that M' has a Galois-invariant realization over an extension of k and we want to show that M has a realization over k . Suppose we have a realization over a Galois extension k' of k . Since all the flats from the original matroid contain different numbers of points, the Galois group does not permute the corresponding lines in the realization. Therefore, the lines and thus also the points from the original matroid M must be defined over k . Therefore, the restriction of this realization gives a realization of M over k , which completes the proof of the lemma. \square

Proof of Theorem 1.1. As in the theorem statement, let X be a scheme of finite type over \mathbb{Z} . We choose an affine open cover of X and let \tilde{X} be the

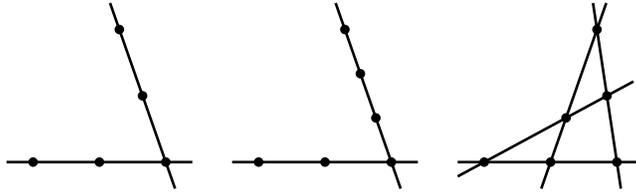


FIGURE 2. Matroids from Theorem 4.1

disjoint union of these affine schemes. By the scheme-theoretic version of Mnev's universality theorem, there is a matroid M whose realization space is isomorphic to an open subset U of $\tilde{X} \times \mathbb{A}^N$ in such a way that U maps surjectively onto X [Laf03, Thm. 1.14]. Now let M' be the matroid as in Lemma 3.10 and we claim that $\Gamma_{M'}$ and $D_{M'}$ have the desired properties for the theorem.

Let k be any infinite field, and then X clearly has a k -point if and only if \tilde{X} has a k -point. Likewise, since k is infinite, any non-empty subset of \mathbb{A}_k^N has a k -point, so U also has a k -point if and only if X has a k -point. By Lemma 3.10, these conditions are equivalent to M' having a Galois-invariant realization over an extension of k . Supposing that X has a k -point and thus M' has a realization over k , then $\Gamma_{M'}$ and $D_{M'}$ have a lifting over $k[[t]]$ by Theorem 3.5. Conversely, if $D_{M'}$ has a lifting over $k[[t]]$, then M' has a Galois-invariant realization over an extension of k by Theorem 3.4, and thus M has a realization over k by Lemma 3.10, so X has a k -point. \square

4. BRILL-NOETHER THEORY

In the theory of limit linear series, a key technique is the observation that that if the moduli space of limit linear series on the degenerate curve has the expected dimension then it lifts to a linear series [EH86, Thm. 3.4]. Here, the expected dimension of limit linear series of degree d and rank r on a curve of genus g is $\rho(g, r, d) = g - (r + 1)(g + r - d)$. It is natural to ask if a tropical analogue of this result is true: if the dimension of the moduli space of divisor classes on a tropical curve of degree d and rank at least r has (local) dimension $\rho(g, r, d)$, then does every such divisor lift?

In this section, we show that the divisors D_M on graphs G_M do not provide a negative answer to this question. We begin with the following classification:

Theorem 4.1. *Let M is a rank 3 simple matroid, with g and d equal to the genus of Γ_M and degree of D_M respectively. If $\rho(g, 2, d) \geq 0$, then M is one of the following matroids:*

- (1) *The one element extension of the uniform matroid $U_{2, n-1}$, with $\rho = n - 2$.*
- (2) *The uniform matroid $U_{4, 3}$, with $\rho = 0$.*
- (3) *The matroid formed by the vectors: $(1, 0, 0)$, $(1, 0, 1)$, $(0, 0, 1)$, $(0, 1, 1)$, $(0, 1, 0)$, with $\rho = 1$.*

- (4) *The matroid in the previous example together with $(1, 0, \lambda)$ for any element λ of the field other than 1 and 0, with $\rho = 0$.*
- (5) *The matroid consisting of the point of intersection between any pair in a collection of 4 generic lines, for which we can take the coordinates to be the vectors from (3) together with $(1, 1, 1)$, with $\rho = 0$.*

Proof. We first compute the invariants for the graph Γ_D and divisor D_M constructed in Section 2. As in that section, we let n be the number of elements of M , ℓ the number of flats and m the number of complete flags. Since Γ_D consists of m edges and $n + \ell$ vertices, it has genus $m - n - \ell + 1$. It is also immediate from its definition that D_M has degree n . Thus, the expected dimension of rank 2 divisors is

$$(3) \quad \rho = m - n - \ell + 1 - 3((m - n - \ell + 1) + 2 - n) = 5n + 2\ell - 2m - 8$$

Now, we consider what happens when we remove a single element e , where e is not contained in all bases. For every flat containing e , we decrease the number of complete flags by 1 if that flat contains at least 3 elements, and if it contains 2 elements, then we decrease the number of flags by 2 and the number of flats by 1. Thus, by (3), ρ drops by $5 - 2f$, where f is the number of flats in M which contain e . Since e must be contained in at least 2 flats, either $M \setminus e$, the matroid formed by removing e has positive ρ or e is contained in exactly 2 flats.

We first consider the latter case, in which e is contained in exactly two flats, which we assume to have cardinality $a + 1$ and $b + 1$ respectively. The integers a and b completely determine the matroid because all the other flats consist of a pair of elements, one from each of these sets. Thus, there are $ab + 2$ flats and $2ab + a + b + 2$ complete flags. By using (3), we get

$$\rho = 5(a + b + 1) + 2(ab + 2) - 2(2ab + a + b + 2) - 8 = -2ab + 3a + 3b - 3.$$

It's possible to check, that up to swapping a and b , the only non-negative values of this equation are for $a = 1$ and b arbitrary or $a = 2$ and b is 2 or 3. These correspond to cases (1), (3), and (4) respectively from the theorem statement.

Now we consider the case that e contained in more than two flats, in which case $M \setminus e$ satisfies $\rho > 0$. By induction on the number of elements, we can assume that $M \setminus e$ is on our list, in which case the possibilities with $\rho > 0$ are (1) and case (3). For the former matroid, if e is contained in a flat of $M \setminus e$, then M is a matroid of the type from the previous paragraph, with a equal to 1 or 2. On the other hand, if e contained only in 2-element flats, then e is contained in $n - 1$ flats, so

$$\rho(M) = \rho(M \setminus e) + 5 - 2(n - 1) = (n - 3) + 7 - 2n = 4 - n.$$

The only possibility is $n = 4$, for which we get (2), the uniform matroid. Finally, if $M \setminus e$ is the matroid in case (3), then the only relevant possibilities are those for which e is contained in at most 3 flats, for which the possible matroids are (4) or (5). \square

Proposition 4.2. *If R is any DVR and M is any of the matroids from Theorem 4.1, then M has a Galois-invariant realization over an extension of R .*

Proof. The matroids (2), (3) and (5) are regular matroids, i.e. realizable over \mathbb{Z} , so they are *a fortiori* realizable over any DVR. Moreover, the other matroids in case (1) and (4) are realizable over R so long as the residue field has at least $n - 2$ and 3 elements respectively. We will show that if the residue field is finite, then the one-element extension of $U_{2,n-1}$ has a Galois-invariant realization over R . The other matroid is similar.

Let M be the one-element extension of $U_{2,n-1}$ and suppose that the residue field k is finite. Then we choose a polynomial with coefficients of degree $n - 1$ in R and whose reduction to k is square-free. Adjoining the roots of this polynomial defines an unramified extension R' of R , and we write a_1, \dots, a_{n-1} for its roots in R . Then, the vectors $(1, a_1, 0), \dots, (1, a_{n-1}, 0), (0, 0, 1)$ give a Galois-invariant realization of M over R' , which is what we wanted to show. \square

REFERENCES

- [AB] Omid Amini and Matthew Baker. Linear series on metrized complexes of algebraic curves. preprint, [arXiv:1204.3508](https://arxiv.org/abs/1204.3508).
- [ABBR] Omid Amini, Matthew Baker, Erwan Brugallé, and Joseph Rabinoff. Lifting harmonic morphisms of tropical curves, metrized complexes, and Berkovich skeleta. preprint, [arXiv:1303.4812](https://arxiv.org/abs/1303.4812).
- [AC13] Omid Amini and Lucia Caporaso. Riemann-Roch theory for weighted graphs and tropical curves. *Adv. Math.*, 240:1–23, 2013.
- [AK06] Federico Ardila and Carly Klivans. The Bergman complex of a matroid and phylogenetic trees. *J. Combin. Theory Ser. B*, 96(1):38–49, 2006.
- [Bak08] Matthew Baker. Specialization of linear systems from curves to graphs. *Algebra & Number Theory*, 2(6):613–653, 2008.
- [BN07] Matt Baker and Serguei Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215:766–788, 2007.
- [BN09] Matt Baker and Serguei Norine. Harmonic morphisms and hyperelliptic graphs. *Int. Math. Res. Not.*, 15:2914–2955, 2009.
- [Dha90] Deepak Dhar. Self-organized critical state of the sandpile automaton models. *Phys. Rev. Lett.*, 64:1613–1616, 1990.
- [EH86] David Eisenbud and Joe Harris. Limit linear series: Basic theory. *Invent. math.*, 85:337–371, 1986.
- [KP11] Eric Katz and Sam Payne. Realization spaces for tropical fans. In *Combinatorial aspects of commutative algebra and algebraic geometry*, volume 6, pages 73–88. Abel Symp., Springer, 2011.
- [Laf03] Laurent Lafforgue. *Chirurgie des grassmanniennes*, volume 19 of *CRM Monograph Series*. American Mathematical Society, 2003.
- [Mne88] Nikolai E. Mnev. *The universality theorems on the classification problem of configuration varieties and convex polytopes varieties*, volume 1346 of *Lect. Notes in Math.*, pages 527–543. Springer, 1988.
- [Oxl92] J. Oxley. *Matroid Theory*. Oxford Univ. Press, 1992.

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, PO Box 208283, NEW HAVEN,
CT 06520

E-mail address: `dustin.cartwright@yale.edu`